

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

Pol Antràs Teresa C. Fort Felix Tintelnot *

May 9, 2017

Abstract

We develop a quantifiable multi-country sourcing model in which firms self-select into importing based on their productivity and country-specific variables. In contrast to canonical export models where firm profits are additively separable across destination markets, global sourcing decisions naturally interact through the firm's cost function. We show that, under an empirically relevant condition, selection into importing exhibits complementarities across source markets. We exploit these complementarities to solve the firm's problem and estimate the model. Comparing counterfactual predictions to reduced-form evidence highlights the importance of interdependencies in firms' sourcing decisions across markets, which generate heterogeneous domestic sourcing responses to trade shocks.

During the last three decades, the world has become increasingly globalized. Dramatic advances in communication, information, and transportation technologies have revolutionized

*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. Pol Antràs: Harvard University, 1805 Cambridge Street, Littauer 207, Cambridge, MA 02138, and NBER (pantras@fas.harvard.edu). Teresa C. Fort: Tuck School of Business at Dartmouth, 100 Tuck Hall, Hanover, NH 03755, and NBER (teresa.fort@dartmouth.edu). Felix Tintelnot: University of Chicago, 1126 East 59th Street, Chicago, IL 60637, and NBER (tintelnot@uchicago.edu). We are grateful to Treb Allen, Isaiah Andrews, Andy Bernard, Emily Blanchard, 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 two anonymous referees for their constructive comments, and to Xiang Ding, BooKang Seol and Linh Vu for excellent research assistance. We have also benefited from very useful feedback from seminar audiences at Aarhus, AEA Meetings in Boston, Barcelona GSE, Bank of Spain, Boston College, Boston University, Brown, Cambridge University, Chicago Booth, Dartmouth, ECARES, the Econometric Society Meeting in Minneapolis, ERWIT in Oslo, Harvard, IMF, John Hopkins SAIS, LSE, Michigan, MIT, UQ à Montreal, National Bank of Belgium, NBER Summer Institute, Northwestern, Princeton, Sciences Po in Paris, SED in Toronto, Stanford, Syracuse, Tsinghua, UBC, UC Berkeley, UC Davis, UC San Diego, Urbana-Champaign, Virginia, and Yale. We thank Jim Davis at the Boston RDC for support with the disclosure process. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

how and where firms produce their goods. Intermediate inputs account for approximately two thirds of international trade (Johnson and Noguera, 2016), and vertical specialization across countries is an important and growing feature of the world economy (Hummels et al., 2001; Hanson et al., 2005). As global value chains rise in importance, a firm’s production is more likely than ever to span multiple countries. There is also mounting evidence that firm-level decisions play a critical role in explaining trade patterns (Bernard et al., 2009), and that they have important ramifications for aggregate productivity, employment, and welfare (Goldberg et al., 2010; Hummels et al., 2014).

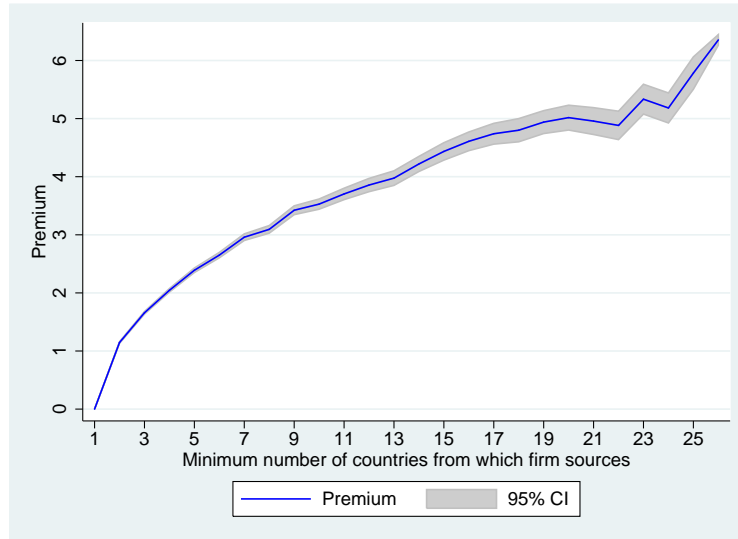
Despite the growing importance of global production sharing, the canonical model of firm-level trade decisions (cf. Melitz, 2003) focuses on exporting rather than importing. Since every international trade transaction involves an exporter and an importer, a natural question is: can one use the structure of the well-known exporting framework to analyze firms’ import decisions? Existing export models cannot be applied directly to analyze foreign sourcing for a simple – yet powerful – reason. While the canonical export model ensures that a firm’s decision to enter each market can be analyzed separately by assuming constant marginal costs, a firm chooses to import precisely because it seeks to lower its marginal costs. In a world in which firm heterogeneity interacts with fixed sourcing costs, the firm’s decision to import from one market will also affect whether it is optimal to import from another market. Foreign sourcing decisions are therefore interdependent across markets, making a model about importing much more complicated to solve theoretically and to estimate empirically.

In this paper, we develop a new framework to analyze firm-level sourcing decisions in a multi-country world. An important focus of the model is on firms’ extensive margin decisions about which products to offshore and the countries from which to purchase them. Bernard et al. (2009) find that these margins account for about 65 percent of the cross-country variation in U.S. imports, and 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 In Figure 1, we extend this evidence to show not only that importers are larger than non-importers, but also that their relative size advantage is increasing in the number of countries from which they source. The figure indicates that firms that 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 importer size advantages are suggestive of sizable country-level fixed costs of sourcing, which limit the ability of small firms to select into importing from a large number of countries.²

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, section C.2).

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

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



Not only do country-level fixed costs of importing appear to be empirically relevant, but they also seem to be heterogeneous across countries. To illustrate this variation, Table 1 shows the number of U.S. firms that import from a country versus total sourcing from that country. The table lists the top ten source 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 show that Canada ranks number one based on the number of U.S. importers and total import value. For most other countries, however, country rank based on the number of importers does not equal the rank based on import values. China is number two for firms but only number three for value; and Mexico, the number two country in terms of value, ranks eighth in terms of number of importers.

The considerable divergence between the intensive and extensive margins presented in Table 1 suggests that countries differ not only in terms of their potential as a marginal cost-reducing source of inputs, but also in terms of the fixed costs firms must incur to import from them. In section 1, we develop a quantifiable multi-country sourcing model that allows for this possibility. Heterogeneous firms self-select into importing based on their productivity and country-specific characteristics (wages, trade costs, and technology). The model delivers a simple closed-form solution for firm profits, in which marginal costs are decreasing in a firm’s *sourcing capability*, which is itself a function of the set of countries from which a firm imports, as well as those countries’ characteristics. Firms can, in principle, buy intermediate inputs from any country

least two countries, etc. The horizontal axis denotes the number of countries from which a firm sources, with 1 corresponding to firms that use only domestic inputs. These 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 section C.3 of 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	37,810	59	145,740	16
China	2	3	21,460	33	121,990	13
Germany	3	5	13,050	20	62,930	7
United Kingdom	4	6	11,450	18	30,750	3
Taiwan	5	11	10,490	16	16,630	2
Italy	6	13	8,500	13	13,230	1
Japan	7	4	7,980	12	112,250	12
Mexico	8	2	7,810	12	125,980	14
France	9	9	6,110	9	22,980	3
Korea, South	10	10	5,610	9	20,390	2

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

in the world, but acquiring the ability to import from a country entails a market-specific fixed cost. As a result, relatively unproductive firms may opt out of importing from high fixed cost countries, even if they are particularly attractive sources of inputs.

In this environment, the optimality of importing from one country generically depends on the other countries from which a firm sources its inputs. This stands in sharp contrast to standard export models, where the assumption of constant marginal costs ensures that the decision to sell in one market is independent of export decisions in other markets. This constant marginal cost assumption is clearly not tenable for sourcing decisions, since the firm chooses to import precisely in order to lower its marginal costs. The resulting interdependence in a firm’s extensive margin import decisions complicates the firm’s problem considerably, as it now involves a combinatorial problem with 2^J possible choices, where J denotes the number of possible source countries.

Despite these complications, we provide the first characterization of the firm’s extensive margin sourcing decisions. First, we show that source countries can be complements or substitutes, depending only on a parametric restriction that relates the elasticity of demand faced by the final-good producer to the dispersion of input productivities across locations. When demand is inelastic or input efficiency differences are small, the addition of a country to a firm’s global sourcing strategy reduces the marginal gain from adding other locations. In such a “substitutes case,” the firm’s optimal choice of countries to include in its sourcing strategy is extremely hard to characterize, both analytically as well as quantitatively. High productivity

firms may opt into countries with high fixed costs but with the potential for high marginal cost savings, thus rendering further marginal cost reductions less beneficial. Although low productivity firms would also like to source from these locations, the high fixed costs may preclude them from doing so. In this scenario, high productivity firms will always source from (weakly) better countries, but they may source from fewer countries than low productivity firms. Global sourcing therefore magnifies any pre-existing differences in underlying firm productivity and increases the skewness in the size distribution of firms, but does not necessarily lead to the hierarchical entry predictions that are well-known for exporting.

Conversely, selection into importing features complementarity across markets when demand is relatively elastic (so profits are particularly responsive to variable cost reductions) and input efficiency levels are relatively heterogeneous across markets (so that the reduction in expected costs from adding an extra country in the set of active locations is relatively high). This case is 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, in such a case, the sourcing strategies of firms follow a strict hierarchical structure in which the number of countries in a firm’s sourcing strategy increases (weakly) with the firm’s core productivity level.³ Crucially, despite this hierarchical sourcing structure, the model also generates the type of discrepancies between the intensive and extensive margins of importing evident in Table 1. We explicitly show that this is not possible in a model that abstracts from country-specific fixed costs.

Our quantitative analysis enables separate identification of the *sourcing potential* of a country – a function of technology, trade costs, and wages capturing the potential of a country as source of marginal cost savings – and the fixed cost of sourcing from that country. We use 2007 Census data on U.S. manufacturers’ mark-ups and import shares to recover the sourcing potential of 66 foreign countries, as well as the average elasticity of demand and dispersion of input productivities faced by U.S. firms. Consistent with the pattern documented in Figure 1, we find robust evidence suggesting that the extensive margin sourcing decisions of U.S. firms are complements. This finding paves the way for an additional methodological contribution of our paper –namely, to solve the firm’s problem and estimate the model structurally. To do so, we apply an iterative algorithm developed by Jia (2008), which exploits the complementarities in the ‘entry’ decisions of firms, and uses lattice theory to reduce the dimensionality of the firm’s optimal sourcing strategy problem. We can therefore estimate the fixed costs of sourcing, which range from a median of 10,000 to 56,000 USD, are around 13 percent lower for countries with a common language, and increase in distance with an elasticity of 0.19. In

³The 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.

line with the premise that countries differ along two dimensions, the relative rankings of these fixed costs are also quite different from the rankings of countries' potential to reduce marginal costs.⁴

The structural estimation of the model is informative not only because it shows the importance of marginal cost savings versus fixed cost heterogeneity across countries, but also because it allows for counterfactual exercises.⁵ We exploit this capability by studying the implications of an increase in China's sourcing potential calibrated to match the observed growth in the share of U.S. firms importing from China between 1997 and 2007. These years are governed both by data availability, and the fact that they span China's accession to the World Trade Organization (WTO). Consistent with other quantitative models of trade, the China shock increases the competitive environment by decreasing the equilibrium industry-level U.S. price index and driving some U.S. final good producers out of the market. Although the net result of these forces is a marked decrease in domestic sourcing (and U.S. employment) in that sector, the net decline masks significant heterogeneity in how the shock affects the sourcing decisions of firms at different points in the size distribution. More specifically, the shock induces a range of U.S. firms to select into sourcing from China, and on average, these firms 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 that does not arise in the absence of fixed costs of offshoring or whenever entry decisions are independent across markets.

To assess the empirical relevance of these channels, we compare the model's counterfactual predictions to the observed changes in U.S. manufacturers' sourcing from the U.S. and third markets between 1997 and 2007. We first show that the same qualitative patterns predicted by the model are evident in the raw data. Firms that begin importing from China over this period grow their domestic and third market sourcing the most, continuing importers have smaller but still positive sourcing changes, and firms that never import from China shrink their domestic sourcing and increase third-market sourcing by a substantially smaller amount than both new and continuing China importers. To ensure that the patterns observed in the raw data are not driven solely by firm-specific demand or productivity shocks, we construct a plausibly exogenous firm-level shock to Chinese sourcing potential in the spirit of Autor et al. (2013) and Hummels et al. (2014). The results show that exogenous increases in firm-level imports from China do not decrease domestic and third market sourcing—as might be expected in a world with no interdependencies in sourcing decisions—but instead are associated with

⁴Building on Jia (2008), in a recent interesting paper Arkolakis and Eckert (2017) develop an algorithm that can solve large combinatorial discrete choice problems in the case in which the firm's discrete choices are substitutes.

⁵This is in contrast to moment inequality methods, which were first adopted in an international trade context by Morales et al. (2014).

increased firm-level sourcing from other markets. We thus provide both structural and reduced-form evidence of the empirical relevance of interdependencies in firms’ extensive margin import decisions.

Our paper contributes to three distinct literatures. First, we add to a large body of theoretical work on foreign sourcing. We follow existing theory that adapts the Melitz (2003) model to characterize heterogeneous firms’ foreign sourcing decisions (e.g., Antràs and Helpman, 2004, 2008), but our framework also shares features with a parallel literature that uses the Eaton and Kortum (2002) model to study offshoring (Rodríguez-Clare, 2010; Garetto, 2013). More specifically, we build on the approach in Tintelnot (2017) of embedding the Eaton and Kortum (2002) stochastic representation of technology into the problem of a firm, though in our context firms choose optimal sourcing rather than final-good production locations. This approach allows us to move beyond the two-country frameworks that have pervaded the literature and develop a tractable multi-country model. A key theoretical insight from our model is that a positive shock to sourcing from one location could lead a firm either to decrease its sourcing from other locations as it substitutes away from them, or instead to grow sufficiently so that it increases its net sourcing from other locations. This prediction is reminiscent of Grossman and Rossi-Hansberg (2008), who show that an offshoring industry may expand domestic employment if a “productivity effect” dominates a “substitution effect.” An important difference is that in our framework, these effects take place within a firm rather than an industry.

Our paper also relates to an extensive empirical literature on offshoring. A number of papers provide reduced form evidence on the determinants of offshoring (Fort, 2017), as well as its impact on firm performance and aggregate productivity (Amiti and Konings, 2007; Goldberg et al., 2010; De Loecker et al., 2016). A related set of papers uses a more structural approach to quantify the effect of importing on firm productivity and prices (Halpern et al., 2015; Gopinath and Neiman, 2014; Blaum et al., 2017). The first part of our estimation provides a similar quantification, implying that a firm sourcing from all foreign countries faces nine percent lower variable costs and achieves 32 percent higher sales than when sourcing exclusively from domestic suppliers.⁶ The most important distinction between those papers and ours is that we provide evidence not just on the intensive margin implications of importing, but also on the firm’s extensive margin sourcing decisions in a multi-country setting with heterogeneous fixed costs across countries. While Blaum et al. (2013) discuss the existence of interdependencies across sourcing decisions in a model with an arbitrary number of countries and inputs, ours is the first paper to characterize the extensive margin of importing in this setting and to solve

⁶Quantitatively, this is lower than the findings of Halpern et al. (2015) 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 productivity of a Hungarian firm by 22 percent. Blaum et al. (2017) obtain even larger cost reduction estimates for some French firms, perhaps due to an alternative interpretation of idiosyncratic differences in sourcing shares (e.g., as measurement error in our context and as structural error in their paper).

the firm’s problem quantitatively.

Finally, we contribute to a growing body of work that analyzes interdependencies in firm-level decisions. Yeaple (2003) and Grossman et al. (2006) first described the inherent difficulties in solving for the extensive margin of imports in a multi-country model with multiple intermediate inputs and heterogeneous fixed costs of sourcing. Those authors obtained partial characterizations of the problem in models with at most three countries and two inputs. We provide the first characterization of the firm’s extensive margin sourcing decision in this setting with multiple inputs and countries, and show how these decisions can be aggregated to describe trade flows across countries. These trade flow equations collapse to the well-known Eaton and Kortum (2002) gravity equation whenever fixed costs are zero (so that there is universal importing), or to the Chaney (2008) gravity equation in the knife-edge case that shuts down interdependencies across markets. In a general setting, our model delivers an extended gravity equation reminiscent of Morales et al. (2014), who estimate a model with interdependencies in firms’ export decisions. That paper uses moment inequalities to partially identify the cost parameters and does not conduct any counterfactual analysis. Tintelnot (2017) solves the optimal plant location problem of multinational firms in a general equilibrium model, however, in a setting with much fewer countries. We overcome the challenges in prior work by combining the theoretical insights on complementarity with Jia’s (2008) algorithm for solving Walmart’s and Kmart’s decisions about whether and where to open new retail establishments. Our paper is the first to adopt this algorithm in an international setting or in a setting with more than two firms.

The rest of the paper is structured as follows. We present the assumptions of our model in section 1 and solve for the equilibrium in section 2. In section 3, we introduce the data and provide descriptive evidence supporting the assumptions underlying our theoretical framework. We estimate the model structurally in section 4, and in section 5, we perform our counterfactual analysis and compare the predictions of the model to reduced-form evidence. Section 6 concludes.

1 Theoretical Framework

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

1.1 Preferences and Endowments

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

$$U_{Mi} = \left(\int_{\omega \in \Omega_i} q_i(\omega)^{(\sigma-1)/\sigma} d\omega \right)^{\sigma/(\sigma-1)}, \quad \sigma > 1, \quad (1)$$

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

$$q_i(\omega) = E_i P_i^{\sigma-1} p_i(\omega)^{-\sigma}, \quad (2)$$

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

$$B_i = \frac{1}{\sigma} \left(\frac{\sigma}{\sigma-1} \right)^{1-\sigma} E_i P_i^{\sigma-1}. \quad (3)$$

There is a unique factor of production, labor, which commands a wage w_i in country i . When we close the model in general equilibrium, we later introduce a freely tradable, 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.

1.2 Technology and Market Structure

There exists a measure N_i of final-good producers in each country $i \in J$, 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. 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. Following Melitz (2003), we assume that firms only learn their productivity φ after incurring an entry cost equal to f_{ei} units of labor in country i . This core productivity is drawn from a country-specific distribution $g_i(\varphi)$, with support in $[\underline{\varphi}_i, \infty)$, and with an associated contin-

uous cumulative distribution $G_i(\varphi)$. For simplicity, we assume that final-good varieties are prohibitively costly to trade across borders.

Intermediates can instead be traded internationally, 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 firm-specific 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 . All intermediates are produced with labor under constant-returns-to-scale technologies. We denote by $a_j(v, \varphi)$ the unit labor requirement associated with the production of firm φ 's intermediate $v \in [0, 1]$ in country $j \in J$.

Although intermediates are produced worldwide, a final-good producer based in country i only acquires the capability to offshore in j after incurring a fixed cost equal to f_{ij} units of labor in country i . 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 *global sourcing strategy* of that firm.

Intermediates are produced by a competitive fringe of suppliers who sell their products at marginal cost.⁷ Shipping intermediates from country j to country i entails iceberg trade costs τ_{ij} . As a result, the cost at which firms from i can procure input v from country j is given by $\tau_{ij} a_j(v, \varphi) w_j$, and the price that firm φ based in country i pays for input v can be denoted by

$$z_i(v, \varphi; \mathcal{J}_i(\varphi)) = \min_{j \in \mathcal{J}_i(\varphi)} \{ \tau_{ij} a_j(v, \varphi) w_j \}. \quad (4)$$

We can then express the marginal cost for firm φ based in country i of producing a unit of a final-good variety as

$$c_i(\varphi) = \frac{1}{\varphi} \left(\int_0^1 z_i(v, \varphi; \mathcal{J}_i(\varphi))^{1-\rho} dv \right)^{1/(1-\rho)}. \quad (5)$$

Building on Eaton and Kortum (2002), we treat the (infinite-dimensional) vectors of firm-specific intermediate input efficiencies $1/a_j(v, \varphi)$ as the realization of an extreme value distribution. More specifically, suppliers in j draw the value of $1/a_j(v, \varphi)$ from the Fréchet distribution

$$\Pr(a_j(v, \varphi) \geq a) = e^{-T_j a^\theta}, \quad \text{with } T_j > 0. \quad (6)$$

These draws are assumed to be independent across locations and inputs. 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 across countries.

⁷Implicitly, we assume that contracts between final-good producers and suppliers are perfectly enforceable, so that the firm-specificity of inputs is irrelevant for the prices at which inputs are transacted.

1.3 Discussion of Assumptions

This completes the description of the key assumptions of the model. A number of dimensions of our setup are worth discussing. First, although we have assumed that inputs are firm-specific, our model is in fact isomorphic to one in which the unit measure of inputs, as well as their associated unit labor requirements $a_j(v, \varphi)$, are identical for all firms and denoted by $a_j(v)$. We emphasize the firm-specificity of inputs to justify why intermediaries (e.g., wholesalers) would not trivially eliminate the need for all firms to incur fixed costs of foreign sourcing. Second, to highlight the importance of importing, we have assumed that final-good varieties cannot be traded across borders. In the Online Appendix (section B.3), we relax this assumption and study the joint determination of the extensive margins of both exports and imports, an approach that has been further pursued by Bernard et al. (2016). Third, our model assumes that all final-good producers combine a measure one of inputs in production. As we demonstrate in the Online Appendix (section B.3), it is simple to generalize our framework to the case in which final-good producers also hire local labor to assemble the bundle of inputs, and in which firms optimally choose the complexity of production, as captured by the measure of inputs used in production (see Acemoglu et al., 2007). The qualitative results of these extensions are analogous to those of our benchmark model, but incorporating these features would significantly complicate the structural estimation. Fourth, tractability concerns also dictate our assumption that wages are pinned down in a non-manufacturing sector, as we discuss at greater length in section 5. Finally, we have introduced an asymmetric market structure in the final- and intermediate-input sectors because this allows our model to nest two key workhorse trade models developed in recent years. It would be feasible to turn the intermediate-input sector into a monopolistically competitive sector with a fixed mass of firms, and the relevant expressions would all be very similar.

2 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 flows across countries.

2.1 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 (5), it is clear that after learning the vector of unit labor requirements for each country $j \in \mathcal{J}_i(\varphi)$, the firm will choose the location of production for each input v that solves $\min_{j \in \mathcal{J}_i(\varphi)} \{\tau_{ij} a_j(v, \varphi) w_j\}$. Using the properties of the Fréchet distribution in (6), 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

$$\chi_{ij}(\varphi) = \frac{T_j (\tau_{ij} w_j)^{-\theta}}{\Theta_i(\varphi)} \quad \text{if } j \in \mathcal{J}_i(\varphi) \quad (7)$$

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

$$\Theta_i(\varphi) \equiv \sum_{k \in \mathcal{J}_i(\varphi)} T_k (\tau_{ik} w_k)^{-\theta}. \quad (8)$$

The term $\Theta_i(\varphi)$ summarizes the *sourcing capability* of firm φ from i . Note that, in equation (7), each country j 's market share in the firm's purchases of intermediates corresponds to this country's contribution to its 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 trade costs when selling to country i will 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 .⁸

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

$$c_i(\varphi) = \frac{1}{\varphi} (\gamma \Theta_i(\varphi))^{-1/\theta}, \quad (9)$$

where $\gamma = [\Gamma(\frac{\theta+1-\rho}{\theta})]^{1/(1-\rho)}$ and Γ is the gamma function.⁹ Note that in light of equation (8), the addition of a new location to the set $\mathcal{J}_i(\varphi)$ increases the sourcing capability of the

⁸It 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 (5). 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.

⁹These derivations are analogous to those performed by Eaton and Kortum (2002) to solve for the aggregate price index in their model. 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.

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 $\mathcal{J}_i(\varphi)$.

Using the demand equation (2) and the derived marginal cost function in (9), we can express the firm's profits conditional on a sourcing strategy $\mathcal{J}_i(\varphi)$ as

$$\pi_i(\varphi) = \varphi^{\sigma-1} (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} B_i - w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij}, \quad (10)$$

where B_i is given in (3). As is clear from equation (10), 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 $w_i f_{ij}$.

It is worth highlighting the connection between our modeling of the gains from importing intermediate inputs and the Armington-style approach that is standard in the literature on importing.¹⁰ More specifically, suppose that all suppliers in a given country produce the same intermediate input using local labor under a constant-returns-to-scale technology featuring a unit labor requirement equal to $(\gamma T_j)^{-1/\theta}$ in each country $j \in J$. Assume, in addition, that inputs are differentiated by country of origin with an elasticity of substitution across inputs from any two locations equal to $1 + \theta$. Finally, as in our framework, assume that in order to import country j 's unique input, final-good producers in i need to incur a fixed costs equal to $w_i f_{ij}$ and iceberg trade costs τ_{ij} . Under these assumptions, it is then straightforward to verify that the resulting firm profits will be identical to those in equation (10) above.

This isomorphism between our model and the love-for-variety approach carries three significant implications. First, it should be clear that the interdependencies in the firm's extensive margin sourcing decisions are also a feature of the Armington-style models that have pervaded the literature on importing. Second, it follows that the results below on the optimal determination of the sourcing strategy $\mathcal{J}_i(\varphi)$, as well as the techniques we develop in section 4 to structurally estimate the model, are also applicable in these types of models. Third, it implies that our model provides an intuitive microfoundation for why being able to import from (several) foreign countries is productivity-enhancing, without resorting to the elusive notion of input differentiation by country of origin. With this in mind, we next turn to an analysis of the optimal sourcing strategy of firms.

¹⁰See, among others, Halpern et al. (2015), Goldberg et al. (2010), and Gopinath and Neiman (2014).

2.2 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 (10). We can alternatively express this problem as

$$\max_{I_{ij} \in \{0,1\}_{j=1}^J} \pi_i(\varphi, I_{i1}, I_{i2}, \dots, I_{iJ}) = \varphi^{\sigma-1} \left(\gamma \sum_{j=1}^J I_{ij} T_j (\tau_{ij} w_j)^{-\theta} \right)^{(\sigma-1)/\theta} B_i - w_i \sum_{j=1}^J I_{ij} f_{ij}, \quad (11)$$

where the indicator variable I_{ij} takes a value of 1 when $j \in \mathcal{J}_i(\varphi)$, and 0 otherwise. The problem in (11) 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. In theory, one could simply calculate firm profits for different combinations of locations and pick the unique strategy yielding the highest level of profits. In practice, however, this would amount to computing profits for 2^J possible strategies, which is clearly infeasible unless one chooses a small enough number J of candidate countries.

Inspection of (11) reveals, however, that the profit function π_i is supermodular in φ and $\Theta_i(\varphi)$, and features increasing differences in (I_{ij}, I_{ik}) for $j, k \in \{1, \dots, J\}$ and $j \neq k$, whenever $(\sigma - 1)/\theta > 1$. These properties of the problem in (11) allow us to establish the following result (the proof is straightforward and is relegated to the Online Appendix):

Proposition 1. The solution $I_{ij}(\varphi) \in \{0, 1\}_{j=1}^J$ to the optimal sourcing problem (11) is such that:

- (a) 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 ;
- (b) if $(\sigma - 1)/\theta \geq 1$, then $\mathcal{J}_i(\varphi_L) \subseteq \mathcal{J}_i(\varphi_H)$ for $\varphi_H \geq \varphi_L$, where $\mathcal{J}_i(\varphi) = \{j : I_{ij}(\varphi) = 1\}$.

Part (a) of Proposition 1 simply states that more productive firms choose a larger sourcing capability—either because they select into more countries or because they select into better countries—thereby magnifying their cost advantage relative to less productive firms. This in turn implies that the equilibrium size distribution of firms will feature more positive skewness than what would be observed without foreign sourcing.

It is important to emphasize that this first result 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 firm might not have an incentive to add further locations to its sourcing strategy. Instead, a low productivity firm from i might not be able to profitably offshore to j^* , but may well find it optimal to source from two foreign countries with associated lower fixed costs.

Part (b) of Proposition 1 states, however, that this possibility can only arise when $(\sigma - 1) / \theta < 1$. When instead $(\sigma - 1) / \theta \geq 1$, the cardinality of the set $\mathcal{J}_i(\varphi)$ is necessarily weakly increasing in φ . Because firm size is increasing in core productivity φ , this prediction is consistent with the upward sloping sales premium documented in Figure 1 in the Introduction. The intuition behind this second result in Proposition 1 rests on the fact that, 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 thus the marginal gain from adding a new location to the set $\mathcal{J}_i(\varphi)$ cannot possibly be reduced by the addition of other countries to the set. This case is more likely to apply whenever demand is elastic and thus profits are particularly responsive to variable cost reductions (high σ), and whenever input efficiency levels are relatively heterogeneous across markets (low θ), so that one achieves a relatively high reduction of costs by adding an extra country into the set of active locations.¹¹

Part (b) of Proposition 1 also has the strong implication that there should be a strict hierarchical order in the extensive margin of offshoring – a ‘pecking order’ which is reminiscent of the one typically obtained in models of exporting with heterogeneous firms, such as Eaton et al. (2011). This prediction is very strong and often violated in the data: it is not uncommon to observe less productive firms sourcing from countries from which more productive firms do not source. Still, in section 3.2 we show that 36 percent of U.S. firms follow the predicted pecking order from the top ten source countries, whereas we would expect only 20 percent to do so if the probabilities to source from individual countries were independent and equal to the share of importers that source from them.

A possible explanation for the violation of a strict hierarchy of import sources is the fact that fixed costs of sourcing might be heterogeneous across firms. With that in mind, our structural estimation in section 4 will incorporate such heterogeneity in fixed costs. In that section, a variant of part (b) of Proposition 1 will be instrumental for reducing the dimensionality of the optimal sourcing problem. In particular, because of increasing differences in the profit function when $\sigma - 1 > \theta$, we can state (see the Online Appendix for a formal proof):

Proposition 2. For all $j \in \{1, \dots, J\}$, define the mapping $V_{i,j}(\varphi, \mathcal{J})$ to take a value of one whenever including country j in the sourcing strategy \mathcal{J} raises firm-level profits $\pi_i(\varphi, \mathcal{J})$, and to take a value of zero otherwise. Then, whenever $(\sigma - 1) / \theta \geq 1$, $V_{i,j}(\varphi, \mathcal{J}') \geq V_{i,j}(\varphi, \mathcal{J})$ for $\mathcal{J} \subseteq \mathcal{J}'$.

The usefulness of this result is best demonstrated with an example. Suppose that one is trying to assess whether a given country j belongs in the firm’s optimal sourcing strategy

¹¹Readers familiar with the work of Eaton and Kortum (2002) might expect that $\theta > \sigma - 1$ is in fact implied by the need for the firm’s 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 9).

$\mathcal{J}_i(\varphi)$. Without guidance from the theory, one would need to compute all 2^J candidate sourcing strategies to answer that question. Proposition 2 implies, however, that if for country j , $V_{i,j}(\varphi, \mathcal{J}) = 1$ when \mathcal{J} is the null set, then j is necessarily in $\mathcal{J}_i(\varphi)$, while if $V_{i,j}(\varphi, \mathcal{J}) = 0$ when \mathcal{J} includes all countries except for j , then j cannot possibly be in $\mathcal{J}_i(\varphi)$. In section 4 we will discuss Jia’s (2008) algorithm, which leverages this logic to devise an iterative algorithm to solve the problem defined in (11) efficiently.

In our above discussion, we have focused on the ‘complements case’ ($\sigma - 1 > \theta$), which allows one to characterize some key properties of the optimal sourcing problem in (11) without any restriction on the relationship between the various countries’ sourcing potentials and fixed costs of sourcing. In the ‘substitutes case’ ($\sigma - 1 < \theta$), this is no longer feasible and one needs to make additional assumptions to obtain a sharp characterization of the firm’s sourcing strategy. For instance, consider a situation in which the fixed costs of offshoring are common for all foreign countries (as in Blaum et al., 2017), so $f_{ij} = f_{ij'}$ for all $j, j' \neq i$. 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, the extensive margin of firms grows 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, however, that this result relies on the assumption of identical offshoring fixed costs across sourcing countries, an assumption that appears particularly unlikely 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 close this section by using the properties of the profit function to discuss comparative statics that apply when holding 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) increases 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 . Third, and perhaps more surprisingly, in the complements case a reduction of *any* τ_{ij} or f_{ij} (weakly) increases firm-level

bilateral input purchases from *all* countries. To see this, note that firm-level intermediate input purchases from any country $j \in \mathcal{J}_i(\varphi)$ are a fraction $(\sigma - 1) \chi_{ij}(\varphi)$ of firm profits, and using (7) and (10), they can thus be expressed as

$$M_{ij}(\varphi) = \begin{cases} (\sigma - 1) B_i \gamma^{(\sigma-1)/\theta} \varphi^{\sigma-1} (\Theta_i(\varphi))^{(\sigma-1-\theta)/\theta} T_j (\tau_{ij} w_j)^{-\theta} & \text{if } j \in \mathcal{J}_i(\varphi) \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

When $(\sigma - 1) / \theta \geq 1$, $M_{ij}(\varphi)$ is thus increasing in all the terms in $\Theta_i(\varphi)$. Intuitively, when demand is sufficiently elastic (i.e., σ is high enough) or the strength of comparative advantage in the intermediate-good sector across countries is sufficiently high (i.e., θ is low enough), the scale effect through the demand response to lower costs dominates the direct substitution effect related to market shares shifting towards the locations whose costs of sourcing have been reduced. It is useful to restate this third result in the following way (see the Online Appendix for a formal proof):

Proposition 3. Holding constant the market demand level B_i , whenever $(\sigma - 1) / \theta \geq 1$, an increase in the sourcing potential $T_j (\tau_{ij} w_j)^{-\theta}$ or a reduction in the fixed cost f_j of any country j , (weakly) increases the input purchases by firms in i not only from j , but also from all other countries.

It should be emphasized that the sharp results above only apply when holding market demand – of which the price index is a key component – fixed. In general equilibrium, these same parameters also affect the level of market demand. As we shall see in our counterfactual exercise in section 5, 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). Despite these nuances, Proposition 3 will still prove to be very useful in interpreting our counterfactual results and in relating them to the observed transformation in the global sourcing practices of U.S. firms over the period 1997-2007.

2.3 Industry and General Equilibrium

Consider now the general equilibrium of the model. As mentioned before, we 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. Given our assumption that final-good producers only observe their productivity after paying the fixed cost of entry, we can use equation (10) to express the free-entry condition in manufacturing as

$$\int_{\tilde{\varphi}_i}^{\infty} \left[\varphi^{\sigma-1} (\gamma \Theta_i(\varphi))^{(\sigma-1)/\theta} B_i - w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} \right] dG_i(\varphi) = w_i f_{ei}. \quad (13)$$

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 (13), 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 Online Appendix we show that one can appeal to monotone comparative statics arguments to prove that:

Proposition 4. Equation (13) 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 (11) 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 (13), by appealing to the marginal cost in (9), to constant-mark-up pricing, and to the fact that spending E_j in manufacturing is a share η of (labor) income. This delivers:

$$N_i = \frac{\eta L_i}{\sigma \left(\int_{\tilde{\varphi}_i}^{\infty} \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} dG_i(\varphi) + f_{ei} \right)}. \quad (14)$$

With this expression in hand, the equilibrium number of active firms is simply given by $N_i [1 - G_i(\tilde{\varphi}_i)]$.¹²

¹²In the Online Appendix (section B.2), we show that in the complements case, and when φ is distributed Pareto with shape parameter κ , we can further reduce equation (14) 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. It is important to stress, however, that this result relies on the

2.4 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 . Using equation (12) and aggregating across firms, we obtain the following expression for aggregate manufacturing imports from country j by firms based in i :

$$M_{ij} = N_i \int_{\tilde{\varphi}_{ij}}^{\infty} M_{ij}(\varphi) dG_i(\varphi) = (\sigma - 1) \gamma^{(\sigma-1)/\theta} N_i B_i T_j (\tau_{ij} w_j)^{-\theta} \Lambda_{ij}, \quad (15)$$

where

$$\Lambda_{ij} = \int_{\tilde{\varphi}_{ij}}^{\infty} I_{ij}(\varphi) (\Theta_i(\varphi))^{(\sigma-1-\theta)/\theta} \varphi^{\sigma-1} dG_i(\varphi). \quad (16)$$

In the second expression, $\tilde{\varphi}_{ij}$ denotes the productivity of the least productive firm from i offshoring to j , while $I_{ij}(\varphi) = 1$ for $j \in \mathcal{J}_i(\varphi)$ and $I_{ij}(\varphi) = 0$ otherwise. We next re-express equation (15) so that it is comparable to gravity equations used in empirical analyses. In particular, plugging the equilibrium values for B_i and N_i in (13) and (14), and rearranging, we obtain

$$M_{ij} = \frac{E_i}{P_i^{1-\sigma}/N_i} \times \frac{Q_j}{\sum_k \frac{E_k}{P_k^{1-\sigma}/N_k} \tau_{kj}^{-\theta} \Lambda_{kj}} \times \tau_{ij}^{-\theta} \times \Lambda_{ij}, \quad (17)$$

where E_i equals country i 's total spending in manufacturing goods (which is a multiple $\sigma/(\sigma-1)$ of country i 's worldwide absorption of intermediate inputs), $Q_j = \sum_k M_{kj}$ denotes the total production of intermediate inputs in country j , and P_i is the ideal manufacturing price index in country i .¹³

Equation (17) resembles a standard gravity equation relating bilateral trade flows to an importer 'fixed effect' (i.e., a term that is common for all exporters holding the importer country constant), an analogous exporter fixed effect, and bilateral iceberg trade barriers τ_{ij} . Notice, however, that equation (17) incorporates an additional term Λ_{ij} that typically varies both across importers and exporters. In fact, the only case in which Λ_{ij} does not vary across exporters is when the fixed costs of offshoring are low enough to ensure that *all* firms acquire the capability to source inputs from *all* countries. In such a case, we have

$$\Lambda_{ij} = \left(\sum_{k \in J} T_k (\tau_{ik} w_k)^{-\theta} \right)^{(\sigma-1-\theta)/\theta} \int_{\tilde{\varphi}_i}^{\infty} \varphi^{\sigma-1} dG_i(\varphi) = \Lambda_i,$$

existence of fixed costs of domestic sourcing which generate a positive measure of inactive firms that do not source any inputs. Because in our empirical work all firms are active and source inputs, we will set $f_{ii} = 0$, and the equilibrium measure of entrants will react to changes in trade costs, wages, and technological parameters.

¹³The ideal manufacturing price index in country i is given by $P_i^{1-\sigma} = N_i \int_{\tilde{\varphi}_i}^{\infty} p_i(\varphi)^{1-\sigma} dG_i(\varphi)$.

and thus Λ_{ij} gets ‘absorbed’ into the importer fixed effect. In this universal importing case, 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} in (7) with $\mathcal{J}_i(\varphi) = J$ for all φ .

When fixed costs of sourcing are large enough to generate selection into importing, 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. As a result, the aggregate elasticity of bilateral trade flows to bilateral trade frictions no longer coincides with the firm-level one, given by θ . In the plausible case in which a reduction in τ_{ij} enhances the extensive margin of imports from country j , the aggregate trade elasticity will thus tend to be higher than θ .

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 $\Theta_i(\varphi)$ and due to industry equilibrium effects. For the special case in which $\sigma - 1 = \theta$, notice however that Λ_{ij} reduces to

$$\Lambda_{ij} = \int_{\tilde{\varphi}_{ij}}^{\infty} \varphi^{\sigma-1} dG_i(\varphi) = \Lambda_{ij}(\tilde{\varphi}_{ij}).$$

Thus, 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 (i.e., a reduction in $\tilde{\varphi}_{ij}$), it is clear that the elasticity of bilateral trade flows with respect to τ_{ij} will now be higher than the firm-level one. Furthermore, as we show in the Online Appendix (section B.2), if we assume that firms draw their core productivity from a Pareto distribution with shape parameter κ (assumed to be higher than $\sigma - 1$ to ensure a finite variance of sales), we can express aggregate manufacturing imports from country j by firms based in i as

$$M_{ij} = \frac{(E_i)^{\kappa/(\sigma-1)}}{\Psi_i} \frac{Q_j}{\sum_k \frac{(E_k)^{\kappa/(\sigma-1)}}{\Psi_k} (\tau_{kj})^{-\theta} (f_{kj})^{1-\kappa/(\sigma-1)}} (\tau_{ij})^{-\kappa} (f_{ij})^{1-\kappa/(\sigma-1)}, \quad (18)$$

where $\Psi_i = f_{ei} \varphi_i^{-\kappa} P_i^{-\kappa} w_i^{\kappa/(\sigma-1)-1} / L_i$. Notice that equation (18) is a well defined gravity equation in which the ‘trade elasticity’ (i.e., the elasticity of trade flows with respect to variable trade costs) can still be recovered from a log-linear specification that includes importer and exporter fixed effects. But notice 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 $\kappa > \sigma - 1 = \theta$).¹⁴

¹⁴It 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

The knife-edge case $\sigma - 1 = \theta$ is useful in illustrating why one should expect the aggregate trade elasticity to be larger than the firm-level one. Yet it masks the fact that whenever $\sigma - 1 \neq \theta$, Λ_{ij} in (16) will be a function of $I_{ij}(\varphi) \Theta_i(\varphi)$ for $\varphi > \tilde{\varphi}_{ij}$, and will thus depend on which other countries are included in the sourcing strategy of firms from i sourcing from j and those other countries' characteristics. In such a case, equation (17) becomes 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 (16)), it appears that the sign of these third-market effects depends crucially on whether $\sigma - 1 > \theta$ or $\sigma - 1 < \theta$. 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.

Interestingly, our model suggests a relatively simple way to control for these extended gravity forces. In particular, defining the importer-specific term $\Xi_i = T_i(\tau_{ii}w_i)^{-\theta} (\sigma - 1) \gamma^{(\sigma-1)/\theta} N_i B_i$, note that we can express

$$\Lambda_{ij} = \frac{1}{\Xi_i} \times (\sigma - 1) N_i B_i \gamma^{(\sigma-1)/\theta} \int_{\tilde{\varphi}_{ij}}^{\infty} I_{ij}(\varphi) \varphi^{\sigma-1} (\Theta_i(\varphi))^{(\sigma-1-\theta)/\theta} T_i(\tau_{ii}w_i)^{-\theta} dG_i(\varphi),$$

where the second term on the right-hand-side corresponds to the *domestic* input purchases aggregated over all firms based in i that import inputs from j . In section 4.2, we show that when including this bilateral aggregate measure of domestic input purchases into a standard gravity specification, the resulting estimate of the trade elasticity θ becomes much lower, in line with the one we estimate at the firm level.

3 Data Sources and Descriptive Evidence

In the theory sections, 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

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 (18) is essentially identical to those obtained by Chaney (2008) or Arkolakis et al. (2008).

the structural estimation, we describe the data used in the paper and provide several novel empirical facts that support the theoretical framework.

3.1 Data Description

The primary data used in the paper are from the U.S. Census Bureau’s 1997 and 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.¹⁵ 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.¹⁶ We also limit the sample to firms with positive sales and employment and exclude all mineral imports from the analysis since they do not represent offshoring. Firms with at least one manufacturing plant account for five percent of firms, 23 percent of employment, 38 percent of sales, and 65 percent of non-mineral imports. In terms of explaining aggregate U.S. sourcing patterns, it is critically important to include manufacturing firms with non-manufacturing 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 in 2007— technology and wages, as well as other controls— from various sources. Country

¹⁵The 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 a more complete depiction of the entire firm compared to studies that rely solely on the CM.

¹⁶We 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, 2015, for details). Unfortunately, there is no practical way to distinguish all FGPs from traditional wholesale establishments. Furthermore, data on 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.

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).

3.2 Descriptive Evidence

We use the 2007 data to assess the model’s assumptions and predictions. First, we provide information on the number of products imported by U.S. firms, and the number of countries from which they source. Second, we show that firms generally source each input from a single location. Finally, we document the extent to which firm sourcing decisions follow a hierarchical pattern.

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. 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. The data show that importers source an average of 12 distinct products from about three foreign countries. The median number of imported products is two, while the 95th percentile is 41. The median number of source countries is two, and the 95th percentile is 11.

One feature of our model is that it delivers a closed-form solution for the share of inputs a firm sources from a particular country. This solution comes from an Eaton and Kortum (2002) selection process in which a firm sources each input from the single, lowest cost location. Table 2 shows that this feature of our model is consistent with the data. The table presents statistics on the firm-level mean, median, and maximum number of countries from which a firm imports a particular product. We report the mean, median, and 95th percentile of these firm-level measures. The median firm imports each distinct product from an average of only one country. The median number of countries per product for firms is always one, 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.¹⁷

¹⁷In the Online Appendix (section C.10) 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 show that this pattern is not driven by sparsity in the data, since the same firm-level statistics on the number of products per country are

Table 2: Firm-level statistics on the number of source countries per imported product

	Firm-level		
	Mean	Median	Max
Mean	1.11	1.03	1.78
Median	1.00	1.00	1.00
95%tile	1.61	1.00	4.00

Notes: Table reports on the number of countries from which a firm imports the same HS10 product.

Before turning to the structural estimation, it is useful to assess the extent to which firms follow a hierarchical pecking order in their sourcing behavior. To do so, we follow Eaton et al. (2011) and 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, and so on. We calculate these statistics irrespective of firm sourcing outside the top ten countries. Columns 1 and 2 in Table 3 show that over 21 thousand firms, or 36 percent of importers, follow a pecking order. To assess the significance of this share, we calculate the share of firms that would follow this hierarchy if firms selected into countries randomly. Specifically, we use the share of importers from country j as the probability that any given firm will source from j , and we assume that each probability is independent. Column 5 shows that fewer than 20 percent of firms would follow a pecking order under random entry—just over half the share observed in the data.

The results here are similar to those in Eaton et al. (2011) where the authors show that 27 percent of French exporters follow a pecking order for the top seven destinations, which is more than double what would be predicted under the same random entry calculation. While our findings are certainly suggestive of a pecking order in which country characteristics make some countries particularly appealing for all U.S. firms, they also point 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.

always greater than 1. We also provide the statistics at the HS6 level, and we show that every statistic on the number of countries from which a firm sources a given product is equal to or lower than the comparable statistic for the number of countries to which a firm exports a given product.

Table 3: U.S. firms importing from strings of top 10 countries

String	Data		Random Entry	
	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.

4 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 source-destination-specific country characteristics. 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. First, 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 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.

4.1 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 (7) in the model that a firm’s share of inputs sourced from country j , χ_{ij} , is simply that country’s contribution to the firm’s sourcing capability, Θ_i^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 (7) by taking logs and normalizing the share of inputs purchased from country j by the firm’s share of domestic inputs leads to

$$\log \chi_{ij}^n - \log \chi_{ii}^n = \log \xi_j + \log \epsilon_j^n, \quad (19)$$

where n denotes the firm. In order to turn the model’s equilibrium condition (7) into an empirical specification, note that this equation includes a firm-country-specific shock ϵ_j^n . When normalizing by the domestic share, we set domestic sourcing potential to 1.

The dependent variable in equation (19) 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, production worker wages, and total 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.

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. It is also consistent with firm-country-specific shocks to the fixed costs of sourcing. In what follows we treat ϵ_j^n as measurement error.¹⁸

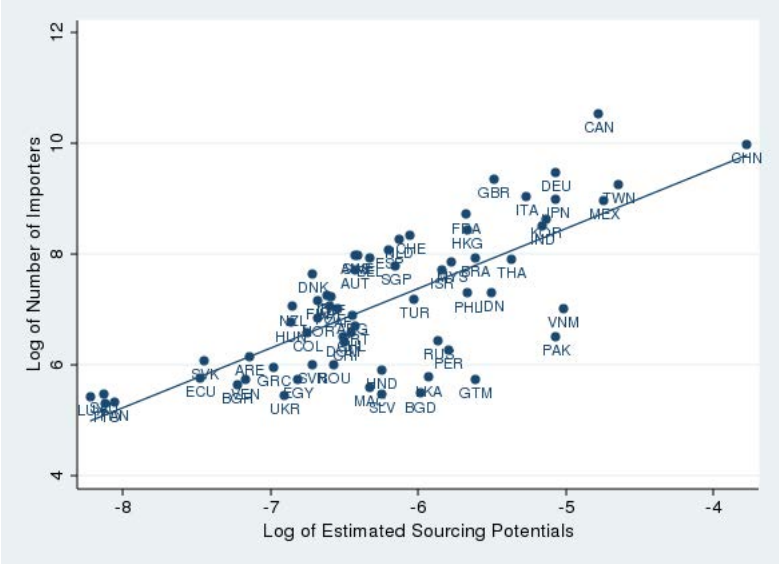
We estimate equation (19) via Ordinary Least Squares (OLS), using country fixed effects to capture the ξ_j terms. The estimated coefficients on these fixed effects represent each country’s sourcing potential, which we note is simply the average share difference by country. By esti-

¹⁸Note 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. Online Appendix Section C.5 discusses the robustness of these estimates.

mating sourcing potentials via OLS, however, we also calculate standard errors, which show that all the sourcing potential fixed effects are significant at the one 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 the number of firms importing from that country. Our parameter estimates suggest that China has the highest sourcing potential for U.S. firms, followed by Canada and Taiwan. More firms import from Germany and the 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. The variation in this figure is helpful for understanding how the structural estimation will identify country fixed costs. For instance, the fact that more firms import from Canada than China, despite the fact that China has the highest estimated sourcing potential, suggests that the fixed cost of importing from Canada are lower than the fixed costs of importing from China.

Figure 2: Country sourcing potential parameters and the extensive margin



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 domestic sourcing potential is normalized to one and the summation of the foreign sourcing potential terms is 0.193, these results imply that the sourcing capability of a firm that purchases inputs from all 67 countries is 19.3 percent larger than that of a firm buying inputs 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 (9). 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.

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

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 findings.¹⁹ Given the potential issues that may affect the accuracy of our estimate of the demand elasticity, we include a sensitivity analysis in section 5.3, in which we consider alternative values for the elasticity of demand, as well as other parameters.

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 , which 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:

$$\begin{aligned} \log \hat{\xi}_j = & \beta_0 + \beta_r \log \text{R\&D}_j + \beta_k \log \text{capital}_j + \beta_F \log \text{number of firms}_j - \theta \log w_j \\ & - \theta (\log \beta_c + \beta_d \log \text{distance}_{ij} + \text{language}_{ij} \log \beta_l + \beta_C \text{control of corruption}_j) + \nu_j, \end{aligned} \quad (20)$$

so that iceberg trade costs are proxied by log distance, common language, and control of corruption.

Equation (20) shows that the parameter θ can be recovered from the estimated coefficient

¹⁹For 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.

on wages. In theory, one could also identify θ using tariffs, but, as we show in the Online Data 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 of labor, we follow Eaton and Kortum (2002) and use a human-capital-adjusted wage. Even adjusting for skill differences across workers, there are other country-level factors that are likely correlated with the average wage, such as infrastructure, that will lead to an upward bias on the wage coefficient. To address this issue, as well the potential for measurement error, we instrument for a country's wage using its population. The first stage results, presented in the Estimation Appendix, show a negative and statistically significant coefficient on the log of country 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. To address that possibility, 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 57 of the 66 foreign countries included in the structural estimation.

The first column of Table 4 presents the results from estimating equation (20) via OLS. Column 2 provides the analogous IV estimates, using population as an instrument for wages. As expected, the IV estimate for θ (1.789) is larger than the OLS estimate. In line with the discussion in section 2.4, 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 3.60 using data on wages. Column 4 shows that we obtain a similar coefficient of 4.544 when estimating the same specification as in equation (20), but using aggregate imports as the dependent variable. In the Online Data Appendix we also provide 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.

Motivated by our discussion at the end of section 2.4, in column 5 we present estimates from a specification that includes the log of domestic input purchases aggregated over all U.S. importers from country j . To the extent that this term controls for the term Λ_{ij} in (15), one would expect this specification to deliver a lower estimate of θ that is in line with the estimate in column 2. Indeed, the resulting estimate of θ (1.268) is lower than the one in column 4 and quite close to our preferred estimate from the firm-level data presented in column 2. This result is reassuring, but should be interpreted with caution for two reasons. First, the orthogonality condition which ensures that our firm-level estimate of θ in column 2 is consistent, does *not* guarantee that the estimate in column 5 is consistent as well (see our Online Data Appendix

Table 4: Estimation of firm and aggregate trade elasticities

	log ξ		log aggregate imports		
	OLS	IV	OLS	IV	IV
log HC adjusted wage	-0.537 (0.184)	-1.789 (0.696)	-0.643 (0.390)	-4.544 (1.844)	-1.268 (0.768)
log distance	-0.341 (0.197)	-0.621 (0.294)	-0.859 (0.418)	-1.733 (0.779)	-0.650 (0.333)
log R&D	0.352 (0.068)	0.524 (0.125)	0.763 (0.144)	1.298 (0.332)	0.251 (0.176)
log capital/worker	-0.184 (0.175)	0.425 (0.390)	-0.264 (0.370)	1.633 (1.033)	0.308 (0.421)
common language	0.105 (0.223)	0.146 (0.289)	0.354 (0.471)	0.479 (0.764)	0.137 (0.317)
control corrupt	0.156 (0.151)	0.621 (0.312)	0.365 (0.319)	1.816 (0.826)	0.414 (0.350)
log no. of firms	0.108 (0.086)	-0.020 (0.130)	0.031 (0.183)	-0.369 (0.345)	-0.062 (0.142)
log domestic purchases					2.392 (0.327)
Constant	-7.250 (0.922)	-11.068 (2.323)	14.499 (1.952)	2.600 (6.156)	-37.389 (6.573)
Observations	57	57	57	57	57

Notes: Standard errors in parentheses. In the IV specifications, the human-capital-adjusted wage is instrumented by population. HC adjusted wage is country wage adjusted for differences in human capital. Domestic purchases is total purchases of U.S. inputs by firms sourcing from a country. First-stage F-statistic on the excluded instrument is 6.49. First-stage regression results are in the Appendix.

for details). Second, and perhaps relatedly, the coefficient we obtain on log domestic purchases is higher than the theoretically-predicted value of 1.²⁰ For these two reasons, we henceforth treat our column 2 coefficient $\theta = 1.789$ as our benchmark estimate of θ .

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 nine percent ($1.193^{(-1/1.789)}$) lower input costs than a firm sourcing purely domestically, and consequently its sales are made around 32 percent ($1.193^{(2.85/1.789)}$) larger by sourcing inputs from all countries.²¹

²⁰When constraining that coefficient to be 1, we estimate a somewhat larger $\theta = 3.175$.

²¹These calculations are based on an ‘average’ firm importing inputs from all countries. In the data, there is of course variation in spending shares across firms, even conditional on their sourcing strategy. In Online Appendix Table C.4, we present statistics on the distribution of total foreign sourcing shares. Using the ‘sufficient statistic’ formula in Blaum et al. (2017) and our estimate of θ , we can compute the cost savings associated with importing for firms in the 90th percentile. These firms import 47 percent of their input purchases, which implies 30 percent cost savings, and a 176 percent increase in their sales due to global sourcing.

Across various specifications, including the 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. Furthermore, for our benchmark estimate $\theta = 1.789$, we can reject the null hypothesis that $\theta > \sigma - 1 = 2.85$ at reasonable significance levels (the p-value for this test is 0.06). As shown in section 1, 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 to solve the firm’s problem numerically and thereby estimate the fixed costs of sourcing from different markets.

4.3 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 as well as on a measure of the source country’s control of corruption. Thus far, the model features fixed costs that differ across countries, but not across firms. This implies that the number of importing firms will be identical to the number of firms that source from the most popular country. In contrast, the data show that around 64,600 firms import, while only about 37,800 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 of sourcing, f_{ij}^n , which are drawn from a lognormal distribution with dispersion parameter β_{disp}^f and scale parameter $\log \beta_c^f + \beta_d^f \log \text{distance}_{ij} + \log \beta_l^f \text{language}_{ij} + \beta_C^f \text{control of corruption}_j$.²² Since by definition active manufacturing firms must use some domestic inputs (at least in terms of production worker services), we cannot identify the fixed costs of domestic sourcing, so we set them equal to zero ($f_{ii}^n = 0$).

In a setting with a large number of countries, the firm faces an enormous discrete choice problem when solving for its optimal sourcing strategy. If there are 66 foreign countries, the firm selects among 2^{66} , which is roughly 10^{20} , possible sourcing strategies. Clearly, calculating the profits for each of these strategies for every firm is infeasible. To reduce the dimensionality of the firm’s problem, we rely on Proposition 2 and adopt an algorithm first developed by Jia (2008). The specifics of the algorithm are as follows. Given a core productivity φ and a guess \mathcal{J} for the firm’s sourcing strategy, \mathcal{J}^n , define the marginal benefit of including country j in

²²Analogous to the core productivity level, φ , the firm learns about its fixed costs of sourcing, f_{ij}^n , after having paid the fixed entry cost, f_{ei} , which is homogeneous among potential firms in country i .

the sourcing strategy \mathcal{J} as:

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

As in Proposition 2, we define a mapping, $V_j^n(\mathcal{J})$ that takes a value of one if this marginal benefit is positive, and takes a value of zero otherwise. Under the empirically relevant condition $\sigma - 1 > \theta$, Proposition 2 shows that this mapping is an increasing function of \mathcal{J} . Jia (2008) shows that when starting from the set $\underline{\mathcal{J}}$ (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 $\overline{\mathcal{J}}$ (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 are not identical (see Appendix Table A.4).²³

Turning to the practical implementation of our structural estimation, we follow Melitz and Redding (2015) and assume that firms' core productivities are distributed Pareto. As in their paper, we set the shape parameter κ of the Pareto distribution equal to 4.25. We do not estimate κ to ease the computational burden associated with estimating non-linearly the other parameters of the model, but we discuss sensitivity results with alternative values for κ in section 5.3. This leaves the following six parameters to estimate: $\delta = \left[B, \beta_c^f, \beta_d^f, \beta_l^f, \beta_C^f, \beta_{\text{disp}}^f \right]$.²⁴ 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

²³In 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.

²⁴We set $\gamma^{1/\theta} \varphi_{\text{US}} = 1$, as it scales input purchases equivalently to an increase in B .

from a standard normal distribution (which, given a parameter guess δ , can be used to calculate the lognormal distributed firm-country specific fixed cost level).²⁵ Note that there is no relationship between the number of simulated firms and the number of actual firms in the data. The model assumes that we have a continuum of firms whose core efficiency, fixed cost draws, and country-specific efficiency shocks follow particular distributions, and we use the simulated firms as evaluation points of these distributions.

We use the simulated firms to construct three sets of moments. The first set of moments includes: a) the share of importers for all manufacturing firms (about a quarter of all firms); and b) the share of importers with firm sales below the median (8.4 percent). This is simply a 2×1 vector, 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 imports from each country. We label this $(J - 1) \times 1$ vector of moments in the data as m_2 and the simulated moment vector as $\hat{m}_2(\delta)$. These two sets of moments are informative about the overall magnitude of the fixed costs of sourcing, as well as on how they vary with distance, language, and control of corruption. In addition, the share of importing firms from the most popular country relative to the total share of importers is indicative of the fixed cost dispersion parameter. The intuition here is that if there were zero dispersion in fixed costs across firms, the total share of importers would be identical to the share of importers from the most popular sourcing country. Similarly, the share of importers among firms with sales below the median firm is informative about the dispersion parameter. For example, even in a model with i.i.d. fixed cost draws across firms, but common draws across countries, an increase in the dispersion of fixed cost draws will lead to more importers in that size bin of firms compared to the share of importers in the full set of firms. Finally, the last moment included in the estimation 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 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 , as B governs the level of input purchases.

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

²⁵We 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.

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

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

$$E[\hat{y}(\delta_0)] = 0. \tag{21}$$

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

$$\hat{\delta} = \arg \min_{\delta} [\hat{y}(\delta)]^\top \mathbf{W} [\hat{y}(\delta)], \tag{22}$$

where \mathbf{W} is a weighting matrix. We weight the moments equally, hence the weighting matrix is the identity matrix.²⁶

The parameter estimates are displayed in Table 5 below. We find that the fixed costs of sourcing are increasing in distance with an elasticity of 0.19, and that sourcing from countries with a common language reduces fixed costs by about 13 percent. The fixed costs of sourcing also seem reasonable in magnitude. The median fixed cost estimate ranges from 10,000 to 56,000 USD, though the assumption of a lognormal distribution means they can be substantially larger for some individual firm-country combinations. Out of the total variance of fixed cost draws across firms and countries, the country-level variation explains about 9 percent of the total variance, with the remainder explained by firm-country-level variation.

Table 5: Estimated parameters

B	β_c^f	β_d^f	β_l^f	β_C^f	β_{disp}^f
0.122	0.022	0.193	0.872	-0.393	0.934
(0.004)	(0.002)	(0.018)	(0.024)	(0.012)	(0.018)

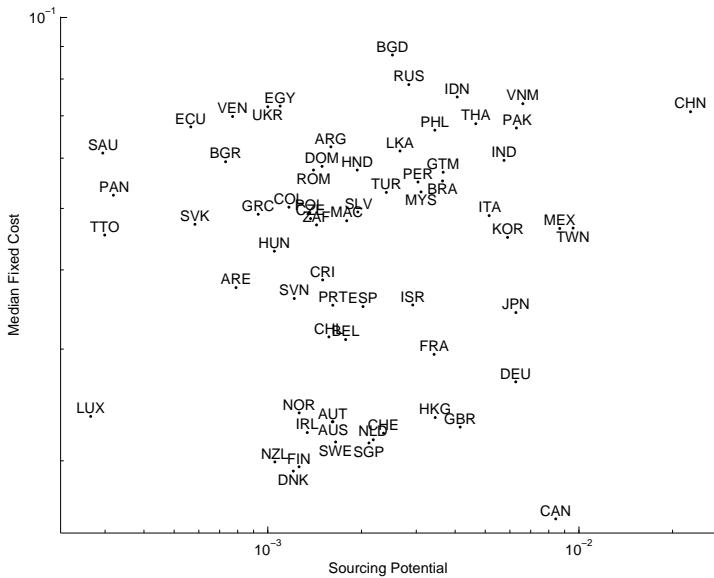
Notes: Table reports coefficients and standard errors from estimating the model via simulated method of moments Standard errors based on 25 bootstrap samples drawn with replacement.

In Figure 3, we show the estimated sourcing potential and median fixed costs by country. It is clear that Canada has one of the highest sourcing potentials, but the lowest fixed costs. Mexico has both higher sourcing potentials and higher fixed costs than Germany and the United Kingdom. These differences in sourcing potentials and fixed costs help reconcile the variation

²⁶We use the identity matrix instead of the inverse of the estimated variance-covariance matrix of the 69 moments as the weighting matrix, since the former leads to a better fit of the import shares of the most popular countries, in particular China, which are most relevant for the counterfactuals below (at the expense of a worse fit of the shares of less popular importing countries).

in a country’s rank in terms of the number of importing firms and the import values displayed in Table 1. They also highlight the importance of heterogeneous fixed costs in matching the model to the data. We conclude this section by describing the estimated model’s fit of the data, and by assessing the need for heterogeneous fixed costs for a good fit.

Figure 3: Estimated sourcing potential and median fixed costs by country

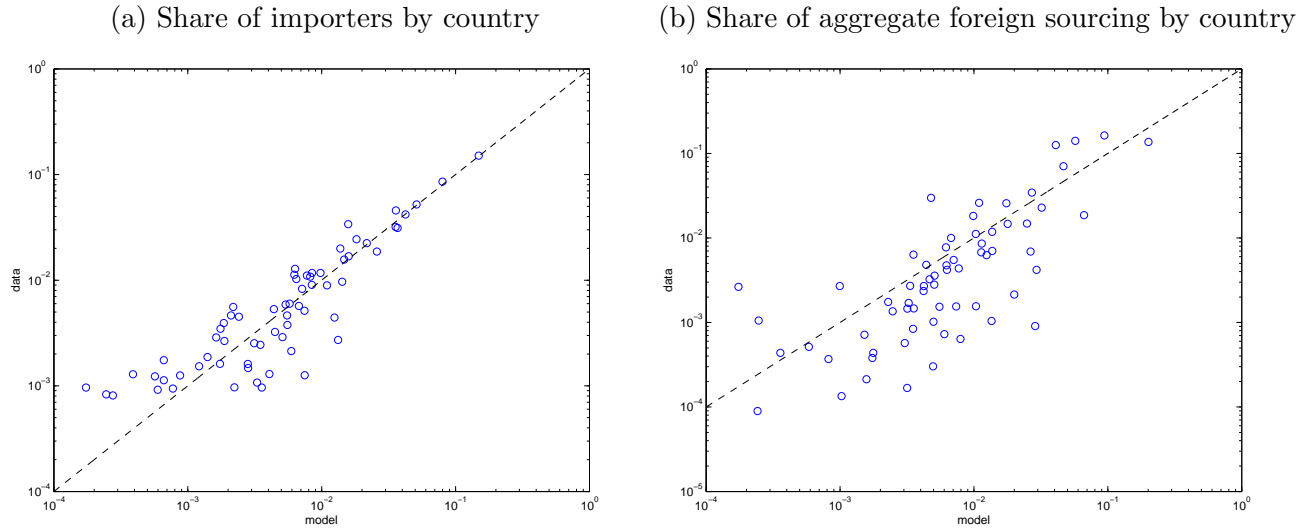


4.4 Fit of the Model

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 data and baseline model, around 26 percent of U.S. firms import (25.8 in the data and 26.8 in the baseline model). For U.S. firms with sales below the median, around 8 percent import (8.5 in the data and 7.3 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 0.98. Figure 4a 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 approximately equal (568,000 in the data and 572,000 in the model). The model also does a good job at matching two sets of moments we did not target directly in the estimation. As shown in Figure 4b, the correlation between the actual and simulated import shares by country is very high (the correlation coefficient is 0.78). In Table 6, we also show that the estimated model does a good job of matching the hierarchal sourcing patterns of firms. While the model does well in all these dimensions, it overpredicts domestic sourcing by about 18.25 percent.²⁷

²⁷This is not apparent from Figure 4b because the figure plots each country’s share in foreign (and not overall) sourcing. There are two possible reasons for this overprediction of the domestic sourcing share. First,

Figure 4: Model fit: share of firms and aggregate sourcing



It is important to emphasize that our model fits the data substantially better than a model with common fixed costs across countries, but with idiosyncratic variation in fixed cost draws across firms. Such a model would be much simpler to estimate, and would not require our adoption of Jia’s (2008) algorithm to solve the firm’s problem. This simplification, however, comes at a large expense. For example, in the common fixed costs model only 17 percent of firms import, as opposed to 26 percent in the baseline model and data (in the common fixed cost model, the share of importing firms equals the share of firms importing from the most popular country). The correlation between the common fixed cost model’s predictions and the data is only 0.68 for the share of importers by country and only 0.67 for import shares by country. Furthermore, while the baseline model does a good job at fitting the hierarchies in firms’ sourcing patterns, the common fixed costs model does poorly, since – given the estimated sourcing potentials – it predicts, for instance, that China, not Canada, is the number one source of inputs.²⁸

as discussed above, we set the fixed costs of domestic sourcing to zero. Second, our sourcing potential estimates, which represent the (geometric) average ratios of firms’ sourcing from a given foreign country relative to the US, may understate the importance of foreign inputs for selected large importers, and therefore lead to an underprediction of aggregate imports (see Online Appendix section C.5 for a detailed discussion of these sourcing potential estimates).

²⁸In Appendix Table A.2 we also illustrate the superior fit of the statistics described in Table 1 by the baseline model compared to the model with common fixed costs. Note that one important point about the comparison of the heterogeneous versus homogeneous fixed cost models is that the former has two sources heterogeneity: country and firm-country-specific variation. A model with common mean fixed costs across countries, but with firm-country-specific variation in the fixed costs, would still require our algorithm in order to solve the firm’s problem, since for a given firm, fixed costs would again be heterogeneous across countries.

Table 6: Hierarchies in sourcing patterns: data and model

String	Data	Baseline model
CA	29.82	28.66
CA-CH	3.67	3.42
CA-CH-DE	0.56	0.58
CA-CH-DE-GB	0.25	0.14
CA-CH-DE-GB-TW	0.13	0.09
CA-CH-DE-GB-TW-IT	0.05	0.02
CA-CH-DE-GB-TW-IT-JP	0.05	0.03
CA-CH-DE-GB-TW-IT-JP-MX	0.08	0.07
CA-CH-DE-GB-TW-IT-JP-MX-FR	0.27	0.12
CA-CH-DE-GB-TW-IT-JP-MX-FR-KR	1.08	0.66
TOTAL Following Pecking Order	36.0	33.8

Notes: This table depicts the percentage of importers following a particular sourcing pattern. The first row shows the percentage of firms only importing from Canada; the second row shows the percentage of firms only importing from Canada and China; and so forth (irrespective of firm sourcing outside these top 10 countries). The ranking of countries is determined by the number of firms sourcing from these countries in the data.

5 Counterfactual: An Increase of China’s Sourcing Potential

In this section, we use the parameter estimates from section 4 to assess how firm-level import decisions, the firm size distribution, and aggregate sourcing by country respond to a shock in China. We focus on China not only because it is one of the biggest U.S. trade partners, but also because its accession to the WTO in 2001 provides an actual shock in the data against which we can compare the model’s predictions. The aggregate U.S. employment effects of this shock have been studied in several contexts (e.g., Autor et al., 2013, 2014; Pierce and Schott, 2016), but with a focus on import competition rather than offshoring opportunities. Our results highlight the empirical relevance of interdependencies inherent in firm-level sourcing decisions, and show that these interdependencies lead to significant heterogeneity in the firm-level implications for domestic employment.

We model the China shock as a change to China’s sourcing potential that is large enough to explain the observed 178 percent increase in the Chinese share of US manufacturer’s imports between 1997 and 2007. Specifically, when multiplying the Chinese sourcing potential estimate for 2007 by a factor of 0.46, and holding all other exogenous variables fixed – but re-solving for

the price index and the mass of firms – we can match the observed increase in aggregate imports from China.²⁹ Our *baseline* specification uses the estimated parameters for 2007, but with the lower 1997 Chinese sourcing potential. We then analyze the changes our model predicts from increasing Chinese sourcing potential to its actual estimated value in 2007. Although it might have been more natural to re-estimate our model for the year 1997, data limitations (particularly the poor coverage of U.S. importers from Canada) preclude us from doing so in a proper manner. 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.³⁰ While firms in our static model are only one-period lived, we can nevertheless compare firms with the same productivity levels and fixed cost draws before and after the China shock.

For the first part of the analysis, we do not take a stance on the underlying cause of the shock to China’s sourcing potential, but when mapping our counterfactuals to reduced-form evidence, we will attempt to isolate a shock to Chinese productivity in the production of intermediate inputs. It is important to emphasize that we are not able to trace the responses of final-good producers in foreign countries to this China shock. As long as wages are pinned down by a non-manufacturing sector, these foreign responses are irrelevant for the quantitative implications of the shock for the sourcing decisions of U.S. final-good producers. Even with endogenous wages, the qualitative implications of the shock would also be unaffected by the decisions of final-good producers abroad. Nevertheless, the aggregate implications of our model for overall sales of U.S. intermediate-input producers and for overall U.S. manufacturing employment could well be affected by these foreign responses.

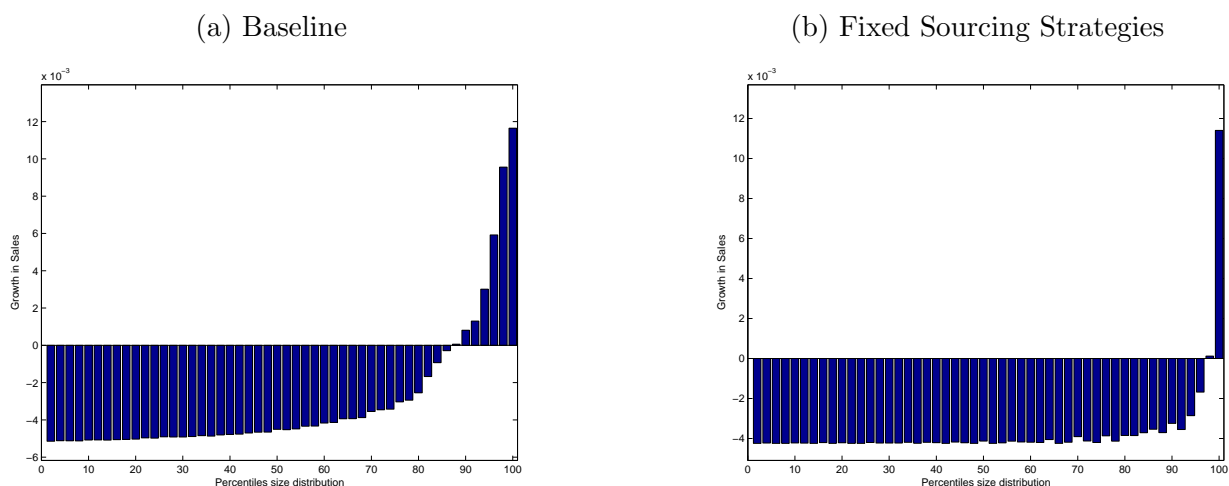
²⁹This counterfactual sourcing potential value still implies that China was the number one country for cost savings in 1997, with Taiwan as the next best country. Of course, the increase in the share of firms sourcing from China may have been triggered by both an increase of the Chinese sourcing potential and a reduction of the fixed costs of sourcing from China. Since one of our goals is to compare our model’s predictions with those from a model with common fixed costs across countries, we focus on a change to only the sourcing potential. In Online Appendix C.8, we alternatively consider a shock only to fixed costs that achieves the same increase in the Chinese import share. If Chinese sourcing potential had not changed, the fixed costs of sourcing from China in 1997 would have needed to be almost nine times larger than their estimated size in 2007. The observed price index change would have been similar to the findings under the sourcing potential shock, though the counterfactual underpredicts the share of firms importing from China in 1997, and obviously it cannot generate an expansion in U.S. and third-market sourcing by continuers (a feature of the data we will document in section 5.4).

³⁰Given our parameter estimates and the counterfactual parameter change, the calculation of the counterfactual price index and mass of firms works as follows. We set total expenditure, E , equal to the the total sales of firms in the data. We then use our estimate of B together with equation (3) and the ideal CES price index associated with (1) to back out the mass of firms, N , in the equilibrium associated with the estimated parameters. Using the free entry condition (13), we set the fixed costs of entry to the average profits of firms implied by the estimated model. In our counterfactual, after changing the Chinese sourcing potential, we solve for the new level of B such that the free entry condition (13), is again satisfied and – given the new level of B – use (14) to determine the new mass of firms. The new equilibrium price index then follows from the standard formula.

5.1 Baseline Predictions

Table 7 (Panel A) documents how the China shock affects sourcing in various markets for different sets of firms. We find that the shock induces 5.3 percent of firms to start importing from China, which is around 80 percent of the actual entry rate observed in the data (6.6 percent). These entrants account for about a third of Chinese imports after the shock. Consistent with the complementarities highlighted by our model, these new China importers increase their sourcing from the U.S. and from third countries by 0.8 and 1.5 percent, respectively. Firms that continue sourcing from China comprise 2.7 percent of firms, and they also slightly increase their domestic and third country sourcing. The expansion of these new and continuing China importers is associated with these firms' being able to sell their products at lower prices, which in turn leads the aggregate price index and the mass of active firms to adjust. As a result, firms for which the shock is not large enough to induce importing from China face tougher competition, which leads their sourcing from the U.S. and other foreign countries to contract by 0.5 and 1.3 percent, respectively. The responses in sourcing from other foreign countries are generally larger than for domestic sourcing, since the former involves both intensive and extensive margin adjustments (i.e., firms selecting into or out of particular countries). The extensive margin adjustments are quantitatively important. The addition of new source countries accounts for 50 percent of the aggregate change in third country sourcing by Entrants, 29 percent of the change by Continuers, and 58 percent of the change by Others. Sourcing responses are also necessarily larger for new China importers, since these firms' extensive margin sourcing change leads them to grow relatively more.³¹

Figure 5: Changes in the size of firms



³¹It is important to note that continuing Chinese importers need not necessarily increase their domestic and third market sourcing. If the price index decline were sufficiently large, the interdependencies in their extensive margin sourcing decisions which lead them to grow could be swamped by the increased competitive pressure that leads all firms to shrink.

Table 7: Third country sourcing effects of Chinese sourcing potential shock

Chinese import status	Change in sourcing from U.S.	Change in sourcing from other countries	Change in sourcing from China	Share of firms	Share of imports from China
<i>Panel A: Baseline model</i>					
Entrants	1.008	1.015	∞	0.053	0.334
Continuers	1.001	1.001	2.148	0.027	0.666
Others	0.995	0.987	-	0.920	0.000
<i>Panel B: Baseline model, fixed sourcing strategies</i>					
Entrants	-	-	-	0.000	0.000
Continuers	1.002	1.002	2.150	0.027	1.000
Others	0.996	0.996	-	0.973	0.000
<i>Panel C: Independent entry decisions model ($\theta = \sigma - 1 = 2.85$)</i>					
Entrants	0.997	0.994	∞	0.053	0.343
Continuers	0.997	0.996	2.116	0.027	0.657
Others	0.997	0.992	-	0.920	0.000
<i>Panel D: Universal importing model (no fixed costs)</i>					
Entrants	-	-	-	0.000	0.000
Continuers	0.987	0.987	2.976	1.000	1.000
Others	-	-	-	0.000	0.000
<i>Panel E: Common fixed costs across countries model (heterogeneous across firms)</i>					
Entrants	1.006	-	∞	0.108	0.328
Continuers	0.999	0.999	2.249	0.066	0.672
Others	0.993	-	-	0.826	0.000

Notes: The table groups firms by Chinese import status. Entrants are those firms (i.e. bundles of productivity levels and fixed cost draws) that begin sourcing from China. Continuers are firms that source from China before and after the shock. Others are firms that do not source from China before or after the shock. As described in the main text, the shock to the Chinese sourcing potential is calibrated so that the Baseline model matches the observed 178 percent increase in the Chinese import share from 1997 to 2007. For the Panels C, D, E we re-calibrate the shock to the Chinese sourcing potential so that the respective models can match the same 178 percent increase in the Chinese import share. Columns 1, 2, and 3 contain the ratio of the total sourcing by each group of firms before and after the shock. Figures in this table do not include changes in sourcing due to changes in the mass of firms after the counterfactual shock. Changes in sourcing due to changes in the overall mass of firms are included in the figures in Table 8.

Figure 5a depicts the percent growth in sales by firm size percentile. Firms are ranked according to their sales before the China shock. Large firms increase their sales by about 1.5 percent. While these magnitudes may 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. Firms below the 86th percentile of the size distribution see their sales shrink on average in response to the China shock.

Overall, aggregate imports from China increase by 222 percent, while aggregate sourcing from the U.S. falls by 0.53 percent (see Table 8).³² The net change in U.S. sourcing masks a substantial amount of domestic churn. As shown in Table 8, new and continuing importers from China increase their domestic sourcing by 4.41 billion USD (or 0.11 percent of total sourcing). Comparing this figure to the total decline in domestic sourcing accounted for by contracting or exiting firms (i.e., 25.04 billion USD) implies that every dollar reduction in domestic sourcing is thus partly offset by a 18 cent increase in domestic sourcing from expanding firms. Since in our model this spending represents jobs in the intermediate goods sector, we predict a substantial amount of employment churn in response to an increase in Chinese sourcing potential. We also find that an increase in China’s sourcing potential lowers the price index in the manufacturing good sector by 0.19 percent. Consumers are therefore better off since they enjoy lower prices for the goods they buy.

Table 8: Gross and net U.S. sourcing effects

	Baseline		Fixed sourcing strategies	
	Difference in sourcing in billion USD	Change in percent of total U.S. sourcing	Difference in sourcing in billion USD	Change in percent of total U.S. sourcing
Increase in domestic sourcing	4.41	0.11	1.97	0.05
Decrease in domestic sourcing	-13.93	-0.36	-12.86	-0.33
by firms that continue to operate				
Decrease in domestic sourcing	-11.11	-0.28	0	0
by firms that shut down				
Total	-20.63	-0.53	-10.88	-0.28

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.

³²The increase in aggregate imports from China is larger than the calibrated 178 increase in the Chinese import share. These numbers differ because an increase in aggregate imports from China raises both the denominator and the numerator of the Chinese import share.

5.2 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 on aggregate trade patterns and welfare, we consider the same shock to Chinese sourcing potential as in Section 5.1 but hold firms' extensive margin sourcing strategies fixed. In the exercise, we recompute the equilibrium price index and the mass of firms. As expected, when firms' sourcing strategies cannot change, the aggregate response to the China shock is substantially smaller. U.S. sourcing decreases by 0.28 percent (see Table 8), and aggregate sourcing from China increases by 115 percent (see Table 7, Panel B), with these responses being about half as large as those under flexible strategies. The implied growth in the Chinese share of US imports between 1997 and 2007 is also much lower than in our baseline model (98% versus 178%).

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 aggregate sourcing from both the U.S. and other countries. In addition, in the baseline analysis the increase in firms' extensive margin sourcing decisions results in higher expenditures on fixed costs. Given constant expenditure in the manufacturing sector, this leads *ceteris paribus* to lower expected profits, which in turn are brought back to zero by a decrease in the number of firms. In contrast, under fixed sourcing strategies, the number of firms remains the same in equilibrium (despite free entry), and consequently there is no decrease in U.S. sourcing from firms that shut down (see Table 8).³³ Firms that source from China slightly increase their sourcing from the U.S. and other countries, while all other firms reduce their sourcing from the U.S. and other countries (see Table 7, Panel B).

The responses of the size distribution are also quite different (see Figure 5b). 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. Under fixed sourcing strategies the price index declines by 0.15 percent, which is about 78 percent of the impact of the shock when firms can adjust their extensive margin decisions.³⁴

5.3 Comparison with Alternative Models and Sensitivity Analysis

As discussed in Section 2.4 when describing the model's predictions for aggregate trade flows, our framework nests two canonical frameworks of the last decade: the Melitz (2003) hetero-

³³The number of firms does not change since total expenditure on manufacturing is constant, and aggregate expenditure on fixed costs is unchanged (due to the lack of extensive margin adjustments), so that the free entry condition is satisfied with the same mass of firms as prior to the counterfactual.

³⁴Given fixed wages, free entry, and a constant share of income spent on manufacturing, the decline in the manufacturing price index captures the full effect of the shock on aggregate real income.

geneous 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. We also compare our baseline model’s predictions to those from a model in which fixed costs are heterogeneous across firms, but common across countries. For each of these alternative models, we re-calibrate the change in the Chinese sourcing potential to match the 178 percent increase in the Chinese share of US imports between 1997 and 2007.³⁵ In the Estimation Appendix, we assess the baseline model’s predictions under a range of alternative parameter values. An important caveat to the model comparisons and robustness analyses 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 first consider a scenario similar to the Chaney (2008) multi-country Melitz model 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-estimate fixed costs under these parameter values in Step 3, leading to a similar model fit as in the baseline. 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 fixed costs, although these are not present in the original Chaney (2008) paper.

The counterfactual predictions from this version of our model in which firms’ import entry decisions are independent across markets are quite different from the baseline results (despite the share of firms in each category being virtually identical in the two counterfactuals). 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), which 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: -0.3 percent (see Table 7, Panel C). Sourcing from third markets also declines. Most notably, the aggregate net effects are exactly equal to the gross effects under these parameter values. One might be concerned that this conclusion rests on our simplifying assumption that wages are unaffected by the China shock. Inspection of equation (12) – after plugging in the restriction $\theta = \sigma - 1$ – reveals, however, that changes in wages would necessarily lead *all* firms to change their sourcing from the U.S. by the same proportion. In sum, when entry decisions are independent,

³⁵In Online Appendix Tables C.6 and C.7, we report various targeted and untargeted moments for each of these variants of the model.

there is no scope for domestic churn and reallocation following the China shock.³⁶

We next 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. Under these conditions, all firms decrease their sourcing from the U.S. and from all other countries except China by 1.3 percent (see Table 7, Panel D). 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.

An important contribution of our paper is to present a new framework for analyzing extensive margin sourcing decisions when the fixed costs to import differ across countries. To assess the importance of this fixed cost heterogeneity, we revisit the case first described in section 4.4 in which fixed costs are identical across countries, but heterogeneous across firms. This model predicts a perfect pecking order of firm sourcing in which firms select into countries based solely on the potential marginal cost savings. Panel E of Table 7 shows that this prediction strongly limits firms' extensive margin sourcing responses. China is the number one sourcing potential country in 2007 and also 1997, even with the lower calibrated potential. As a result, the common fixed cost model predicts that firms that do not import from China will never find it profitable to import from any other country. This is evident in the Others row, which shows that non-China importers do not change their sourcing from other countries. In addition, the Entrants row also shows that new China importers do not change their third-market sourcing. Under the common fixed cost assumption, new China importers cannot have intensive margin changes in third market sourcing (since they did not import from anywhere in 1997), and the degree of complementarity we estimate is not powerful enough to induce them to add third markets. The latter stands in sharp contrast to our baseline model, where half of the change in Entrants' third-market sourcing was driven by extensive margin changes. In the next section, we will also show that Entrants' extensive margin changes are also key features in the actual US data.³⁷

A key takeaway from these counterfactual exercises is that the interdependencies in firms' extensive margin sourcing decisions lead to significant differences between the gross and net changes caused by the shock. Although the China shock is associated with a sizable decrease in domestic sourcing in the U.S., some firms increase their domestic input purchases considerably

³⁶In a Melitz-Chaney model of exporting, domestic sales of all surviving firms change by the same percentage; larger firms grow and small firms shrink only because the latter only sell domestically, while the former also export. Focusing on sales in a particular market (not just the domestic one), all firms scale up or down proportionally.

³⁷Another difference between Panel E and our baseline results in panel A, is that continuing China importers are predicted to shrink in both domestic and third markets. On the other hand, the common fixed costs model is successful in delivering gross increases and decreases in sourcing from U.S. suppliers, though to a lesser extent than the baseline model.

following the shock. More generally, the third-market effects – that is how a shock to one country (in this case China) affects sourcing from other countries – are quite different for both firm and aggregate outcomes in the presence of interdependencies in the extensive margin.

Before turning to a comparison of our counterfactual predictions to actual data, we test the sensitivity of our results to specific values of the key structural parameters. To do so, we re-estimate the fixed costs of sourcing using 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 summarize the main results here, with detailed results presented in Appendix Table A.3. The estimates for the effects of language and control of corruption 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 decreasing 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 dollar decrease in U.S. sourcing, a 38 cent 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.

5.4 Comparing Counterfactual Predictions to Actual Changes

We conclude this section by comparing the counterfactual predictions from the model to actual changes in U.S. firms' sourcing from 1997 to 2007. We analyze changes from 1997 to 2007 since data on firm sales and input purchases are available only in years ending in 2 and 7, and these years conveniently span China's accession to the WTO. Table 9 presents an analog to Table 7 in which we use the data, deflated to 1997 dollars, to calculate the relative changes in domestic sourcing, sourcing from other countries, and sourcing from China. We use an unbalanced panel of firms in this exercise to ensure that the observed patterns are not driven by selection of more productive firms that are more likely to survive over this period.³⁸

While the counterfactual predicts that firms will grow or shrink their sourcing by small amounts, Table 9 shows that in the data, all types of firms changed their sourcing significantly. This is not surprising since the counterfactual holds all else constant, while in reality many factors change over the ten-year period. Nevertheless, one can still compare relative changes in sourcing across Chinese import status to assess whether the qualitative predictions of the

³⁸This selection does not occur in the model since the core productivities of firms are unchanged by the counterfactual. The share of new China importers (Entrants) in the unbalanced panel is eight percent. This is slightly higher than the 6.6 percent entry rate mentioned earlier, which is the difference between the 1997 share of China importers (1.9 percent) and the 2007 share of imports (8.5 percent). These numbers do not perfectly coincide because, in reality, not all firms that import from China in 1997 continue to do so in 2007, and there some firms that enter and exit over the period also import from China.

model are evident in the data. Consistent with the counterfactual predictions, firms that begin sourcing from China grow both their domestic sourcing and their sourcing from other countries relatively more than continuing China importers and non-China importers. Table 9 also shows that continuing importers grow their domestic sourcing and sourcing from other countries more than non-importers. In fact, as in the counterfactual predictions, aggregate domestic sourcing by firms that never source from China shrinks over this period. The main qualitative difference from the model is that sourcing from third markets by non-China importers grows, but this growth is substantially smaller relative to foreign sourcing growth of both new and continuing China importers. It is interesting to note that the shares of imports from China accounted for by entrants and continuers are almost identical in the data and model.

Table 9: Observed changes in third country sourcing from 1997 - 2007 by firms' 2007 Chinese import status

Chinese import status	Change in sourcing from U.S.	Change in sourcing from other countries	Change in sourcing from China	Share of firms	Share of imports from China
Entrants	2.68	5.93	-	0.08	0.37
Exiters	0.15	0.02	0.00	0.00	0.00
Continuers	1.03	1.32	5.75	0.01	0.63
Others	0.74	1.03	-	0.91	0.00

Notes: This table is based on an unbalanced panel of manufacturing firms from 1997 and 2007. Exiters (entrants) are those firms that stop (begin) sourcing from China. Continuers are firms that source from China in 1997 and 2007. Columns 1, 2, and 3 contain the ratio of total sourcing in 2007 deflated dollars relative to 1997 dollars by each group of firms. The share of firms in column 4 is the number of 2007 firms in each China import status category over total active firms in 2007.

The fact that both the model and data show China importers growing their domestic and third market sourcing relatively more than non-China importers provides support for the empirical relevance of the interdependencies highlighted in the model. As shown in both Panels D and E of Table 9, we would not expect to see increased domestic and third market sourcing by U.S. firms sourcing from China if entry decisions were independent across markets, or if there were no fixed costs of sourcing. While the common fixed cost model would still lead to increased domestic and other sourcing by new China importers, it also makes the extreme prediction that all new importers would necessarily source from China –a prediction clearly at odds with the empirical evidence presented in Table 1.

A potential concern with this comparison to the data is firms may change their Chinese imports in response to positive (or negative) demand, productivity, or technology shocks over the period. To make a cleaner comparison of the counterfactual predictions of our model to the

observed evolution of sourcing in the data, we exploit the significant productivity growth within China and its accession to the WTO in 2001 to construct a firm-specific, plausibly exogenous shock to Chinese sourcing potential. In the spirit of Autor et al. (2013) and Hummels et al. (2014), we instrument for changes in U.S. firms' imports from China using changes in Chinese export shares to other developed countries in a firm's 1997 input industries.³⁹ We then estimate the relationship between predicted changes in sourcing from China and firm sourcing from domestic and third markets according to

$$\Delta y_n = \beta_0 + \beta_{Ch} \Delta China_n + \varepsilon_n, \quad (23)$$

where $\Delta China_n = \frac{Imports_n^{Ch2007} - Imports_n^{Ch1997}}{(Imports_n^{Ch2007} + Imports_n^{Ch1997})/2}$ is a Davis-Haltiwanger-Schuh (DHS) growth rate of firm n 's imports from China that captures changes in firms' intensive and extensive margin sourcing. We estimate the relationship between China sourcing and five dependent variables (all represented by a DHS growth rate unless otherwise noted): domestic input purchases, imports from other countries, the log difference in the number of countries from which the firm sources (excluding China but including the U.S.), and firm-level employment.⁴⁰ All dollar values are deflated using NBER industry deflators.

Table 10 presents results from estimating equation (23) via OLS and two-stage least squares for the balanced panel of manufacturing firms present in both 1997 and 2007.⁴¹ Standard errors are clustered at the industry level to match the same level of aggregation as the China shock variable. The OLS estimates presented in columns 1-4 are all significant at the one percent level and suggest that a ten percentage point increase in firm-level sourcing from China was associated with a 0.6 point increase in domestic sourcing, a 2.6 percent increase in the number of countries from which the firm sources, and a 3.6 point increase in sourcing from other foreign countries. Columns 5-7 show that IV estimates are larger and still significant

³⁹Unlike Autor et al. (2013) who use changes in the levels of Chinese imports across industries in other high-income countries as an instrument, we use changes in Chinese market shares in the original 15 European Union countries. We choose market shares rather than levels to address the potential for correlated demand shocks between those countries and the U.S., which are a more serious concern in our framework. All our results are also robust to using the same high-income countries as in ADH, but our first stage is stronger using EU15 countries, especially in the robustness test in which we instrument for both firm-level sourcing and import competition. Additional details are in the Estimation and Online Data Appendices.

⁴⁰We use DHS growth rates as a simple way to account for both intensive and extensive margin adjustments. Our results are qualitatively similar when using an indicator equal to one for firms that begin sourcing from China over the period. They are also similar when using log differences in firms' domestic input use and employment as outcome variables. We do not obtain significant results for the log difference in firms' third-market sourcing – a result consistent with the model's predictions about the importance of changes in firms' extensive margin sourcing decisions.

⁴¹Here we focus on a balanced panel since we require a 1997 firm industry to construct our instrument. The results are significantly stronger if we used an unbalanced panel and assigned shocks to new firms based on their initial industry. We do not present those results since they may be biased by firms' endogenous entry and exit decisions.

at the one percent level.⁴² Given the considerable focus in the literature on the employment effects of Chinese imports, we also estimate the relationship between changes in China sourcing and firm-level employment. While the OLS estimate shows that 10 point increase in the growth rate of Chinese imports is associated with a one point increase in the firm’s U.S. employment, the IV estimate is negative but statistically insignificant. This result is still consistent with our framework since the model predictions relate to changes in a firm’s total domestic input/employment use, and not to employment at the *firm* itself. The bottom panel of Table 10 also presents first-stage statistics, which show a positive and significant coefficient on the instrument and a Kleibergen-Paap F-statistic of 28.25.

Table 10: Estimates of the impact of the China shock on firm-level sourcing

Dependent variable is firm-level change from 1997 to 2007 in:

	Domestic inputs	No. of countries	Foreign inputs	Firm empl.	Domestic inputs	No. of countries	Foreign inputs	Firm empl.
	OLS				IV			
China, DHS	0.064 (0.010)	0.255 (0.007)	0.362 (0.013)	0.097 (0.007)	0.758 (0.214)	0.551 (0.080)	0.670 (0.198)	-0.092 (0.162)
Constant	0.054 (0.019)	0.144 (0.013)	0.315 (0.026)	-0.075 (0.014)	-0.054 (0.039)	0.098 (0.017)	0.267 (0.044)	-0.046 (0.032)
N	127,400	127,400	127,400	127,400	127,400	127,400	127,400	127,400
First Stage Statistics	Coeff (se) 2.685 (0.505)				KP F stat 28.25			

Notes: All variables are changes or growth rates from 1997 to 2007. China, DHS is a Davis-Haltiwanger-Schuh growth rate in firm imports from China. Domestic inputs, foreign inputs, and firm employment are a DHS growth rate. No. of countries is the log difference in the number of countries (excluding China, but including the U.S.) from which the firm sources inputs. Foreign inputs exclude China. Standard errors are in parentheses and clustered by 439 NAICS industries. In the IV specifications, firm-level sourcing from China is instrumented by the change in Chinese market share in EU15 countries of a weighted average of the firm’s inputs. KP F-stat is the Kleibergen Paap F-statistic. N is rounded for disclosure avoidance.

There are two potential concerns with this reduced form analysis. First, the exclusion restriction might be violated if industries in which China gained market share in EU15 countries are also industries that faced greater competition from China in final-good markets. Although we instrument with a shock to a firm’s inputs rather than outputs, input and output shocks may still be correlated. In the Online Data Appendix, we show that the results are robust to controlling for import penetration from China, and to instrumenting for import penetration from China with a shock to a firm’s output industries. Another concern with our approach is that the exclusion restriction could be violated because the multilateral nature of the China shock may have indirectly affected firm sourcing from third markets by affecting input suppliers in those markets.⁴³ Unfortunately there is no clean way to control for shocks to a firm’s

⁴²The difference in magnitude between our OLS and IV estimates are similar to those in Hummels et al. (2014), and consistent with offshoring being a firm-level response to negative shocks, as documented in Monarch et al. (2017) and Bernard et al. (2017).

⁴³To the extent that the China shock displaced input suppliers either in the U.S. or other foreign countries, these indirect effects might bias our estimates of β_{Ch} down (e.g. Hanson and Robertson, 2010). Alternatively,

suppliers in these data, but this could be a fruitful avenue to explore with richer information on supply-side firm-level networks.

The results in Tables 9 and 10 provide strong support for the empirical relevance of our theoretical mechanism. First, they show that firms sourcing from China increase their domestic and third market sourcing substantially more than firms that do not import from China. This is inconsistent with the predicted responses in a world with independent entry decisions since, as shown in Panel D of Table 7, under this scenario all firms *decrease* their domestic sourcing by the same amount. Second, they show that increased domestic and foreign sourcing occurs not just across, but also within firms. Third, they rule out firm-specific demand or productivity shocks as drivers of these changes. Most importantly, the results show that increased firm-level imports from China do not decrease domestic and third market sourcing –as might be expected in a world with no interdependencies in sourcing decisions– but instead are associated with increased firm-level sourcing from other markets.

6 Conclusion

This paper provides a new framework in which to analyze the global sourcing decisions of firms in a multi-country world where 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. These special cases highlight the fact that in a general setting, sourcing decisions interact through the cost function, so that determining the extensive margin of importing must involve 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.

A key contribution of this paper is to overcome these challenges by showing that –under a simple parametric restriction that is consistent with the data– a firm’s decision to source from one country is complementary to its decision to source from other countries. These complementarities in a firm’s extensive margin sourcing decisions underpin our new methodological approach to reduce the dimensionality of the firm’s problem and solve it. By extending the pioneering work of Jia (2008) to a multi-firm environment, we 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

the China shock may have led to wage decreases or productivity increases among suppliers that in turn led U.S. firms to increase sourcing from them. To the extent that these changes disproportionately affect industries in which a firm’s inputs were shocked (e.g., Torres-Ruiz and Utar, 2013), our estimates might be biased up.

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 the shocked country may expand, while firms not sourcing from that country shrink. We show that a ‘China shock’ calibrated to match the growth in U.S. foreign sourcing from China between 1997 and 2007 generates qualitative effects consistent with the observed evolution of U.S. firms’ domestic and third-market sourcing.

Our theoretical framework is necessarily stylized, but it can flexibly accommodate various extensions. As shown in the Online Appendix and further explored in Bernard et al. (2016), it can easily be extended to include the joint determination of the extensive margin of importing *and* exporting. It is also straightforward to incorporate fixed costs of sourcing costs at the input level (see the Online Appendix) and also at the supplier level. This latter approach is explored by Bernard et al. (2015) and Furusawa et al. (2016) in their studies of buyer-seller relationships in Japan. Similarly, we have abstracted from the type of contractual frictions inherent in global sourcing transactions, but as outlined in Antràs (2016), 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 in the presence of demand linkages across markets.

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* 2016. Princeton University Press.
- **and Elhanan Helpman**, “Global Sourcing,” *Journal of Political Economy*, 2004, *112* (3), 552–580.
- **and —**, “Contractual Frictions and Global Sourcing,” in “The Organization of Firms in a Global Economy,” Cambridge, MA: Harvard University Press, 2008.
- Arkolakis, Costas and Fabian Eckert**, “Combinatorial Discrete Choice,” Working Paper 2017.
- **, Arnaud Costinot, and Andrés Rodríguez-Clare**, “New Trade Models, Same Old Gains?,” *American Economic Review*, February 2012, *102* (1), 94–130.
- **, 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, David Dorn, Gordon H. Hanson, and Jae Song**, “Trade Adjustment: Worker-level Evidence,” *The Quarterly Journal of Economics*, 2014, *129*, 199–1860.
- 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**, “Factoryless Goods Producing Firms,” *American Economic Review Papers and Proceedings*, May 2015, *105* (5), 518–523.
- **, Andreas Moxnes, and Yukiko U. Saito**, “Production Networks, Geography and Firm Performance,” Working Paper 21082, National Bureau of Economic Research April 2015.
- 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.
- **, —, —, and —**, “Global Firms,” paper commissioned by the *Journal of Economic Literature*, 2016.
- Bernard, Andrew B., Teresa C. Fort, Valerie Smeets, and Frederic Warzynski**, “Offshoring and Reorganization,” Technical Report, Tuck School of Business 2017.

- 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.
- , **Claire LeLarge, and Michael Peters**, “The Gains from Input Trade With Heterogeneous Importers,” *American Economic Journal: Macro*, 2017.
- Broda, Christian and David E. Weinstein**, “Globalization and the Gains From Variety,” *The Quarterly Journal of Economics*, May 2006, *121* (2), 541–585.
- 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.
- De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik**, “Prices, Markups, and Trade Reform,” *Econometrica*, March 2016, *84* (2), 445–510.
- Eaton, Jonathan and Samuel Kortum**, “Technology, Geography, and Trade,” *Econometrica*, 2002, *70* (5), 1741–1779.
- , – , and **Francis Kramarz**, “An Anatomy of International Trade: Evidence From French Firms,” *Econometrica*, 2011, *79* (5), 1453–1498.
- 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,” *Review of Economic Studies*, 2017, *84* (2), 650–687.
- and **Shawn Klimek**, “The Effects of Industry Classification Changes on US Employment Composition,” Technical Report, Tuck School of Business 2016.
- Furusawa, Taiji, Tomohiko Inui, Keiko Ito, and Heiwai Tang**, “Global Sourcing and Domestic Production Networks,” Technical Report, mimeo, John Hopkins University 2016.
- 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*, March 2014, *104* (3), 793–831.
- Grossman, Gene M. and Esteban Rossi-Hansberg**, “Trading Tasks: A Simple Theory of Offshoring,” *American Economic Review*, 2008, *98* (5), 1978–97.

- 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,” *American Economic Review*, December 2015, *105* (12), 3660–3703.
- Hanson, Gordon H. and Raymond Robertson**, “China and the Manufacturing Exports of Other Developing Countries,” in Robert Feenstra and Shang Jin Wei, eds., *China’s Growing Role in World Trade*, University of Chicago Press and the NBER, 2010, pp. 137–159.
- , **Raymond J. Mataloni, and Matthew J. Slaughter**, “Vertical Production Networks in Multinational Firms,” *Review of Economics and Statistics*, 2005, *87* (4), 664–678.
- 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.
- Hummels, David, Jun Ishii, and Kei-Mu Yi**, “The Nature and Growth of Vertical Specialization in World Trade,” *Journal of International Economics*, 2001, *54*, 75–96.
- , **Rasmus Jorgensen, Jakob Munch, and Chong Xiang**, “The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data,” *American Economic Review*, 2014, *104*, 1597–1629.
- 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**, “A Portrait of Trade in Value Added over Four Decades,” *Review of Economics and Statistics*, 2016, *forthcoming*.
- Melitz, Marc J.**, “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 2003, *71* (6), 1695–1725.
- **and Stephen J. Redding**, “New Trade Models, New Welfare Implications,” *American Economic Review*, March 2015, *105* (3), 1105–1146.
- 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, Jooyoun Park, and Jagadeesh Sivadasan**, “Domestic Gains from Offshoring? Evidence from TAA-linked U.S. Microdata,” *Journal of International Economics*, 2017, *105*, 150–173.
- 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.
- Pierce, Justin R. and Peter K. Schott**, “A Concordance between Ten-Digit U.S. Harmonized System Codes and SIC/NAICS Product Classes and Industries,” *Journal of Economic and Social Measurement*, 2012, *37*, 61–96.
- **and** — , “The Surprisingly Swift Decline of U.S. Manufacturing Employment,” *American Economic Review*, 2016, *106* (7), 1632–1662.
- 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,” *The Quarterly Journal of Economics*, 2017, *132* (1), 157.
- Torres-Ruiz, Luis and Hale Utar**, “International Competition and Industrial Evolution: Evidence from the Impact of Chinese Competition on Mexican Maquiladoras,” *Journal of Development Economics*, 2013, *105*, 267–287.
- 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.

A Estimation Appendix

A.1 Measuring Firm-Level Offshoring Shares

We measure a firm’s total inputs using production worker wages from the Census of Manufactures (CM), total cost of materials from the Economic Censuses of manufacturing, construction, and mining, and merchandise purchases from the Census of Wholesale. Inputs from any foreign country j are simply the firm’s total imports from j . Domestic inputs are the difference between total inputs and imports. A firm’s share of inputs from country j , χ_{ij} , is computed as imports from j divided by total input purchases. Additional details are in described in the Online Data Appendix.

A.2 Estimation of the Trade Elasticity

Table A.1 presents the first stage regressions for the IV estimates of θ , where we instrument for country wages using population. As expected, the estimated coefficient on population is negative and significant at the five percent level. The F-statistics for the excluded instrument is 6.49. In unreported results (available in our online replication files), we verify that the Anderson-Rubin F test and χ^2 test of significance of the endogenous regressors are statistically significant at the five percent level.

Table A.1: First stage regressions for trade elasticity estimates

Dependent variable is log HC adj. wage	
log population	-0.32 (0.12)
log distance	-0.20 (0.14)
log R&D	0.32 (0.09)
log KL	0.14 (0.17)
Common language	-0.05 (0.17)
Control of corruption	0.22 (0.12)
log no. of firms	-0.04 (0.07)
Constant	-1.48 (0.81)
R^2	.883
Observations	57

Notes: First stage regressions for the IV estimates presented in Table 4. Wage is the log of the human capital-adjusted wage. Population is the excluded instrument.

A.3 Estimation Results and Counterfactual Predictions

In this Appendix, we provide three additional tables related to our structural estimation and counterfactuals. Table A.2 compares the share of imports and importers in the data (as in Table 1), in the baseline model, and in the model with common fixed costs across countries. Table A.3 reports our estimation and counterfactual results for various alternative parameter values. Table A.4 contains details on the performance of our application of Jia’s (2008) algorithm.

Table A.2: Share of imports and importers: data and models

	Share of Importers			Share of Imports		
	Data	Baseline Model	C.F.C. Model	Data	Baseline Model	C. F. C. Model
Canada	0.585	0.554	0.304	0.163	0.094	0.055
China	0.333	0.297	1.000	0.137	0.201	0.244
Germany	0.201	0.191	0.207	0.070	0.047	0.036
UK	0.178	0.134	0.117	0.034	0.027	0.019
Taiwan	0.163	0.158	0.356	0.019	0.066	0.067
Italy	0.132	0.059	0.159	0.015	0.025	0.026
Japan	0.124	0.133	0.207	0.126	0.041	0.036
Mexico	0.121	0.137	0.315	0.141	0.057	0.058
France	0.094	0.068	0.090	0.026	0.017	0.014
South Korea	0.087	0.081	0.191	0.023	0.032	0.032

Notes: This Table depicts for the 10 most popular importing countries the share of importers that buy from that country and the share of import volume.

A.4 China Shock Measure for Reduced Form Estimates

We instrument for changes in firm-level sourcing from China using changes in a novel measure of Chinese comparative advantage in a firm’s inputs. Specifically, we measure Chinese comparative advantage in the inputs of industry h and year t as

$$China_{ht}^{input} = \sum_{m \in h} s_{mh} \left(\frac{EU15imports_{mt}^{China}}{EU15imports_{mt}^{World/US}} \right),$$

where s_{mh} is the expenditure share of inputs from industry m in industry h from the 1997 BEA input-output table. The terms in parentheses are China’s share of imports in the original 15 European Union countries in industry m and year t , excluding imports from the U.S., measured using bilateral trade data from the UN Comtrade database.⁴⁴ We use Chinese market shares in these countries since it is unlikely that demand or supply shocks for U.S. firms will drive changes in their aggregate import shares. Since firms often span multiple industries, we assign changes in these input shocks as the weighted average of firm manufacturing sales across industries in 1997. All time-series variation in this shock is therefore driven by changes in Chinese market shares in EU15 countries in a firm’s inputs.

Our approach is similar to Hummels et al. (2014), who construct firm-level shocks to Danish importers’ sourcing decisions by focusing on transport cost shocks to the set of countries and products the firm imported in a pre-period. We do not use pre-period imports since it is critical for our instrument to identify firm-level extensive margin decisions to start importing from China. The spirit of our identification strategy is most similar to Autor et al. (2013), who instrument for Chinese imports per worker in the U.S. using Chinese exports to eight high-income countries. There are two important distinctions between our approaches. First, we construct shocks to inputs, while their focus is on shocks to final goods. Second, we use changes in Chinese market shares, rather than changes in the levels of Chinese exports. Using Chinese market shares helps to address a potential concern with their identification strategy—namely that correlated industry demand shocks in other high-income countries and the U.S. may have increased imports from China in both places. Unlike the Autor et

⁴⁴The EU15 countries are Austria, Belgium, Luxembourg, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain, Sweden, the U.K. In addition to excluding U.S. imports in the denominator, we also exclude all trade among EU15 countries in the calculation.

Table A.3: 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	$\theta = 1.3$	$\theta = 2$	$\theta = 2.85$	$\kappa = 3$	$\kappa = 4$	$\kappa = 5$	$\sigma = 2.79$	$\sigma = 5$
Parameter estimates									
B	0.122	0.121	0.123	0.124	0.103	0.118	0.132	0.237	0.072
β_c^f	0.023	0.030	0.020	0.014	0.040	0.024	0.020	0.017	0.039
β_d^f	0.192	0.209	0.188	0.187	0.414	0.224	0.118	0.002	0.421
β_l^f	0.871	0.939	0.868	0.872	0.778	0.856	0.901	0.944	0.768
β_C^f	-0.394	-0.434	-0.395	-0.394	-0.393	-0.389	-0.415	-0.468	-0.393
β_{disp}^f	0.939	0.935	0.932	0.927	1.312	0.989	0.819	0.633	1.318
<i>Median fixed cost range across countries</i>									
Minimum (in tsd USD)	10	14	9	7	19	11	9	5	19
Maximum (in tsd USD)	56	82	50	35	175	65	42	29	175
Counterfactual predictions									
Price index change (in percent)	-0.189	-0.268	-0.167	-0.113	-0.430	-0.215	-0.134	-0.072	-0.408
Change in third country sourcing by new importers from China (in percent)	1.5	3.6	0.9	-0.6	0.1	1.3	2.0	-0.3	1.5
Gross increase in U.S. sourcing (in percent of total U.S. sourcing)	0.113	0.345	0.053	0.000	0.011	0.090	0.141	0.000	0.096
Net change in U.S. sourcing (in percent of total U.S. sourcing)	-0.527	-0.554	-0.519	-0.497	-0.854	-0.571	-0.433	-0.307	-0.885

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 .

al. (2013) measure, our instrument does not capture any industry growth, but instead relies solely on changes in Chinese imports' relative importance in an industry. We use EU15 countries since we do not face the same data constraints as ADH who need trade data back to 1990. All our results are robust to using the ADH countries, but our first stage statistics are somewhat weaker, especially for the results in Online Appendix Table C.11, where we instrument for both firm-level Chinese imports and Chinese import penetration in a firm's industry.

As discussed in section 2, changes in market demand (B_i) may have a significant impact in industry equilibrium. A key component of market demand is the price index, which was significantly affected by changes in Chinese productivity. To the extent possible, we control for price changes by deflating firm inputs and imports from other countries using the NBER industry deflators. The first-stage estimates and additional details on the variables' construction are in the Online Data Appendix.

Table A.4: Cardinality of differences in bounds

Cardinality of differences in bounds		0	1	2	3	4	5	6	7	8	9	10 - 25	≥ 26
Variation in θ	Baseline	11220190126	0	918094	84695	6693	376	16	0	0	0	0	0
	$\theta = 1.3$	12359440019	0	3682595	600585	98164	15680	2474	444	26	12	1	0
	$\theta = 2$	12073844987	0	515622	36813	2342	233	3	0	0	0	0	0
	$\theta = 2.85$	12897360000	0	0	0	0	0	0	0	0	0	0	0
Variation in κ	$\kappa = 3$	12343521037	0	801134	72030	5290	477	25	5	2	0	0	0
	$\kappa = 4$	12786062392	0	1034458	94304	7843	902	99	2	0	0	0	0
	$\kappa = 5$	11792382315	0	1107735	101502	7671	719	49	9	0	0	0	0
Variation in σ	$\sigma = 2.79$	13145760000	0	0	0	0	0	0	0	0	0	0	0
	$\sigma = 5$	12304230446	0	2885568	476715	72738	12206	1947	353	27	0	0	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.