

TWO-SIDED HETEROGENEITY AND TRADE

Andrew B. Bernard, Andreas Moxnes, and Karen Helene Ulltveit-Moe*

Abstract—This paper develops a multicountry model of international trade that provides a simple microfoundation for buyer-seller relationships in trade. We explore a rich data set that identifies buyers and sellers in trade and establish a set of basic facts that guide the development of the theoretical model. We use predictions of the model to examine the role of buyer heterogeneity in a market for firm-level adjustments to trade shocks, as well as to quantitatively evaluate how firms' marginal costs depend on access to suppliers in foreign markets.

I. Introduction

GLOBAL trade is the sum of millions of transactions between individual buyers (importers) and sellers (exporters). Microlevel data have traditionally revealed exports of individual firms, summed across all buyers or, conversely, imports of individual firms, summed across all sellers. Naturally, theories of international trade have also focused on firms on either side of the market—exporters in Melitz (2003) or importers in Antràs, Fort, and Tintelnot (2017). In this paper, we explore the individual matches between exporters and importers and examine the consequences of this microstructure on firm-level and aggregate outcomes. In doing so, we build a model of international trade where exporters and importers are put on an equal footing.

We have access to a rich data set for Norwegian firms where the identities of both the exporter and the importer are known and where a firm's annual export transactions can be linked to specific buyers in every destination country, and each firm's annual import transactions can be linked to specific suppliers in every source country. This allows us to establish a set of basic facts about sellers and buyers across markets that guide the development of a parsimonious multicountry theoretical model with two-sided heterogeneity.

In the model, exporters vary in their efficiency in producing differentiated intermediate goods and pay a relation-specific fixed cost to match with each buyer. These fixed costs can be related to bureaucratic procedures, contract agreements, and the customization of output to the requirements of particular buyers. Importers bundle inputs into a

final product with heterogeneity in efficiency. Due to the presence of the relation-specific cost, not every exporter sells to every buyer in a market. Highly productive exporters reach many customers, and their marginal customer is small; highly productive importers purchase from many sellers, and their marginal supplier is small. This setup delivers parsimonious expressions for both upstream firms' exports and downstream firms' imports, which in equilibrium may differ because a seller can match to multiple buyers and a buyer can match to multiple suppliers. Buyer-seller matches are therefore entirely explained by selection based on heterogeneity and fixed costs. These represent the simplest possible ingredients of a model that are needed in order to explain broad features of the buyer-seller data.

Our theoretical modeling of the two-sided nature of trade brings several new insights. At the firm level, trade integration lowers marginal costs among downstream firms by reducing the cost of inputs and facilitating more matches between input suppliers and final goods producers. The importance of intermediate inputs for productivity growth has strong empirical support. Gopinath and Neiman (2014) find that a collapse in imports leads to a fall in productivity among Argentinian firms during the 2001–2002 crisis, while Amiti and Konings (2007), Goldberg et al. (2010), and Khandelwal and Topalova (2011) all find that declines in input tariffs are associated with sizable measured productivity gains. The model can generate firm-level responses to trade cost shocks that are consistent with the empirical evidence. Our work highlights that measured firm-level productivity gains not only arise from falling costs or access to higher-quality inputs but also from gaining access to new suppliers.

At the macrolevel, global trade will depend on the magnitude of relation-specific costs: lower relation-specific costs facilitate more matches between buyers and sellers, therefore generating more trade between nations as well as improving consumer welfare. In the aggregate, the model also retains the properties of one-sided models, as it gives us a simple gravity equation of bilateral trade flows as well as the same welfare results as in Arkolakis, Costinot, and Rodriguez-Clare (2012). In that sense, our model nests previous work while featuring a richer microfoundation.

We explore various empirical applications of the model starting with predictions for firm-level exports. According to the model, lower variable trade costs in a destination country will lead to higher firm-level export growth when buyers in that market are less dispersed in terms of their productivity. When buyers are more similar, an exporter will find many new profitable matches, whereas if buyers are dispersed, only a few more matches will become profitable. In other words, the customer-extensive margin response will be strong when buyer heterogeneity is low. We

Received for publication February 12, 2016. Revision accepted for publication May 1, 2017. Editor: Amit K. Khandelwal.

* Bernard: Tuck School of Business at Dartmouth, CEPR, and NBER; Andreas Moxnes: University of Oslo and CEPR; Karen Helene Ulltveit-Moe: University of Oslo and CEPR.

Thanks go to Richard Baldwin, Arnaud Costinot, Dave Donaldson, Adam Kleinbaum, Ben Mandel, Kjetil Storesletten, and Tony Venables, as well as seminar participants at Columbia, DINR, ERWIT 2013, Michigan, MIT, NBER, NY Fed, Princeton, and Stanford for helpful comments. We thank Angela Gu for excellent research assistance. A special thanks to the efforts of Statistics Norway for undertaking the identification of buyers and linking the transactions. A.M. is grateful for financial support from the Nelson A. Rockefeller Center for Public Policy and the Social Sciences at Dartmouth College and the Center for Global Business and Government at Tuck.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/rest_a_00721.

develop a theory-consistent sufficient statistic for unobservable trade costs and test this prediction based on our rich data set on buyers and sellers by exploiting variation in import shares across industries and countries over time. We find strong empirical support for the prediction from the model. An implication of our work is therefore that characteristics on the importer side (such as buyer heterogeneity) matter for firm-level adjustment dynamics. The firm-level export response after a change to trade policy, exchange rate movements, or other kinds of shocks will vary across countries depending on the characteristics of the importers.

Second, based on the predictions of the model, we develop an empirical methodology to evaluate downstream firms' marginal cost response when foreign market access is changing due to a fall in trade barriers or a reduction in the pool of potential suppliers. We show that a sufficient statistic for a firm's change in marginal costs is the level of, and the change in, intermediate import shares and the trade elasticity. Evaluating the impact of the 2008–2009 trade collapse on firms' production costs, we find that worsened market access during the trade collapse had a substantial negative impact on production costs, especially for downstream firms with high *ex ante* exposure to international markets. The empirical exercise also allows us to assess the fit of the model and evaluate the relative importance of the supplier margin. Overall, the model does well in explaining the fall in the number of buyer-seller connections during the trade collapse.

This paper is related to several new streams of research on firms in international trade. Importing firms have been the subject of work documenting their performance and characteristics. Bernard, Jensen, and Schott (2009), Castellani, Serti, and Tomasi (2010), and Muuls and Pisu (2009) show that the heterogeneity of importing firms rivals that of exporters for the United States, Italy, and Belgium, respectively. Amiti and Konings (2007), Halpern, Koren, and Szeidl (2011), and Bøler, Moxnes, and Ulltveit-Moe (2015) relate the importing activity of manufacturing firms to increases in productivity. In recent work, Blaum, Lelarge, and Peters (forthcoming) develop a model of firm-level imports and show, as we do, that a firm's marginal costs depend on the share of intermediates sourced domestically as well as the trade elasticity. They generalize this result and show that this holds for a wide class of models, while our framework emphasizes the two-sided nature of trade—that one firm's exports is another firm's imports.

Rauch (1999), Rauch and Watson (2004), Antràs and Costinot (2011), and Petropoulou (2011) consider exporter-importer linkages, while Chaney (2014) develops a search-based model of trade where firms must match with a contact in order to export to a destination. These studies adopt a search and matching approach to linking importers and exporters, while we abstract from these mechanisms and instead focus on the implications of buyer heterogeneity for trade.

Our work is also related to the literature on exports and heterogeneous trade costs initiated by Arkolakis (2010,

2011). In these papers, the exporter faces a rising marginal cost of reaching additional (homogeneous) customers. In our framework, buyers themselves are heterogeneous in their expenditures, but in equilibrium, exporting firms face rising costs per unit of exports as they reach smaller importers.

Our paper is most closely related to the nascent literature using matched importer-exporter data. Blum, Claro, and Horstmann (2010, 2012) examine characteristics of trade transactions for the exporter-importer pairs of Chile-Colombia and Argentina-Chile, while Eaton et al. (2014) consider exports of Colombian firms to specific importing firms in the United States. Blum et al. (2010, 2012) find, as we do, that small exporters typically sell to large importers, and small importers buy from large exporters. Their focus is on the role of import intermediaries in linking small exporters and small customers. Eaton et al. (2014) develop a model of search and learning to explain the dynamic pattern of entry and survival by Colombian exporters and to differentiate between the costs of finding new buyers and maintaining relationships with existing ones. Monarch (2013) estimates switching costs using a panel of U.S. importers and Chinese exporters, and Dragusanu (2014) explores how the matching process varies across the supply chain using U.S.-Indian data. Sugita, Teshima, and Siera (2014) study matching patterns in U.S.-Mexico trade, while Benguria (2014) estimates a trade model with search costs using matched French-Colombian data. In contrast to those papers but similar to Carballo, Ottaviano, and Martincus (2013), we focus on the role of importer heterogeneity across destinations. Carballo et al. (2013) focus on the distribution of export sales across buyers within a product-country, while we study the implications of importer heterogeneity on exporting firms' responses to exogenous shocks to trade barriers and the role of buyer-seller matches in the marginal cost of importers.

The rest of the paper is structured as follows. In section II, we document the main data set and present a set of facts on the role of buyers in trade, the heterogeneity of buyers and sellers, and their bilateral relationships. In section III, we develop a multicountry trade model with heterogeneous sellers and buyers, which is guided by the basic facts in section II. Section IV tests the predictions of the model with respect to the impact of trade cost shocks and the role of importer heterogeneity on firm-level performance and adjustment. Section V develops an empirical methodology to quantify the role of market access to intermediates and number of supplier connections in explaining the impact of a supply shock on downstream firms' marginal cost. Section VI concludes.

II. Exporters and Importers

A. Data

The main data set employed in this paper is based on Norwegian transaction-level customs data from 2004–2012. The

TABLE 1.—THE MARGINS OF TRADE, 2006

	Sellers	Products	Buyers	Density	Intensive
Exports (log)	0.57*** (0.02)	0.53*** (0.02)	0.61*** (0.02)	-1.05*** (0.04)	0.32*** (0.02)
<i>N</i>	205	205	205	205	205
<i>R</i> ²	0.86	0.85	0.81	0.81	0.50

Given that OLS is a linear estimator and its residuals have an expected value of 0, the coefficients for each set of regressions sum to unity, with each coefficient representing the share of overall variation in trade explained by the respective margin. Significance of the robust standard error in parentheses: ****p* < 0.01.

data have the usual features of transaction-level trade data in that it is possible to create annual flows of exports by product, destination, and year for all Norwegian exporters. In addition, these data have information on the identity of the buyer for every transaction in every destination market. As a result, we are able to see exports of each seller at the level of the buyer-product-destination-year.¹ Our data include the universe of Norwegian nonoil merchandise exports, and we observe export value and quantity. In 2005 total Norwegian nonoil merchandise exports amounted to US\$41 billion, equal to approximately 18% of mainland Norway GDP (GDP excluding the oil and gas sector). The firm-level evidence from Norwegian nonoil exports looks remarkably similar to that of other developed countries (see Cebeci et al., 2012; Irarrazabal, Moxnes, & Opmolla, 2013; and Mayer & Ottaviano, 2008). Table 5 in section A in the online appendix reports the top five exported products from mainland Norway.

B. Basic Facts

This section explores the matched buyer-seller data for Norwegian exporters. We establish the relevance of the buyer dimension as a margin of trade and document a set of facts on the heterogeneity of buyers and sellers and their relationships. We let these facts guide our model of international trade and subsequent empirical specifications.

Fact 1: The buyer margin explains a large fraction of the variation in aggregate trade. We decompose total exports to country *j*, x_j , into the product of the number of unique exporting firms, *f*; the number of exported products, *p*; the number of unique buyers (importers), *b*; the density of trade, *d*, the fraction of all possible exporter-product-buyer combinations for country *j* for which trade is positive; and the average value of exports, \bar{x} . Hence, $x_j = f_j p_j b_j d_j \bar{x}_j$, where $d_j = o_j / (f_j p_j b_j)$, o_j is the number of exporter-product-buyer observations for which trade with country *j* is positive and $\bar{x}_j = x_j / o_j$ is average value per exporter-product-buyer. We regress the logarithm of each component on the logarithm of total exports to a given market in 2006, for example, $\ln f_j$, against $\ln x_j$. The results, shown in table 1, confirm and extend previous findings on the importance of the extensive and intensive margins of trade. While it has been shown in

¹ Statistics Norway identifies buyers using the raw transaction-level records; however, they aggregate the data to the annual level before allowing external access to the data.

a variety of contexts that the numbers of exporting firms and exported products increase as total exports to a destination increase, our results show the comparable importance of the number of importing buyers in total exports. In fact, the buyer margin is as large as or larger than the firm or product margins.²

It is well documented that the total value of exports, the number of exporting firms, and the number of exported products are systematically related to destination market characteristics such as GDP and distance. Looking within the firm across markets, we study how the buyer margin responds to these gravity variables. We find that a firm's number of customers is significantly higher in larger markets and smaller in remote markets, that is, importers per exporter vary systematically with GDP and distance (see table 6 and figure 3 in the online appendix).³

Fact 2: The populations of both sellers and buyers of Norwegian exports are characterized by extreme concentration. The top 10% of exporters to an OECD country typically account for more than 90% of aggregate exports to that destination. At the same time, the top 10% of buyers from an OECD country are as dominant and also account for more than 90% of aggregate purchases (see table 7 in the online appendix).⁴ Although a handful of exporters and importers account for a large share of aggregate trade, these large firms are matching with many partners. One-to-one matches represent 9.5% of all exporter-importer connections but account for only 4.6% of aggregate trade (see table 8 in the online appendix). Many-to-many matches, those where both exporter and importer have multiple connections, make up almost two-thirds of aggregate trade. These facts motivate us to develop a model allowing for suppliers to match with several customers and buyers to match with multiple sellers. Using trade data for Chile and Colombia, as well as Argentina and Chile, Blum et al. (2012) similarly point to the dominance of large exporter-large importer matches among the total number of trading pairs. However, the theoretical model they develop fails to capture this feature of the data.

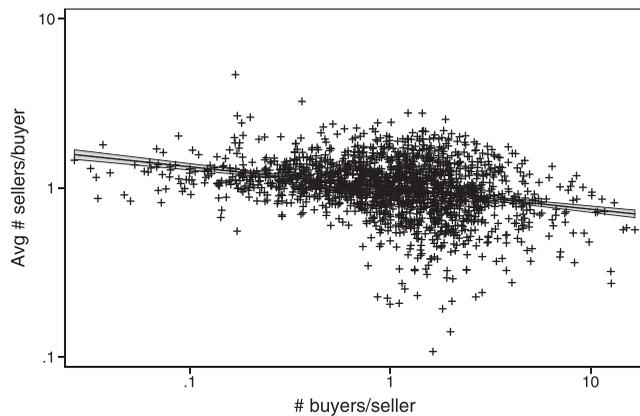
Fact 3: The distributions of buyers per exporter and exporters per buyer are characterized by many firms with few connections and a few firms with many connections. We find that the shapes of the distributions are remarkably similar across markets and illustrate this in figures 4 and 5 in the online appendix, plotting the results for China, the United States, and Sweden. The distributions appear to be largely consistent with a Pareto distribution as the CDFs are close to linear except in the tails. The Pareto fails to capture the

² Using data for Costa Rica, Uruguay, and Ecuador, Carballo et al. (2013) also find support for the role of the buyer margin in explaining the variation in trade. Their findings on the relative importance of buyers versus firms and products mirror our results.

³ The response of the buyer margin to gravity variables has been confirmed by Carballo et al. (2013).

⁴ This concentration of imports and exports in a small set of firms is similar to that found by Bernard et al. (2009) for the United States, Eaton et al. (2014) for Colombian exporters, and Mayer and Ottaviano (2008) for other European countries.

FIGURE 1.—MATCHING BUYERS AND SELLERS ACROSS MARKETS, 2006



The figure shows all possible values of the number of buyers per exporter in a given market, a_i , on the x-axis, and the average number of Norwegian connections among these buyers, $b_j(a_i)$, on the y-axis. Both variables are de-meant, and axes are in logs. The interpretation of a point with the coordinates (10, 0.1) is that an exporter with ten times more customers than is typical for that market has customers there with on average 1/10 the typical number of Norwegian suppliers.

discreteness of the actual empirical distribution (the number of customers per exporter is discrete), but we view the Pareto as a continuous approximation of the discrete case.⁵

Fact 4: Within a market, exporters with more customers have higher total sales, but the distribution of exports across customers does not vary systematically with the number of customers. Not surprisingly, firms with more buyers typically export more. The average firm with ten customers in a destination exports more than ten times as much as a firm with only one customer. However, firm-level exports to the median buyer are roughly constant, or even declining slightly, so that better-connected sellers are not selling more to their median buyer in a destination compared to less well-connected sellers. In our theoretical model, the variation in firm sales in a market is driven by the extensive margin of the number of customers.

Fact 5: There is negative degree assortativity among sellers and buyers. We characterize sellers according to their number of buyers and buyers according to their number of sellers. We find that the better connected a seller is, the less well connected is its average buyer. Figure 1 provides an overview of seller-buyer relationships. The slope of the fitted regression line is -0.13 , so a 10% increase in number of customers is associated with a 1.3% decline in average connections among the customers.⁶ In recent work by Bernard, Moxnes, and Saito (forthcoming), negative degree assortativity is also found for buyer-seller links among Japanese firms. Their

⁵These results are largely consistent with the findings by Blum et al. (2010, 2012) and Carballo et al. (2013).

⁶Using the median number of connections instead of the average number of connections as the dependent variable also generates a significant and negative slope coefficient. Estimating the relationship separately for each country, instead of pooling all countries, produces a negative assortativity coefficient for 89% of the countries we have sufficient data for (defined as countries with ten or more observations in the regression). In section B.8 in the online appendix, we show that the elasticity is informative of a structural parameter of the model.

Japanese data set covers close to the universe of domestic buyer-seller links and therefore contains information about the full set of buyer linkages. Our results are also in line with those found by Blum et al. (2010) in their analysis of Chilean-Argentinian trade.

Negative degree assortativity does not mean that well-connected exporters sell only to less-connected buyers; instead, it suggests that well-connected exporters typically sell to both well-connected buyers and less-connected buyers, whereas less-connected exporters typically only sell to well-connected buyers (see figure 8 in the online appendix for an illustration).⁷

Fact 6: Firms tend to follow a hierarchical pecking order in their choice of connections. We investigate the pervasiveness of buyer hierarchies following a procedure similar to Eaton, Kortum, and Kramarz (2011). First, we rank every buyer in a market according to the number of Norwegian connections of that buyer, r_b (country subscripts suppressed). The probability of connecting to a buyer, ρ_{r_b} , is b 's number of connections relative to the number of firms exporting to that market. Under independence, the probability of connecting only to the most connected buyer is $p_1 = \rho_1 \prod_{r=2}^B (1 - \rho_r)$, where B is the total number of buyers in the market. The probability of connecting only to the most and second-most connected buyers is $p_2 = \rho_1 \rho_2 \prod_{r=3}^B (1 - \rho_r)$, and so on. The likelihood of following the hierarchy under independence is therefore $\sum_{i=1}^B p_i$. We compare the likelihood of following this hierarchy under independence relative to what we find in the data for each country in our data set. For the vast majority of countries, there are more firms following the hierarchy relative to the statistical benchmark (the observations are above the 45-degree line). The data therefore refute the statistical benchmark of independence (see figure 9 in the online appendix for an illustration).

Section D in the online appendix presents a range of robustness checks for the facts presented. We also provide external validity by showing results for a different country, Colombia. We find that the basic facts also hold in the Colombian data. Finally, one may question if the basic facts

⁷Degree assortativity is a meaningful measure only in economic environments with many-to-many matching. Moreover, negative degree assortativity can coexist with positive assortative matching on the intensive (export value) margin. Sugita et al. (2014) study one-to-one matches in Mexico-U.S. trade and find evidence that more capable sellers typically match with more capable buyers; Dragusanu (2014) and Benguria (2014) also find evidence of positive assortativity on the intensive margin. In fact, this would also be the outcome of a one-to-one matching version of our model because the profits of a match are supermodular in seller and buyer efficiency (see section B.3 in the online appendix). Section B.8 in the online appendix provides additional evidence on intensive margin assortativity in our data. We note that social networks typically feature positive degree assortativity, that is, highly connected nodes tend to attach to other highly connected nodes, while negative correlations are usually found in technical networks such as servers on the Internet. In the friendship network among prison inmates considered by Jackson and Rogers (2007), the correlation between a node's in-degree (incoming connections) and the average in-degree of its neighbors is 0.58. The correlation in our data is -0.31 . Serrano and Boguna (2003) find evidence of negative assortativity in the network of trading countries—highly connected countries, in terms of trading partners, tend to attach to less connected countries.

we have presented can also be generated from a simple stochastic process where buyers and sellers meet randomly. We investigate this in the online appendix section C, where we simulate a balls and bins model of trade similar to Armenter and Koren (2014). The main finding is that a random model fails to explain key empirical characteristics of exporter-importer connections.

III. A Model with Two-Sided Heterogeneity

In this section, we develop a multicountry trade model that provides a microfoundation for buyer-seller relationships and allows us to examine the role of buyer heterogeneity and buyer-seller links for firm-level adjustments. As in Melitz (2003), firms (sellers) within narrowly defined industries produce with different efficiencies. We think of these firms as producers of intermediates as in Ethier (1979). Departing from Melitz (2003), we assume that intermediates are purchased by final goods producers (buyers or customers) that bundle inputs into final goods that in turn are sold to consumers. Final goods producers also produce with different efficiencies, giving rise to heterogeneity in their firm size, as well as a sorting pattern between sellers and buyers in equilibrium.

A. Setup

Each country i is endowed with L_i workers, and the labor market is characterized by perfect competition, so that wages are identical across sectors and workers. In each country, there are three sectors of production: a homogeneous goods sector characterized by perfect competition, a traded intermediates sector, and a nontraded final goods sector; the two last sectors are characterized by monopolistic competition. Workers are employed in the production of the homogeneous goods, as well as the production of the intermediates.⁸ The homogeneous good is freely traded and is produced under constant returns to scale with one hour of labor producing w_i units of the homogeneous good. Normalizing the price of this good to 1 sets the wage rate in country i to w_i .

Consumers. Consumers derive utility from consumption of the homogeneous good and a continuum of differentiated final goods. Specifically, upper-level utility is Cobb-Douglas between the homogeneous good and an aggregate differentiated good with a differentiated good expenditure share μ , and lower-level utility is a constant elasticity of substitution (CES) across differentiated final goods with an elasticity of substitution $\sigma > 1$.

Intermediates. Intermediates are produced using only labor by a continuum of firms, each producing one variety of the differentiated input. Firms are heterogeneous in productivity z , and firms' productivity is a random draw from a

⁸ Adding workers to the final goods sector would add more complexity to the model without generating new insights.

Pareto distribution with support $[z_L, \infty)$ and shape parameter $\gamma > \sigma - 1$, so that $F(z) = 1 - (z_L/z)^\gamma$. Lowercase symbols refer to intermediate producers, whereas uppercase symbols refer to final goods producers.

Final goods producers. Final goods are produced by a continuum of firms, each producing one variety of the final good. Their production technology is CES over all intermediate inputs available to them,

$$Z(v) \left(\int_{\Xi_j(v)} c(\omega)^{(\sigma-1)/\sigma} d\omega \right)^{\sigma/(\sigma-1)},$$

where productivity for firm v is denoted by $Z(v)$, which is drawn from the Pareto distribution $G(Z) = 1 - Z^{-\Gamma}$ with support $[1, \infty)$. $c(\omega)$ represents purchases of intermediate variety ω , and $\Xi_j(v)$ is the set of varieties available for firm v in country j . To simplify the notation, the elasticity of substitution among intermediates is identical to the elasticity of substitution among final goods, both denoted by σ . This restriction does not significantly affect the qualitative results of the paper. We also impose $\Gamma > \gamma$, which ensures that the price index for final goods is finite (see section B.2 in the online appendix).

Relationship-specific investments. Intermediate producers sell to an endogenous measure of final goods producers, and they incur a match-specific fixed cost for each buyer they choose to sell to. Hence, the act of meeting a buyer and setting up a supplier contract is associated with a cost that is not proportional to the value of the buyer-seller transaction. These costs may typically be related to bureaucratic procedures, contract agreements, and costs associated with sellers customizing their output to the requirements of particular buyers.⁹ Formally, we model this as a match-specific fixed cost, f_{ij} , paid by the seller in terms of labor, and it may vary according to seller country i and buyer country j . Consequently, buyer-seller links are the result of producers of intermediates that endogenously choose their set of customers.

The total mass of buyers and sellers, N_i and n_i , in each country i is proportional to total income Y_i , so there are more firms in larger economies. As there is no free entry, the production of intermediates and final goods leaves rents. We follow Chaney (2008) and assume that consumers in each country derive income not only from labor but also from the dividends of a global mutual fund. Each consumer owns w_i shares of the fund, and profits are redistributed to them in units of the numeraire good. Total worker income in country

⁹ Kang, Mahoney, and Tan (2009) provide examples of such relationship-specific investments and analyze under what circumstances firms are more likely to make these types of investments. For example, a newly adopted just-in-time (JIT) business model by Dell required that its suppliers prepare at least three month's buffering in stock. However, Dell did not offer any guarantee on purchasing volumes due to high uncertainty in final product markets.

i, Y_i , is then $w_i (1 + \psi) L_i$, where ψ is the dividend per share of the global mutual fund.¹⁰

Variable trade barriers. Intermediates are traded internationally, and firms face standard iceberg trade costs $\tau_{ij} \geq 1$, so that τ_{ij} must be shipped from country i in order for one unit to arrive in country j .¹¹

Sorting functions. Due to the presence of the match-specific fixed cost, a given seller in i will find it optimal to sell only to buyers in j with productivity higher than a lower bound \underline{Z}_{ij} . Hence, we introduce the equilibrium sorting function $\underline{Z}_{ij}(z)$, which is the lowest possible productivity level Z of a buyer in j that generates a profitable match for a seller in i with productivity z . We solve for $\underline{Z}_{ij}(z)$ in section III C. Symmetrically, we define $\underline{z}_{ij}(Z)$ as the lowest efficiency for a seller that generates a profitable match for a buyer in country j with productivity Z . By construction, $\underline{z}_{ij}(Z)$ is the inverse of $\underline{Z}_{ij}(z)$, that is, $Z = \underline{Z}_{ij}(\underline{z}_{ij}(Z))$.

Pricing. As intermediates and final goods markets are characterized by monopolistic competition, prices are a constant markup over marginal costs. For intermediate producers, this yields a pricing rule $p_{ij} = \bar{m} \tau_{ij} w_i / z$, where $\bar{m} \equiv \sigma / (\sigma - 1)$ is the markup.¹² For final goods, the pricing rule becomes $P_j = \bar{m} q_j(Z) / Z$, where $q_j(Z)$ is the ideal price index for intermediate inputs facing a final goods producer with productivity Z in market j . The restriction of identical elasticities of substitution across final and intermediate goods also implies that the markup \bar{m} is the same in both sectors. Using the Pareto assumption for seller productivity z , the price index on inputs facing a final goods producer with productivity Z can be written as

$$q_j(Z)^{1-\sigma} = \frac{\gamma z_L^\gamma}{\gamma_2} \sum_k n_k (\bar{m} \tau_{kj} w_k)^{1-\sigma} \underline{z}_{kj}(Z)^{-\gamma_2}, \quad (1)$$

where $\gamma_2 \equiv \gamma - (\sigma - 1)$.

Exports of intermediates. Given the production function of final goods producers specified above and conditional on a match (z, Z) , firm-level intermediate exports from country i to j are

¹⁰ Section B.9 in the online appendix develops an extension of the model where the number of buyers N_i is determined by a free entry condition. In that case, the number of buyers N_i is indeed proportional to country income Y_i . Introducing free entry on the seller side is more complex, as there is no closed-form solution for the number of sellers in a market n_i .

¹¹ We normalize $\tau_{ii} = 1$ and impose the common triangular inequality, $\tau_{ik} \leq \tau_{ij} \tau_{jk} \forall i, j, k$.

¹² Because marginal costs are constant, the optimization problem of finding the optimal price and the optimal measure of buyers simplifies to standard constant markup pricing and a separate problem of finding the optimal measure of buyers. We abstract from variable markups in our model, although in our data, unit values vary with respect to transaction size and destination. We leave this for future research.

$$r_{ij}(z, Z) = \left(\frac{p_{ij}(z)}{q_j(Z)} \right)^{1-\sigma} E_j(Z), \quad (2)$$

where $E_j(Z)$ is total spending on intermediates by a final goods producer with productivity Z in market j . The specific form of $E_j(Z)$ depends on the equilibrium sorting pattern in the economy (see section III C and the online appendix sections B.1 and B.2).

B. A Limiting Case

Because the lower support of the seller productivity distribution is z_L , a buyer (final goods producer) can potentially meet every seller (intermediate goods producer) in the economy. An implication is that we have two types of buyers: (a) buyers that match with a subset of the sellers and (b) buyers that match with every seller. Case (a) is characterized by $\underline{z}_{ij}(Z) > z_L$, while case (b) is characterized by $\underline{z}_{ij}(Z) \leq z_L$.

The discontinuity of the Pareto distribution at z_L is inconvenient, as the sorting function $\underline{z}_{ij}(Z)$ will be nonsmooth (not continuously differentiable) and important relationships will not have closed-form solutions. Henceforth, we choose to work with a particular limiting economy. Specifically, we let $z_L \rightarrow 0$, so that even the most productive buyer is not large enough to match with the smallest seller. In addition, we assume that the measure of sellers is an inverse function of the productivity lower bound, $n_i = z_L^{-\gamma} n'_i$, where n'_i is the normalized measure of sellers. Therefore, a lower productivity threshold is associated with more potential firms.¹³ When z_L declines, a given seller is more likely to have lower productivity, but there are also more sellers, so that the number of sellers in a given country with productivity z or higher remains constant. In equilibrium, the two forces exactly cancel out, so that the sorting patterns, as well as expressions for trade flows and other equilibrium objects, are well defined.

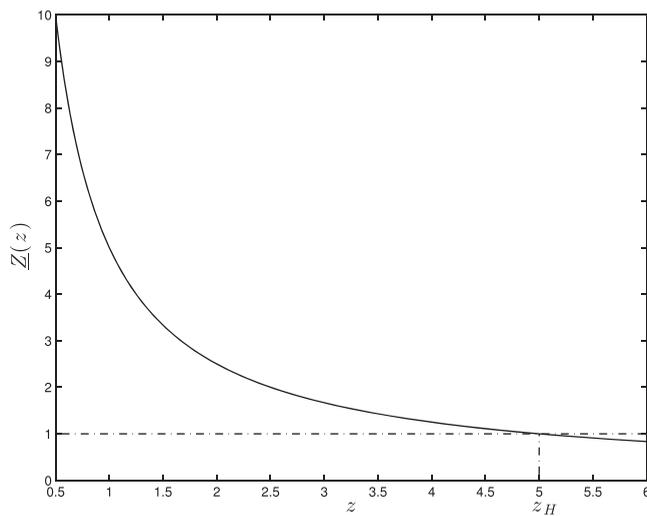
The support of the buyer distribution is $[1, \infty)$, which means that a highly productive seller can potentially meet every buyer in the market. This discontinuity is analytically tractable, so we allow for this to occur in equilibrium. We denote the productivity of the marginal seller that meets every buyer $z_H \equiv \underline{z}_{ij}$, equation (1). Hence, sellers with $z \geq z_H$ meet every buyer in the market.

C. Equilibrium Sorting

Based on the setup presented in section III A, we now pose the question: For a given seller of intermediates in country i , what is the optimal number of buyers to match with in market j ? An intermediate firm's net profits from a (z, Z) match is $\Pi_{ij}(z, Z) = r_{ij}(z, Z) / \sigma - w_i f_{ij}$. Given the optimal price from section III A, the matching problem of the firm is equivalent to determining $\underline{Z}_{ij}(z)$, the lowest-productivity buyer that generates a profitable match for a seller with productivity z . Hence, we find $\underline{Z}_{ij}(z)$ by solving for $\Pi_{ij}(z, Z) = 0$.

¹³ n'_i is constant as $z_L \rightarrow 0$. The normalization is similar to Oberfield (forthcoming).

FIGURE 2.—MATCHING FUNCTION



Inserting the demand equation (2) and a firm’s optimal price, we can express $\underline{Z}_{ij}(z)$ implicitly as

$$q_j(\underline{Z})^{\sigma-1} E_j(\underline{Z}) = \sigma w_i f_{ij} (\bar{m} \tau_{ij} w_i)^{\sigma-1} z^{1-\sigma}. \tag{3}$$

A complication is that the price index is also a function of the unknown $\underline{z}_{ij}(Z)$ and, furthermore, that total spending on intermediates, $E_j(Z)$, is unknown and depends on the equilibrium sorting pattern. In online appendix sections B.1 and B.2, we show that we can start with a guess of the functional forms for $\underline{z}_{ij}(Z)$ and $E_j(Z)$, derive the equilibrium, and then confirm that the functional forms are indeed valid. The solution to the sorting function is

$$\underline{Z}_{ij}(z) = \frac{\tau_{ij} w_i \Omega_j}{z} (w_i f_{ij})^{1/(\sigma-1)}, \tag{4}$$

where

$$\Omega_j = \left(\frac{\sigma}{\kappa_3} \frac{\gamma}{\gamma_2} \sum_k Y_k (\tau_{kj} w_k)^{-\gamma} (w_k f_{kj})^{-\gamma_2/(\sigma-1)} \right)^{1/\gamma}, \tag{5}$$

and κ_3 is a constant ($\kappa_3 \equiv \mu(\Gamma - \gamma)/\Gamma$). These expressions are valid under any distribution for buyer productivity; it is not necessary to assume Pareto-distributed buyer productivity to derive this particular result.

We plot the matching function $\underline{Z}_{ij}(z)$ in figure 2.¹⁴ $\underline{Z}_{ij}(z)$ is downward sloping in z , so more efficient sellers match with less efficient buyers on the margin. The point z_H on the horizontal axis denotes the cutoff productivity where a seller matches with every buyer. A firm with efficiency z matches with lower-efficiency buyers whenever variable or fixed trade costs (τ_{ij} and f_{ij}) are lower (the curve in figure 2

¹⁴The figure is based on parameter values $\tau_{ij} w_i \Omega_j (w_i f_{ij})^{1/(\sigma-1)} \times (Y_j/N_j)^{-1/\gamma} = 5$.

shifts toward the origin). Higher wages in country i mean that exporters (from i) cannot profitably match with lower-efficiency buyers. Conversely, higher GDP in the destination market, Y_j , increases the range of profitable matches.

The model is multicountry in that matching costs, variable trade costs, and wages in third countries affect the buyer cutoff between i and j . A firm from i matches with a greater range of (lower-efficiency) buyers in j when trade costs from third countries to j are higher (market access to j , Ω_j , is lower). This occurs because the downstream firms’ price index on inputs, $q_j(Z)$, is decreasing in market access (see equation [19] in online appendix B.1). Ω_j in equation (5) therefore has a similar interpretation as the multilateral resistance variables in Anderson and van Wincoop (2004) and Eaton and Kortum (2002). A key difference, although, is that Ω_j endogenizes the density of matching patterns through the fixed matching costs f_{ij} .

Highly productive downstream firms also will have a lower input price index— $q_j(Z)$ is decreasing in Z . Hence, all else equal, a given seller will face tougher competition when selling to a high-productivity buyer (which will, in equilibrium, have many suppliers).

D. Firm-Level Exports and Imports

Having determined the equilibrium sorting function between intermediate and final goods producers, we can now derive equilibrium expressions for firm-level exports and imports and decompose trade into the extensive margin in terms of number of buyers (suppliers) and the intensive margin in terms of sales per buyer (supplier).

Firm-level exports. Using equations (1), (2), and (4), for a given firm with productivity $z < z_H$, firm-level intermediate exports to market j are

$$r_{ij}^{TOT}(z) = \kappa_1 Y_j (w_i f_{ij})^{1-\Gamma/(\sigma-1)} \left(\frac{z}{\tau_{ij} w_i \Omega_j} \right)^\Gamma, \tag{6}$$

where κ_1 is a constant ($\kappa_1 \equiv \sigma\Gamma/[\Gamma - (\sigma - 1)]$).¹⁵ The $z > z_H$ case is in our context less interesting because the seller will match with every buyer, and the expression for firm-level trade therefore resembles the case with no buyer heterogeneity. It is also straightforward to determine marginal exports—exports to the least productive buyer. Using the fact that $\pi_{ij}(z, Z) = r_{ij}(z, Z)/\sigma - w_i f_{ij}$, we get

$$r_{ij}(z, \underline{Z}_{ij}(z)) = \sigma w_i f_{ij}. \tag{7}$$

Hence, marginal exports are entirely pinned down by the relation-specific fixed cost. We can also derive the optimal measure of buyers in an export market j for an upstream firm

¹⁵ See online appendix section B.3 for the corresponding expression for firms with $z \geq z_H$.

with productivity $z < z_H$ in country i (see online appendix B.3), which yields

$$b_{ij}(z) = Y_j (w_i f_{ij})^{-\Gamma/(\sigma-1)} \left(\frac{z}{\tau_{ij} w_i \Omega_j} \right)^\Gamma. \quad (8)$$

We emphasize two properties of these results. First, a firm will sell more in larger markets (higher Y_j), but the marginal export flow (i.e., a firm's transaction to the smallest buyer), is unaffected by market size because the marginal transaction is pinned down by the magnitude of the relation-specific fixed cost.¹⁶ Second, the elasticity of exports and the number of buyers with respect to variable trade barriers equals Γ , the shape parameter of the buyer productivity distribution. Hence, a lower degree of buyer heterogeneity (higher Γ) amplifies the negative impact of higher variable trade costs for both exports and the number of customers. This is in contrast to models with no buyer heterogeneity, where the firm-level trade elasticity is determined by the elasticity of substitution, σ (e.g., in Melitz, 2003). The intuition is that in markets with low heterogeneity (high Γ), there are many potential buyers that a seller can form profitable matches with after, for example, a decline in trade barriers. Consequently, trade liberalization in a destination market with low heterogeneity among importers translates into higher export growth than in a market with high heterogeneity among importers.

Proposition 1. *For $z < z_H$, the elasticity of firm-level exports with respect to variable trade costs equals Γ , the Pareto shape coefficient for buyer productivity.*

A potential concern is that this result is not robust to other distributional assumptions. Section B.5 in the online appendix derives general expressions for the firm-level trade elasticity given any distribution for buyer productivity. We show that the qualitative result that the elasticity is higher in markets with less buyer dispersion continues to hold for many commonly used distributions (log normal, exponential, and Fréchet).

Firm-level imports. The model also delivers parsimonious expressions for a downstream firm's intermediate imports as well as a firm's measure of suppliers. Section B.3 in the online appendix shows that intermediate imports from country i to a downstream firm in j are

$$R_{ij}^{TOT}(Z) = \kappa_4 Y_i (w_i f_{ij})^{1-\gamma/(\sigma-1)} \left(\frac{Z}{\tau_{ij} w_i \Omega_j} \right)^\gamma, \quad (9)$$

while the measure of suppliers is

$$S_{ij}(Z) = Y_i (w_i f_{ij})^{-\gamma/(\sigma-1)} \left(\frac{Z}{\tau_{ij} w_i \Omega_j} \right)^\gamma. \quad (10)$$

¹⁶ Also, a higher match cost f_{ij} dampens both firm exports and the number of buyers because $1 - \Gamma/(\sigma - 1) < 0$, given the previous restrictions that $\gamma - (\sigma - 1) > 0$ and $\Gamma > \gamma$.

At the firm level, an upstream firm's exports to country j , r_{ij}^{TOT} , are not identical to a downstream firm's imports from i , R_{ij}^{TOT} . At the aggregate level, of course, total export revenue must equal total import costs between i and j .

In the model, falling trade barriers or a greater number of potential suppliers lower marginal costs among downstream firms by reducing the cost of inputs and facilitating more matches between input and final goods producers. Specifically, as shown in section B.1 in the online appendix, the price index for intermediates for a downstream firm in j is given by

$$q_j(Z)^{1-\sigma} = Z^{\gamma_2} \frac{\bar{m}^{1-\sigma} \kappa_3}{\sigma} \Omega_j^{\sigma-1}, \quad (11)$$

that is, the marginal cost of a final goods producer in country j is inversely proportional to the market access term Ω_j . We summarize this in the following proposition:

Proposition 2. *A downstream firm's marginal costs are inversely proportional to the market access term Ω_j .*

This result follows directly from the sorting function described in equations (4) and (5). Hence, proposition 2 holds for any distribution of buyer productivity, not just Pareto.

The importance of intermediate inputs for productivity growth has strong empirical support. Gopinath and Neiman (2014) find a large productivity decline due to an input cost shock during the 2001–2002 Argentinian crisis, while Amiti and Konings (2007), Goldberg et al. (2010), and Khandelwal and Topalova (2011) all find that declines in input tariffs are associated with sizable measured productivity gains. Hence, the model generates firm-level responses to trade cost shocks that are consistent with the empirical evidence. Moreover, our theoretical results show that measured productivity gains can arise not only from falling costs or access to higher-quality inputs but also from being able to connect to new suppliers.

While our model is able to explain a range of new facts at the microlevel, it produces results on aggregate trade and welfare similar to models with one-sided heterogeneity (see online appendix section B.4 for details and for discussion of the difference between trade elasticities at the firm and aggregate level).

E. Linking Facts and Theory

We have pointed out that our theory is guided by the basic facts on buyer-seller relationships presented in section IIB. This section revisits the basic facts and examines the extent to which the model fits them.

According to fact 1 and table 6 in the online appendix, a firm's number of customers is increasing in GDP and decreasing in distance. As displayed in equation (8), according to the model, the number of buyers per firm increases

with market size and falls with trade costs, with elasticities 1 and $-\Gamma$, respectively.

The distribution of firm-level exports $r_{ij}^{TOT}(z)$, imports $R_{ij}^{TOT}(Z)$, the number of customers per exporter $b_{ij}(z)$ and the number of exporters per customer $S_{ij}(Z)$ are all Pareto, broadly consistent with facts 2 and 3.¹⁷

Fact 4 states that while total firm-level exports are increasing in the number of customers, the distribution of exports across buyers is roughly invariant to the firm's number of customers. In our model, the within-firm sales distribution is (see online appendix section B.7)

$$\Pr[r_{ij} < r_0 | z] = 1 - \left(\frac{\sigma w_i f_{ij}}{r_0} \right)^{\Gamma/(\sigma-1)},$$

so that all exporters to a market j have the same Pareto distribution of sales across buyers.

According to fact 5, highly connected exporters to market j have, on average, customers with few connections to Norwegian exporters. In the model, among exporters from i with b_{ij} customers in j , the average number of connections in i among these customers is (see online appendix section B.8):

$$\tilde{S}_{ij}(b_{ij}) = \frac{\Gamma}{\Gamma - \gamma} \left(\frac{b_{ij}}{b_{ij}(1)} \right)^{-\gamma/\Gamma}.$$

Hence, the elasticity is negative with a slope coefficient $-\gamma/\Gamma$.

IV. Firm-Level Adjustment to Trade Shocks

Proposition 1 states that the firm-level trade elasticity with respect to variable trade barriers is higher when importer productivity is less dispersed. In this section, we aim to test this main prediction of the model, that is, the role of buyer heterogeneity.

A sufficient statistic. An empirical challenge is that we do not directly observe either variable trade barriers τ_{ij} or the market access term Ω_j . We solve this by obtaining a sufficient statistic based on the predictions of the model. We proceed as follows. From equation (27) in the online appendix, we know that the aggregate trade share is

$$\pi_{ij} = Y_i (w_i f_{ij})^{1-\gamma/(\sigma-1)} (\Omega_j \tau_{ij} w_i)^{-\gamma}.$$

Solving this for $\Omega_j \tau_{ij} w_i$ and inserting it back into the expression for firm-level exports in equation (6) gives us

$$r_{ij}^{TOT}(z) = \kappa_1 Y_j Y_i^{-\Gamma/\gamma} (w_i f_{ij})^{1-\Gamma/\gamma} \pi_{ij}^{\Gamma/\gamma} z^\Gamma. \quad (12)$$

Hence, the unobserved variable trade cost and market access terms are replaced by the observable trade share π_{ij} .

¹⁷ The distributions of buyers per seller and sellers per buyer in the model are exactly Pareto, while those in the data approximate a Pareto except in the tails. Adding random matching to the model would allow the theoretical CDFs to more closely align with the empirical CDFs.

Empirical specification. We take the logs of equation (12); add subscripts m , k , and t to denote firm, industry, and year, respectively, and remove subscript i as Norway is always the source country in our data. Furthermore, we add a subscript j to the importer heterogeneity term Γ , as we want to use differences in importer heterogeneity as a source of identification. This gives us

$$\ln x_{mjkt} = \alpha_{mj} + \delta_{jt} + \ln Y_{jkt} + \frac{\Gamma_j}{\gamma} \ln \pi_{jkt},$$

where $\ln x_{mjkt}$ denotes a firm-level export variable, α_{mj} is a firm-country fixed effect, that captures time-invariant firm-country-specific factors such as idiosyncratic demand across destinations, and δ_{jt} is a destination-year fixed effect that captures time-varying country-wide shocks such as the real exchange rate or changes in relation-specific costs. We choose to work with empirical specifications exploiting industry-level variation (subscript k) because this allows us to include country-year fixed effects. This is potentially important because those fixed effects will absorb various factors that may be correlated with the trade shares π_{jkt} . In the robustness section below, we also experiment with other combinations of fixed effects.

We do not have sufficient variation in the Norwegian data to estimate every single measure of buyer dispersion Γ_j across markets. Instead we choose to calculate Γ_j using an international cross-country database (see the next section) on firm size distribution. Specifically, we estimate

$$\ln x_{mjkt} = \alpha_{mj} + \delta_{jt} + \beta_1 \ln Y_{jkt} + \beta_2 \ln \pi_{jkt} + \beta_3 \ln \pi_{jkt} \times \Gamma_j + \epsilon_{mjkt}, \quad (13)$$

where we have added an error term ϵ_{mjkt} . The error term is likely to include measurement error in log exports and buyer dispersion, as well as remaining unobserved factors that also determine firm-level exports, such as idiosyncratic trade frictions or demand shocks that are not captured by the industry-level π_{jkt} and Y_{jkt} .¹⁸ Because $\partial \ln x_{mjkt} / \partial \ln \pi_{jkt} = \beta_2 + \Gamma_j \beta_3$, the prediction of our model is that $\beta_3 > 0$, so that the elasticity is higher in markets with less importer dispersion.

Instrumental variable approach. A concern is that changes in the trade shares π_{jkt} are endogenous. For example, high productivity growth among one or several Norwegian firms could increase Norway's total market share, creating a causal relationship from firm-level export growth to the aggregate trade share. We deal with this by using the remaining Nordic countries' (Denmark, Finland, Iceland,

¹⁸ Also, note that the inclusion of firm-country fixed effects means that identification is coming only from within firm-destination changes in the variables. This implies that firms that export to a destination j only in one period (singletons) are dropped from the estimating sample. Our identifying assumption is therefore that the impact of trade barriers among singletons is identical to the impact of trade barriers among continuing exporters.

and Sweden) trade share, $\pi_{Nordic,jkt}$, as an instrument for Norway's trade share, π_{jkt} . Because of geographical and cultural proximity, as well as substantial economic integration among the Nordic countries, their trade shares are highly positively correlated (see section IIA). The exclusion restriction is that changes in the Nordic market share do not have a direct impact on Norwegian firm-level exports. Although we cannot completely rule this out, we find it unlikely because the Nordic market shares are typically very small in other countries (see section IVA). We estimate the model by 2SLS using $\ln \pi_{Nordic,jkt}$ and $\ln \pi_{Nordic,jkt} \times \Gamma_j$ as instruments for $\ln \pi_{jkt}$ and $\ln \pi_{jkt} \times \Gamma_j$, respectively.

Identification. Identification comes from comparing within firm-destination export growth across industries and firms, while controlling for country-specific trends. Our approach resembles a triple-differences model as we compare growth in exports across both industries and firms. Specifically, for two firms A and B and two industries 1 and 2, the β 's are identified by firm A's exports growth in country-industry jk relative to (a) its own export growth in industry 2 and (b) other firms' export growth in industry 2.¹⁹

A. Measures of Dispersion

To test our hypothesis, we require data on the degree of firm heterogeneity among importers located in different countries. Ideally, in line with our theoretical model, we would want a measure of buyer productivity dispersion in different markets. A close proxy for this is a measure of dispersion in firm size.²⁰ We therefore use data on the firm size distribution in different countries to calculate two measures of dispersion; a Pareto slope coefficient (Γ_j^1) and the standard deviation of log employment (Γ_j^2).²¹

Our preferred data source is Bureau van Dijk's Orbis database. Orbis has information on over 100 million private companies across the world.²² Orbis does not cover all firms and, especially among smaller firms, sampling may vary across countries. We therefore calculate dispersion based on the population of firms with more than fifty employees.²³ We calculate our two measures of dispersion for all countries with 1,000 or more Orbis firms. In total, this gives us information on buyer dispersion in 48 countries, covering 89% of Norwegian exports (based on 2006 values). There

¹⁹ The fixed effects α_{mn} and δ_{jt} are differenced out for $\Delta \ln y_{mjk't} - \Delta \ln y_{mjk't}$ and $\Delta \ln y_{mjk't} - \Delta \ln y_{m'jk't}$ where $k' \neq k$ and $m' \neq m$.

²⁰ The relationship between productivity and size has also been documented in a set of studies for many countries (see Bartelsman, Scarpetta, and Schivardi, 2013 for recent evidence). Helpman, Melitz, and Yeaple (2004) also use the firm size distribution as a proxy for firm-level heterogeneity.

²¹ We regress the empirical $1 - CDF$ on firm employment, both in logs, for each destination market. The resulting slope coefficient is the negative of the Pareto slope coefficient.

²² See <http://www.bvdinfo.com/Products/Company-Information/International/ORBIS.aspx> and Alfaro and Chen (2013) for a thorough discussion of the coverage of the database.

²³ Varying this size threshold has a negligible effect on our estimates of dispersion.

is substantial variation across countries. Dispersion in Russia is much lower than dispersion in Germany and Sweden. Also, the standard errors associated with the Pareto coefficient estimates are typically very small, suggesting that the Pareto distribution fits the empirical firm-size distribution quite well.²⁴

B. Construction of Variables

Our sufficient statistic approach requires data on Norway's trade share, π_{jkt} , and the Nordic countries' trade share, $\pi_{Nordic,jkt}$. Moreover, we need data on country income Y_{jkt} . Within the context of the theoretical model, the correct proxy for Y_{jkt} is absorption. Hence, we construct Y_{jkt} as output minus exports plus imports from UNIDO's Industrial Demand-Supply Balance Database (IDSB), which provides nominal output, total imports, and exports at the four-digit level of ISIC revision 3, for 127 manufacturing sectors and 121 countries over the sample period 2004 to 2012. In addition, our approach requires bilateral trade data by ISIC sector. We convert six-digit Harmonized System bilateral trade data to ISIC revision 3 by utilizing a concordance from the World Bank.²⁵ The trade shares are then calculated as $\pi_{jkt} = X_{NOjkt}/Y_{jkt}$ and $\pi_{Nordic,jkt} = X_{Nordic,jkt}/Y_{jkt}$ where X_{NOjkt} and $X_{Nordic,jkt}$ are trade from Norway and the remaining Nordic countries, respectively. The mean (median) trade share of Norway in 2004 was 0.21% (0.004%). There is a strong positive correlation between π_{jkt} and $\pi_{Nordic,jkt}$ in the data. Figure 11 in the online appendix shows a local polynomial regression of π_{jkt} on $\pi_{Nordic,jkt}$ (in logs) in 2004, where the market shares are measured relative to the mean log market share in country j .²⁶ We observe that even after controlling for the overall market share of Norway in country j , there is a positive and significant relationship.

C. Results

The 2SLS results from estimating the specification in equation (13) are shown in table 2. Columns 1 and 2 use total firm-level exports as the dependent variable, while columns 3 and 4 use the firm-level number of buyers (both in logs). The first two columns use the Pareto coefficient as the measure of firm-level heterogeneity, and the two last columns use the standard deviation of log employment.

We find that the export elasticity is significantly dampened in markets with more heterogeneity, consistent with the predictions of our model. The elasticity for the number of buyers is also consistent with the model, although the magnitude of the estimate is smaller than for the export elasticity. The coefficients for the interaction term are positive rather

²⁴ Figure 10 in the online appendix shows the resulting Pareto coefficients. Results on standard errors are available on request.

²⁵ Specifically, we use the COMTRADE/BACI trade database from CEPII and the WITS concordance from http://wits.worldbank.org/product_concordance.html.

²⁶ This is identical to including country fixed effects in a regression. The correlation is similar in other years.

TABLE 2.—MARKET ACCESS AND HETEROGENEITY: 2SLS ESTIMATES

	Exports (1)		Number of Buyers (2)		Exports (3)		Number of Buyers (4)	
Y_{jkt}	.18***	(.01)	.05***	(.00)	.18***	(.01)	.05***	(.00)
π_{jkt}	.30***	(.01)	.07***	(.00)	.33***	(.01)	.08***	(.00)
$\pi_{jkt} \times \Gamma_j^1$ (Pareto)	.07***	(.01)	.01**	(.00)				
$\pi_{jkt} \times \Gamma_j^2$ (SD)					-.10***	(.01)	-.01***	(.00)
Firm-country FE	Yes		Yes		Yes		Yes	
Country-year FE	Yes		Yes		Yes		Yes	
N	264,544		264,544		264,544		264,544	

Robust standard errors clustered by firm in parentheses: *** $p < 0.01$, ** $p < 0.05$. All variables in logs. Y_{jkt} and π_{jkt} are absorption and Norwegian market share in country-industry jk , respectively. π_{jkt} and $\pi_{jkt} \times \Gamma_j^1$ are instrumented with $\pi_{Nordic,jkt}$ and $\pi_{Nordic,jkt} \times \Gamma_j^1$ respectively, where $\pi_{Nordic,jkt}$ is the Nordic (excluding Norway) market share in country-industry jk .

TABLE 3.—2SLS ESTIMATES: ROBUSTNESS I

	(1)	(2)	(3)	(4)	(5)	(6)
	Exports	Number of Buyers	Exports	Number of Buyers	Marginal Buyer	Median Buyer
Y_{jt}	.18*** (.01)	.05*** (.00)	.24*** (.01)	.08*** (.00)	.02 (.02)	.08*** (.02)
π_{jt}	.27*** (.01)	.07*** (.00)	.33*** (.02)	.11*** (.00)	.00 (.04)	.12*** (.04)
$\pi_{jkt} \times \Gamma_j^1$ (Pareto)	.03*** (.01)	.00 (.00)	.03*** (.01)	.00 (.00)	.05 (.03)	.10*** (.03)
GDP/capita interaction	Yes	Yes	No	No	No	No
Firm-country FE	Yes	Yes	No	No	Yes	Yes
Country-year FE	Yes	Yes	No	No	Yes	Yes
Firm-country-year FE	No	No	Yes	Yes	No	No
2-digit industry FE	No	No	Yes	Yes	No	No
N	264,544	264,544	264,544	264,544	14,551	14,551

Robust standard errors in parentheses, clustered by firm: *** $p < 0.01$. All variables in logs. Y_{jkt} and π_{jkt} are absorption and Norwegian market share in country-industry jk , respectively. Γ_j^1 is the Pareto shape parameter. π_{jkt} and $\pi_{jkt} \times \Gamma_j^1$ are instrumented with $\pi_{Nordic,jkt}$ and $\pi_{Nordic,jkt} \times \Gamma_j^1$ respectively, where $\pi_{Nordic,jkt}$ is the Nordic (excluding Norway) market share in country-industry jk . Only exporters with > 5 buyers in columns (5) and (6).

than negative in columns 1 and 2 since the Pareto coefficient is inversely related to dispersion. The magnitudes are also economically significant: increasing the Pareto coefficient by 1 standard deviation raises the elasticity, $\beta_2 + \beta_3 \Gamma_j$, by 33%, suggesting that firm heterogeneity is quantitatively important for our understanding of firm-level trade adjustment.²⁷

We report OLS and first-stage results in table 10 in the online appendix. The OLS estimates are close to the IV estimates. The first-stage results confirm that the market shares among other Nordic countries are strongly associated with Norway's market share in country j (see also figure 11 in the online appendix).

The model predicts that the trade elasticity of exports to variable trade barriers is identical to the elasticity of the number of customers to variable trade barriers (see equations [6] and [8]), while the empirical results show that the export elasticity is stronger than the customer elasticity. One possible explanation for these discrepancies is that we are testing the predictions of the model using within-firm changes in a market over time, while the model is about cross-firm variation in a market at a point in time. Actual matching costs may have both sunk and fixed components.

²⁷ Γ_j is normalized with mean 0 and standard deviation 1; hence an increase of 1 standard deviation increases the elasticity from β_2 to $\beta_2 + \beta_3$. Inserting the numbers from the table, we get $(\beta_2 + \beta_3) / \beta_2 \approx 4/3$.

Robustness. A potential concern is that buyer dispersion may be correlated with other factors that also affect the trade elasticity; for example, both buyer dispersion and trade elasticities may be different in low-income countries. We address this issue by adding an interaction term between GDP per capita and the trade share π_{jkt} . The 2SLS results are reported in columns 1 and 2 in table 3. Overall the results are relatively close to the baseline case in table 2.²⁸

We also experiment with a different set of fixed effects. In columns 3 and 4 in table 3, we replace the firm-country and destination-year fixed effects with firm-destination-year and two-digit ISIC industry fixed effects, essentially exploiting variation only within a single firm-destination pair, across various sectors. This reduces the magnitude of the interaction term somewhat, but it is still significant and positive.

The exclusion restriction would no longer hold if, for example, firm-level exports from Norway cause a change in the Nordic trade share $\pi_{Nordic,jkt}$. We therefore introduce an alternative instrument to test the sensitivity of our results. Specifically, we calculate the Norwegian import share, defined as imports in industry k from country j at time t relative to total imports in that industry, $\pi_{NO,jkt}^{Imports}$. We then use $\ln \pi_{NO,jkt}^{Imports}$ and $\ln \pi_{NO,jkt}^{Imports} \times \Gamma_j$ as instruments for $\ln \pi_{jkt}$ and

²⁸ We also report the correlation matrix for the Pareto coefficient and other variables in table 9 in the online appendix.

In $\pi_{jkt} \times \Gamma_j$, respectively.²⁹ In our data, $\pi_{NO,jkt}^{Imports}$ and $\pi_{NO,jkt}$ are positively correlated, perhaps because trade barriers are to some extent symmetric. The exclusion restriction in this case is that a change in the import share from a given country does not directly cause a change in firm-level exports (other than through the impact through $\pi_{NO,jkt}$). Columns 1 and 2 of table 11 in the online appendix show that this alternative instrument gives us qualitatively the same results, although the coefficient magnitudes differ somewhat compared to the baseline.

A possible concern with this alternative instrument is that firms may be importing inputs from industry-country pair jk , which may have a direct impact on exports to the same industry-country pair. A simple way to check for this is to count how often firm-level exports to jk are associated with firm-level imports from jk . We find that across all countries and ISIC industries, the average share of firms both exporting to and importing from jk relative to the number of firms exporting to jk is 5.2% (and the median is 1.0%). We therefore conclude that this is a relatively minor concern in our data set.

Another concern is that our measure of dispersion from Orbis may not completely capture the degree of heterogeneity among buyers in that country. We deal with this by introducing an alternative measure of dispersion from the Exporter Dynamics Database (EDD) from the World Bank (year 2006). For each exporting country and year, EDD contains information about the coefficient of variation of log exports across the full population of exporters in a country. It is well known that the propensity to import is much higher for exporters than nonexporters (Bernard et al., 2007). Therefore, dispersion in exports from EDD may be more closely related to our Γ_j compared to the baseline measure of dispersion from Orbis. A caveat with EDD is that the sample of countries is smaller than with Orbis (39 versus 48 countries). Columns 3 and 4 of table 11 in the online appendix show the results. The interaction term for the coefficient of variation is negative and significant, indicating that more dispersion in a market lowers the elasticity.

As an additional robustness check, we test a second prediction from the model: that a firm's exports to her marginal (smallest) buyer are unaffected by both market size and trade costs. Exports to the marginal buyer are pinned down by the magnitude of the relation-specific cost. We estimate equation (13) by 2SLS using the value of a firm's marginal export and exports to the firm's median buyer as dependent variables. According to our theory, the coefficients for absorption Y_{jkt} and market access π_{jkt} should be 0 when the dependent variable is exports to the marginal or median buyer. The results largely confirm the predictions from the model (see table 3). The marginal export flow appears to be unrelated to market

size and access. However, exports to the median buyer are increasing in market size and market access.³⁰

Finally, in our model, Γ_j is time invariant, whereas one may speculate that our proxies for dispersion may change over time, and perhaps even more so if trade openness π_{jkt} is changing. Again using EDD data, we can test for this by calculating the CV year by year. Regressing CV_{jt} on CV_{jt-1} gives us a coefficient of 0.95 (SE 0.02), showing a very high degree of persistence in dispersion over time.

In sum, we confirm one of the main predictions of the model: improvement in market access results in higher export growth to countries where firms are less heterogeneous.

V. The Role of Supply Shocks

The empirical analysis focused on the prediction of our model with respect to the role of buyer heterogeneity. In this section, we turn to the other central feature of the model: the link between downstream firms' foreign market access, their number of suppliers, and marginal costs. According to the model, a downstream firm's production costs and measure of suppliers will depend on its access to foreign markets—the number of potential suppliers there as well as trade costs from supplier to buyer (proposition 2). This section asks two main questions. First, are these economic mechanisms quantitatively important in explaining changes in downstream firms' production costs? We will use the 2008–2009 trade collapse as a natural experiment, a period where intermediate imports and the number of suppliers declined substantially (see sections VA and VB). Second, we will use the same trade collapse to study how well the model can explain the observed drop in supplier connections (the out-of-sample fit of the model).

The forces driving the trade collapse are complex (see Eaton et al., 2013). Here we ask how much one particular channel, worsening market access to suppliers (e.g., due to increased trade costs, f_{ij} and τ_{ij} , and/or fewer potential suppliers, n_i) affected firms' production costs and supplier network. While there is little doubt that the crisis caused the exit of many firms worldwide, there is also evidence of increased trade barriers in the aftermath of the collapse in 2008 (see Evenett, 2009, and Kee, Neagu, & Nicita, 2013). Our approach, however, does not rule out other explanations for the trade collapse. Rather, our quantitative framework aims at isolating the part of the trade collapse that was due to a change in market access to suppliers and thus a loss of buyer-seller relationships.

A. Data

This quantitative exercise requires data on firms' imports across suppliers and source countries, as well as data on firms' total purchases of intermediate goods. In this part

²⁹ Some country-pairs have imports close to 0 or very few firms importing in that pair. We eliminate those cases by restricting the analysis to industry-country pairs with $\pi_{NO,jkt}^{Imports} > .05$.

³⁰ In the min and median exports regressions, we use only firms with more than five customers.

of the paper, we therefore use customs data on imports that have an identical structure to the export data described above. In addition, we match the import data to balance sheet data for manufacturing firms, which includes a variable for total intermediate purchases. The balance sheet data are from Statistics Norway's Capital database, an annual unbalanced panel of all nonoil manufacturing joint-stock firms. It includes approximately 8,000 firms per year, which is roughly 90% of all manufacturing firms.³¹

In our data, the adding and dropping of suppliers constitute a significant share of aggregate import growth every year. We decompose aggregate imports into three margins: firms that enter or exit importing, firms that enter or exit a relationship with a supplier, and continuing relationships. Specifically, the change in aggregate imports can be calculated as

$$\Delta x_t = \sum_{f \in N} x_{ft} - \sum_{f \in E} x_{ft-1} + \sum_{f \in C} \Delta x_{ft},$$

where f indexes firms, N is the set of new firms entering importing, E is the set of existing firms exiting importing, and C is the set of firms continuing to import. Furthermore,

$$\Delta x_{ft} = \sum_{j \in A_f} x_{fjt} - \sum_{j \in D_f} x_{fjt-1} + \sum_{j \in C_f} \Delta x_{fjt}, \quad (14)$$

where j indexes supplier relationships, A_f is the set of suppliers added by firm f , D_f is the set of suppliers dropped by firm f , and C_f is the set of suppliers continued by firm f . In table 12 in the online appendix, we present decompositions based on our sample. The global downturn during the Great Recession of 2008–2009 hit Norwegian trade hard as well. Total intermediate imports fell by 16% from 2008 to 2009. Roughly one-fifth of the decline was due to buyer-supplier relationships breaking up. There is also significant churning in buyer-supplier relationships every year; for example, from 2008 to 2009, gross retired supplier relationships alone contributed to a 17% drop in aggregate imports.

B. Methodology

The data at hand allow us to estimate the change in a downstream firm's market access, Ω_j , which is inversely proportional to their marginal costs (see proposition 2). We show that we can estimate market access using standard firm-level import data and without imposing any assumptions on model parameters except the trade elasticity γ .

Following Dekle, Eaton, and Kortum (2007), we solve the model in changes. When we use equation (5), the change in the market access term Ω_j is

$$\hat{\Omega}_{mj} \equiv \left(\sum_i \pi_{mij} \hat{\rho}_{ij} \right)^{1/\gamma}, \quad (15)$$

³¹ Statistics Norway's capital database is described in Raknerud, Ronnigen, and Skjerpen (2004).

where \hat{x} denotes the annual change x_t/x_{t-1} and ρ_{ij} is a composite index of costs associated with sourcing from location i , $\hat{\rho}_{ij} \equiv \hat{Y}_i (\hat{\tau}_{ij} \hat{w}_i)^{-\gamma} (\hat{w}_i \hat{f}_{ij})^{1-\gamma/(\sigma-1)}$. Henceforth, we use the terminology *sourcing costs* for ρ_{ij} . Finally, π_{mij} is firm m 's trade share in $t-1$, $\pi_{mij} \equiv X_{mijt-1} / \sum_k X_{mkjt-1}$. We have added a firm subscript m to the market access term Ω_{jm} because at the firm level, ex ante trade shares π_{mij} vary across firms.³²

Using equation (9) and equation (21) in the online appendix, the change in a downstream firm's import share from i is

$$\hat{\pi}_{mij} \equiv \frac{\hat{R}_{ij}^{TOT}(Z)}{\hat{E}_j(Z)} = \hat{\rho}_{ij} \hat{\Omega}_{mj}^{-\gamma}. \quad (16)$$

Using the import share π_{mij} instead of the value of imports R_{ij}^{TOT} is useful because it allows us to eliminate a firm's unobserved productivity Z (which appear in R_{ij}^{TOT} ; see equation [9]), thus isolating sourcing costs ρ_{ij} . By the same logic, using a firm's import share, π_{mij} , instead of import value, R_{ij}^{TOT} , eliminates a firm's unobserved demand shocks, so that our estimates are not contaminated by the drop in demand during the 2008–2009 trade collapse. Intuitively, equations (15) and (16) make it clear that one can use data on the change in the import share to obtain information about the change in sourcing costs. This allows us to calculate the change in market access, $\hat{\Omega}_{mj}$, a weighted average of sourcing costs, using ex ante import shares π_{mij} as weights. Note that the assumption that the number of firms is proportional to output is innocuous. The sufficient statistic approach ensures that we simply identify the total change in sourcing costs $\hat{\rho}_{ij}$. No assumption about the determination of the subcomponents of ρ_{ij} (n_i or Y_i , w_i and so on) is needed.

Fixed-point procedure. There is no closed-form solution for $\hat{\Omega}_{mj}$ because $\hat{\Omega}_{mj}$ and $\hat{\rho}_{ij}$ are nonlinear functions of each other. Hence, we solve numerically for $\hat{\Omega}_{mj}$ using the following fixed-point procedure.

- Step 1: Choose initial values for $\hat{\rho}_{ij}$.
- Step 2: Solve for $\hat{\Omega}_{mj}^y$ for firm m , using equation (15) and ex ante trade shares π_{mij} for firm m .
- Step 3: From equation (16), calculate $\hat{\rho}_{ij} = \hat{\Omega}_{mj}^y \hat{\pi}_{mij}$.

The resulting sourcing cost $\hat{\rho}_{ij}$ varies across firms because of measurement error and firm-country specific shocks. We eliminate this noise by taking the median of $\hat{\rho}_{ij}$ across firms. We return to step 2 if the difference between the current and previous $\hat{\rho}_{ij}$ is large, and we stop if the difference is sufficiently small. The fixed-point procedure converges quickly. The choice of starting values $\hat{\rho}_{ij}$ has no impact on the solution.

³² In the model, import shares do not vary across downstream firms. One could add firm-country specific shocks to the relation-specific fixed cost that would bring the model closer to the data.

TABLE 4.—A SUPPLY SHOCK: THE TRADE COLLAPSE

	Median	Mean	Weighted Mean	SD
<i>Data</i>				
$\ln \hat{\pi}_{mij}$	-.101	-.212	-.212	1.096
$\ln \hat{S}_{mij}$	0	-.083	-.080	.546
$\ln \hat{S}_{mij}, \geq 2$ suppliers	-.154	-.212	-.164	.523
<i>Model</i>				
$\ln \hat{\Omega}_{mj}^{\gamma}$	-.013	-.027	-.034	.033
$\ln \hat{\pi}_{mij}$	-.114	-.110	-.106	.056
$\ln \hat{S}_{mij}$	-.114	-.110	-.106	.056
$\ln \hat{S}_{mij}, \geq 2$ suppliers	-.110	-.106	-.117	.053

2008 to 2009 changes. Observations for 3,306 firms and 30 countries. Firm revenue is used as weights in weighted mean calculations. $\hat{\Omega}_{mj}^{\gamma}$ is change in market access for firm m , $\hat{\pi}_{mij}$ is change in the import share from i for firm m , and \hat{S}_{mij} is change in the measure of suppliers from i for firm m .

Normalization. We can identify $\hat{\rho}_{ij}$ only up to a constant because for given m and j , one of the i elements in the vector $\hat{\pi}_{mij}$ is linearly dependent on the other elements. We normalize the change in domestic sourcing cost to one, $\hat{\rho}_{1j} = 1$ where $i = 1$ is the domestic market.

Knowing the change in sourcing costs $\hat{\rho}_{ij}$, one needs only one model parameter, the import elasticity γ , to calculate the firm-level change in marginal costs from equation (15). The change in marginal costs varies across firms because their ex ante trade shares π_{ijm} differ: some firms are using imported inputs intensively and other firms are not.

C. Results and Out-of-Sample Fit

Table 4 provides an overview of observed changes in import shares and supplier connections and the results from the quantitative exercise. We calculate firm-specific import shares, π_{mij} , and the related 2008 to 2009 change, $\hat{\pi}_{mij}$, for all source countries, i , including Norway itself. We restrict the data in two ways. First, firms with no foreign sourcing are dropped, as their $\hat{\Omega}_{mj}$ is normalized to 1 (see the previous section). Second, we focus on the set of importing countries where a firm has positive imports in both 2008 and 2009. This is necessary because $\hat{\pi}_{mij}$ is not defined if a firm adds or drops a sourcing market. Third, while the raw data include imports from every possible source, we limit the analysis to the top thirty source countries in terms of total import value. Focusing on the top thirty sources ensures that we typically have more than one firm with suppliers from a given source country. With these restrictions, aggregate imports in our final sample account for 78% of total imports in the manufacturing sector.

Turning to the results, we find that the mean market access ($\log \hat{\Omega}_{mj}^{\gamma}$), across all importers, fell by 2.7%, while the weighted mean, using firm revenue as weights, fell by 3.4%. This translates into substantial cost increases among the importing firms. With a trade elasticity of $\gamma = 4$, the weighted mean increase in marginal costs is roughly 1% ($(1 - 0.034)^{1/\gamma}$).³³

³³ Recall that the change in domestic sourcing costs is normalized to 0. Hence, we only identify changes in marginal costs coming from changes in foreign market access.

There is also great dispersion across firms. For firms importing small amounts, there is almost no change in $\hat{\Omega}_{mj}$. For big importers, however, the decline in $\hat{\Omega}_{mj}$ is much larger; the decline for the first decile of $\ln \hat{\Omega}_{mj}^{\gamma}$ is -0.08 , or a 2% increase in marginal costs given $\gamma = 4$. Our results suggest that the trade collapse had a relatively large negative impact on production costs among importing firms and was driven by changes in sourcing costs. Recall that these marginal cost estimates are not contaminated by the demand side of the economy, because the demand side was effectively differenced out in the quantitative procedure.

Model fit. Table 4 shows the median, mean, and weighted mean change in firms' import shares ($\ln \hat{\pi}_{mij}$) across all firm-country pairs, in the data and in the model. Import shares fell in both the data and the model, and the model's median change is very close to the data. This is as expected because we used data on changes in trade shares to calculate $\hat{\rho}_{ij}$ and $\hat{\Omega}_{mj}$. Moreover, we evaluate the out-of-sample fit in terms of changes in buyer-seller linkages in the economy (a moment we did not target). In the model, the change in the firm-level measure of suppliers is given by

$$\hat{S}_{mij} = \frac{\hat{\rho}_{ij}}{\hat{w}_i \hat{f}_{ij}} \hat{\Omega}_{mj}^{-\gamma} \hat{Z}^{\gamma}.$$

We calculate \hat{S}_{mij} for each firm using our estimates of $\hat{\rho}_{ij}$ and $\hat{\Omega}_{mj}^{\gamma}$, while keeping other factors fixed (productivity Z and relation-specific costs $w_i f_{ij}$). The model captures the decline in supplier connections well; the model generates an average 11% fall in the number of supplier connections, while the actual average decline was 8%.

The fit for the median $\ln \hat{S}_{mij}$ is poor, as the median log change in the data is 0. However, this occurs because S in the data is an integer and cannot take a value lower than 1. If we take the median of $\ln \hat{S}_{mij}$ across firm-country pairs with two or more suppliers, we find a median decline in suppliers of 15%—slightly more than the model prediction.

Looking across source countries, we also find a positive correlation between the average response in the data and in the model. Figure 12 in the online appendix shows the true and predicted mean for the change in the number of suppliers, $\ln \hat{S}_{mij}$, across all firms with one or more connections for each source country. The drop in supplier connections was particularly strong in Canada and Taiwan in both the model and data. For other countries, the fit is worse. The observed decline in the number of Japanese suppliers was relatively large, whereas in the model, it was not. We conclude that the model is able to quantitatively replicate the role of diminishing buyer-seller connections, one important margin for trade adjustments during the 2008–2009 trade collapse.

VI. Conclusion

We use highly disaggregated trade transaction data from Norway to explore the role of buyers and buyer-seller

relationships in international trade. We present a series of basic facts about buyer-seller relationships in international trade that point to a distinct role of buyers in explaining variation in trade, extreme concentration of exports across both sellers and buyers and Pareto-shaped distributions of buyers per exporter and sellers per importer. We find that large exporters reach more customers, but exports to the median customer are not increasing with the number of customers within a destination, and there is negative degree assortivity in the exporter-importer matches. In other words, large exporters on average reach importers that buy from a relatively smaller number of Norwegian firms.

Guided by these facts, we develop a parsimonious multicountry model of heterogeneous exporters and importers where matches are subject to a relation-specific fixed cost. A central feature of the model is that lower variable trade costs will lead to higher export growth when buyers in the destination market are less dispersed in terms of their productivity. When buyers are more similar, an exporter will find many new profitable matches, whereas if buyers are dispersed, only a few more matches will become profitable. In other words, the customer extensive margin response will be strong when buyer heterogeneity is small. We test this prediction by exploiting variation in import shares across industries and countries over time and find strong empirical support.

The theoretical model also shows that firms' access to suppliers is important for firm performance and marginal costs. To evaluate the role of these features in explaining variation in trade, we develop an empirical methodology to estimate downstream firms' marginal cost response due to exogenous shocks to foreign market access. We show that a sufficient statistic for a firm's change in marginal costs is the level of, and the change in, intermediate import shares and the trade elasticity. The methodology is subsequently applied to evaluate the impact of the 2008–2009 trade collapse on firms' production costs. Our results indicate that worsened market access during the trade collapse had a significantly negative impact on production costs, and especially so for downstream firms that were ex ante highly exposed to international markets. The quantitative exercise shows that the model matches well the fall in the number of buyer-seller matches observed during the trade collapse.

Our results suggest that buyer heterogeneity and buyer-seller connections are important in understanding firm-level trade, as well as fluctuations in marginal costs and measured productivity. Future research might fruitfully focus on the growth and stability of exporter-importer relationships, as well as the sources of heterogeneity across sellers and buyers.

REFERENCES

Alfaro, Laura, and Maggie X. Chen, "Market Reallocation and Knowledge Spillover: The Gains from Multinational Production," Harvard Business School working paper (2013).

- Amiti, Mary and Jozef Konings, "Trade Liberalization, Intermediate Inputs and Productivity: Evidence from Indonesia," *American Economic Review* 97 (2007), 1611–1638.
- Anderson, James E., and Eric van Wincoop, "Trade Costs," *Journal of Economic Literature* 42 (2004), 691–751.
- Antràs, Pol, and Arnaud Costinot, "Intermediated Trade," *Quarterly Journal of Economics* 126 (2011), 1319–1374.
- Antràs, Pol, Teresa C. Fort, and Felix Tintelnot, "The Margins of Global Sourcing: Theory and Evidence from U.S. Firms," *American Economic Review* 107 (2017), 2514–2564.
- Arkolakis, Costas, "Market Access Costs and the New Consumers Margin in International Trade," *Journal of Political Economy* 118 (2010), 1151–1199.
- , "A Unified Theory of Firm Selection and Growth," NBER working paper 17553 (2011).
- Arkolakis, Costas, Arnaud Costinot, and Andres Rodriguez-Clare, "New Trade Models, Same Old Gains?" *American Economic Review* 102 (2012), 94–130.
- Armenter, Roc, and Miklos Koren, "A Balls-and-Bins Model of Trade," *American Economic Review* 104 (2014), 2127–2151.
- Bartelsman, Eric J., Stefano Scarpetta, and Fabiano Schivardi, "Cross-Country Differences in Productivity: The Role of Allocation and Selection," *American Economic Review* 103 (2013), 305–334.
- Benguria, Felipe, "Production and Distribution of International Trade: Evidence from Matched Exporter-Importer Data," technical report (2014).
- Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott, "Firms in International Trade," *Journal of Economic Perspectives* 21 (2007), 105–130.
- Bernard, Andrew B., J. Bradford Jensen, and Peter K. Schott, "Importers, Exporters, and Multinationals: A Portrait of Firms in the U.S. That Trade Goods," in T. Dunne, J. B. Jensen, and M. Roberts, eds., *Producer Dynamics: New Evidence from Micro Data* (Chicago: University of Chicago Press, 2009).
- Bernard, Andrew B., Andreas Moxnes, and Yukiko Saito, "Production Networks, Geography and Firm Performance," *Journal of Political Economy* (forthcoming).
- Blaum, Joaquin, Claire Lelarge, and Michael Peters, "The Gains from Input Trade with Heterogeneous Importers," *American Economic Journal: Macroeconomics* (forthcoming).
- Blum, Bernardo S., Sebastian Claro, and Ignatius Horstmann, "Facts and Figures on Intermediated Trade," *American Economic Review* 100 (2010), 419–423.
- , "Import Intermediaries and Trade Costs: Theory and Evidence," University of Toronto mimeograph (2012).
- Bøler, Esther Ann, Andreas Moxnes, and Karen Helene Ulltveit-Moe, "R&D, International Sourcing and the Joint Impact on Firm Performance," *American Economic Review* 105 (2015), 3704–3739.
- Carballo, Jeronimo, Gianmarco I. P. Ottaviano, and Christian Volpe Martincus, "The Buyer Margins of Firms' Exports," CEPR discussion paper 9584 (2013).
- Castellani, Davide, Francesco Serti, and Chiara Tomasi, "Firms in International Trade: Importers' and Exporters' Heterogeneity in Italian Manufacturing Industry," *World Economy* 33 (2010), 424–457.
- Cebeci, Tolga, Ana Fernandes, Caroline Freund, and Martha Denisse Pierola, "Exporter Dynamics Database," World Bank policy research working paper 6229 (2012).
- Chaney, Thomas, "Distorted Gravity: the Intensive and Extensive Margins of International Trade," *American Economic Review* 98 (2008), 1707–1721.
- , "The Network Structure of International Trade," *American Economic Review* 104 (2014), 3600–3634.
- Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Unbalanced Trade," *American Economic Review* 97 (2007), 351.
- Dragusanu, Raluca, "Firm-to-Firm Matching along the Global Supply Chain," Harvard University working paper (2014).
- Eaton, Jonathan, and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica* 70 (2002), 1741–1779.
- Eaton, Jonathan, Marcela Eslava, David Jenkins, C. J. Krizan, and James Tybout, "A Search and Learning Model of Export Dynamics," Society for Economic Dynamics 2015 meeting paper (2014).
- Eaton, Jonathan, Samuel S. Kortum, and Francis Kramarz, "An Anatomy of International Trade: Evidence from French Firms," *Econometrica* 79 (2011), 1453–1498.

- Eaton, Jonathan, Sam Kortum, Brent Neiman, and John Romalis, "Trade and the Global Recession," NBER working paper 16666 (2013).
- Ethier, W., "Internationally Decreasing Costs and World Trade," *Journal of International Economics* 9:1 (1979), 1–24.
- Evenett, Simon, *Broken Promises: The Second GTA Report* CEPR, 2009.
- Goldberg, Pinelopi K., Amit K. Khandelwal, Nina Pavcnik, and Petia Topalova, "Imported Intermediate Inputs and Domestic Product Growth: Evidence from India," *Quarterly Journal of Economics* 125 (2010), 1727–1767.
- Gopinath, Gina, and Brent Neiman, "Trade Adjustment and Productivity in Large Crises," *American Economic Review* 104 (2014), 793–831.
- Halpern, Laszlo, Miklos Koren, and Adam Szeidl, "Imported Inputs and Productivity," CeFIG working paper, Central European University (2011).
- Helpman, Elhanan, Marc J. Melitz, and Stephen R. Yeaple, "Export versus FDI with Heterogeneous Firms," *American Economic Review* 94 (2004), 300–316.
- Irrazabal, Alfonso, Andreas Moxnes, and Luca David Oromolla, "The Margins of Multinational Production and the Role of Intrafirm Trade," *Journal of Political Economy* 121 (2013), 74–126.
- Jackson, Matthew O., and Brian W. Rogers, "Meeting Strangers and Friends of Friends: How Random Are Social Networks?" *American Economic Review* 97 (2007), 890–915.
- Kang, M.-P., J. Mahoney, and D. Tan, "Why Firms Make Unilateral Investments Specific to Other Firms: The Case of OEM Suppliers," *Strategic Management Journal* 30 (2009), 117–135.
- Kee, Hiau Looi, Cristina Neagu, and Alessandro Nicita, "Is Protectionism on the Rise? Assessing National Trade Policies during the Crisis of 2008," *this REVIEW* 95 (2013), 342–346.
- Khandelwal, Amit K., and Petia Topalova, "Trade Liberalization and Firm Productivity: The Case of India," *this REVIEW* 93 (2011), 995–1009.
- Mayer, Thierry, and Gianmarco Ottaviano, "The Happy Few: The Internationalisation of European Firms," *Intereconomics: Review of European Economic Policy* 43 (2008), 135–148.
- Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica* 71 (2003), 1695–1725.
- Monarch, Ryan, "It's Not You, It's Me: Breakups in U.S.-China Trade Relationships," Center for Economic Studies working paper CES-WP-14-08 (2013).
- Muuls, Mirabelle, and Mauro Pisu, "Imports and Exports at the Level of the Firm: Evidence from Belgium," *World Economy* 32 (2009), 692–734.
- Oberfield, Ezra, "Business Networks, Production Chains, and Productivity: A Theory of Input-Output Architecture," *Econometrica* (forthcoming).
- Petropoulou, Dimitra, "Information Costs, Networks and Intermediation in International Trade," Globalization and Monetary Policy Institute working paper 76, Federal Reserve Bank of Dallas (2011).
- Raknerud, Arvid, Dag Ronningen, and Terje Skjerpen, "Documentation of the Capital Database," *Statistics Norway Documents* 2004/16 (2004), 1–13.
- Rauch, James E., "Networks versus Markets in International Trade," *Journal of International Economics* 48 (1999), 7–35.
- Rauch, James E., and Joel Watson, "Network Intermediaries in International Trade," *Journal of Economics & Management Strategy* 13 (2004), 69–93.
- Serrano, Ma Angeles, and Marian Boguna, "Topology of the World Trade Web," *Physical Review* 68 (2003), 015101.
- Sugita, Yoichi, Kensuke Teshima, and Enrique Siera, "Assortative Matching of Exporters and Importers," Research Institute of Economy, Trade and Industry discussion paper 17016 (2014).