

TWO-SIDED HETEROGENEITY AND TRADE*

Andrew B. Bernard[†]

Tuck School of Business at Dartmouth, CEPR & NBER

Andreas Moxnes[‡]

University of Oslo, CEPR & NBER

Karen Helene Ulltveit-Moe[§]

University of Oslo & CEPR

August 2014

Abstract

This paper explores the consequences of joint heterogeneity on the supply side and the demand side in international trade using a dataset from Norway where exporters and importers are identified in each transaction. The buyer-seller data reveal new facts on the distributions of buyers per exporter and exporters per buyer, the matching among sellers and buyers, and the variation of buyer dispersion across destinations. The paper develops a trade model with heterogeneous importers and heterogeneous exporters where matches incur a relation-specific fixed cost. The model generates new testable predictions emphasizing the importance of buyer-seller relationships in explaining trade patterns.

Keywords: Heterogeneous firms, exporters, importers, trade elasticity

JEL codes: F10, F12, F14.

*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. Moxnes is grateful for financial support from The Nelson A. Rockefeller Center for Public Policy and the Social Sciences at Dartmouth College.

[†]100 Tuck Hall, Hanover, NH 03755, USA, *tel:* +1 603 646 0302, *email:* andrew.b.bernard@tuck.dartmouth.edu

[‡]PO Box 1095 Blindern, 0317 Oslo, Norway, *tel:* +47 22857048, *email:* andreamo@econ.uio.no

[§]PO Box 1095 Blindern 0317 Oslo Norway *tel:* +47 22857240, *email:* k.h.ulltveit-moe@econ.uio.no

1 Introduction

Empirical studies of firms within industries consistently report substantial heterogeneity in measures of performance such as size and productivity. The importance of such heterogeneity for aggregate and firm-level export outcomes is well established. More recently, researchers have found comparable variation in size and performance across importers (Bernard et al., 2009). However, there has been far less work on the role of demand side (importer) heterogeneity in international trade.¹ This paper explores the interaction of exporter and importer heterogeneity and the consequences for firm-level and aggregate trade and productivity.

The paper makes use of a novel dataset that links annual Norwegian non-oil export transactions with every importer in every country. We establish a set of basic facts about sellers and buyers across markets and develop a parsimonious theoretical model with two-sided heterogeneity. The model is able to match the basic facts and generates additional testable implications about the role of buyer heterogeneity in international trade. A key theoretical and empirical finding is that buyer-side heterogeneity plays an important role in explaining the response of exports to aggregate shocks and trade liberalization.

In our data, the identities of both the exporter and the importer are known. We can link a firm's annual export transactions to specific buyers in every destination country and, at the same time, examine all of an importer's transactions with Norwegian firms. We establish a set of facts that guide the development of the model. First, the populations of sellers and buyers of Norwegian exports are both characterized by extreme concentration, mirroring the findings in Bernard et al. (2009). Although a handful of firms account for a large share of aggregate trade, one-to-one matches are typically not important in the aggregate. Hence, a model allowing for many-to-many matches is needed in order to explain the aggregate facts. Second, the distributions of buyers per exporter and exporters per buyer approximately follow a power-law over a wide range of magnitudes: many firms have only one connection, but the big firms typically have tens or hundreds of connections. Third, there is negative degree assortitivity among sellers and buyers, meaning that relatively well-connected exporters on average sell to relatively less-connected importers. This is driven by a selection effect, where well-connected exporters not only sell to well-connected importers, but also to less-connected importers, whereas less-connected exporters are unlikely to match with less-connected importers. Fourth, firms tend to follow a hierarchical pecking order in their choice of connections: an exporter selling to the second ($k+1^{\text{th}}$) most connected buyer has a higher likelihood of selling to the (k^{th}) most connected buyer relative to the prediction of a random matching model. This mirrors the finding in Eaton et al. (2011), but at a lower level of aggregation. We check the external validity of our results on trade networks using import data from Colombia that has similar buyer-seller information to

¹Exceptions are Blum et al. (2010) and Blum et al. (2012), Carballo et al. (2013) and Eaton et al. (2012) who examine exporter-importer pairs for individual pairs of countries.

that in the Norwegian data. We find that the basic characteristics of exporter-importer relationships are confirmed by the Colombian data (Appendix I).

We develop a framework to match the basic facts on trade networks by building a multi-country model of international trade with heterogeneity among importers as well as exporters. 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.

The model offers a number of firm-level and macro insights. At the firm-level, falling trade barriers lower marginal costs among final goods producing firms by reducing the cost of inputs and by facilitating more matches between input suppliers and final goods producers. The importance of intermediate inputs for productivity growth has strong empirical support; 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 can generate firm-level responses to trade cost shocks that are consistent with the empirical evidence. Moreover, our work suggests 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 macro level, the gains from trade are identical to the class of models considered in Arkolakis et al. (2012): welfare depends only on two sufficient statistics, the share of expenditure on domestic goods and the elasticity of trade with respect to variable trade costs. However, our work suggests that estimating the the elasticity of trade to variable costs is challenging. To the extent that variable trade costs and relation-specific costs covary in the data, and relation-specific costs are not properly controlled for, the estimated variable trade elasticity will be biased.² Finally, the model shows that relation-specific costs shape aggregate outcomes: lower relation-specific costs facilitate more matches between buyers and sellers, therefore generating more trade between nations as well as improving consumer welfare.

Beyond matching the basic facts, the theoretical model generates three main testable implications. First, a demand shock in a destination market has no impact on a firm's exports to its marginal customer in that market. This occurs because the marginal transaction is determined only by the relation-specific cost. Second, the change in a firm's exports following a demand shock in

²For example, the EU's single market has eliminated tariffs on most goods as well as focusing on the reduction of non-tariff barriers. The harmonization of product regulation, rules relating to company law and corporate governance aims to help companies operate throughout the EU on the basis of a single set of rules, see http://ec.europa.eu/internal_market/top_layer/business-environment/index_en.htm.

the destination country depends on the extent of buyer heterogeneity in that market. Specifically, the trade elasticity is higher in markets with less dispersion of buyer efficiency. This occurs because exporters will meet more buyers when these firms are less heterogeneous, i.e. the extensive margin response is stronger. Third, dispersion in exports across firms in a destination market is inversely related to dispersion in buyer productivity in that market. Exports are more dispersed in markets with less buyer dispersion. The intuition is that if dispersion among buyers is high, then imports are concentrated in a few large buyers, and even small and low productivity exporters will sell to them, thus compressing the exports distribution.

We find empirical support for all three predictions from the model. A positive demand shock has no impact on exports to the marginal buyer, whereas the number of buyers and total firm-level exports increase. The firm-level elasticity of exports (and buyers) with respect to a demand shock is higher in countries with less dispersion in buyer productivity. Finally, using a differences-in-differences estimator, we find that exports to country-product pairs are less dispersed in markets with more buyer dispersion.

An implication of our work is that the variance of demand matters for how responsive firm's trade flows are to changes in trade policy, exchange rate movements, and other types of shocks. Previous research has shown that dispersion in firm size and productivity differs both across regions and over time due to policy-induced distortions (Bartelsman et al., 2013, Braguinsky et al., 2011, Garicano et al., 2013 and Hsieh and Klenow, 2009). Our findings may add to the understanding of the impact of policy changes on international trade. More broadly, our framework enhances the understanding of how relation-specific costs shape international trade both at the micro and macro level.

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 et al. (2009), Castellani et al. (2010) and Muuls and Pisu (2009) show that the heterogeneity of importing firms rivals that of exporters for the US, Italy and Belgium respectively. Amiti and Konings (2007), Halpern et al. (2011) and Boler et al. (2012) relate the importing activity of manufacturing firms to increases in productivity.

Papers by Rauch (1999), Rauch and Watson (2004), Antràs and Costinot (2011) and Petropoulou (2011) consider exporter-importer linkages. Chaney (2011) also has a search-based model of trade where firms must match with a contact in order to export to a destination. These papers adopt a search and matching approach to linking importers and exporters, while in this paper we abstract from these mechanisms and instead focus on the implications of buyer heterogeneity for international 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 et al. (2010) and Blum et al. (2012) examine characteristics of trade transactions for the exporter-importer pairs of Chile-Colombia and Argentina-Chile while Eaton et al. (2012) consider exports of Colombian firms to specific importing firms in the United States. Blum et al. (2010) and Blum et al. (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. (2012) 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 to maintaining relationships with existing ones. In contrast to those papers but similar to Carballo et al. (2013), we focus on the role of importer heterogeneity across destinations. Carballo et al. (2013) focus on export margins across goods, countries and buyers, while we study the implications of importer heterogeneity on exporting firms' responses to exogenous shocks to trade barriers and demand. Monarch (2013) estimates switching costs using a panel of U.S importers and Chinese exporters, while Dragusanu (2014) explores how the matching process varies across the supply chain using U.S.-Indian data. Sugita et al. (2014) study matching patterns in U.S.-Mexico trade while Benguria (2014) estimates a trade model with search costs using matched French-Colombian data.

The rest of the paper is structured as follows. In Section 2 we describe our main data source, while in Section 3 we document a set of facts on the role of buyers in trade, the heterogeneity of buyers and sellers, and their bilateral relationships which will guide our theoretical model and subsequent empirical specification. In Section 4 we develop a multi-country trade model with networks of heterogeneous sellers and buyers, while in Section 5 we test the empirical predictions of the model. Section 6 concludes.

2 Data

The main data set employed in this paper is based on Norwegian transaction-level customs data from 2005-2010. The 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. However, in addition, this data has 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.³ Our data include the universe of Norwegian non-oil merchandise

³Statistics 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.

exports, and we observe export value and quantity. In 2005 total Norwegian non-oil merchandise exports amounted to US\$41 Billion, equal to approximately 18 percent of Mainland Norway GDP (GDP excluding the oil and gas sector). Norwegian merchandise exports constituted around 1/3 of total Norwegian exports in 2005, and were undertaken by 18,219 sellers who sold 5,154 products to 81,362 buyers across 205 destinations.⁴

The firm-level evidence from Norwegian non-oil exports looks remarkably similar compared to other developed countries, see e.g. Cebeci et al. (2012), Alfonso Irarrazabal and Opromolla (2013) and Mayer and Ottaviano (2008). Tables 1 and 2 report the top 5 exported products from Mainland Norway. Ordered by export value, exports of intermediate inputs (metals, fertilizers), capital equipment (vessels) and food (fish) constitute the largest shares of exports. Ordered by the number of exporters, differentiated products such as machinery and various parts and components rank the highest.

The empirical analysis in Section 5 also rely on international trade data from COMTRADE and the World Bank’s Enterprise Surveys, the World Bank’s Exporter Dynamics database as well as the Bureau van Dijk’s ORBIS database to calculate measures of buyer heterogeneity across export destinations. Furthermore, we use Colombian buyer-seller data to test the external validity of our results. See Appendix Section I for a description.

3 Exporters and Importers

3.1 Basic Facts

This section explores the matched exporter-importer data. 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. To examine the role of buyers in the variation of exports across countries, we decompose total exports to country j , x_j , into the product of the number of exporting firms, f , the number of exported products, p , the number of buyers (importers), b , the density of trade, d , i.e. 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, e.g.

⁴In the same year, merchandise exports in the U.S. and UK was 60 and 50 percent of total exports respectively.

$\ln f_j$, against $\ln x_j$. Given that OLS is a linear estimator and its residuals have an expected value of zero, 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. The results, shown in Table 3, confirm and extend previous findings on the importance of the extensive and intensive margins of trade. While it has been shown in a variety of contexts that the number of exporting firms and products increases 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 or larger than the firm or product margins.

It is well documented that the total value of exports, the number of exporting firms and the number of exported products are all systematically related to destination market characteristics such as GDP and distance. Looking within the firm across markets, we show how the buyer margin responds to these standard gravity variables by regressing a firm's number of customers on a firm fixed effect, distance and GDP in the destination market (all in logs). The results in Table 4 column 2 show that a firm's number of customers is significantly higher in larger markets and smaller in remote markets, i.e. importers per exporter vary systematically with GDP and distance.⁵


Fact 2: The populations of sellers and buyers of Norwegian exports are both characterized by extreme concentration. The top 10 percent of sellers account for 98 percent of Norwegian aggregate exports. At the same time, the top 10 percent of buyers are almost as dominant and account for 96 percent of the purchases of Norwegian exports (Table 5). 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 are typically not important in the aggregate. Table 6 shows that one-to-one matches represent 9.5 percent of all exporter-importer connections but account for only 4.6 percent of aggregate trade. Many-to-many matches, i.e. 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.

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 plot the number of buyers of each exporting firm in a particular market against the fraction of exporters selling in the market who sell to at least that many buyers. We find that the distributions are remarkably similar across markets, Figure 1 plots the results for China, the US and Sweden.⁶ The average number of buyers per seller is 4.5 in the U.S. and 3.6 in China and Sweden (see Table 5). The distributions

⁵We also use total firm-level exports and average firm-level exports per buyer as dependent variables in columns 1 and 3.

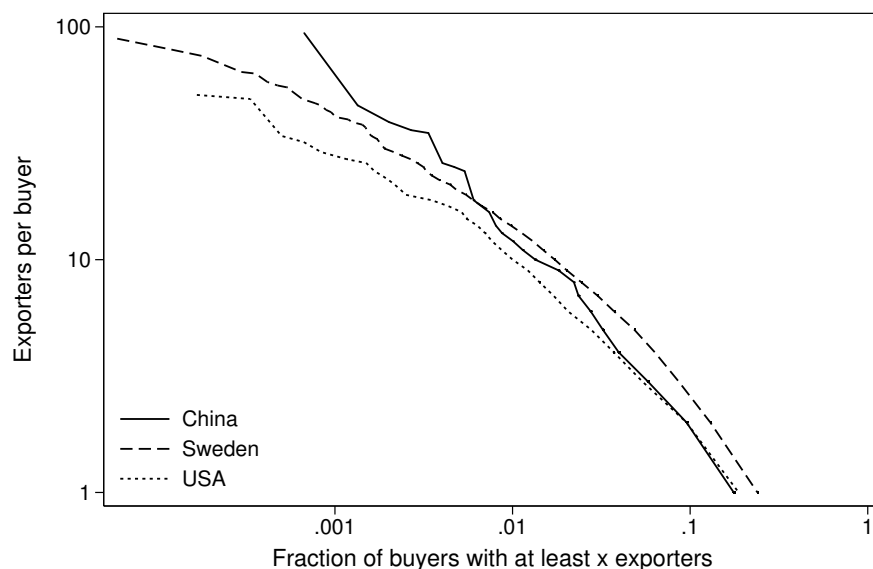
⁶To interpret Figure 1 as the empirical CDF, let x_j^ρ be the ρ th percentile of the number of buyers per exporter in market j . We can then write $\Pr[X \leq x_j^\rho] = \rho$. If the distribution is Pareto with shape parameter a and location parameter x_0 , we have $1 - (x_0/x_j^\rho)^a = \rho$, and taking logs this gives us $\ln x_j^\rho = \ln x_0 - \frac{1}{a} \ln(1 - \rho)$. Hence, the slope in Figure 1 is $-1/a$.

Figure 1: Distribution of the number of buyers per exporter.




Note: 2006 data, log scale. The estimated slope coefficients are -1.02 (s.e. 0.010) for China, -1.02 (s.e. 0.002) for Sweden and -1.13 (s.e. 0.005) for the U.S. The distribution is Pareto if the slope is constant. The slope coefficient equals the negative of the inverse of the Pareto shape parameter ($-1/a$, see footnote 7).

Figure 2: Distribution of the number of exporters per buyer.



Note: 2006 data, log scale. The estimated slope coefficients are -0.92 (s.e. 0.002) for China, -0.88 (s.e. 0.001) for Sweden and -0.80 (s.e. 0.001) for the U.S. The distribution is Pareto if the slope is constant. The slope coefficient equals the negative of the inverse of the Pareto shape parameter ($-1/a$, see footnote 3).

Figure 3: Number of buyers & firm-level exports.




Note: 2006 data. The Figure shows the fitted line from a kernel-weighted local polynomial regression of firm-destination log exports on firm-destination log number of customers. Axes scales are in logs. Exports are normalized, see footnote 7.

appear to be largely consistent with a Pareto distribution as the cdfs are close to linear except in the tails of the distribution. Note that the Pareto fails to capture the 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.

We also plot the number of exporters per buyer in a particular market against the fraction of buyers in this market who buys from at least that many exporters (see Figure 2). Again the distributions are approximately Pareto, except in the tails, with many buyers having a few suppliers, and a few buyers with many suppliers. The average number of exporters per buyer in China, Sweden and the US is 1.7, 1.9 and 1.6, respectively.

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. Figure 3 plots the relationship between a firm's number of customers on the horizontal axis and its total exports on the vertical axis using log scales. The solid line is the fit from a kernel-weighted local polynomial regression, and the gray area is the 95 percent confidence interval. We pool all destination countries and normalize exports such that average exports for one-customer firms in

Figure 4: Number of buyers & within-firm dispersion in exports.




Note: 2006 data. The Figure shows the fitted lines from kernel-weighted local polynomial regressions of the x 'th percentile of within-firm-destination log exports on firm-destination log number of customers. Axes scales in logs. Exports are normalized, see footnote 7.

each destination equal 1.⁷ Not surprisingly, firms with more buyers typically export more. The average firm with 10 customers in a destination exports more than 10 times as much as a firm with only one customer.

In Figure 4, we examine how the distribution of exports across buyers varies with the number of buyers. The plot shows the fitted lines from polynomial regressions of the 10th, median and 90th percentile of firm-level log exports (across buyers) and the log number of customers using log scales. We focus on firms with 10 or more customers because the 10th and 90th percentiles are not well defined for firms with fewer than 10 buyers. Again, we pool all destinations and normalize exports such that average exports for one-customer firms are 1. Firm-level exports to the median buyer are roughly constant, so that better-connected sellers are not selling more to their median buyer in a destination compared to less well-connected sellers. The 10th and 90th percentiles are also relatively flat. Dispersion in firm-level exports (across buyers), measured as the difference between the 90th and 10th percentiles, is constant for firms with more than 10 buyers.

⁷The unit of observation is a firm-destination. Log exports are expressed relative to average log exports for one-customer firms, $\ln Exports_{mj} - \ln \overline{ExportsOCF_j}$, where $\ln Exports_{mj}$ is log exports from seller m to market j and $\ln \overline{ExportsOCF_j}$ is average log exports for one-customer firms in market j . This normalization is similar to removing country fixed effects from export flows. Furthermore it ensures that the values on the vertical axis are expressed relative to one-customer firms.

Figure 5: Matching buyers and sellers across markets.




Note: 2006 data. The Figure shows all possible values of the number of buyers per Norwegian firm in a given market j , a_j , on the x-axis, and the average number of Norwegian connections among these buyers, $b_j(a_j)$, on the y-axis. Axes scales are in logs. Both variables are demeaned, i.e. we show $b_j(a_j)/\bar{b}_j(a_j)$, where $\bar{b}_j(a_j)$ is the average number of Norwegian connections among all buyers in market j . The fitted regression line and 95% confidence intervals are denoted by the solid line and gray area. The slope coefficient is -0.13 (s.e. 0.01).

Fact 5: There is negative degree assortivity 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, the less well-connected is its average buyer. Figure 5 provides an overview of seller-buyer relationships. The Figure shows all possible values of the number of buyers per Norwegian firm in a given market, a_j , on the x-axis, and the average number of Norwegian connections among these buyers, $b_j(a_j)$, on the y-axis. Both variables are demeaned and axes are in logs.⁸ The interpretation of a point with the coordinates (10,0.1) is that the customers of Norwegian exporters in a market with 10 times more customers than average have 1/10th the average number of Norwegian suppliers. The slope of the fitted regression line is -0.13, so a 10 percent increase in number of customers is associated with a 1.3 percent decline in average connections among the customers.⁹ In recent work by Bernard et al. (2014), negative degree assortivity is also found

⁸This Figure shows $b_j(a_j)/\bar{b}_j(a_j)$, where $\bar{b}_j(a_j)$ is the average number of Norwegian connections among all buyers in j .

⁹Using 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 assortivity coefficient for 89 percent of the countries we have

Figure 6: Matching buyers and sellers.



Note: 2006 data. Destination market is Sweden. Each bar represents a group of exporters. The groups are (i) Firms with 1 connection, (ii) 2-3, (iii) 4-10 and (iv) 11+ connections. For each group, we plot the share of buyers that have 1, 2-3, 4-10, 11+ connections. For example, the left bar shows that among exporters with 1 connection, roughly 30 percent of these connections are made with buyers that also have 1 connection.

for buyer-seller links among Japanese firms. The Japanese dataset covers close to the universe of domestic buyer-seller links and therefore contains information about the full set of buyer linkages (not only the linkages going back to the source market, as in the current paper).

Note that negative degree assortivity does not mean that well-connected exporters only sell 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. This is illustrated in Figure 6. We divide firms into groups with 1 connection, 2-3, 4-10 and 11+ connections in the largest export market, Sweden.¹⁰ For each group, we then calculate the share of customers that have 1 Norwegian connection, 2-3, 4-10 and 11+ Norwegian connections. The far left bar shows that among exporters with 1 Swedish connection, around 30 percent of the total number of matches are made with buyers with 1 Norwegian connection. The far right bar shows that among exporters with 11+ Swedish connections, almost half of the number

sufficient data for (defined as countries with 10 or more observations in the regression). In appendix F, we show that the elasticity is informative of a structural parameter of the model.

¹⁰The median, 75th percentile and 90th percentile number of number of customers per exporter is 1, 3 and 7 respectively. Patterns for other markets are broadly similar.

of matches made are with buyers with 1 Norwegian connection. Hence, more popular exporters are much more exposed to single-connection buyers.

Note that degree assortivity is only a meaningful measure in economic environments with many-to-many matching. Moreover, negative degree assortivity can coexist with positive assortative matching on the intensive (export value) margin. For example, 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.¹¹ 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 Appendix C.

Interestingly, social networks typically feature positive degree assortivity, 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 (Jackson and Rogers, 2007).¹²

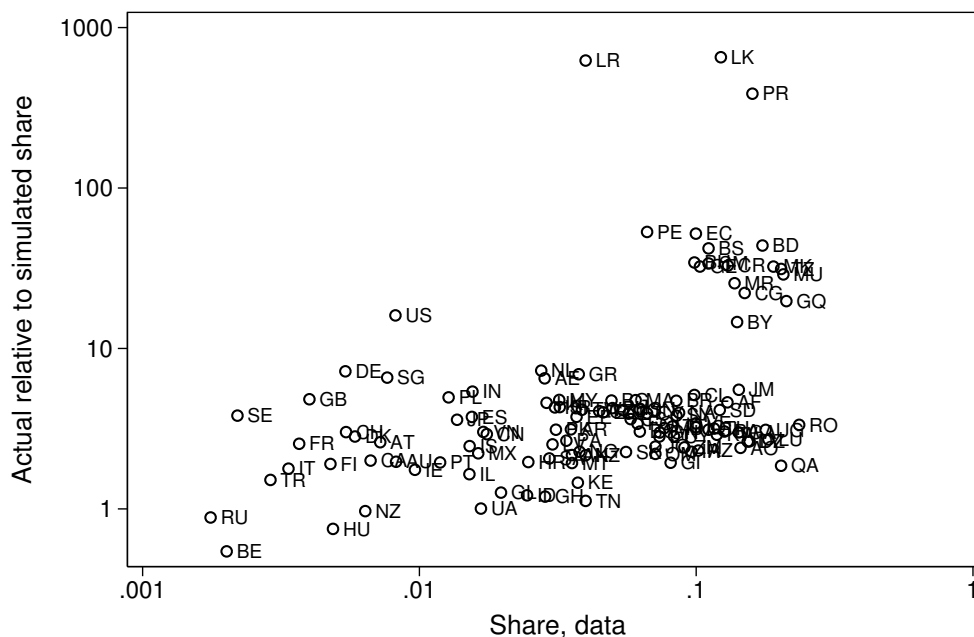
Fact 6: Firms tend to follow a hierarchical pecking order in their choice of connections. One feature of the model presented in the next section is that firms obey a hierarchy, or so called pecking order. Specifically, an exporter selling to the second ($k+1^{\text{th}}$) most connected buyer will also sell to the (k^{th}) most connected buyer. This mirrors the property of other models with heterogeneity in productivity and fixed costs, but at a different level of aggregation. We investigate the pervasiveness of hierarchies following a procedure similar to Eaton et al. (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 buyer 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 dataset.

Figure 7 shows the actual shares of firms following the hierarchy on the horizontal axis and the actual relative to the simulated shares under the assumption of independence on the vertical axis. For the vast majority of countries, there are more firms following the hierarchy relative to the statistical benchmark (the value on the vertical axis is higher than one for most countries). According to our model, all firms follow a strict hierarchy, which is clearly refuted by the data. Extending the model with randomness in relation-specific costs or revenues would allow for deviations from the strict hierarchy in the model.

¹¹Dragusanu (2014) and Benguria (2014) also find evidence of positive assortivity on the intensive margin.

¹²In 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 sorting in the network of trading *countries*; i.e.

Figure 7: Pecking order hierarchy across buyers.



Note: 2006 data. All destination markets with more than 20 sellers and buyers are included. The shares on the horizontal axis represent the number of exporters selling only to the top buyer, the top and second top buyer, and so on, relative to the number of exporters in that destination. The vertical axis represents the actual shares relative to the simulated shares under the assumption that connection probabilities are independent ($\sum_{i=1}^B p_i$). Axes on log scales.

3.2 Robustness

The basic facts presented here show empirical regularities between buyers and sellers irrespective of which product is traded. Firms with many customers are typically firms selling many products. This might suggest a framework where firms meet new buyers by expanding product scope rather than by overcoming match-specific fixed costs which is the mechanism in the theoretical model in the next section. A simple way to control for the product dimension is to re-calculate the facts with the firm-product instead of the firm as the unit of analysis.¹³ The qualitative evidence from the facts reported above remains robust to this change. For example, the distribution of the number of buyers per firm-product combination is approximately Pareto (Fact 3) and firm-products selling to many customers match on average with less connected buyers (Fact 5). These findings suggest that the basic facts cannot be explained by variation in the product dimension alone.

Our theoretical model is based on the assumption that intermediate goods are differentiated products, whereas products in the data are a mix of homogeneous and differentiated goods. We

highly connected countries, in terms of trading partners, tend to attach to less connected countries.

¹³A product is defined as a HS1996 6 digit code. Results available upon request.

therefore re-calculate the facts above for differentiated products only. Specifically, we drop all products that are classified as “reference priced” or “goods traded on an organized exchange” according to the Rauch classification.¹⁴ The qualitative evidence from the facts section remains robust to this change. A different concern is that the data includes both arm’s length trade and intra-firm trade, whereas our model is about arm’s length trade exclusively. We therefore drop all Norwegian multinationals from the dataset and recalculate the facts.¹⁵ Again, the evidence is robust to this change.

The data used in this paper is the universe of non-oil merchandise exports. A subset of the exporters are outside manufacturing and a potential concern is that the model is less relevant for trade intermediaries. We match the customs data to the manufacturing census, which allows us to remove exporters outside manufacturing. The qualitative evidence from the facts reported above remains robust to this change.¹⁶

An additional concern is that Norway may somehow be unusual and the facts are not found elsewhere. In Appendix I, we test the external validity of our results using import data from Colombia that has the similar buyer-seller information to that in the Norwegian data. We find that the basic facts also hold in the Colombian data..

Finally, one may question if the basic facts presented above can be generated from a simple stochastic process where buyers and sellers meet randomly. If so, a theory for the relationship between exporters and importers may seem superfluous. We investigate this in Appendix Section H, where we simulate a balls and bins model of trade similar to Armenter and Koren (2013). The main finding is that a random model fails to explain key empirical characteristics of exporter-importer connections.

4 A Trade Model with Two-Sided Heterogeneity

In this section, we develop a multi-country trade model with networks of heterogeneous sellers and buyers. 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) who 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. The key ingredient in our model

¹⁴The Rauch classification is concorded from SITC rev. 2 to 6 digit HS 1996 using conversion tables from the UN (<http://unstats.un.org/unsd/trade>).

¹⁵The trade transactions themselves are not identified as intra-firm or arm’s length. Norwegian multinationals account for 38 percent of the total value of Norwegian exports.

¹⁶The export value for non-manufacturing firms is 9 percent relative to total exports in 2006. Detailed results available upon request.

is heterogeneity in efficiency that in turn gives rise to heterogeneity in size both among sellers and buyers. However, two-sided heterogeneity in size could potentially also arise from other sources, e.g. differences in endowments among buyers and differences in quality among sellers. The significant testable implications from such alternative models would not depart much from the current setup.

We let the model be guided by the descriptive evidence and basic facts on sellers and buyers and their relationships as presented above. In particular, buyer and seller productivities are Pareto distributed, which gives rise to high levels of concentration in trade both on the supply and demand side, as well as Pareto distributed degree distributions (number of customers per firm and number of firms per customer), consistent with Facts 2 and 3. Due to the presence of a buyer-seller match-specific fixed cost, more efficient exporters connect with more buyers, consistent with Fact 4. This in turn leads to negative sorting, so that well-connected exporters on average connect to customers that are less well-connected, consistent with Fact 6.

4.1 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 good sector characterized by perfect competition, a traded intermediates sector and a non-traded final goods sector; the two last sectors are characterized by monopolistic competition. Workers are employed in the production of the homogeneous good as well as the production of the intermediates.¹⁷ 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 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 Pareto distribution with support $[z_L, \infty)$ and shape parameter $\gamma > \sigma - 1$, so that $F(z) = 1 - (z_L/z)^\gamma$. As a notational convention, lower case symbols refer to intermediate producers whereas upper case 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

¹⁷Adding workers to the final goods sector would only add more complexity to the model, without generating new insights.

to them,

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

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 $\Omega_j(v)$ 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 Appendix B).

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 the search for suppliers, bureaucratic procedures, contract agreements and costs associated with sellers customizing their output to the requirements of particular buyers.¹⁸ 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, production networks are the result of intermediate firms that endogenously choose their set of customers.

There are exogenous measures of buyers and sellers, N_i and n_i , in each country i . 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 i , Y_i , is then $w_i(1 + \psi)L_i$, where ψ is the dividend per share of the global mutual fund.

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 .¹⁹

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 4.3. Symmetrically, we define $\underline{z}_{ij}(Z)$ as the lowest efficiency for a seller that

¹⁸Kang et al. (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 months buffering in stock. However, Dell did not offer any guarantee on purchasing volumes due to high uncertainty in final product markets.

¹⁹We normalize $\tau_{ii} = 1$ and impose the common triangular inequality, $\tau_{ik} \leq \tau_{ij}\tau_{jk} \forall i, j, k$.

generates a profitable match for a buyer in country j with productivity Z . By construction, $z_{ij}(Z)$ is the inverse of $\underline{Z}_{ij}(z)$, i.e. $Z = \underline{Z}_{ij}(z_{ij}(Z))$.

Pricing. As intermediates and final goods markets are characterized by monopolistic competition, prices are a constant mark-up 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 mark-up.²⁰ 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 . Note that the restriction of identical elasticities of substitution across final and intermediate goods also implies that the mark-up \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

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

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

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

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 4.3 and Appendices A-B.

4.2 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: (i) buyers that match with a subset of the sellers, and (ii) buyers that match with every seller. Case (i) is characterized by $z_{ij}(Z) > z_L$, while case (ii) is characterized by $z_{ij}(Z) \leq z_L$.

The discontinuity of the Pareto distribution at z_L is inconvenient, as the sorting function $z_{ij}(Z)$ will be non-smooth (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

²⁰Because marginal costs are constant, the optimization problem of the firm of finding the optimal price and the optimal measure of buyers simplifies to standard constant mark-up pricing and a separate problem of finding the optimal measure of buyers.

lower productivity threshold is associated with more potential firms.²¹ 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 and 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 $z_H \equiv \underline{z}_{ij}(1)$. Hence, sellers with $z \geq z_H$ meet every buyer in the market.

4.3 Equilibrium Sorting

Based on the setup presented in Section 4.1, 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 4.1, 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 is willing to sell to. Hence, we find $\underline{Z}_{ij}(z)$ by solving for $\pi_{ij}(z, Z) = 0$. Inserting the demand equation (2) and a firm's optimal price, we can express $\underline{Z}_{ij}(z)$ implicitly as

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

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 Appendices A-B, 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:²²

$$\underline{Z}_{ij}(z) = \frac{\tau_{ij} w_i \Omega_j}{z} (w_i f_{ij})^{1/(\sigma-1)} \left(\frac{Y_j}{N_j} \right)^{-1/\gamma}, \quad (4)$$

where

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

and κ_3 is a constant.²³ We plot the matching function $\underline{Z}_{ij}(z)$ in Figure 8.²⁴ $\underline{Z}_{ij}(z)$ is downward sloping in z , so more efficient sellers match with less efficient buyers on the margin. A firm with


²¹ n'_i is constant as $z_L \rightarrow 0$. The normalization is similar to Oberfeld (2013).

²²The sorting function in equation (4) is valid under any distribution for buyer productivity, i.e. it is not necessary to assume Pareto buyer productivity to derive this particular result.

²³ $\kappa_3 = \mu(\Gamma - \gamma) / \Gamma$.

²⁴The Figure is based on parameter values $\tau_{ij} w_i \Omega_j (w_i f_{ij})^{1/(\sigma-1)} (Y_j / N_j)^{-1/\gamma} = 5$.

Figure 8: Matching function.



efficiency z matches with lower efficiency buyers whenever variable or fixed trade costs (τ_{ij} and f_{ij}) are lower (the curve in Figure 8 shifts towards 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 multi-country in that matching costs, variable trade costs, and wages in other source countries affect the buyer cutoff. A firm matches with a greater range of (lower efficiency) buyers when trade costs from third countries to j are higher (market access to j , Ω_j , is lower). Ω_j in equation (5) therefore has a similar interpretation as the multilateral resistance variable in Anderson and van Wincoop (2004).²⁵

4.4 Export Margins and Buyer Dispersion

Having determined the equilibrium sorting function between intermediate and final goods producers, we can now derive equilibrium expressions for firm-level trade and decompose trade into the extensive margin in terms of number of buyers and the intensive margin in terms of sales per buyer, leading to additional testable implications of the model.

Firm-level exports

Using (2), for a given firm with productivity $z < z_H$, we can express total firm-level intermediate exports, from country i to j across all the buyers with which the firm has matched as $r_{ij}^{TOT}(z) =$

²⁵The point z_H on the horizontal axis denotes the cutoff productivity where a seller matches with every buyer.

$N_j \int_{\underline{Z}_{ij}(z)} r_{ij}(z, Z) dG(Z)$. In Appendix C, we show that firm-level exports to market j are

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

where κ_1 is a constant.²⁶ The corresponding expression for firms with $z \geq z_H$ is shown in Appendix C. 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. The sorting function also allows us to determine marginal exports, i.e. exports to the least productive buyer. We insert equation (4) into (18) which yields

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

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 a firm with productivity $z < z_H$ in country i (see Appendix C), which yields

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

We emphasize two properties of these results. First, the elasticity of exports and of 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 trade elasticity is determined by the elasticity of substitution, σ (see Krugman (1980)). Also note that, as expected, a higher match cost f_{ij} dampens both firm exports and the number of buyers.²⁷

The second key property of these results is that the elasticity of exports and of the number of buyers with respect to demand in the destination market, Y_j , is determined by the ratio of buyer to seller heterogeneity, Γ/γ . 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 a positive shift in buyer expenditure. Consequently, a positive demand shock in a market with low heterogeneity among buyers translates into more exports than in a market with high heterogeneity among buyers. We summarize these findings in the following proposition.

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. The elasticity of firm-level exports with respect to destination country demand, Y_j , equals Γ/γ , the ratio of the buyer to seller productivity Pareto shape coefficient.*

²⁶ $\kappa_1 \equiv \sigma\Gamma/[\Gamma - (\sigma - 1)]$.

²⁷The elasticity of exports with respect to f_{ij} is $1 - \Gamma/(\sigma - 1)$, which is negative given the previous restrictions that (i) $\gamma - (\sigma - 1) > 0$ and $\Gamma > \gamma$.

In Section 5, we empirically test this prediction of the model by exploiting cross-country differences in the degree of firm size heterogeneity.

The Export Distribution

In a model without buyer heterogeneity, the export distribution inherits the properties of the productivity distribution, and with Pareto distributed productivities, the shape coefficient for the export distribution is simply $\gamma/(\sigma - 1)$. In our model with buyer heterogeneity, dispersion in the export distribution is determined by seller heterogeneity *relative to* buyer heterogeneity. To see this, for $z < z_H$ firms we calculate

$$\Pr [r_{ij}^{TOT}(z) < r_0^{TOT}] = 1 - (r_{ij}^{TOT}(z_L)/r_0^{TOT})^{\gamma/\Gamma}.$$

We summarize this in the following proposition:

Proposition 2. *For $z < z_H$, the distribution of firm-level exports from country i to country j is Pareto with shape parameter γ/Γ . Hence, while more heterogeneity in seller productivity translates into more heterogeneity in export sales, more heterogeneity in buyer productivity leads to less heterogeneity in export sales.*

The intuition for this result is the following. If buyer expenditure is highly dispersed, then purchases are concentrated in a few large buyers and most exporters will sell to them. This tends to dampen the dispersion in the number of buyers reached by different exporters. On the other hand, if buyer expenditure is less dispersed, then there are fewer large buyers in the market, and consequently higher dispersion in the number of buyers reached by different exporters.

An implication of our work is therefore that buyer dispersion plays a role in shaping the sales distribution, and consequently the firm size distribution, in a market. As documented by Luttmer, 2007 and Axtell, 2001, the Pareto distribution is a good approximation of the U.S. firm size distribution, although the results here raise the question of whether this is due to underlying the productivity distribution of sellers or buyers. Our results also add to the debate on firm-level heterogeneity and misallocation of resources (see e.g. Hsieh and Klenow (2009)). Hence, the variation in the strength of the link between productivity and size across countries, industries and over time reported by Bartelsman et al. (2013) may not only be the result of policy-induced distortions, but also due to differences in buyer distributions across markets.

Firm-level performance

In the model, falling trade barriers lower marginal costs among final goods firms by reducing the cost of inputs and by facilitating more matches between input and final goods producers. Specifically, as shown in Appendix A equation (15), the marginal cost of a final goods producer in country j

is inversely proportional to the market access term Ω_j . The importance of intermediate inputs for productivity growth has strong empirical support, e.g. 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 can generate firm-level responses to trade cost shocks that are consistent with the empirical evidence. Moreover, our work suggests that measured productivity gains can arise not only from falling costs or access to higher quality inputs, but also from gaining access to new suppliers.

4.5 Aggregate Relationships

We now proceed to derive expressions for total trade and welfare. Aggregate trade from i to j is

$$X_{ij} = n_i N_j \int_1 \int_{z_{ij}(Z)} r_{ij}(z, Z) dF(z) dG(Z).$$

Solving the integrals, the trade share $X_{ij} / \sum_k X_{kj}$ is²⁸

$$\frac{X_{ij}}{\sum_k X_{kj}} = \frac{n'_i (w_i f_{ij})^{1-\gamma/(\sigma-1)} (\tau_{ij} w_i)^{-\gamma}}{\sum_k n'_k (w_k f_{kj})^{1-\gamma/(\sigma-1)} (\tau_{kj} w_k)^{-\gamma}}. \quad (9)$$

We emphasize two implications for aggregate trade. First, the relation-specific cost f_{ij} dampens aggregate trade with a partial elasticity $1 - \gamma / (\sigma - 1) < 0$. Hence, the presence of the relation-specific cost has macro implications for trade flows. Second, the partial aggregate trade elasticity with respect to variable trade barriers, $\partial \ln X_{ij} / \partial \ln \tau_{ij}$, is $-\gamma$, the Pareto coefficient for seller productivity. This result mirrors the finding in models with one-sided heterogeneity, e.g. Eaton et al. (2011). Hence, our model produces similar macro trade elasticities compared to models with one-sided heterogeneity while being able to explain a range of new facts at the micro level. It may seem surprising that the aggregate trade elasticity is γ , given that the firm-level elasticity is Γ . This occurs because the aggregate elasticity is the weighted average of firm-level elasticities for $z < z_H$ firms and $z \geq z_H$ firms. These elasticities are Γ and $\sigma - 1$ respectively (see Appendix C). In equilibrium, the weighted average of the two is γ .²⁹

Real wages in our model are

$$\frac{w_j}{Q_j} = \kappa_6 \left(n'_j N_j \right)^{1/\gamma} \left(\frac{f_{jj}}{L_j} \right)^{-\gamma_2 / [\gamma(\sigma-1)]} \frac{\pi_{jj}^{-1/\gamma}}{\tau_{jj}}, \quad (10)$$

where κ_6 is a constant (see Appendix D).³⁰ Higher spending on home goods (higher π_{jj}) lowers real wages with an elasticity $1/\gamma$, mirroring the finding in Arkolakis et al. (2012). A potential

²⁸We can alternatively write $X_{ij} = \kappa_5 n'_i Y_j (w_i f_{ij})^{1-\gamma/(\sigma-1)} (\tau_{ij} w_i \Omega_j)^{-\gamma}$ where $\kappa_5 = \Gamma \sigma \gamma / [\gamma_2 (\Gamma - \gamma)]$.

²⁹Aggregate trade can alternatively be written $X_{ij} = n_i \int_{z_L}^{z_H} r_{ij}^{TOT}(z) dF(z) + n_i \int_{z_H}^{\infty} \tilde{r}_{ij}^{TOT}(z) dF(z)$, where $\tilde{r}_{ij}^{TOT}(z)$ is exports for $z > z_H$ firms (see Appendix C). Solving the two integrals yields exactly the same expression for X_{ij} as the equation above.

³⁰ $\kappa_6 = \left(\frac{\sigma}{\kappa_3} \frac{\gamma}{\gamma_2} \right)^{1/\gamma} \left(\bar{m}^{2(1-\sigma)} \frac{\mu}{\sigma} \right)^{1/(\sigma-1)} (1 + \psi)^{-1/\gamma + 1/(\sigma-1)}$.

complication in our context, however, is that the elasticity γ refers to the elasticity of trade with respect to variable trade barriers τ_{ij} , see equation (9). To the extent that variable trade barriers and relation-specific costs f_{ij} co-vary, it may be empirically challenging to identify γ from the data.³¹ The expression for w_j/Q_j also reveals a novel role for market size in our model. Holding the number of firms (varieties) fixed, a larger market L_j increases real wages with an elasticity $\gamma_2/[\gamma(\sigma-1)] < 1$. This occurs because larger final goods firms get access to more intermediate inputs, which in turn lowers the price index of intermediate goods, $q_i(Z)$.

4.6 Linking Facts and Theory

In presenting the model we pointed out that our theory was guided by the basic facts on buyer-seller relationships presented in Section 3.1. Before turning to testing empirical implications of the model, we revisit the basic facts and examine the extent to which the model fits them. As shown in Proposition 2, the distribution of firm-level exports from i to j is Pareto, consistent with Fact 2. Appendix C shows that the distribution of purchases by firms located in j buying from i is also Pareto, giving rise to a high degree of concentration in trade on the buyer side. Our model also has Pareto distributions of buyers per seller and sellers per buyer, consistent with Fact 3 (see expressions for $b_{ij}(z)$ and $L_{ij}(Z)$ in Appendix C).³² 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 Appendix E)

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

so that all exporters to a market j have the same Pareto distribution of sales across buyers. Fact 5 shows that highly connected exporters to market j have, on average, customers that have 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 Appendix F):

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

Hence, the elasticity is negative with a slope coefficient $-\gamma/\Gamma$. Fact 6 shows that firms are more likely to follow a hierarchy in choosing connections relative to a statistical benchmark. According to the model, firms would follow a strict hierarchy, so that the shares of firms following the hierarchy in Figure 7 would be one.

³¹For example, if f_{ij} and τ_{ij} are positively correlated and f_{ij} is not controlled for, then the estimate of γ will be biased upwards and the gains from trade will be biased towards zero.

³²The 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 some random matching to the model would allow the theoretical cdfs to more closely align with the empirical cdfs.

We also revisit the empirical relationship between the margins of trade and market characteristics and link them to the model. According to Fact 1 and Table 4 a firm’s number of customers is increasing in GDP and decreasing in distance. As displayed in equation (8), the model predicts that the number of buyers per firm increases with market size and falls with trade costs, with elasticities Γ/γ and $-\Gamma$ respectively.

5 Empirical Implications

In this section, we test three main predictions of the model developed above that emphasize the importance of buyer heterogeneity in explaining trade patterns. The first prediction is that a positive demand shock (an increase in Y_j) facing firm m should raise firm-level exports, but the marginal export flow, i.e. the firm’s transaction to the smallest buyer, should remain unchanged as the marginal transaction is pinned down by the magnitude of the relation-specific fixed cost. The second prediction is that a similar-sized positive demand shock facing firm m across different destinations should translate into relatively higher sales in markets with less heterogeneity, as stated in Proposition 1. The third prediction is that heterogeneity in sales across exporters is not only driven by heterogeneity in exporter productivity, but inversely related to importer heterogeneity, as stated in Proposition 2.

5.1 A Measure of Demand

We start by calculating a measure of firm-destination specific demand. The objective is to create a variable that proxies for market size in the destination country, Y_j in equation (6). In addition, we would like the variable to be firm-specific, so that we can control for market-wide factors that may also impact sales by including fixed effects that vary at the destination level by year.

We therefore choose to proxy for the demand facing Norwegian firm m in destination country j for all its exported products by calculating total imports in j of those products from sources other than Norway. Given the small market share of Norwegian firms in most destinations, this measure should be exogenous with respect to firm m ’s exports. We proceed by using product-level (HS6 digit) trade data from COMTRADE and denote total imports of product p to country j at time t from all sources except Norway as I_{pjt} .³³ The firm-level demand shock d_{mjt} in market j at time t is then defined as the unweighted average of imports for the products that firm m is exporting

$$d_{mjt} = \frac{1}{N_m} \sum_{p \in \Omega_m} \ln I_{pjt},$$

where Ω_m is the set of products firm m is exporting (to any country in any year), and N_m is the

³³We use CEPII’s BACI database using the HS 1996 revision.

number of products firm m is exporting.³⁴ We also investigate the robustness of our results to other specifications of demand. These are discussed in Section 5.3.

5.2 Demand Shocks and the Marginal Buyer

According to the model (see Section 4.4), a positive demand shock in market j will increase total firm-level exports and the number of buyers, but will have no impact on sales to the marginal buyer. This occurs because the gross profits associated with the marginal buyer exactly equals the buyer-seller match fixed cost. To test this prediction we let equations (6)-(8) guide us, and estimate

$$\ln y_{mjt} = \alpha_{mt} + \beta_{jt} + \eta \ln d_{mjt} + \epsilon_{mjt}, \quad (11)$$

where y_{mjt} is an outcome variable for firm m in market j at time t and d_{mjt} is the demand shock facing firm m in market j . We include both firm-year (α_{mt}) and country-year (β_{jt}) fixed effects, allowing for changes in time-varying firm-specific factors such as productivity, and time-varying market-wide shocks, e.g. the real exchange rate. We estimate the model for total firm-level exports ($\sum_b y_{mbjt}$), number of buyers, the firm's marginal export ($\min_b y_{mbjt}$ where b denotes buyer), and exports to the firm's median buyer ($median_b y_{mbjt}$).

Identification then comes from comparing growth in exports within the same firm across markets, while controlling for country-specific trends. Our approach resembles a triple differences model as we compare growth in exports both across markets and across firms. Specifically, for two firms A and B and two markets 1 and 2, η is identified by the difference in firm A's exports growth to markets 1 and 2, relative to the difference in firm B's exports growth in markets 1 and 2.³⁵

The results largely confirm the predictions from the model. Table 7 shows that total exports and the number of buyers per firm (columns 1 and 2) are positively and significantly related to positive demand shocks in the destination country. As predicted by the model, positive demand shocks have no impact on the marginal export flow (column 3).

However, exports to the median buyer (column 4) are increasing in firm-level demand shocks while the model predicts that the distribution of exports across buyers would be unchanged.³⁶ In addition, the model predicts that the elasticity of exports to a demand shock is identical to the elasticity of the number of customers to a demand shock, 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

³⁴ Ω_m is the same in all destinations and in all years, so that firm behavior across time and countries does not change the set. A few importer-product pairs are missing in one or more years, these pairs are dropped.

³⁵The fixed effects α_{mt} and β_{jt} are differenced out for $\Delta \ln y_{mjt} - \Delta \ln y_{mj-1,t} - (\Delta \ln y_{jt} - \Delta \ln y_{j-1,t})$.

³⁶In the min and median exports regressions (columns (3) and (4)), we only use firms with more than 5 customers. The sample is also restricted to countries with information about firm size dispersion from the World Bank Enterprise Surveys, so that the sample size is identical to the sample size in the regressions in Section 5.3. Results based on the entire sample are not significantly different.

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. Another possible reason for this discrepancy and the positive coefficient for exports to the median buyer is that the empirical productivity distributions of buyers and sellers may deviate from the assumed Pareto shape.

5.3 Demand Shocks and Importer Heterogeneity

One of the main features of the theoretical framework is the role of buyer-side heterogeneity in determining the response of firm exports to demand shocks, i.e. that the demand shock elasticity is greater in markets with less buyer heterogeneity. Hence, we would expect a similar-sized demand shock facing firm m across different destinations to translate into relatively greater changes in sales for markets with less heterogeneity, as stated in Proposition 1. We test this prediction by amending the model in equation (11) by including an interaction term for buyer dispersion, allowing us to check whether the demand elasticities are higher in markets with less heterogeneity. Specifically, we estimate

$$\ln y_{mjt} = \alpha_{mt} + \beta_{jt} + \eta_1 \ln d_{mjt} + \eta_2 \ln d_{mjt} \times \Theta_j + \epsilon_{mjt}, \quad (12)$$

where Θ_j is a measure of buyer dispersion in destination market j .

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.³⁷ We therefore use data on the firm size distribution from the World Bank’s Enterprise Surveys, and calculate a Pareto slope coefficient (Θ^1), the 90/10 percentile ratio (Θ^2), and the standard deviation of log employment for each country (Θ^3).³⁸ The Enterprise Surveys are firm-level surveys of a representative sample of an economy’s private sector (manufacturing and services) including companies in the formal sector with 5 or more employees.³⁹ The survey aims to achieve cross-country comparisons so that our dispersion measures should not be contaminated by differences in sampling design.

The results from estimating the specification in (12) are shown in Table 8. We find that the elasticity for both export value and the number of buyers is significantly dampened in markets with more heterogeneity, consistent with the predictions of our model. Note that the coefficients for the interaction term are positive rather than negative in columns (1) and (2) since the Pareto coefficient

³⁷The relationship between productivity and size has also been documented in a set of studies for many of countries (see e.g. Bartelsman et al. (2013) for recent evidence).

³⁸We calculate the Pareto slope coefficient by regressing 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.

³⁹The survey covers 87 countries, mostly developing countries. In 2006 these countries received 29 percent of Norwegian exports. We drop countries where the survey has fewer than 100 observations per country. These countries are: Brazil, Eritrea, Guyana, Jamaica, Lebanon, Lesotho, Montenegro, Oman and Turkey.

is inversely related to dispersion. The magnitudes are also economically significant: Moving from the 25th to the 75th percentile of the Pareto coefficient Θ^1 increases the demand elasticity, $\eta_1 + \eta_2 \Theta^1$, by 11 percent, suggesting that demand-side factors are quantitatively important for our understanding of trade elasticities.⁴⁰

Robustness

In this section, we perform a number of robustness checks. First, a concern is that Norwegian exports to countries included in the Enterprise Surveys only amount to 1/3 of total exports. We therefore check the robustness of our results by using alternative data sources on dispersion, allowing us to include other destination countries in the sample.

An alternative data source providing information on firm size dispersion is Bureau van Dijk's Orbis database, which has information on over 100 million private companies across the world.⁴¹ Unfortunately, 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 50 employees. We calculate Pareto coefficients for firm employment, as in the baseline case, for all countries with 1000 or more Orbis firms. In total, this gives us information on buyer dispersion for 48 countries, covering 89 percent of Norwegian exports (based on 2006 values).⁴² The estimates in columns (1) and (2) of Table 9 show that using Orbis produces remarkably similar results to those reported for the baseline case in Table 8 even though the sample of countries (and firms) is quite different.

The World Bank's Exporter Dynamics database provides data on exports for 39 countries.⁴³ Unfortunately, the Exporter Dynamics database does not include firm-level information on firm size or on exports, but it does provide the mean and standard deviation of exports across firms. This allows us to calculate the coefficient of variation for exporters for all 39 countries, which we use as our measure buyer dispersion. A potential concern is that this measure of buyer dispersion is inferred from the exports distribution. However, as our buyers are importers, and as importers themselves

⁴⁰The 25th and 75th percentiles of Θ^1 are 0.58 and 0.80, so that the demand elasticities are 0.41 and 0.46 respectively.

⁴¹<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.

⁴²The 48 countries are Argentina, Austria, Australia, Bosnia and Herzegovina, Belgium, Bulgaria, Brazil, Belarus, Canada, Switzerland, Germany, Denmark, Estonia, Egypt, Spain, Finland, France, United Kingdom, Greece, Hong Kong, Croatia, Hungary, Ireland, India, Italy, Japan, Korea, Sri Lanka, Lithuania, Latvia, Morocco, Macedonia, Mexico, Netherlands, Peru, Poland, Portugal, Romania, Serbia, Russia, Sweden, Slovenia, Slovakia, Tunisia, Turkey, Ukraine, United States and South Africa.

⁴³See Cebeci et al. (2012) for details on the data set. In 2006, the countries for which the database provide information received 20 percent of Norwegian exports. The countries included are Albania, Bangladesh, Belgium, Burkina Faso, Bulgaria, Brazil, Botswana, Chile, Cameroon, Costa Rica, Dominican Republic, Ecuador, Estonia, Egypt, Spain, Guatemala, Iran, Jordan, Kenya, Cambodia, Laos, Morocco, Macedonia, Mali, Mauritius, Malawi, Mexico, Nicaragua, New Zealand, Peru, Pakistan, Sweden, Senegal, El Salvador, Turkey, Tanzania, Yemen and South Africa.

tend to be exporters (Bernard et al., 2007), there should be a strong positive correlation between imports and exports dispersion. In fact, we can estimate this correlation using the Norwegian data, and we do indeed find a strong positive correlation. We refer the reader to Appendix G for more information. We estimate equation (12) using the the calculated coefficient of variation (Θ^5). Columns (3) and (4) in Table 9 show that the same pattern of significance holds in this case, although the magnitudes are not directly comparable due to the different measures of dispersion.

A second concern is that buyer dispersion may be correlated with other factors that also affect the demand elasticity; for example both buyer dispersion and demand elasticities may be different in low-income countries. We address this issue by purging GDP per capita from our Pareto shape coefficient Θ^1 . Specifically we regress Θ^1 on GDP per capita and use the fitted residual, Θ^6 . The results are reported in columns (5) and (6) in Table 9. Overall the results are very similar to the baseline case in Table 8. A third concern is that the demand shock variable d_{mjt} may suffer from measurement error, as imports may not fully capture demand facing Norwegian firms. As a simple test, we instead replace d_{mjt} with GDP_{jt} as our proxy for demand. In this case, we cannot include country-year fixed effects but we do include country fixed effects and a real exchange rate control variable. The results in columns (7) and (8) in Table 9 show the same pattern as in the baseline case, although the standard errors are somewhat higher.

In sum, we confirm one of the main predictions of the model: Export markets with more homogeneous buyer distributions have greater elasticities for both exports and the number of buyers than do markets with more heterogeneous firm distributions.

5.4 Sorting among Importers and Exporters

According to Proposition 2, more importer heterogeneity is associated with less exporter heterogeneity. We test this prediction exploiting variation in dispersion across countries and industries in our data. Specifically, we ask whether increased buyer dispersion in a market is correlated with less dispersion among Norwegian exporters serving that market.

We proceed by defining a market as a country-industry combination, where an industry is defined as a unique 2 digit HS code. To measure buyer dispersion in a given market we again use the World Bank's Exporter Dynamics Database. The database contains information about dispersion in exports for different countries and HS 2 digit industries. We proceed as in the previous section, and calculate the coefficient of variation Θ_{jp}^7 , for each country j and each 2-digit HS industry p . Moreover, we calculate corresponding measures of dispersion for Norwegian exporters, y_{jp} , serving market j for industry p . The model we estimate is then

$$y_{jp} = \alpha_j + \delta_p + \mu\Theta_{jp}^7 + \epsilon_{jp},$$

where α_j and δ_p are country and industry fixed effects, and all variables are measured in logs.

This is a differences-in-differences model, where identification comes from comparing differences in dispersion across industries within a country (first difference) across different countries (second difference). Country-specific variation in dispersion will be differenced out by α_j , while industry-specific variation in dispersion will be differenced out by δ_p .

There are two potential concerns with the chosen approach. First, buyer dispersion is inferred from the value of exports, but as discussed above (see Section 5.3 and Appendix G), import dispersion is highly correlated with export dispersion. Second, as dispersion is measured per industry, we implicitly make the assumption that the buyers of goods in an industry (e.g. beverages) are themselves exporting in the same industry. Although this certainly does not hold perfectly, we know from input-output tables that the sourcing matrix is dominated by the diagonal, i.e. that industries tend to source a significant share of their intermediates from themselves.⁴⁴

Table 10 presents the results. As our dispersion measures can be calculated for each year, each column shows the results for a different annual cross-section. For a industry-destination pair to be included in the sample, we must choose a threshold for how many Norwegian firms are exporting to jp and how many foreign firms are exporting from jp . We set a threshold of 30 or more firms in the top panel and 50 or more firms in the bottom panel.

Focusing on the top panel, in all years except 2006, the estimates show that more buyer dispersion is significantly associated with less seller dispersion. The magnitudes are also quantitatively large; a one percent increase in buyer dispersion leads to a 0.3-0.7 decrease in seller dispersion. The bottom panel shows the results in the 50 firm threshold case. This decreases the number of country-industry pairs included in the sample and increases the standard errors, but the magnitudes are largely unchanged.

6 Conclusion

We use highly disaggregated trade transaction data from Norway to explore the role of buyers (importers) in international trade. We find that the extensive margin of the number of buyers plays an important role in explaining variation in exports in the aggregate and at the firm level. The importer margin is comparable in magnitude to previously documented extensive margins of trade of exporters, destinations and products.

We introduce a series of basic facts about buyer-seller relationships in international trade which point to extreme concentration of exports across both sellers and buyers, distinct differences in the degree of dispersion of buyer expenditures across destinations, 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

⁴⁴See the discussion of input-output linkages in Caliendo and Parro (2012).

destination, and that there is negative degree assortivity in the exporter-importer matches. In other words, large exporters on average reach importers who buy from a relatively smaller number of Norwegian firms.

Guided by these facts, we develop a parsimonious multi-country model of heterogeneous exporters and importers where matches are subject to a relation-specific fixed cost. This framework yields interesting new testable implications on the importance of buyer-side heterogeneity that are largely confirmed by our empirical analysis. An increase in foreign demand increases firm-level exports but the marginal export flow does not change as it is pinned down by the magnitude of the relation-specific fixed cost. The response of firm-level exports to comparable demand shocks across destinations varies systematically with the dispersion of expenditures. Specifically, the export response is amplified in destinations with less buyer dispersion. Finally, we provide evidence supporting the theoretical prediction that more buyer dispersion in a market is associated with less dispersion in exports to that market.

The results suggest that demand-side characteristics play an important role in determining the firms' export response to shocks, and more broadly that relation-specific costs help us understand the micro and macro structure of international trade. Future research might fruitfully focus on the growth and stability of these exporter-importer networks as well as the sources of heterogeneity in buyer expenditure itself.

References

- Alfaro, L. and M. X. Chen (2013). Market reallocation and knowledge spillover: The gains from multinational production. Harvard business school working paper, Harvard Business School. 41
- Alfonso Irarrazabal, A. M. and L. D. Opmolla (2013). The margins of multinational production and the role of intrafirm trade. *Journal of Political Economy* 121(1), 74–126. 2
- Amiti, M. and J. Konings (2007). Trade liberalization, intermediate inputs and productivity: Evidence from indonesia. *American Economic Review* 97(5), 1611–1638. 1, 4.4
- Anderson, J. E. and E. van Wincoop (2004, September). Trade costs. *Journal of Economic Literature* 42(3), 691–751. 4.3
- Antràs, P. and A. Costinot (2011). Intermediated trade. *The Quarterly Journal of Economics* 126(3), 1319–1374. 1
- Arkolakis, C. (2010). Market access costs and the new consumers margin in international trade. *Journal of Political Economy* 118(6), 1151–1199. 1
- Arkolakis, C. (2011). A unified theory of firm selection and growth. Working paper, NBER Working Paper 17553. 1
- Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012). New trade models, same old gains? *American Economic Review* 102(1), 94–130. 1, 4.5
- Armenter, R. and M. Koren (2013). A balls-and-bins model of trade. *forthcoming in American Economic Review*. 3.2
- Axtell, R. L. (2001). Zipf distribution of us firm sizes. *Science* 293, 1818–1820. 4.4
- Bartelsman, E. J., S. Scarpetta, and F. Schivardi (2013, February). Cross-country differences in productivity: The role of allocation and selection. *American Economic Review* 103(1), 305–334. 1, 4.4, 37
- Benguria, F. (2014). Production and distribution of international trade: Evidence from matched exporter-importer data. Technical report. 1, 11
- Bernard, A. B. and S. Dhingra (2014). Contracting and the division of the gains from trade. Working paper, Tuck School of Business. 47
- Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2007, Summer). Firms in international trade. *Journal of Economic Perspectives* 21(3), 105–130. 5.3

- Bernard, A. B., J. B. Jensen, and P. K. Schott (2009). Importers, exporters, and multinationals: A portrait of firms in the U.S. that trade goods. In T. Dunne, J. Jensen, and M. Roberts (Eds.), *Producer dynamics: New Evidence from Micro Data*. University of Chicago Press. 1
- Bernard, A. B., A. Moxnes, and Y. Saito (2014). Geography and firm performance in the Japanese production network. Discussion Paper 14-E-033, RIETI. 3.1
- Blum, B. S., S. Claro, and I. Horstmann (2010). Facts and figures on intermediated trade. *American Economic Review* 100(2), 419–23. 1, 1
- Blum, B. S., S. Claro, and I. J. Horstmann (2012). Import intermediaries and trade costs: Theory and evidence. mimeo, University of Toronto. 1, 1
- Boler, E. A., A. Moxnes, and K. H. Ulltveit-Moe (2012). Technological change, international sourcing and the joint impact on productivity. CEPR Discussion Paper 8884. 1
- Braguinsky, S., L. G. Branstetter, and A. Regateiro (2011). The incredible shrinking Portuguese firm. NBER Working Papers 17265, National Bureau of Economic Research. 1
- Caliendo, L. and F. Parro (2012). Estimates of the trade and welfare effects of NAFTA. NBER Working Papers 18508, National Bureau of Economic Research. 44
- Carballo, J., G. I. Ottaviano, and C. V. Martinicus (2013). The buyer margins of firms' exports. CEPR Discussion Paper 9584. 1, 1
- Castellani, D., F. Serti, and C. Tomasi (2010, March). Firms in international trade: Importers' and exporters' heterogeneity in Italian manufacturing industry. *The World Economy* 33(3), 424–457. 1
- Cebeci, T., A. Fernandes, C. Freund, and M. D. Pierola (2012). Exporter dynamics database. World Bank Policy Research Working Paper 6229, World Bank. 2, 43
- Chaney, T. (2008). Distorted gravity: the intensive and extensive margins of international trade. *American Economic Review* 98(4), 1707–1721. 4.1
- Chaney, T. (forthcoming). The network structure of international trade. *American Economic Review*. 1
- Dragusanu, R. (2014). Firm-to-firm matching along the global supply chain. 1, 11
- Eaton, J., M. Eslava, D. Jenkins, C. J. Krizan, and J. Tybout (2012). A search and learning model of export dynamics. working paper. 1, 1

- Eaton, J., S. S. Kortum, and F. Kramarz (2011, September). An anatomy of international trade: Evidence from French firms. *Econometrica* 79(5), 1453–1498. 1, 3.1, 4.5
- Garicano, L., C. LeLarge, and J. van Reenen (2013). Firm size distortions and the productivity distribution: Evidence from France. NBER Working Papers 18841, National Bureau of Economic Research. 1
- Goldberg, P. K., A. K. Khandelwal, N. Pavcnik, and P. Topalova (2010). Imported intermediate inputs and domestic product growth: Evidence from India. *The Quarterly Journal of Economics* 125(4), 1727–1767. 1, 4.4
- Halpern, L., M. Koren, and A. Szeidl (2011). Imported inputs and productivity. Cefig working paper, Central European University. 1
- Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing tfp in China and India. *Quarterly Journal of Economics* 124, 1403–1448. 1, 4.4
- Jackson, M. O. and B. W. Rogers (2007). Meeting strangers and friends of friends: How random are social networks? *American Economic Review* 97(3), 890–915. 3.1, 12
- Kang, M.-P., J. Mahoney, and D. Tan (2009). Why firms make unilateral investments specific to other firms: the case of OEM suppliers. *Strategic Management Journal* 30, 117–135. 18
- Khandelwal, A. K. and P. Topalova (2011). Trade liberalization and firm productivity: The case of India. *Review of Economics and Statistics* 93(3), 995–1009. 1, 4.4
- Krugman, P. (1980). Scale economies, product differentiation and the pattern of trade. *American Economic Review* 70(5), 950–959. 4.4
- Luttmer, E. G. J. (2007). Selection, growth and the size distribution of firms. *The Quarterly Journal of Economics* 122(3), 1103–1144. 4.4
- Mayer, T. and G. Ottaviano (2008, May). The happy few: The internationalisation of European firms. *Intereconomics: Review of European Economic Policy* 43(3), 135–148. 2
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6), 1695–1725. 4
- Monarch, R. (2013). It’s not you, it’s me: Breakups in U.S.-China trade relationships. 1
- Muuls, M. and M. Pisu (2009, May). Imports and exports at the level of the firm: Evidence from Belgium. *The World Economy* 32(5), 692–734. 1

- Oberfeld, E. (2013). Business networks, production chains, and productivity: A theory of input-output architecture. Working paper. 21
- Petropoulou, D. (2011). Information costs, networks and intermediation in international trade. Globalization and Monetary Policy Institute Working Paper 76, Federal Reserve Bank of Dallas. 1
- Rauch, J. E. (1999, June). Networks versus markets in international trade. *Journal of International Economics* 48(1), 7–35. 1
- Rauch, J. E. and J. Watson (2004). Network intermediaries in international trade. *Journal of Economics & Management Strategy* 13(1), 69–93. 1
- Serrano, M. A. and M. Boguna (2003). Topology of the world trade web. *The Accounting Review* 68(1). 12
- Sugita, Y., K. Teshima, and E. Siera (2014). Assortative matching of exporters and importers. 1, 3.1

Table 1: Top HS 8 digit products by # exporters.

HS code	Description	Share of exporters, %
84799090	Subgroup of: 847990 Parts of machines and mechanical appliances n.e.s.	9.1
84733000	Parts and accessories for automatic data-processing machines or for other machines of heading 8471, n.e.s.	7.6
73269000	Articles of iron or steel, n.e.s. (excl. cast articles or articles of iron or steel wire)	5.8
39269098	Subgroup of: 392690 Articles of plastics or other materials of headings 3901 to 3914, for civil aircraft, n.e.s	4.9
84099909	Subgroup of: 840999 Parts suitable for use solely or principally with compression-ignition internal combustion piston engine "diesel or semi-diesel engine", n.e.s	4.4

Note: 2006 data. HS codes refer to 2006 edition. Oil and gas exports excluded (HS 27x products).

Table 2: Top HS 8 digit products by value.

HS code	Description	Share of value, %
76012001	Subgroup of: 760120 Unwrought aluminium alloys	9.9
03021201	Subgroup of: 030212 Fresh or chilled Pacific salmon	5.1
75021000	Nickel, not alloyed, unwrought	4.8
89069009	Subgroup of: 890690 Vessels, incl. lifeboats (excl. warships, rowing boats and other vessels of heading 8901 to 8905 and vessels for breaking up)	1.3
31052000	Mineral or chemical fertilisers containing the three fertilising elements nitrogen, phosphorus and potassium (excl. those in pellet or similar forms,	1.2

Note: 2006 data. HS codes refer to 2006 edition. Oil and gas exports excluded (HS 27x products).

Table 3: The margins of trade.

VARIABLES	(1) Sellers	(2) Products	(3) Buyers	(4) Density	(5) Intensive
Exports (log)	0.57 ^a (0.02)	0.53 ^a (0.02)	0.61 ^a (0.02)	-1.05 ^a (0.04)	0.32 ^a (0.02)
N	205	205	205	205	205
R ²	0.86	0.85	0.81	0.81	0.50

Note: We decompose total exports to country j , x_j , into the product of the number of trading firms, f , the number of traded products, p , the number of buyers, b , the density of trade, d , i.e. the fraction of all possible firm-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 firm-product-buyer observations for which trade with country j is positive and $\bar{x}_j = x_j / o_j$ is average exports per firm-product-buyer. We regress the logarithm of each component on the logarithm of total exports to a given market in 2006, $\ln f_j$ against $\ln x_j$. Robust standard errors in parentheses. ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$.

Table 4: Within-firm Gravity.

VARIABLES	(1)	(2)	(3)
	Exports	# Buyers	Exports/Buyer
Distance	-0.48 ^a	-0.31 ^a	-0.17 ^a
GDP	0.23 ^a	0.13 ^a	0.10 ^a
Firm FE	Yes	Yes	Yes
N	53,269	53,269	53,269
R^2	0.06	0.15	0.01

Note: 2006 data. Robust standard errors in parentheses, clustered by firm. ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$. All variables in logs.

Table 5: Descriptive statistics.

	Overall	Sweden	Germany	US	China	OECD	non-OECD
Number of exporters	18,219	8,614	4,067	2,088	725	1,588.2	98.2
Number of buyers	81,362	16,822	9,627	5,992	1,489	3,055.6	144.5
Buyers/exporter, mean	9.0	3.6	3.6	4.5	3.6	2.7	1.6
Buyers/exporter, median	2	1	1	1	1	1	1
Exporters/buyer, mean	2.0	1.9	1.5	1.6	1.7	1.5	1.2
Exporters/buyer, median	1	1	1	1	1	1	1
Share trade, top 10% sellers	.98	.94	.97	.96	.86	.90	.75
Share trade, top 10% buyers	.96	.95	.95	.97	.89	.89	.73
Log max/median exports	13.0	10.7	11.4	11.2	7.9	8.7	4.6
Log max/median imports	12.2	10.8	10.8	11.7	8.4	8.4	4.6
Share in total NO exports, %	100	11.3	9.6	8.8	2.1	81.6	18.4

Note: 2006 data. The overall column refers to outcomes unconditional on destination country. OECD and non-OECD are the unweighted means of outcomes for all countries in the two groups. Log max/median exports (imports) is the log ratio of the largest exporter (importer), in terms of trade value, relative to the median exporter (importer).

Table 6: Types of matches, %.

	(1)	(2)	(3)	(4)
	One-to-one	Many-to-one	One-to-many	Many-to-many
Share of value, %	4.6	26.9	4.9	63.6
Share of counts, %	9.5	40.1	11.0	39.4

Note: 2006 data. Column (1) refers to matches between exporters (E) and importers (I) where both have one connection in a market, column (2) refers to matches where the E has many connections and the I has one, columns (3) refers to matches where the E has one connection and the I has many, column (4) refers to matches where both E and I have many connections. The unit of observation is firm-destination, e.g. an exporter with one customer in two destinations is counted as a single-customer exporter. The first row shows the trade value for each group relative to total trade. The second row shows the number of matches in the group relative to the total number of matches.

Table 7: Firm responses to demand shocks.

	(1)	(2)	(3)	(4)
	Exports	# Buyers	Marginal buyer	Median buyer
d_{mjt}	.43 ^a	.14 ^a	.00	.45 ^a
	(.02)	(.01)	(.07)	(.05)
Country-year FE	Yes	Yes	Yes	Yes
Firm-year FE	Yes	Yes	Yes	Yes
N	105,756	105,756	8,106	8,106
Firms-years	44,068	44,068	4,055	4,055
Destinations	75	75	57	57

Note: Robust standard errors in parentheses, clustered by firm-year. ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$. All variables in logs. The dep. variables in columns (3) and (4) are the minimum (median) export value for a firm, across its buyers; $\min_b y_{mbjt}$ and $\text{median}_b y_{mbjt}$. Sample is restricted to countries with information about dispersion from the World Bank Enterprise Surveys. Only exporters with > 5 buyers in columns (3) and (4).

Table 8: Demand shocks and heterogeneity.

	(1)	(2)	(3)	(4)	(5)	(6)
	Exports	# Buyers	Exports	# Buyers	Exports	# Buyers
d_{mjt}	.30 ^a (.05)	.04 ^b (.02)	.60 ^a (.07)	.27 ^a (.03)	.70 ^a (.08)	.30 ^a (.03)
$d_{int} \times \Theta^1$ (Pareto)	.20 ^a (.08)	.15 ^a (.03)				
$d_{int} \times \Theta^2$ (P90/10)			-.04 ^b (.02)	-.03 ^a (.01)		
$d_{int} \times \Theta^3$ (Std. Dev.)					-.18 ^a (.05)	-.11 ^a (.02)
Country-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	105,756	105,756	105,756	105,756	105,756	105,756
Firms-years	44,068	44,068	44,068	44,068	44,068	44,068
Destinations	75	75	75	75	75	75

Note: Robust standard errors in parentheses, clustered by firm-year. ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$. All variables in logs. Θ^1 , Θ^2 and Θ^3 denote the interaction between the demand shock d_{int} and the Pareto shape parameter, the log firm size 90/10 percentile ratio, and the standard deviation of log employment, respectively.

Table 9: Robustness: Demand shocks and heterogeneity.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Exports	# Buyers	Exports	# Buyers	Exports	# Buyers	Exports	# Buyers
d_{mjt}	.61 ^a (.11)	.27 ^a (.05)	.22 ^a (.04)	.03 ^c (.02)	.39 ^a (.01)	.17 ^a (.01)	.06 (.26)	-.07 (.08)
$d_{mjt} \times \Theta^4$ (Pareto Orbis)	.23 ^a (.04)	.16 ^a (.02)						
$d_{mjt} \times \Theta^5$ (CV WBED)			-.11 ^a (.04)	-.06 ^a (.02)				
$d_{mjt} \times \Theta^6$ (Alt Pareto WBES)					.27 ^a (.08)	.21 ^a (.03)		
$d_{mjt}^{GDP} \times \Theta^1$ (Pareto WBES)							.50 ^c (.29)	.17 ^c (.09)
Country FE	No	No	No	No	No	No	Yes	Yes
Country-year FE	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Firm-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	296,045	296,045	90,951	90,951	103,716	103,716	102,856	102,856
Firms-years	100,895	100,895	58,939	58,939	43,757	43,757	43,732	43,732
Destinations	50	50	37	37	74	74	75	75

Note: Robust standard errors in parentheses, clustered by firm-year. ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$. All variables in logs. Θ^4 is the Pareto coefficient from Orbis data, see main text; Θ^5 denotes the log coefficient of variation obtained from the World Bank's Exporter Dynamics Database (WBED). Θ^6 is the residual from regressing the Pareto shape coefficient from the World Bank Enterprise Survey (WBES), Θ^1 , on log GDP/capita. d_{mjt}^{GDP} denotes the alternative demand measure based on country GDP.

Table 10: Buyer versus seller heterogeneity.

	Threshold	2005	2006	2007	2008	2009	2010
N Destinations	30 firms	-.37 ^a	-.15	-.63 ^b	-.40 ^c	-.68 ^a	-.36 ^a
		(.11)	(.14)	(.22)	(.22)	(.15)	(.09)
		101	107	52	56	47	23
N Destinations	50 firms	-.16	-.13	-.77 ^c	-.55 ^c	-.89 ^b	-.49 ^a
		(.21)	(.29)	(.39)	(.28)	(.37)	(.04)
		72	73	28	29	27	12
Country FE		Yes	Yes	Yes	Yes	Yes	Yes
Industry FE (HS2)		Yes	Yes	Yes	Yes	Yes	Yes

Note: The dependent variable is the log coefficient of variation for Norwegian exports to a industry-destination pair. The independent variable is the log coefficient of variation for foreign exports from a industry-destination pair (WBED data). Robust standard errors in parentheses, clustered by country. ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$. Each column represents a regression for a particular year. Threshold=30 firms uses a threshold of 30 or more buyers and sellers per country-industry, while threshold=50 uses a threshold of 50 or more buyers and sellers per country-industry.

Appendix

A Equilibrium Sorting

The solution to the sorting function is:

$$\underline{z}_{ij}(Z) = \frac{\tau_{ij} w_i \Omega_j}{Z} (w_i f_{ij})^{1/(\sigma-1)} \left(\frac{Y_j}{N_j} \right)^{-1/\gamma}$$

Proof. Equation (3) implicitly defines the $\underline{z}_{ij}(Z)$ function. We start with the guess $\underline{z}_{ij}(Z) = S_{ij} Z^s$ and the inverse $\underline{Z}_{ij}(z) = (z/S_{ij})^{1/s}$, where S_{ij} and s are unknowns. Furthermore, the relationship between E and Z is not yet determined, but we start with a guess $E_j(Z) = \kappa_3 (Y_j/N_j) Z^\gamma$, where κ_3 is a constant term, and show in Section B that this is consistent with the equilibrium. Inserting these expressions, as well as the price index (equation (1)), into equation (3) yields

$$\begin{aligned} \frac{1}{\sum_k n_k (\bar{m} \tau_{kj} w_k)^{1-\sigma} (S_{kj} \underline{Z}^s)^{-\gamma_2}} &= \frac{\sigma w_i f_{ij} \gamma z_L^\gamma}{E_j(\underline{Z}) \gamma_2} (\bar{m} \tau_{ij} w_i)^{\sigma-1} z^{1-\sigma} \\ \frac{\underline{Z}^{s\gamma_2+\gamma}}{\sum_k n_k (\bar{m} \tau_{kj} w_k)^{1-\sigma} S_{kj}^{-\gamma_2}} &= \frac{\sigma w_i f_{ij} \gamma}{\kappa_3 (Y_j/N_j) \gamma_2} z_L^\gamma (\bar{m} \tau_{ij} w_i)^{\sigma-1} z^{1-\sigma}. \end{aligned}$$

Hence,

$$\frac{1}{s} = \frac{1-\sigma}{s(\gamma_2 + \gamma/s)} \iff \frac{1}{s} = -1,$$

and

$$\begin{aligned} \left(\frac{1}{S_{ij}} \right)^{1/s} &= \left[\frac{\sigma w_i f_{ij} \gamma z_L^\gamma}{\kappa_3 (Y_j/N_j) \gamma_2} (\bar{m} \tau_{ij} w_i)^{\sigma-1} \sum_k n_k (\bar{m} \tau_{kj} w_k)^{1-\sigma} S_{kj}^{-\gamma_2} \right]^{1/(s\gamma_2+\gamma)} \iff \\ S_{ij} &= \left[\frac{\sigma w_i f_{ij} \gamma z_L^\gamma}{\kappa_3 (Y_j/N_j) \gamma_2} (\tau_{ij} w_i)^{\sigma-1} \sum_k n_k (\tau_{kj} w_k)^{1-\sigma} S_{kj}^{-\gamma_2} \right]^{1/(\sigma-1)}. \end{aligned} \quad (13)$$

In sum, the cutoff is

$$\underline{z}_{ij}(Z) = \frac{S_{ij}}{Z}. \quad (14)$$

□

We proceed by solving for S_{ij} and q_j . Inserting the expression for the cutoff (equation (14)) into the price index in equation (1) yields

$$q_j(Z)^{1-\sigma} = Z^{\gamma_2} \bar{m}^{1-\sigma} \frac{\gamma z_L^\gamma}{\gamma_2} \sum_k n_k (\tau_{kj} w_k)^{1-\sigma} S_{kj}^{-\gamma_2}.$$

Inserting the expression for S_{kj} from equation (13) then yields

$$q_j(Z)^{1-\sigma} = Z^{\gamma_2} \bar{m}^{1-\sigma} \frac{\kappa_3}{\sigma w_i f_{ij} N_j} \frac{Y_j}{\left(\frac{S_{ij}}{\tau_{ij} w_i} \right)^{\sigma-1}}.$$

This must hold for all i , so

$$(w_i f_{ij})^{-1/(\sigma-1)} \frac{S_{ij}}{\tau_{ij} w_i} = (w_k f_{kj})^{-1/(\sigma-1)} \frac{S_{kj}}{\tau_{kj} w_k}.$$

By exploiting this fact, we can transform the expression for S_{ij} ,

$$\begin{aligned} S_{ij}^{\sigma-1} &= (\tau_{ij} w_i)^{\sigma-1} \frac{\sigma w_i f_{ij}}{\kappa_3 (Y_j/N_j)} \frac{\gamma z_L^\gamma}{\gamma_2} \sum_k n_k (\tau_{kj} w_k)^{1-\sigma} (\tau_{kj} w_k)^{-\gamma_2} (w_k f_{kj})^{-\gamma_2/(\sigma-1)} \left((w_k f_{kj})^{-1/(\sigma-1)} \frac{S_{kj}}{\tau_{kj} w_k} \right)^{-\gamma_2} \\ &= (\tau_{ij} w_i)^{\sigma-1} \frac{\sigma w_i f_{ij}}{\kappa_3 (Y_j/N_j)} \frac{\gamma z_L^\gamma}{\gamma_2} \left((w_i f_{ij})^{-1/(\sigma-1)} \frac{S_{ij}}{\tau_{ij} w_i} \right)^{-\gamma_2} \sum_k n_k (\tau_{kj} w_k)^{-\gamma} (w_k f_{kj})^{-\gamma_2/(\sigma-1)} \iff \\ S_{ij}^\gamma &= (\tau_{ij} w_i)^\gamma \frac{\sigma}{\kappa_3 (Y_j/N_j)} (w_i f_{ij})^{\gamma/(\sigma-1)} \frac{\gamma z_L^\gamma}{\gamma_2} \sum_k n_k (\tau_{kj} w_k)^{-\gamma} (w_k f_{kj})^{-\gamma_2/(\sigma-1)} \iff \\ S_{ij} &= \tau_{ij} w_i (w_i f_{ij})^{1/(\sigma-1)} (Y_j/N_j)^{-1/\gamma} z_L \left(\frac{\sigma}{\kappa_3} \frac{\gamma}{\gamma_2} \sum_k n_k (\tau_{kj} w_k)^{-\gamma} (w_k f_{kj})^{-\gamma_2/(\sigma-1)} \right)^{1/\gamma}. \end{aligned}$$

We define

$$\Omega_j \equiv \kappa_2 \left(\sum_k n'_k (\tau_{kj} w_k)^{-\gamma} (w_k f_{kj})^{-\gamma_2/(\sigma-1)} \right)^{1/\gamma},$$

where $\kappa_2 = \left(\frac{\sigma}{\kappa_3} \frac{\gamma}{\gamma_2} \right)^{1/\gamma}$ and given the normalization $n_i = z_L^{-\gamma} n'_i$, we get the closed form solution for the sorting function,

$$z_{ij}(Z) = \frac{\tau_{ij} w_i \Omega_j}{Z} (w_i f_{ij})^{1/(\sigma-1)} \left(\frac{Y_j}{N_j} \right)^{-1/\gamma}.$$

Note that we can now write the price index as

$$\begin{aligned} q_j(Z)^{1-\sigma} &= Z^{\gamma_2} \bar{m}^{1-\sigma} \frac{\kappa_3}{\sigma w_i f_{ij}} \frac{Y_j}{N_j} \left(\frac{S_{ij}}{\tau_{ij} w_i} \right)^{\sigma-1} \\ &= Z^{\gamma_2} \bar{m}^{1-\sigma} \frac{\kappa_3}{\sigma w_i f_{ij}} \frac{Y_j}{N_j} \left(\frac{\tau_{ij} w_i (w_i f_{ij})^{1/(\sigma-1)} (Y_j/N_j)^{-1/\gamma} \Omega_j}{\tau_{ij} w_i} \right)^{\sigma-1} \\ &= Z^{\gamma_2} \frac{\bar{m}^{1-\sigma} \kappa_3}{\sigma} \left(\frac{Y_j}{N_j} \right)^{\gamma_2/\gamma} \Omega_j^{\sigma-1}. \end{aligned} \tag{15}$$

B Final Goods Producers Expenditure on Intermediates and Productivity

In this section, we derive the equilibrium relationship between final goods expenditure E and productivity Z . Revenue for a final goods producer is

$$R_i = \left(\frac{P_i}{Q_i} \right)^{1-\sigma} \mu Y_i = \left(\frac{\bar{m} q_i(Z)}{Z Q_i} \right)^{1-\sigma} \mu Y_i,$$

where $P_i = \bar{m}q_i(Z)/Z$ is the price charged and Q_i is the CES price index for final goods. The price index for final goods is

$$\begin{aligned}
 Q_i^{1-\sigma} &= N_i \int_1^\infty P_i(Z)^{1-\sigma} dG(Z) \\
 &= N_i \int_1^\infty (\bar{m}q_i(Z)/Z)^{1-\sigma} dG(Z) \\
 &= N_i \frac{\bar{m}^{2(1-\sigma)\kappa_3}}{\sigma} \frac{\Gamma}{\Gamma-\gamma} \left(\frac{Y_i}{N_i}\right)^{\gamma_2/\gamma} \Omega_i^{\sigma-1}.
 \end{aligned} \tag{16}$$

Rewriting revenue as a function of E and inserting the equilibrium expressions for $q_i(Z)$ and Q_i yields

$$\begin{aligned}
 \bar{m}E_i &= \left(\frac{\bar{m}q_i(Z)}{ZQ_i}\right)^{1-\sigma} \mu Y_i \\
 &= \bar{m}^{1-\sigma} Z^{\sigma-1} \frac{Z^{\gamma_2} \bar{m}^{1-\sigma} \kappa_3}{\sigma} \left(\frac{Y_i}{N_i}\right)^{\gamma_2/\gamma} \Omega_i^{\sigma-1} \mu Y_i \iff \\
 &\quad \frac{\bar{m}^{2(1-\sigma)\kappa_3}}{\sigma} \frac{\Gamma}{\Gamma-\gamma} N_i \left(\frac{Y_i}{N_i}\right)^{\gamma_2/\gamma} \Omega_i^{\sigma-1} \\
 E_i(Z) &= \kappa_3 \frac{Y_i}{N_i} Z^\gamma,
 \end{aligned} \tag{17}$$

where $\kappa_3 = \mu(\Gamma - \gamma)/\Gamma$. Hence, total spending on intermediates is increasing in productivity with an elasticity γ . Note that the expression for $E_i(Z)$ is the same as the one we started with in Section F.

C Firm-level Trade

Using equations (2) and (1), as well as the sorting function $\underline{Z}_{ij}(z)$, sales for a (z, Z) match are

$$r_{ij}(z, Z) = \left(\frac{p_{ij}(z)}{q_j(Z)}\right)^{1-\sigma} E_j(Z) = \sigma \left(\frac{zZ}{\tau_{ij}w_i\Omega_j}\right)^{\sigma-1} \left(\frac{Y_j}{N_j}\right)^{(\sigma-1)/\gamma}. \tag{18}$$

Note that revenue is supermodular in (z, Z) : $\partial^2 r / \partial z \partial Z > 0$. Buyer productivity is distributed Pareto, $G(Z) = 1 - Z^{-\Gamma}$. For firms with $z < \underline{z}_{ij}(Z_L) \equiv z_H$, total firm-level exports to country j are

$$\begin{aligned}
 r_{ij}^{TOT}(z) &= N_j \int_{\underline{Z}_{ij}(z)} r_{ij}(z, Z) dG(Z) \\
 &= \kappa_1 N_j (w_i f_{ij})^{1-\Gamma/(\sigma-1)} \left(\frac{z}{\tau_{ij}w_i\Omega_j}\right)^\Gamma \left(\frac{Y_j}{N_j}\right)^{\Gamma/\gamma},
 \end{aligned} \tag{19}$$

where we defined $\kappa_1 \equiv \sigma\Gamma/[\Gamma - (\sigma - 1)]$. We can alternatively express revenue as a function of the hurdle $\underline{Z}_{ij}(z)$, which yields

$$r_{ij}^{TOT}(z) = \kappa_1 N_j w_i f_{ij} \underline{Z}_{ij}(z)^{-\Gamma}.$$

For firms with $z \geq z_H$, total firm-level exports are

$$\begin{aligned}\tilde{r}_{ij}^{TOT}(z) &= N_j \int_{Z_L} r_{ij}(z, Z) dG(Z) \\ &= \kappa_1 N_j \left(\frac{z}{\tau_{ij} w_i \Omega_j} \right)^{\sigma-1} \left(\frac{Y_j}{N_j} \right)^{(\sigma-1)/\gamma}.\end{aligned}$$

Using the sorting function, we can also derive the measure of buyers in country j for a firm in country i with productivity $z < z_H$,

$$\begin{aligned}b_{ij}(z) &= N_j \int_{\underline{Z}_{ij}(z)} dG(Z) \\ &= N_j (w_i f_{ij})^{-\Gamma/(\sigma-1)} \left(\frac{z}{\tau_{ij} w_i \Omega_j} \right)^\Gamma \left(\frac{Y_j}{N_j} \right)^{\Gamma/\gamma}.\end{aligned}\quad (20)$$

Given that z is distributed Pareto, the distribution of customers per firm (out-degree distribution) is also Pareto.

Knowing firm-level exports from equation (19) as well as the number of buyers from equation (20), the firm's average exports is given by

$$\frac{r_{ij}^{TOT}(z)}{b_{ij}(z)} = \kappa_1 w_i f_{ij}.\quad (21)$$

Inversely, we calculate purchases from i of a final goods firm Z located in j . This is

$$\begin{aligned}R_{ij}^{TOT}(Z) &= n_i \int_{\underline{z}_{ij}(Z)} r_{ij}(z, Z) dF(z) \\ &= \kappa_4 n'_i (w_i f_{ij})^{1-\gamma/(\sigma-1)} \left(\frac{Z}{\tau_{ij} w_i \Omega_j} \right)^\gamma \frac{Y_j}{N_j},\end{aligned}$$

where $\kappa_4 = \sigma\gamma/[\gamma - (\sigma - 1)]$. The firm-level measure of sellers for a buyer located in j with productivity Z is

$$L_{ij}(Z) = n_i \int_{\underline{z}_{ij}(Z)} dF(z) = n'_i (w_i f_{ij})^{-\gamma/(\sigma-1)} \left(\frac{Z}{\tau_{ij} w_i \Omega_j} \right)^\gamma \frac{Y_j}{N_j}.\quad (22)$$

Hence, given that Z is distributed Pareto, both the distribution of purchases R_{ij}^{TOT} and the distribution of number of sellers per buyer $L_{ij}(Z)$ (indegree distribution) are Pareto. These results are symmetric to the findings on the seller side.

Finally, equilibrium firm-level profits for intermediate producers with productivity $z < z_H$ is given by

$$\begin{aligned}\pi_{ij}(z) &= \frac{r_{ij}^{TOT}(z)}{\sigma} - w_i f_{ij} b_{ij}(z) \\ &= \left(\frac{\kappa_1}{\sigma} - 1 \right) N_j (w_i f_{ij})^{1-\Gamma/(\sigma-1)} \left(\frac{z}{\tau_{ij} w_i \Omega_j} \right)^\Gamma \left(\frac{Y_j}{N_j} \right)^{\Gamma/\gamma}.\end{aligned}$$

D Welfare

As shown in equation (16), the price index on final goods is

$$Q_i^{1-\sigma} = \bar{m}^{2(1-\sigma)} \frac{\mu}{\sigma} N_i \left(\frac{Y_i}{N_i} \right)^{\gamma_2/\gamma} \Omega_i^{\sigma-1}.$$

Using the expression for the trade share in equation (9), we can rewrite Ω_i as

$$\Omega_i = \left(\frac{\sigma}{\kappa_3} \frac{\gamma}{\gamma_2} n'_j (w_i f_{ij})^{1-\gamma/(\sigma-1)} \right)^{1/\gamma} \pi_{jj}^{-1/\gamma} \frac{1}{\tau_{jj} w_j}.$$

Inserting this back into the price index Q_i and rearranging yields the real wage

$$\frac{w_j}{Q_j} = \kappa_6 \left(n'_j N_j \right)^{1/\gamma} \left(\frac{f_{jj}}{L_j} \right)^{1/\gamma - 1/(\sigma-1)} \frac{\pi_{jj}^{-1/\gamma}}{\tau_{jj}},$$

where κ_6 is a constant.⁴⁵ Holding n'_j , N_j , f_{jj} , τ_{jj} and L_j constant, the change in real wages going from autarky (A) to trade (T) with an import share $1 - \pi_{jj}$ is

$$\frac{(w_j/Q_j)^T}{(w_j/Q_j)^A} = \pi_{jj}^{-1/\gamma}.$$

E The Within-Firm Export Distribution

Using the expression for sales for a given (z, Z) match in equation (18) as well as the sorting function $\underline{Z}_{ij}(z)$, the distribution of exports across buyers for a seller with productivity z is

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

Hence, within-firm sales is distributed Pareto with shape coefficient $\Gamma/(\sigma-1)$. Note that the distribution is identical for every exporter in i selling to j .

F Sorting

Using the Norwegian trade data, Figure 5 shows the empirical relationship between a firm's number of customers in destination j and average number of connections to Norwegian exporters among its customers, i.e. the correlation between the degree of a node and the average degree of its neighbors. In this section, we derive the corresponding relationship in the model.

Using equations (22) and (4), the number of connections for the marginal customer of a firm with productivity z is $L_{ij}(\underline{Z}_{ij}(z)) = n'_i z^{-\gamma}$. Using equation (20), we can rewrite this as

$$\underline{L}_{ij}(b_{ij}) = n'_i N_j (w_i f_{ij})^{-\gamma/(\sigma-1)} (\tau_{ij} w_i \Omega_j)^{-\gamma} \frac{Y_j}{N_j} b_{ij}^{-\gamma/\Gamma},$$

⁴⁵ $\kappa_6 = \left(\frac{\sigma}{\kappa_3} \frac{\gamma}{\gamma_2} \right)^{1/\gamma} \left(\bar{m}^{2(1-\sigma)} \frac{\mu}{\sigma} \right)^{1/(\sigma-1)} (1 + \psi)^{-1/\gamma + 1/(\sigma-1)}$.

which relates a firm's number of customers b_{ij} to the number of connections for the firm's marginal customer, \underline{L}_{ij} .

In the data, we explore the average number of connections among all the firm's customers, not just the marginal one. The average number of connections among the customers of a firm with productivity z is

$$\begin{aligned}\hat{L}_{ij}(z) &= \frac{1}{1 - G(\underline{Z}_{ij}(z))} \int_{\underline{Z}_{ij}(z)} L_{ij}(Z) dG(Z) \\ &= \frac{\Gamma}{\Gamma - \gamma} n'_i z^{-\gamma}.\end{aligned}$$

The average number of connections among the customers of a firm with b_{ij} customers is then

$$\hat{L}_{ij}(b_{ij}) = \frac{\Gamma}{\Gamma - \gamma} n'_i \left(\frac{b_{ij}}{b_{ij}(1)} \right)^{-\gamma/\Gamma}.$$

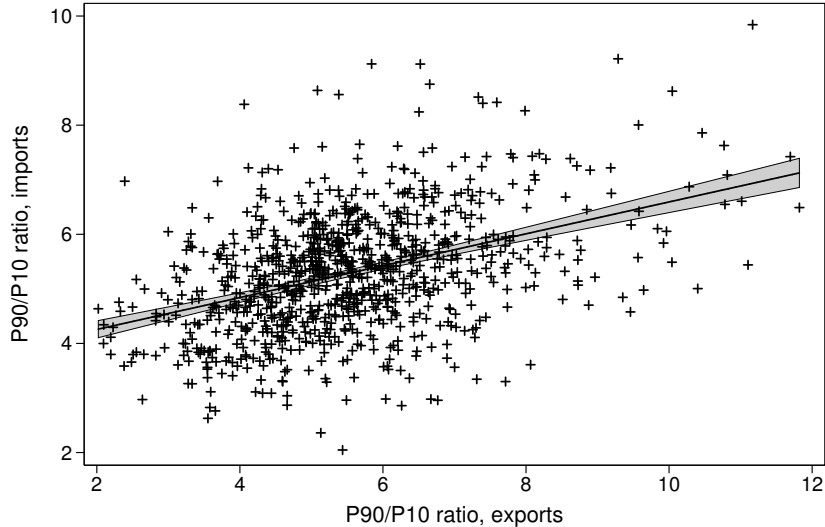
Hence, the elasticity of \hat{L}_{ij} with respect to b_{ij} is $-\gamma/\Gamma$.

G Dispersion in exports and imports

In Section 5.4, we test the hypothesis that imports dispersion is negatively correlated with exports dispersion. As imports dispersion is not directly observed, we instead use exports dispersion from the World Banks Exporter Dynamics database (WBED) as a proxy for imports dispersion. The robustness check in Section 5.3 also uses the WBED data in the same way.

In this Section, we estimate the correlation between exports and imports dispersion using the Norwegian data. For the 2004 cross-section, we observe both export and import values by firm, product and year. We proceed as follows. First, the data is aggregated to the HS 2 digit level, as in Section 5.4. Second, the exports and imports log 90/10 percentile ratios are calculated for each product-destination combination. In Figure 9, we plot the resulting scatter for every product-destination pair with more than 10 firms present. Choosing a different threshold has a negligible impact on the results. The correlation is positive and significant, and the estimated slope coefficient is 0.29 (s.e. 0.02). This suggests that the WBED data should proxy imports dispersion reasonably well.

Figure 9: Heterogeneity of importer expenditure across markets.



Note: 2004 data. The Figure shows log 90/10 percentile ratios for imports and exports for product-destination pairs with more than 10 firms present. The fitted regression line and 95% confidence intervals are denoted by the solid line and gray area. The slope coefficient is 0.29 (s.e. 0.02).

H A Random Matching Model

In this section, we ask to what extent a random matching model can replicate the basic facts presented in the main text. The main finding is that a random model fails to explain key empirical

facts.

We model the matching process as a balls-and-bins model, similar to Armenter and Koren (2013). There are B buyers, S sellers and n balls. The number of bins is SB , the total number of possible buyer-seller combinations, and we index each bin by sb . The probability that a given ball lands in bin sb is given by the bin size s_{sb} , with $0 < s_{sb} \leq 1$ and $\sum_s^S \sum_b^B s_{sb} = 1$. We assume that $s_{sb} = s_s s_b$, so that the buyer match probability (s_b) and seller match probability (s_s) are independent. Trade from seller s to buyer b is the total number of balls landing in bin sb , which we denote by r_{sb} . A buyer-seller match is denoted by $m_{sb} = 1 [r_{sb} > 0]$.

Parameters and simulation. We simulate the random model as follows. Focusing on Norway's largest export destination, Sweden, we set B and S equal to the number of buyers in Sweden and exporters to Sweden (see Table 5). The number of balls, n , equals the total number of connections made (24,400). The match probabilities s_s correspond to each seller's number of customers relative to the total number of connections made; s_b correspond to each buyer's number of suppliers relative to the total number of connections made.

Results. We focus on the key relationships described in the main text; (i) degree distributions,⁴⁶ (ii) number of connections versus total sales and within-firm sales dispersion and (iii) assortivity in in-degree and average out-degree of the nodes in:

(i) We plot the simulated degree distributions in Figure 10, in the same way as in the main text. Given that the match probabilities s_b and s_s are taken from the actual data, it is not surprising that the simulated degree distributions resemble the actual distributions in Figures 1 and 2.

(ii) The relationship between the number of customers and total exports per seller is plotted in the left panel of Figure 11. The relationship is positive and log linear. The right panel plots the number of customers on the horizontal axis and the value of 10th, 50th and 90th percentile of buyer-seller transactions (within firm) on the vertical axis. In contrast to the actual data and our main model (see Figure 4), the large majority of firms sell the same amount to each buyer; hence both the 10th and the 90th percentile cluster at $r_{sb} = 1$. For the firms with dispersion in sales, the magnitude of dispersion is small, with the 90th percentile not exceeding $r_{sb} = 2$.

(iii) Figure 12 plots the relationship between out-degree and mean in-degree (and the opposite), as illustrated in the main text in Figure 5. The relationship is essentially flat, so that the contacts of more popular sellers are on average similar to the contacts of less popular sellers. This is also at odds with the data and our main model.

In sum, the random matching model is not able to reproduce all the basic facts from the data.

⁴⁶The degree of a node in a network is the number of connections it has to other nodes, while the degree distribution is the probability distribution of these degrees over the whole network.




Figure 10: Distribution of out-degree and in-degree.




Figure 11: Firm-level total exports and within-firm dispersion in exports.

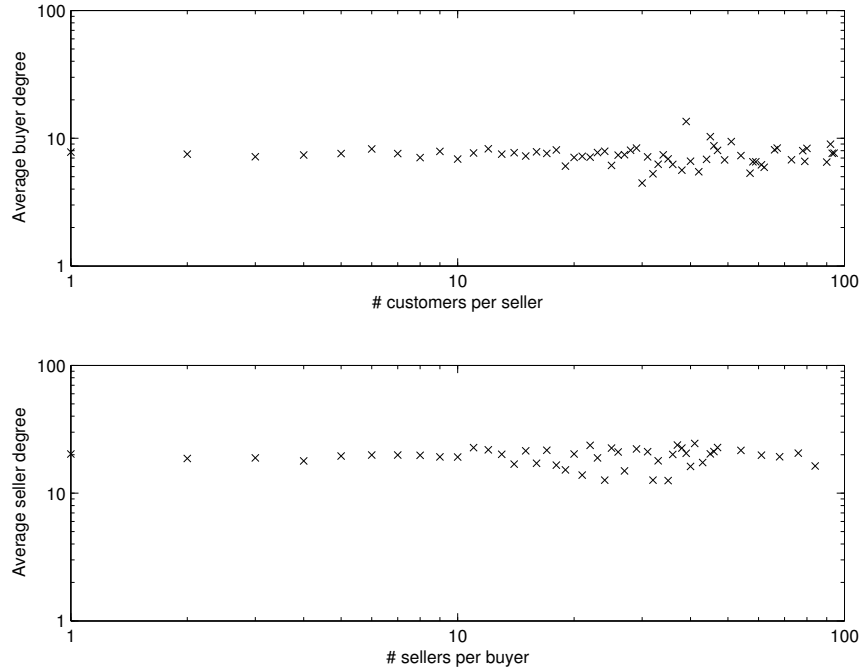


Figure 12: Degree and average degree of customers/suppliers.

I Basic Facts Revisited

This section presents descriptive evidence on buyer-seller relationships using trade data from a different country, Colombia. We show that the basic facts from Section 3 also hold in the Colombian data.

The data set includes all Colombian import transactions in 2011 as assembled by ImportGenius.⁴⁷ As in the Norwegian data, we can identify every domestic buyer (importer) and foreign sellers (exporters) in all source countries. However unlike the Norwegian data, transactions must be matched to firms (either exporters or importers) using raw names and thus are potentially subject to more error than the comparable Norwegian data. However, there is no reason to believe the noise in the data is systematic and thus we are comfortable using the data as a robustness check. Note that since we only have import data from Colombia, the roles of buyers and sellers are reversed compared to the Norwegian data, i.e. in the descriptive evidence that follows, an exporter represents a foreign firm exporting to Colombia, and an importer denotes a Colombian firm purchasing from abroad.

We reproduce the same facts as in the Norwegian data. Table 11 reports exporter and importer concentration for all imports and imports from Colombia's largest sourcing markets in 2011, U.S., China and Mexico. Both sellers and buyers of Colombian imports are characterized by extreme

⁴⁷The data are available at <http://importgenius.com>. See Bernard and Dhingra (2014) for details on the data construction.

Table 11: Descriptive statistics: Colombian Imports.

	Overall	U.S.	China	Mexico
Number of exporters	95,185	28,926	32,677	5,349
Number of buyers	34,166	15,047	15,445	5,050
Share trade, top 10% sellers	.90	.93	.84	.96
Share trade, top 10% buyers	.93	.93	.87	.93
Share in total CO imports, %	100	26.2	15.5	11.4

Note: 2011 data. The overall column refers to outcomes unconditional on importer country.

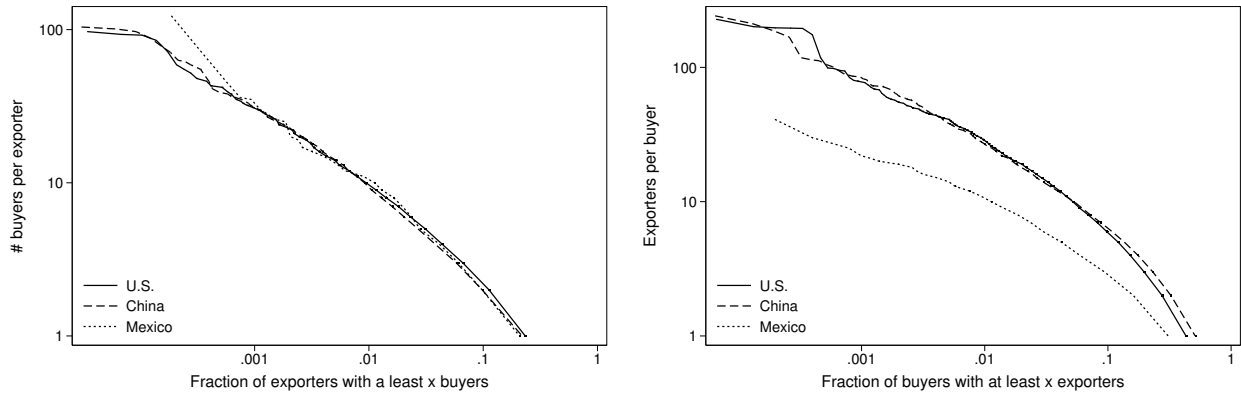
Table 12: Types of matches, % : Colombia.

	(1) One-to-one	(2) Many-to-one	(3) One-to-many	(4) Many-to-many
Share of value, %	4.9	36.4	7.6	51.1
Share of counts, %	15.8	36.5	12.8	34.9

Note: 2011 data. See Table 6 footnote.


concentration, mirroring the finding in Table 5 (basic fact 2). Figure 13 confirms that the degree distributions in Colombia are close to Pareto, mirroring the finding in Figures 1 and 2 in the main text. Moreover, Table 12 shows that one-to-one matches are relatively unimportant in total imports (basic fact 3). Figures 14 and 15 show that while more connected exporters typically sell more, the within-firm distribution of sales is relatively constant, mirroring the finding in Figures 3 and 4 (basic fact 4). Figure 16 illustrates that more popular exporters on average match to less connected importers, mirroring the finding in Figure 5 (basic fact 5). Figure 17 shows that exporters to Colombia are more likely obey the hierarchy among Colombian importers relative to what a random matching model would predict, mirroring the finding in Figure 7 (basic fact 6).

Figure 13: Distribution of # buyers per exporter (left) and exporters per buyer (right): Colombia.




Note: 2011 data. Buyers per exporter: The estimated slope coefficients are -0.74 (s.e. 0.0004) for U.S., -0.78 (s.e. 0.001) for China and -0.78 (s.e. 0.001) for Mexico. Exporters per buyer: The estimated slope coefficients are -0.99 (s.e. 0.002) for U.S., -0.74 (s.e. 0.002) for China and -0.74 (s.e. 0.002) for Mexico.

Figure 14: Number of buyers & firm-level exports: Colombia.




Note: 2011 data. See Figure 3 footnote.

Figure 15: Number of buyers & within-firm dispersion in exports: Colombia.




Note: 2011 data. See Figure 4 footnote.

Figure 16: Matching buyers and sellers across markets: Colombia.



Note: 2011 data. The linear regression slope is -0.14 (s.e. 0.01). See Figure 5 footnote.

Figure 17: Firms obeying a pecking order hierarchy: Colombia.



Note: 2011 data. See Figure 7 footnote.