

NBER WORKING PAPER SERIES

TWO-SIDED SEARCH IN INTERNATIONAL MARKETS

Jonathan Eaton
David Jinkins
James R. Tybout
Daniel Xu

Working Paper 29684
<http://www.nber.org/papers/w29684>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2022

We gratefully acknowledge support from the National Science Foundation (Grant SES-1426645) and the United States Census Bureau. This research was performed at a Federal Statistical Research Data Center under FSRDC Project number 1109. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release (DRB approval requests 5488-2017-0109, 5623-2017-0228, and 5952-2017-0808). Any opinions, findings, conclusions and recommendations expressed are those of the authors and do not necessarily represent the views of the NSF or the U.S. Census Bureau. We thank Costas Arkolakis, Andrew Bernard, Joris Pinske, and participants in numerous seminars for many useful comments on earlier drafts. We also wish to thank Pietra Rivoli and Julia Hughes (President, U.S. Fashion Industry Association) for very helpful discussions. Finally, we are grateful to Yingyan Zhao and Zi Wang for excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Jonathan Eaton, David Jinkins, James R. Tybout, and Daniel Xu. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Two-Sided Search in International Markets

Jonathan Eaton, David Jinkins, James R. Tybout, and Daniel Xu

NBER Working Paper No. 29684

January 2022

JEL No. F12,F14

ABSTRACT

We develop a dynamic model of international business-to-business transactions in which sellers and buyers search for each other, with the probability of a match depending on both individual and aggregate search effort. Fit to customs records on U.S. apparel imports, the model captures key cross-sectional and dynamic features of international buyer-seller relationships. We use the model to make several quantitative inferences. First, we calculate the search costs borne by heterogeneous importers and exporters. Second, we provide a structural interpretation for the life cycles of importers and exporters as they endogenously acquire and lose foreign business partners. Third, we pursue counterfactuals that approximate the phaseout of the Agreement on Textiles and Clothing (the “China shock”) and the IT revolution. Lower search costs can significantly improve consumer welfare, but at the expense of importer profits. On the other hand, an increase in the population of foreign exporters can congest matching to the extent of dampening or even reversing the gains consumers enjoy from access to extra varieties and more retailers.

Jonathan Eaton
Department of Economics
Pennsylvania State University
303 Kern Graduate Building
University Park, PA 16801
and NBER
jxe22@psu.edu

James R. Tybout
Department of Economics
Pennsylvania State University
517 Kern Graduate Building
University Park, PA 16802
and NBER
jtybout@psu.edu

David Jinkins
Copenhagen Business School
Room: POR/16.A-2.66
Solbjerg Pl. 3
2000 Frederiksberg
Denmark
david.jinkins@gmail.com

Daniel Xu
Department of Economics
Duke University
213 Social Science Bldg
419 Chapel Drive
Box 90097
Durham, NC 27708-0097
and NBER
daniel.xu@duke.edu

1 Overview

To break into a foreign market, either as an exporter or as an importer, a firm must find foreign business partners. And since most international partnerships are short-lived, a trading firm must continually seek new connections to maintain or expand its presence in a foreign market. The resulting patterns of international buyer-seller connections are surprisingly fluid, and largely determine the dynamics of firm-level trade flows.

To study these patterns we build a model of two-sided search and matching which we use to quantify search costs and their implications for trade dynamics and welfare. We apply the model to imports of consumer goods, incorporating three types of agents: foreign exporters, domestic retailers, and domestic consumers. Heterogeneous exporters and retailers search for each other, taking stock of their current situation and the aggregate search behavior of other agents. The matches that result determine which goods each retailer carries. Consumers then choose where to shop, buying the individual goods that retailers offer. When an exporter and a retailer form a new business relationship, forward-looking Nash bargaining between them determines the wholesale price. Overall, the model connects the dynamic formation of trade relationships to evolving retailer varieties, prices, and consumer welfare.

Fit to customs records on U.S. apparel imports, our model allows us to evaluate the impacts of market-wide policy and technological changes on network structure, trade, and welfare. First, to approximate the 2005 phaseout of the Agreement on Textiles and Clothing (ATC, formerly the Multi Fibre Arrangement), we explore the effects of changing the number and mix of potential foreign apparel suppliers accessing the U.S. market. Holding the quality mix of potential foreign suppliers constant, an increase in their numbers sufficient to match the observed 2005 growth rate in *active* suppliers has virtually no effect on overall U.S. welfare. Interestingly, the congestion in matching generated by these added suppliers leads high-quality sellers to reduce their search effort disproportionately. The resulting product lines offered by the typical retailer consequently have less consumer appeal, offsetting the standard love-of-variety benefit that comes from enlarging the population of active exporters. An implication is that market integration can reduce welfare in trade models that combine heterogeneous sellers with search and matching frictions.

A second, related, experiment changes the quality mix of sellers accessing the U.S. market,

a likely outcome of the ATC phaseout and Chinese policy reforms (Khandelwal et al., 2013; Bai et al., 2017). Other things equal, we find that increasing the fraction of high-quality potential suppliers improves consumer welfare, despite reducing the number of exporters actively serving the U.S. market. Here improvements in the quality of the typical retailer’s offerings more than offset a negative variety effect. In contrast, improving the relative access of low-quality sellers to the U.S. market can reduce consumer welfare through congestion effects on matching. Improving the access of these low-quality sellers also favors the largest retailers while reducing the quality of their offerings. In this sense our model may help explain the rise of “big box” retailing.

Third, we perform an experiment to approximate firms’ increased market visibility resulting from advances in information and communication technology (ICT) in the early 2000s. A reduction in search costs consistent with the increase in matching between 1998 and 2004 increases consumer welfare significantly, but at the expense of the profits of both retailers and suppliers. The simple reason is that increased search effort resulting from better search technology means more entry, which spreads consumer spending more thinly across active firms on both sides of the market.

Besides providing a laboratory for quantitative analysis, our model speaks to the distribution of match rents between retailers and exporters, and how these rents drive their spending on search for new partners. In our baseline estimation we find that retailers and suppliers spend about the same amount on search in the aggregate. But since there are far more active suppliers than retailers, retailers spend less on search per match. Also, even though high-productivity active sellers are outnumbered by low-productivity active sellers one to four, their spending on search is more than six times the spending of their low-productivity competitors.

Finally, the model relates firm dynamics to the acquisition and dissolution of business partnerships, which in turn relate to search costs and firm competition. Firms’ search intensities vary over their life cycles. As retailers and exporters accumulate more connections, they develop market visibility, making it easier for them to find new business partners. But the cost of replacing expiring relationships, which is convex in the number of replacements, puts an eventual brake on their number of partners.

Our paper relates to a wide variety of earlier contributions. First, it builds on the literature addressing firm-level export dynamics, recently reviewed by Alessandria et al. (2020). Like several other papers, we emphasize customer accumulation as a primary driver of exporter growth (Albornoz et al., 2012; Drozd and Nosal, 2012; Eaton et al., 2021; Chaney, 2014; Carballo et al., 2018; Piveteau, 2021; Fitzgerald et al., 2019).¹ We expand on these papers by treating both exporters and importers as heterogeneous and by modeling search decisions on both sides. In our model, the formation of new relationships depends on equilibrium market tightness. As a result, the accumulation of each exporter’s “customer capital” depends not only on its own effort but also on endogenous market congestion.² In fact, endogenous market congestion plays an important role in our results.

A second literature on exporter dynamics focuses on firms’ life-cycles and their relation to the “fat” tails that typically characterize firm-size distributions. Some studies have generated these tails through stochastic shocks to firm productivity or demand (Luttmer, 2007, 2011; Arkolakis, 2016; Gumpert et al., 2020). Instead, as in Eaton et al. (2021), we generate these tails with a search cost function that incorporates “visibility effects.” That is, we allow firms with a large market presence to find new business partners with relative ease. Our approach explains well both the size distribution of firms and the matrix of transition probabilities across sizes measured in number of business partners.

A third literature focuses on transnational firm-to-firm trading patterns and the question of who matches with whom (Rauch, 2001; Rauch and Trindade, 2002; Sugita et al., forthcoming; Bernard et al., 2018, forthcoming; Eaton et al., 2018; Benguria, 2021). Our model assumes that matching is random, so we rule out assortative matching *ex ante*.³ However, we do allow the mix of active market participants to respond endogenously on both sides of the market. Thus we treat the realized mixes of business partners for importers and exporters as equilibrium objects.⁴

¹Recent contributions focusing on the accumulation of customers in a domestic context include Foster et al. (2016) and Gourio and Rudanko (2014).

²To focus on equilibrium search and matching, we abstract from learning, say by exporters about their own ability or about a destination market. Papers that consider exporter or importer learning effects include Rauch and Watson (2003); Eaton et al. (2021); Timoshenko (2015); Li (2018); Berman et al. (2019); Allen (2014), and Steinwender (2018).

³We present evidence that random matching is a reasonable approximation to the data.

⁴Also see Antras and Chor (2021) for a survey of the large recent literature on value chains and firm-to-firm

Finally, since our firms deal in clothing, our paper relates to a substantial literature on global apparel markets. Two papers are of particular relevance. First, Cahal et al. (2021) analyze the strategies apparel importers pursue when building their portfolio of foreign suppliers. Second, Khandelwal et al. (2013) demonstrate how the relaxation of China’s export licensing regime improved the quality of its apparel exports. We address these issues in the context of a dynamic structural model, sacrificing some of the nuances captured by these papers in order to focus on market equilibrium in the presence of endogenous search efforts by heterogeneous agents.

2 Data and Stylized Facts

While nothing in our framework restricts its application to international trade, we study the network of U.S. apparel importers (buyers) and their foreign suppliers (sellers) for several reasons. First, since most buyers are either wholesale or retail firms, we can keep the buyer side of the market relatively simple. In particular, the revenue function of such firms is nearly separable across categories of consumer goods, so that we can approximate their payoffs with a simple functional form. Second, we can observe cross-border transactions in customs records, while data on domestic firm-to-firm transactions are difficult to come by for the United States. Accordingly, we choose an industry in which domestic supply plays a relatively minor role. Finally, the U.S. apparel market has changed dramatically over the past 30 years, with major new sources of merchandise having emerged abroad and with the phaseout of quantitative restrictions on imports. These developments have changed the network structure of the market in ways that we can compare with our model’s predictions.

Before describing the details of our model, we review some aggregate patterns in U.S. apparel trade over the past 25 years and some micro features of the associated buyer-seller network. Some of these network features have been documented for other markets in the emerging literature on firm-to-firm trade, which includes studies on the United States and Colombia (Eaton et al., 2008, 2021; Bernard et al., 2018), Chile (Blum et al., 2010), Mexico (Sugita et al., forthcoming), Norway (Bernard et al., 2018), and Ireland (Fitzgerald et al.,

production networks.

2019).⁵ Other network features have received less attention, particularly those concerning match count transitions.

2.1 Data Sources

Our quantitative analysis is based largely on customs records from the U.S. Census Bureau. These data describe all merchandise shipments into the United States during the period 1996 to 2011.⁶ Each record includes a ten-digit Harmonized Schedule (HS) product code, shipment value, shipment quantity, entry or exit port, date of transaction, mode of transportation (land, sea, air), and the domestic firm’s identification code (i.e., its employer identification number or EIN). Critically for our study, each record also includes a string identifier based on the name and address of the foreign firm that is party to the transaction. This identifier allows us to track buyer-seller pairs through time.

The name and address of a given exporter may be recorded differently for different shipments, generating noise in this identifier. This noise leads to overstatement in the number of exporters and business relationships, as well as in the rates of relationship turnover. Nonetheless, as discussed in Appendix A, the patterns we report are qualitatively robust to this source of noise so long as we exclude single-shipment matches.⁷

2.2 Aggregate trends in apparel trade

Figure 1 shows that, after 2000, imports rapidly displaced domestic production as the primary source of apparel for U.S. consumers. The import penetration rate rises from approximately 30 percent in 1992 to around 80 percent in 2007.⁸ These trends reflect the emergence of China and other developing economies as exporters and the phasing out of the ATC in 2005.

⁵Bernard and Moxnes (2018) review many of the stylized facts in the firm-to-firm trade literature.

⁶We end our sample period in 2011 because this was the most recent year for which data were available at the shipment level from the U.S. Census Bureau when we began our analysis.

⁷Kamal and Monarch (2017) and Krizan et al. (2020) analyze the extent and implications of imperfect importer-exporter matching using U.S. customs records.

⁸Domestic consumption is the gross value of domestic apparel production plus apparel imports, less apparel exports. The value of domestic production is downloaded from the Bureau of Economic Analysis. Trade aggregates are from the WTO. The dip in both consumption and imports around 2009 reflects the financial crisis.

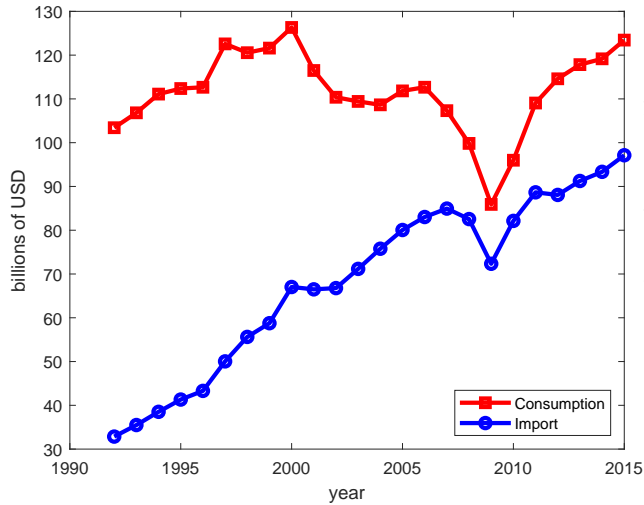


Figure 1: U.S. apparel consumption/imports




Figure 2: Number of sellers, 1996-2011

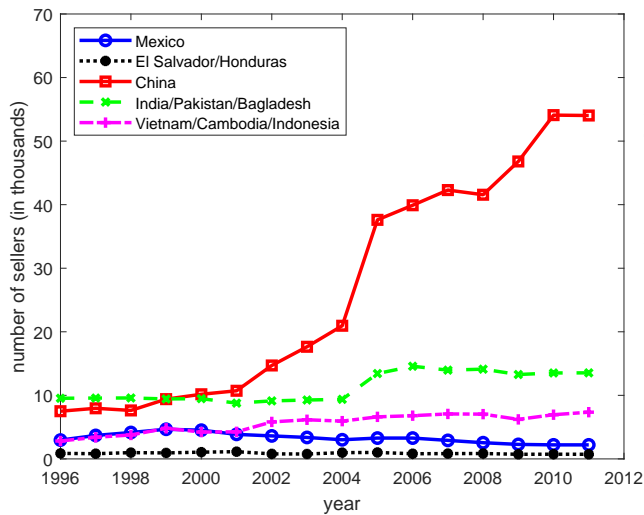


Figure 3: Number of sellers by country, 1996-2011




Figure 4: Value of imports by country, 1996-2011

As imports have come to dominate the domestic market, the number of firms exporting to the United States has steadily grown. Figure 2 shows the number of foreign suppliers making shipments to the U.S. Breaking suppliers down by country, Figures 3 and 4 show that China dominates this growth. India, Pakistan, Bangladesh, and Vietnam gained market share after 2005 as well, while other countries held stable or lost ground. Overall, these patterns suggest that the number and mix of exporters serving the U.S. apparel market were heavily impacted by trade policy reforms and external shocks. Our model is designed to shed light on the quantitative and welfare implications of these developments.

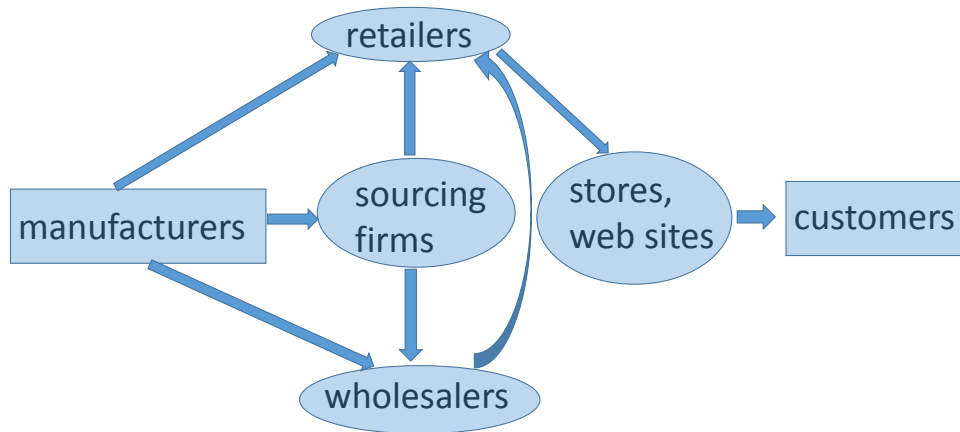


Figure 5: Industry structure

2.3 Firm-level features of the apparel market

To motivate the key components of our structural model, we now turn to the firm-level activities underlying these aggregates.

2.3.1 Types of players

We can roughly divide the players in the market into four categories, depicted in Figure 5: manufacturers, sourcing firms (match-makers), wholesalers (including branded importers), and general merchandise retailers.⁹ At one end of the chain are manufacturers who produce the apparel. At the other end are the consumers who ultimately wear it. In between are intermediaries in several different categories.

Manufacturers sell their output either to general merchandise retailers or to wholesalers. General merchandise retailers include big box stores such as Walmart and Target, as well as department stores such as Macy’s and Nordstrom’s. Firms in this group sell directly to consumers. Among wholesalers we include small-scale designers as well as large apparel firms such as Ralph Lauren, Gap, Land’s End, VF Corp, and Hanes.¹⁰ Such firms may sell directly to consumers as well but also sell to retailers.

⁹Plunkett-Analytics (2015) and Gereffi and Memedovic (2003) provide related classifications.

¹⁰VF Corp owns JanSport, The North Face, Timberland, Lee, Wrangler, and Nautica.

Complicating the picture further is that some connections between foreign manufacturers and U.S. importers of either type are brokered by sourcing firms that provide buyer-seller match-making, design, and other services. Examples include the Gulati Group, Apparel Sourcing Group, Inc., Li & Fung, and W. E. Connor.

The lines between the different types of agents are fuzzy, as it's not unusual for a firm to engage in more than one activity. For example, The Gulati Group also does some clothing manufacturing, and Hanes some manufacturing facilities. Also, in addition to selling their merchandise to department store chains and big box stores, some branded importers such as Ralph Lauren and VF Corp engage directly in retail sales through their web sites or brick-and-mortar stores.¹¹

For our analysis we approximate this structure by partitioning the players in the market into three mutually unaffiliated types: manufacturers, intermediaries, and consumers. Hence we ignore wholesale and retail firms that own their production facilities. We gloss over the distinction between sourcing using in-house staff versus using a sourcing firm. Finally, we treat the sales of intermediaries to other intermediaries the same as sales directly to final consumers, thus eliminating potential double marginalization.

How much distance does this simplification put between our model and the actual structure of the market? First, while a few branded importers such as Hanes own some production facilities, the vast majority do not (Plunkett-Analytics, 2015). The small fraction of apparel imports classified as affiliated trade (Figure 6) reflects this lack of vertical integration. Moreover, arm's-length relationships constitute virtually all of the growth in matches.

Second, it doesn't appear that apparel importers rely heavily on sourcing firms to match with foreign manufacturers. Small-scale operations often get started by attending trade fairs such as "Apparel Sourcing USA" or "Sourcing at Magic," which bring them face to face with foreign manufacturers.¹² McFarlan et al. (2012) report that, in 2004, 9 of the 10 largest

¹¹Census classifies *establishments* according to their main activity, such as wholesaling or retailing. Customs records report the importing *firm*. A firm may own multiple establishments in different categories.

¹²"Apparel Sourcing USA ... offers apparel brands, retailers, wholesalers and independent design firms a dedicated sourcing marketplace for finding the best international apparel manufacturers." <http://www.apparelsourcingshow.com/newyork/en/for-attendees/about-International-Apparel-Sourcing-Show.html>. "Sourcing at Magic" advertises on Facebook as "The largest fashion sourcing event in North America offering one-stop shopping for the entire apparel, footwear and accessories supply chain." The event website <http://10times.com/sourcing-at-magic> provides a partial list of attendees, which includes




Figure 6: Number of buyers, related party versus arm’s length trade

apparel retailers in the U.S. ran their own sourcing offices.¹³ In a recent report, the U.S. Fashion Industry Association surveyed 30 executives representing various segments of the apparel importing market (Lu, 2016).¹⁴ Among this group, 78 percent indicated that they “direct source from a selected supplier and mill matrix using [their] own designs and selecting fabric from the mill resource.” In contrast, only 41 percent indicated that they engaged a third party to source production.

Third, interviews with industry experts suggest that the branded importers set similar retail prices for their products whether they market them directly or through general retail outlets, consistent with the broader finding that 72 percent of on-line prices are identical to the prices charged by brick-and-mortar stores for the same products (Cavallo, 2017). Hence branded importers don’t seem to price discriminate across outlets. An interpretation is that they simply outsource their brick-and-mortar sales operations while retaining control of pricing.¹⁵

many representatives of apparel manufacturers located in South Asia.

¹³Kohl’s was an exception. More recently, some of these retailers have augmented their internal sourcing efforts with the services of sourcing companies. For example, WalMart signed a 6-year deal with Li & Fung in 2010 (McFarlan et al., 2012).

¹⁴Almost all of these executives represented large firms. Among them, 77 percent self-identified as retailers, 69 percent identified as branded importers, 69 percent identified as importer/wholesalers, and 27 percent identified as manufacturer/suppliers. (Percentages do not sum to 100 because most firms engage in more than one activity.)

¹⁵Online apparel sales accounted for roughly 87 percent of total apparel sales in the U.S. during 2014, up from 38 percent in 2003 (statista.com, 2022).

2.3.2 Network dynamics

To keep up with evolving fashions, retailers and branded importers source new products frequently. Access to real-time scanner data on sales has accelerated this development. Retailers are abandoning their practice of replenishing inventory each of the four seasons in favor of high-frequency design innovation requiring small-batch just-in-time production (McFarlan et al. (2012); Taplin (2014)).

In some cases retailers procure new production runs via long-standing relationships with manufacturers, but most buyer-seller partnerships are short-lived. Terry (2008) reports that “Apparel companies’ relationships with contract manufacturers in low-cost countries have historically been transient. Deals sometimes last only a few months as brands continuously pursue the lowest cost. On average, one-third to three-quarters of an apparel company’s contractor portfolio turns over every year.” This frequency is consistent with the annual match separation hazard of 0.81 we estimate from our data.¹⁶

A large amount of this relationship turnover is associated with product turnover. Distinguishing products at the HS6 level, we find that, for the average buyer, 52 percent of current year apparel imports are products that it neither carried last year *nor* sourced from one of last year’s suppliers.¹⁷ Goods that it imported last year but sourced from a different supplier accounted for another 14 percent of sales. Only 25 percent of imports were through product-supplier pairs carried forward from the previous year. These patterns are consistent with the view that apparel importers devote considerable effort to finding new suppliers and that changes in matches are closely linked to changes in varieties.

Given the short duration of a typical match, there’s a lot of year to year fluctuation in the number of foreign business partners that individual buyers and sellers deal with. Table 1 and Table 2, report annual transition rates for sellers’ buyer counts and buyers’ seller counts respectively. Several patterns emerge. First, corroborating our descriptive narratives, there is a substantial amount of churning of business relationships in our data. This is reflected in the small diagonal terms of the transition matrix for both sellers and buyers. Second, the

¹⁶We estimate this hazard by regressing the log of the fraction of matches surviving t years on t . The R^2 for this regression is 0.98, implying that match longevity is well approximated by the Poisson distribution.

¹⁷An example of a product-supplier pair would be “T-shirts, singlets and other vests; of textile materials (other than cotton), knitted or crocheted” from Ocean Sky Apparel in El Salvador.

transition matrix also exhibits a more pronounced downward adjustment probability when the number of relationships are relatively small. For instance in Table 1, when a seller has 1 – 9 buyers, it often has more than 60% probability of losing the number of relationships. Similar patterns hold for sellers per buyer transition in Table 2. In our model, these patterns naturally emerge from the heterogeneous search efforts as well as congestion in the matching market. Finally, even firms with many connections run some risk of dropping to zero. These events are primarily due to their exit—a possibility we will incorporate into our model.

Table 1: Year-to-year transition rates: buyers per seller*

year t, year t+1	0	1	2	3	4	5	6	7	8	9	≥ 10
1	0.65	0.27	0.05	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	0.32	0.31	0.21	0.09	0.03	0.02	0.01	0.00	0.00	0.00	0.00
3	0.19	0.22	0.23	0.17	0.09	0.05	0.02	0.01	0.01	0.00	0.01
4	0.13	0.15	0.18	0.18	0.14	0.09	0.05	0.03	0.02	0.01	0.02
5	0.10	0.10	0.13	0.16	0.16	0.12	0.08	0.05	0.03	0.02	0.04
6	0.08	0.07	0.10	0.13	0.14	0.13	0.11	0.08	0.05	0.03	0.07
7	0.07	0.06	0.08	0.09	0.12	0.13	0.12	0.10	0.07	0.05	0.11
8	0.07	0.05	0.05	0.07	0.10	0.11	0.11	0.11	0.09	0.07	0.16
9	0.06	0.05	0.05	0.06	0.08	0.09	0.10	0.10	0.10	0.08	0.24
≥ 10	0.05	0.03	0.03	0.03	0.04	0.04	0.05	0.05	0.06	0.06	0.56

*Based on monthly U.S. customs records. Buyer-seller pairs are considered to be matched during the period between their first observed shipment and their last observed shipment. Matches that generate shipments in the first sample year (1996) are treated as active from the beginning of the sample and matches that generate shipments in the last sample year (2011) are treated as active through the end of the sample. Buyer-seller pairs that generate a single shipment are not considered to have matched.

2.3.3 Degree distributions

Consistent with the tendency to lose business partners in Tables 1 and 2, it’s unusual for a firm to sustain large portfolios of foreign partners. Table 3 reports the frequency distributions of sellers per buyer and buyers per seller. Note that the vast majority of the firms have less than 10 partners in our data (86% of the buyers and 99% of the sellers)¹⁸. Nonetheless, both degree distributions come close to obeying a power law, implying that a small fraction of firms attain very large sizes.¹⁹

¹⁸The seller’s degree distribution is limited to its U.S. partners, so this pattern is not surprising.

Table 2: Year-to-year transition rates: sellers per buyer*

year t, year t+1	0	1	2	3	4	5	6	7	8	9	≥ 10
1	0.58	0.26	0.09	0.04	0.02	0.01	0.01	0.00	0.00	0.00	0.01
2	0.34	0.24	0.19	0.10	0.05	0.03	0.02	0.01	0.01	0.00	0.02
3	0.25	0.16	0.18	0.14	0.09	0.06	0.03	0.02	0.02	0.01	0.03
4	0.21	0.11	0.14	0.14	0.13	0.08	0.06	0.04	0.03	0.02	0.06
5	0.19	0.07	0.10	0.12	0.12	0.11	0.07	0.06	0.04	0.03	0.09
6	0.17	0.06	0.08	0.09	0.11	0.11	0.09	0.07	0.05	0.04	0.13
7	0.16	0.05	0.05	0.07	0.09	0.10	0.09	0.09	0.06	0.06	0.19
8	0.15	0.04	0.05	0.06	0.07	0.08	0.08	0.08	0.07	0.06	0.25
9	0.15	0.03	0.03	0.04	0.06	0.07	0.08	0.08	0.07	0.07	0.32
≥ 10	0.12	0.01	0.01	0.01	0.02	0.02	0.02	0.02	0.03	0.03	0.71

*Based on monthly U.S. customs records. Buyer-seller pairs are considered to be matched during the period between their first observed shipment and their last observed shipment. Matches that generate shipments in the first sample year (1996) are treated as active from the beginning of the sample and matches that generate shipments in the last sample year (2011) are treated as active through the end of the sample. Buyer-seller pairs that generate a single shipment are not considered to have matched.

2.3.4 Seller heterogeneity within buyers

In addition to cross-firm heterogeneity in partner counts, we observe substantial heterogeneity across sellers in the purchases of a given buyer. Table 4 reports the average share of the largest supplier to buyers with a single supplier, two suppliers, and so on. Apparel importers typically have a dominant supplier whose share in total imports drops only modestly as the number of suppliers increases. These dominant supplier market shares help us identify the distribution of seller types, as we discuss in section 4.2 below.

2.3.5 Search costs

The frequent dissolution of buyer-seller partnerships suggests that the cost of maintaining a network of business connections is high, regardless of whether a firm uses its own sourcing agents, a third-party sourcing firm, or some combination.²⁰ What form do these costs take?

¹⁹This power-law feature appears in data from other countries, including Colombia (Eaton et al., 2008, 2021; Bernard et al., 2018) and Norway (Bernard et al., 2018). In our sample both distributions remain roughly Pareto over time, but the shape parameter for buyers per seller rises from 1.99 in 2000 to 2.84 in 2011, reflecting an increase in the relative number of sellers to buyers.

²⁰WalMart’s sourcing budget was \$10 billion circa 2011 (McFarlan et al., 2012), while its gross income was \$110 billion in 2012 (downloaded December 27, 2016 from <http://www.marketwatch.com/investing/stock/wmt/financials>). Neither figure is specific to apparel.

Table 3: Firm Distributions by Partner Counts, 2011

x	Frac. buyer with at most x sellers	Frac. seller with at most x buyers
1	0.407	0.798
2	0.554	0.951
3	0.645	0.970
4	0.709	0.980
5	0.743	0.987
6	0.780	0.991
7	0.808	0.993
8	0.823	0.995
9	0.837	0.996
10	0.855	0.997

Table 4: Within-buyer Share of Largest Seller

Number of sellers per buyer	Average share, largest seller
1	1.0
2	0.763
3	0.660
4	0.600
5	0.553

On the buyers’ side, a case study of U.S. apparel import intermediaries quoted one respondent on the importance of visiting manufacturers’ factories and learning their capabilities: “You know, [go] into the factory and see what they’re making for other people, or what their lines do, and then basically [give] them that type of products. . . . [T]o go to somebody who makes cotton underpants, and give them synthetic with charms, it’s not the right thing to do because they’re not gonna be the best of that” (Ha-Brookshire and Dyer, 2008). Buyers also wish to avoid factories that fall short in terms of shop floor safety, child labor standards, and environmental impact.²¹ Since each importer has its own standards regarding acceptable practice, the industry norm is for each firm to perform an audit of each factory it deals with before placing any orders.²²

We know less about search efforts on the manufacturers’ side. As mentioned above,

²¹In a 2016 survey of U.S. apparel importers, “33 percent rated ‘unmet social and environmental compliance’ as having a high or very high impact on their supply chain, much higher than concerns for other supply chain risks such as ‘labor disputes,’ ‘political unrest,’ and ‘lack of resources to manage supply chain risks.’ ” (Lu, 2016).

²²This observation is based on a telephone interview with the president of the U.S. Fashion Industry Association, December 14, 2016.

some manufacturers attend trade fairs. Interviews with manufacturers of plastics products in Colombia suggest that the costs of finding foreign buyers can include maintenance of an appealing website in English, web searches for firms abroad that buy one’s type of product, maintenance of a marketing staff, and maintenance of sales offices in destination markets (Dominguez et al., 2021).

In summary, US apparel imports have risen substantially, with steady increases in the numbers of sellers and buyers. Concurrent changes in the trade policy environment such as China’s entry into the WTO and the phaseout of the ATC affected the composition of sellers in this market. Underlying these aggregate changes are the fluid creation and destruction of business relationships between heterogeneous importers and exporters. Our model of two-sided search accounts for these developments in an equilibrium framework. We use the model as a lab to investigate the effects of trade policies and a general reduction of search costs facilitated, say, by the IT revolution.

3 A model of buyer-seller networks

We now develop a continuous-time two-sided search model. As discussed above, we keep this exercise tractable by treating the importing agents in Figure 5 as a single type of firm. The resulting structure, depicted in Figure 7, thus involves three types of agents: sellers (exporting manufacturers), buyers (importing retailers), and consumers.

Consumers acquire goods exclusively through retailers, who offer different but possibly overlapping menus of products, depending upon their current suppliers. A consumer derives utility not only from the products acquired at a retailer, but from amenities the retailer itself may offer, such as a convenient location, pleasant ambience, or attentive service. How a consumer allocates spending across retailers reflects preferences both for the products offered and retailer amenities.²³

Since our focus is on matching between manufacturers and retailers, we assume that consumers choose among retailers with full information about their amenities and product lines.

²³The market demand functions implied by our model can represent an aggregation of consumer-specific nested-logit demand functions, with each consumer purchasing a single unit of the most-preferred product (Verboven,1996).

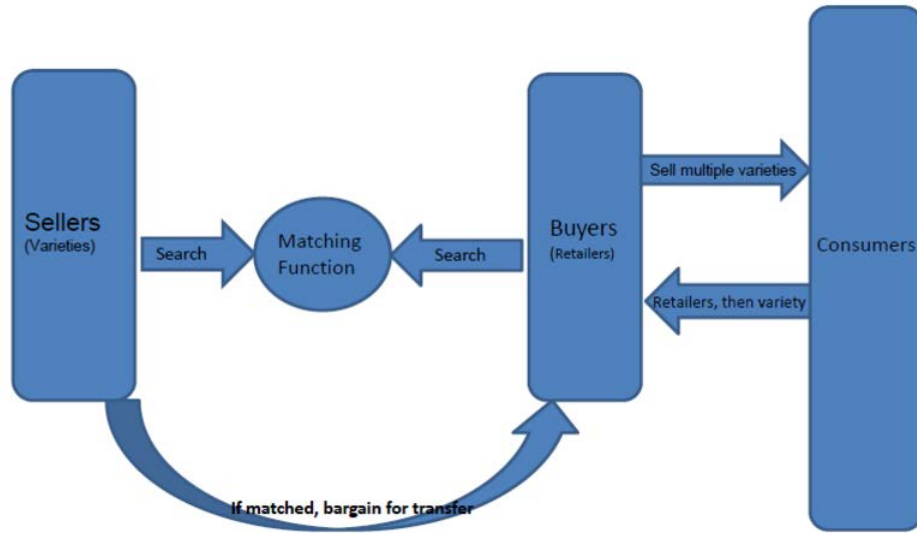


Figure 7: Model diagram

In the wholesale market, however, a buyer or seller can match only by investing in costly search.

For buyers and sellers to be willing to spend on searching for each other, their meetings must create rents. They bargain continuously and bilaterally over these rents, with the expected outcomes determining the returns to search for each.

Other things equal, more intense search increases the hazard of finding a new partner. But the hazard depends on two other factors as well. One is market tightness. When many buyers search for new suppliers, while few suppliers are searching for new buyers, matching hazards will tend to be low for buyers and high for suppliers. Our matching function, which we adopt from the labor-search literature, specifies how search intensities on each side of the market determine aggregate market tightness. Second, how easily agents find each other depends upon their previous successes. We allow agents who have already accumulated a large portfolio of business partners to find it easier to locate still more. This feature of our model, taken from Eaton et al. (2021), helps us capture the “fat-tailed” distributions of buyers across sellers and sellers across buyers discussed above.

3.1 The Retail Market

We now turn to our specification of the retail sector.

Preferences and pricing: As in Atkin et al. (2018) and Hottman et al. (2016), we start from a nested CES demand structure in which consumers have preferences over retailers, and within retailers, over products. Specifically, assume the retail market is populated by a continuum of stores. The utility that a store provides a consumer depends both on its amenities and the set of products it offers, which need not be disjoint across stores. Indexing stores by y and products (which correspond to exporting firms) by x , consumer preferences over retailers are

$$C = \left[\int_{y \in Y} (\mu_y C_y)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

where,

$$C_y = \left[\sum_{x \in J_y} (\xi_x C_x)^{\frac{\alpha-1}{\alpha}} \right]^{\frac{\alpha}{\alpha-1}},$$

J_y is the set of products offered at store y , while μ_y and ξ_x are exogenous parameters reflecting the appeal of retailer y and product x , respectively.²⁴

The price index for retailer y is thus

$$p_y = \left[\sum_{x \in J_y} \left(\frac{p_{xy}}{\xi_x} \right)^{1-\alpha} \right]^{\frac{1}{1-\alpha}}$$

while the price index for retailers as a group is

$$P = \left[\int_y \left(\frac{p_y}{\mu_y} \right)^{1-\eta} \right]^{\frac{1}{1-\eta}}.$$

Search frictions prevent retailers from instantaneously adjusting their product offerings.

²⁴The parameter μ_y could reflect, among other things, the availability of products outside those provided by our exporters at retailer y . Atkin et al. (2018) use the same nesting structure although, unlike us, they do not limit their analysis to a single product group. Alternative nesting structures are possible. In particular, consumers might have preferences over bundles of types of goods, each of which is a CES aggregation over the bundles available from alternative retailers. That is, consumers first allocate spending across product categories, then across retailers in each category. Which specification is preferable depends upon the importance of transport and shopping time costs to consumers.

Rather, at each moment they set their own retail prices taking current offerings and other retailers' prices as given. The first-order conditions for store y 's prices satisfy

$$q_{xy} + \sum_{x' \in J_y} \frac{\partial q_{x'y}}{\partial p_{xy}} (p_{x'y} - c_{x'y}) = 0 \quad \forall x \in J_y, \quad (1)$$

where $c_{x'y}$ is the marginal cost of supplying product x' to consumers through retailer y , including both the manufacturing and shipping costs incurred by the producer of x' and the retailing costs incurred by y .²⁵

Operating profits: Equation (1) implies the standard result that the within-retailer cannibalization effect exactly offsets the cross-store substitution effect, so the mark-up is simply (Atkeson and Burstein, 2008; Hottman et al., 2016):

$$\frac{p_{xy} - c_{xy}}{p_{xy}} = \frac{1}{\eta}. \quad (2)$$

The instantaneous profit flow jointly generated by retailer y and its suppliers is²⁶

$$\pi_y^T = \frac{E}{\eta P^{1-\eta}} \left[\sum_{x \in J_y} \left(\frac{\eta}{\eta - 1} \right)^{1-\alpha} \tilde{c}_{xy}^{1-\alpha} \right]^{\frac{1-\eta}{1-\alpha}} \mu_y^{\eta-1}, \quad (3)$$

where $\tilde{c}_{xy} = \frac{c_{xy}}{\xi_x}$ is the quality-adjusted marginal cost incurred by buyer-seller pair $x - y$ per unit supplied in the retail market.

3.2 The Wholesale Market and Payoff Functions

We now turn to the flow payoffs for buyers (importers) and sellers (foreign suppliers) in the wholesale market.

Buyer-seller transfers: Suppose there are I intrinsic types of buyers indexed by $i \in \{1, 2, \dots, I\}$, so that if buyer y is a type- i retailer, $\mu_y = \mu_i$. Similarly, suppose there are J intrinsic types of sellers indexed by $j \in \{1, 2, \dots, J\}$, so that if seller x is type- j exporter,

²⁵Buyer-seller pairs set retail prices to maximize the joint value of the surplus generated by their business relationship. Hence c_{xy} is the cost for the *pair* to offer the product in the retail market. We thus assume that bargaining between seller and buyer is efficient, eliminating problems of double marginalization.

²⁶See appendix B for details.

matches between this seller and a type- i buyer generate a quality-adjusted marginal cost of \tilde{c}_{ij} . Finally, let $\mathbf{s} = [s_1, s_2, \dots, s_J]$ be a vector of counts of the number of sellers of each type currently matched to a particular buyer, and let $\mathbf{b} = [b_1, \dots, b_I]$ be a vector of counts of the number of buyers of each type currently matched to a particular seller. Then, by equation (3), the gross surplus flow accruing to a type- i buyer and its portfolio of suppliers \mathbf{s} is:

$$\pi_i^T(\mathbf{s}) = \frac{E}{\eta P^{1-\eta}} \left[\sum_j \left(\frac{\eta}{\eta-1} \right)^{1-\alpha} s_j \tilde{c}_{ij}^{1-\alpha} \right]^{\frac{\eta-1}{\alpha-1}} \mu_i^{\eta-1} \quad (4)$$

Note that when the elasticity of substitution across products exceeds the elasticity of substitution across retailers ($\alpha > \eta > 1$), this surplus exhibits diminishing returns with respect to the number of suppliers of any type. That is, buyers who add additional sellers reduce total surplus per supplier.

How does this flow of surplus get divided between a particular buyer and her portfolio of sellers? We assume that the surplus associated with a particular buyer-seller match is divided up according to the ‘‘Rolodex’’ bargaining protocol, as described by Brugemann et al. (2019).²⁷ As demonstrated in Appendix C, at each moment the profit transferred to each seller of type j is

$$\tau_{ji}(\mathbf{s}) \approx \beta \frac{\partial \pi_i^T(\mathbf{s})}{\partial s_j} = \frac{\beta}{\alpha-1} \left(\frac{\eta}{\eta-1} \right)^{-\eta} \frac{E}{P^{1-\eta}} \left[\sum_{j'} s_{j'} \tilde{c}_{ij'}^{1-\alpha} \right]^{\frac{\alpha-\eta}{1-\alpha}} \tilde{c}_{ij}^{1-\alpha} \mu_i^{\eta-1} \quad (5)$$

where $\beta \in [0, 1]$ is a parameter measuring the bargaining strength of the seller, and the equality is approximate because we’ve used a derivative to describe a discrete one-unit change in s_j .

Our nested CES demand system parametrically fixes gross surpluses, variable costs, and, given a bargaining weight, transfers as fractions of consumer expenditures.²⁸ Specifically, by equation (2), a relation between a type- i buyer and a type- j seller that generates $E_{ij}(\mathbf{s}) =$

²⁷The Rolodex bargaining game is a modified version of the Stole and Zwiebel (1996) bargaining protocol, under which buyers bargain continuously with each of the sellers they’re matched with, treating each as the marginal supplier. The Brugemann et al. (2019) modification is the following: ‘‘when a [seller] rejects a counteroffer from the firm, [it] moves to the end of the queue and the firm enters a bargaining session with the seller who, among those without agreement, is now at the top of the queue.’’ The equilibrium payoffs in this game coincide with the Shapley values.

²⁸Alvarez et al. (2020) show that these properties of our demand system can be relaxed by allowing for increasing marginal costs among suppliers.

$\left(\frac{P_{ij}}{\xi_j}\right)^{1-\alpha} \left(\frac{P_i(\mathbf{s})}{\mu_i}\right)^{\frac{1-\eta}{1-\alpha}} P^{\eta-1} E$ in revenues will yield $\pi_{ij}^T = \frac{E_{ij}(\mathbf{s})}{\eta}$ in surplus, and variable production costs $\left(\frac{\eta-1}{\eta}\right) E_{ij}(\mathbf{s})$. Therefore, the total surplus generated by any type- i buyer in state \mathbf{s} is simply $\pi_i^T = \frac{E_i(\mathbf{s})}{\eta}$, where $E_i(\mathbf{s}) = \sum_{j=1}^J s_j E_{ij}(\mathbf{s})$ is the total amount of consumer expenditure they attract. Also, a type- i buyer in state \mathbf{s} transfers surplus $\sum_{j=1}^J s_j \tau_{ji}(\mathbf{s}) = \beta \left(\frac{\eta-1}{\alpha-1}\right) \frac{E_i(\mathbf{s})}{\eta}$ to her suppliers, keeping $\pi_i^B(\mathbf{s}) = \left[1 - \beta \left(\frac{\eta-1}{\alpha-1}\right)\right] \frac{E_i(\mathbf{s})}{\eta}$ in gross profits (before search costs) for herself.

3.3 Search and Matching

We now turn to how buyers and sellers search and meet each other.

3.3.1 Market Aggregates and Market Slackness

How easily an agent matches with a partner depends on her “visibility.” For any two agents on the same side of the market, the ratio of their visibilities is also the ratio of their hazards for meeting a new business partner.

Let $M_i^B(\mathbf{s})$ be the measure of type- i buyers with seller portfolio \mathbf{s} , and specify these buyers’ visibility as:

$$H_i^B(\mathbf{s}) = \sigma_i^B(\mathbf{s}) M_i^B(\mathbf{s})$$

where $\sigma_i^B(\mathbf{s})$ is the search intensity of any one of these buyers. Aggregating over all I types of buyers and all possible seller portfolios $\mathbf{s} \in \mathbb{S}$, the overall visibility of buyers is:

$$H^B = \sum_{i=1}^I \sum_{\mathbf{s} \in \mathbb{S}} H_i^B(\mathbf{s})$$

Analogously, let $M_j^S(\mathbf{b})$ be the measure of sellers of type j with buyer portfolio $\mathbf{b} \in \mathbb{B}$. If each of these sellers searches with intensity $\sigma_j^S(\mathbf{b})$, this group’s visibility is:

$$H_j^S(\mathbf{b}) = \sigma_j^S(\mathbf{b}) M_j^S(\mathbf{b}),$$

and the overall visibility of sellers is:

$$H^S = \sum_{j=1}^J \sum_{\mathbf{b} \in \mathbb{B}} H_j^S(\mathbf{b})$$

Following much of the labor search literature, we assume a matching function that is homogeneous of degree one in the visibility of buyers and sellers. Specifically, we assume that the measure of matches per unit time is (Petrongolo and Pissarides, 2001):²⁹

$$X = f(H^S, H^B) = H^B \left[1 - \left(1 - \frac{1}{H^B} \right)^{H^S} \right] \approx H^B \left[1 - e^{-H^S/H^B} \right] \quad (6)$$

From buyers' perspective, we define market slackness as:

$$\theta^B = \frac{f(H^S, H^B)}{H^B} \quad (7)$$

The larger θ^B , the more matches take place per unit of buyer visibility. Likewise, market slackness from sellers' perspective is:

$$\theta^S = \frac{f(H^S, H^B)}{H^S}. \quad (8)$$

The share of matches involving buyers of type i with match portfolio \mathbf{s} is

$$v_i^B(\mathbf{s}) = \frac{H_i^B(\mathbf{s})}{H^B},$$

and the share of matches involving sellers of type j with match portfolio \mathbf{b} is

$$v_j^S(\mathbf{b}) = \frac{H_j^S(\mathbf{b})}{H^S}.$$

²⁹We have also experimented with the specification

$$x = \frac{H^B H^S}{\left[(H^B)^\phi + (H^S)^\phi \right]^{1/\phi}}$$

3.3.2 Optimal search

To derive the policy functions $\sigma_i^B(\mathbf{s})$ and $\sigma_j^S(\mathbf{b})$ we adopt the search cost function used in Eaton et al. (2021), which expresses the cost of search for any buyer or seller as a function of its search effort σ^A and the number of its current business partners, n^A :³⁰

$$k(\sigma^A, n^A) = \frac{k_0 (\sigma^A)^2}{(n^A + 1)^{\gamma^A}}, \quad A \in \{B, S\}. \quad (9)$$

Here $n^B = \sum_{j=1}^J s_j$ and $n^S = \sum_{i=1}^I b_i$. This functional form, which generalizes Arkolakis (2010), implies that the cost of search is convex in search intensity and that

$$\lim_{\sigma^A \rightarrow 0} k(\sigma^A, n^A) = 0.$$

The parameter γ^A determines the how having more connections affects search costs, capturing the relative strengths of greater “visibility” from having more current partners ($\gamma^A > 0$) or of “fishing out” effects that make it increasingly difficult to find new partners when the easy opportunities have already been exploited ($\gamma^A < 0$).

Buyer’s problem: The flow value of a type- i buyer who is currently in state \mathbf{s} is:

$$\begin{aligned} \rho V_i^B(\mathbf{s}) &= \pi_i^B(\mathbf{s}) - k_s^B(\sigma_i^B(\mathbf{s})) + \sigma_i^B(\mathbf{s})\theta^B \sum_{j=1}^J v_j^S [V_i^B(\mathbf{s} + \mathbf{1}_j) - V_i^B(\mathbf{s})] \\ &\quad + \delta \sum_{j=1}^J s_j [V_i^B(\mathbf{s} - \mathbf{1}_j) - V_i^B(\mathbf{s})] \end{aligned}$$

where $\mathbf{1}_j$ is a $J \times 1$ vector with j^{th} element 1 and 0’s elsewhere and $v_j^S = \sum_{\mathbf{b} \in \mathbb{B}} v_j^S(\mathbf{b})$ is the probability that the next seller the buyer meets will be type- j . The seller reaps profit flow $\pi_i^B(\mathbf{s}) - k_s^B(\sigma_i^B(\mathbf{s}))$ until the next event occurs. With hazard $s_j\delta$ this event is the exogenous termination of a type- j seller relationship, and with hazard $\sigma_i^B(\mathbf{s})\theta^B v_j^S$ it’s a new match with a type- j seller.³¹ The optimal search policy for type- i buyers with \mathbf{s} sellers, $\sigma_i^B(\mathbf{s})$, therefore

³⁰More general functions need not impose a quadratic form or might allow the scalar k_0 to differ between buyers and sellers. Identifying additional parameters would require more data than we have. Nonetheless, the shapes of degree distributions did allow us to separately identify γ^B and γ^S , as discussed in section 4.2 below.

³¹The model can be extended to allow the match δ separation hazard to depend upon the type of seller.

satisfies

$$\frac{\partial k^B(\sigma_i^B, \mathbf{s})}{\partial \sigma_i^B} = \theta^B \sum_{j=1}^J v_j^S [V_i^B(\mathbf{s} + \mathbf{1}_j) - V_i^B(\mathbf{s})]. \quad (10)$$

Sellers' problem: Similarly, the value to a type- j seller of being matched with a type- i buyer in state \mathbf{s} is:

$$\begin{aligned} (\rho + \delta)V_{ji}^S(\mathbf{s}) &= \tau_{ji}(\mathbf{s}) + \sigma_i^B(\mathbf{s})\theta^B \sum_{k=1}^J v_k^S [V_{ji}^S(\mathbf{s} + \mathbf{1}_k) - V_{ji}^S(\mathbf{s})] \\ &\quad + \delta \sum_{k=1}^K (s_k - \mathbf{1}_{k=j}) [V_{ji}^S(\mathbf{s} - \mathbf{1}_k) - V_{ji}^S(\mathbf{s})] \end{aligned}$$

Intuitively, a business relationship with a type- i buyer who has \mathbf{s} suppliers will terminate with exogenous hazard δ , become a relationship with a type- i buyer who has $\mathbf{s} + \mathbf{1}_k$ suppliers with hazard $\sigma_i^B(\mathbf{s})\theta^B v_k^S$, and become a relationship with a type- i buyer who has $\mathbf{s} - \mathbf{1}_k$ suppliers with hazard $(s_k - \mathbf{1}_{k=j})\delta$.³² Given our random search assumption, the *ex ante* expected value of a new business relationship for a type- j seller is:

$$V_j^S = \sum_i \sum_{\mathbf{s} \in \mathbb{S}} v_i^B(\mathbf{s}) V_{ji}^S(\mathbf{s}).$$

So the optimal search intensity for any seller with a total of b active buyers satisfies:

$$\frac{\partial k^S(\sigma^S, \mathbf{b})}{\partial \sigma^S} = \theta^S V_j^S. \quad (11)$$

3.3.3 The Steady State

Given that all relationships end with exogenous hazard δ , the equations of motion for our measure of type- i buyers with \mathbf{s} sellers is:

³²The factor $(s_k - \mathbf{1}_{k=j})$ adjusts for the risk of being dropped when exogenous separations occur. Suppose a particular type- j seller, call it x , is partnered with a particular buyer, y . If buyer y loses one of its type- j sellers, it will be seller x with probability $1/s_j$.

$$\begin{aligned} \dot{M}_i^B(\mathbf{s}) &= \sum_j [\sigma_i^B(\mathbf{s} - \mathbf{1}_j)\theta^B v_j^S M_i^B(\mathbf{s} - \mathbf{1}_j) + \delta(s_j + 1)M_i^B(\mathbf{s} + \mathbf{1}_j)] \\ &\quad - [\sigma_i^B(\mathbf{s})\theta^B M_i^B(\mathbf{s}) + \delta n^B(\mathbf{s})M_i^B(\mathbf{s})], \\ \mathbf{s} &\in \mathbb{S}; i = 1, \dots, I. \end{aligned} \quad (12)$$

This group gains a member whenever any of the $M_i^B(\mathbf{s} - \mathbf{1}_j)$ buyers in state $\mathbf{s} - \mathbf{1}_j$ gains a type- j supplier, which occurs with hazard $\sigma_i^B(\mathbf{s} - \mathbf{1}_j)\theta^B v_j^S$. Similarly, it gains a member whenever any of the $M_i^B(\mathbf{s} + \mathbf{1}_j)$ buyers in state $\mathbf{s} + \mathbf{1}_j$ loses a type- j supplier because of exogenous attrition, which occurs with hazard $\delta(s_j + 1)$. By analogous logic, the group loses existing members that either add a supplier (with hazard $\sigma_i^B(\mathbf{s})\theta^B$) or lose one (with hazard $\delta n^B(\mathbf{s}) = \delta \sum_j s_j$). Finally, the measure of buyers of type i with $s = 0$ sellers evolves according to:

$$\dot{M}_i^B(\mathbf{0}) = \delta \sum_j M_i^B(\mathbf{1}_j) - \sigma_i^B(\mathbf{0})\theta^B M_i^B(\mathbf{0}) \quad i = 1, \dots, I \quad (13)$$

Replacing \mathbb{S} with \mathbb{B} , \mathbf{s} with \mathbf{b} , and i with j in equations (12) and (13) gives the corresponding equations of motion for measures $M_j^S(\mathbf{b})$ of sellers.

To characterize the steady state of this system we set $\dot{M}_i^B(\mathbf{s}) = \dot{M}_j^S(\mathbf{b}) = 0$ and solve the system of $I \cdot (\|\mathbb{S}\| + 1) + J \cdot (\|\mathbb{B}\| + 1)$ equations implied by both versions of (12) and (13) for buyers and sellers.³³ In doing so we, treat the measures of each intrinsic type as exogenous constants and impose the adding-up constraints:

$$M_i^B = \sum_{\mathbf{s} \in \mathbb{S}} M_i^B(\mathbf{s}) \quad (14)$$

$$M_j^S = \sum_{\mathbf{b} \in \mathbb{B}} M_j^S(\mathbf{b}), \quad (15)$$

Because the units of measurement for these objects are arbitrary, we impose $\sum M_i^B = 1$, and we treat the ratio of total sellers to total buyers as a parameter to be estimated. Given the policy functions, this normalization is sufficient to pin down the ratio of search efforts on each

³³Solving for transition dynamics is feasible but more involved. We limit our attention here to steady states.

side of the market, H^B/H^S , which in turn determines the market tightness indices θ^B and θ^S .

4 Fitting the model to data

We can separate the parameters of our model into those that determine payoff flows and those that, given the payoff parameters, determine search policies. We describe our approach to estimating each in turn, report our results, and assess the quality of the fit.

4.1 Transfer function parameters

Equation (5) provides the basis for estimating the key payoff parameters of our model, but several transformations are necessary in order to bring it to the data.

First, the transfer of surplus to the seller $\tau_{ji}(\mathbf{s})$ is not observed, since payments to the seller in the data also include compensation for production costs. To convert equation (5) to an expression for total (observed) export payments, we assume that some fraction λ of the variable costs $c_{ij}q_{ji}$ incurred in an $i - j$ partnership is attributable to the seller. Total export payments (surplus plus cost reimbursement) can then be written as:

$$r_{ji}(\mathbf{s}) = \frac{E}{P^{1-\eta}} \left(\frac{\eta}{\eta-1} \right)^{-\eta} \left[\sum_{j'=1}^J s_{j'} \tilde{c}_{ij'}^{1-\alpha} \right]^{\frac{\alpha-\eta}{1-\alpha}} \tilde{c}_{ij}^{1-\alpha} \mu_i^{\eta-1} \left[\frac{\beta}{\alpha-1} + \lambda \right]. \quad (16)$$

Note that equation (16) retains the form of equation (5) because surpluses and variable production costs are proportional to one another, as discussed in section 3.2.

Next we eliminate the (unobserved) sum in square brackets by using the within-buyer i revenue share of a type- j seller:

$$h_{j|i} = \frac{\tilde{c}_{ij}^{1-\alpha}}{\sum_{j'=1}^J s_{j'} \tilde{c}_{ij'}^{1-\alpha}} \quad (17)$$

We can then rewrite equation (16) in terms of observables and fixed effects as:

$$r_{ji}(\mathbf{s}) = (h_{j|i})^{\frac{\alpha-\eta}{\alpha-1}} \frac{E}{P^{1-\eta}} \left(\frac{\eta}{\eta-1} \right)^{-\eta} \left(\frac{\mu_i}{\tilde{c}_{ij}} \right)^{\eta-1} \left[\frac{\beta}{\alpha-1} + \lambda \right] \quad (18)$$

Defining the coefficient

$$\tilde{\alpha} = \frac{\alpha - \eta}{\alpha - 1},$$

and the match effect

$$v_{ij} = \ln \left\{ \left(\frac{\eta}{\eta - 1} \right)^{-\eta} \left[\frac{\beta}{\alpha - 1} + \lambda \right] \right\} + \ln \left(\frac{\mu_i}{\tilde{c}_{ij}} \right)^{\eta-1},$$

equation 18, in logs, becomes our estimating equation:

$$\ln r_{jit} = \tilde{\alpha} \ln h_{j|i,t} + v_{ij} + d_t + \varepsilon_{jit} \quad (19)$$

where time effects d_t absorb temporal variation in $\ln(E/P^{1-\eta})$, and ε_{jit} is a time-varying match-specific shock capturing factors outside our model.

What might such factors be? In our model random matching is the only driver of within-match variation in $h_{j|i}$. But match-specific demand or marginal cost shocks may also drive variation in r_{jit} . Moreover, to the extent that buyers also source merchandise from domestic suppliers, there is likely to be measurement error in $h_{j|i,t}$ that isn't completely absorbed by the fixed effects. Finally, intertemporal variation in sellers' shares of match-specific marginal costs (λ) might also introduce noise into the relationship between exporter revenues and their shares in buyer payments.

Such sources of transitory shocks to r_{ijt} may induce covariation in $h_{j|i,t}$, while measurement error in $h_{j|i,t}$ can induce attenuation bias. To address such problems we use an instrumental variable approach to estimating $\tilde{\alpha}$. To construct our instrument, we exploit information on the rivals of seller j who also supply buyer i in period t . Specifically, we use the cross-buyer average log sales of these rival firms to buyers *other than* buyer i . This variable, denoted $\ln \bar{r}_{-j|-i,t}^c$, should reflect the marginal costs of the rival sellers and thus be correlated with j 's share of buyer i 's imports. Yet, since their sales to buyer i are excluded from the average, it is arguably uncorrelated with idiosyncratic shocks to buyer i . Similarly, since all of seller j 's sales are excluded, $\ln \bar{r}_{-j|-i,t}^c$ should be orthogonal to idiosyncratic shocks that affect seller j . A variant of this instrument, which we denote $\ln \bar{r}_{-j,-i}^c$ averages over rivals' sellers to not- i buyers in *all* years we observe, not just year t .

Table 5: Transfer function estimates*

	1st Stage IV-FE (dependent variable: $\ln h_{j i,t}$)		2nd stage IV-FE (dependent variable: $\ln r_{jit}$)	
	(1)	(2)	(3)	(4)
$\tilde{\alpha}$	–	–	0.1937 (0.0083)	0.0725 (0.0108)
$\ln \bar{r}_{-j -i,t}^c$	–0.2388 (0.0024)	–	–	–
$\ln \bar{r}_{-j -i}^c$	–	–0.3810 (0.0044)	–	–
instrument	–	–	$\ln \bar{r}_{-j -i,t}^c$	$\ln \bar{r}_{-j -i}^c$
match effects	yes	yes	yes	yes
year effects	yes	yes	yes	yes
R ²	0.8896	0.9045	0.8615	0.8314
obs.	771,200	1,256,000	771,200	1,256,000

*Alternative instrumental variable estimates of equation 19 appear in columns (3) and (4). They are based on predicted values of $\ln h_{j|i,t}$ from the regressions reported in columns (1) and (2), respectively. The instrument in column (1), $\ln \bar{r}_{-j|-i,t}^c$, is the average current period log imports of buyers other than j from non- i suppliers to j . The instrument in column (2), $\ln \bar{r}_{-j|-i}^c$, is average log sales in *all* periods by these non- i suppliers of j to non- j importers.

Both instruments should be negatively correlated with $\ln h_{j|i,t}$ because rivals that sell large quantities to non- i buyers are likely to sell large quantities to buyer i too, driving down j 's share of i 's total imports. This is indeed the case, as can be seen in the first two columns of Table 5. The results for the two instruments are qualitatively similar, though the instrument based on rival sales in all years, $\ln \bar{r}_{-j,-i}^c$, yields a stronger negative relationship. Note that these regressions control for fixed match effects, and thus isolate variation in $\ln h_{j|i,t}$ over the course of seller j 's relationship with buyer i , as rival sellers come and go.

The last two columns of Table 5 report IV estimates of $\tilde{\alpha}$. Column (3) corresponds to the instrument based exclusively on rivals' contemporaneous sales, $\ln \bar{r}_{-j|-i,t}^c$, and column (4) corresponds to the instrument based on rival sales in all years they are observed, $\ln \bar{r}_{-j,-i}^c$. The results imply that there are modestly diminishing returns to adding additional sellers of any type, or put differently, the elasticity of substitution across varieties within a store (α) exceeds the elasticity of substitution across stores (η).

Since α and η are not separately identified we take $\alpha = 4.35$ from Hottman et al.

(2016) and use our estimate $\tilde{\alpha} = 0.193$ to calculate $\eta = 3.68$. Also, to infer $var(\ln \mu_i)$ and $var(\ln \xi_{ij})$ from our fixed effects estimates, we assume that all variation in marginal costs \tilde{c}_j is due to product quality variation ($var(\ln \tilde{c}_j) = var(\ln(c_{jb}/\xi_{jb})) = var(\ln \xi_j)$). We use our assumption of random matching to write the variance of the match effect as $var(v_{ij}) = (\eta - 1)[var(\ln \mu_i) + var(\ln \xi_j)]$. Then projecting the estimated match effects on buyer effects and seller effects, we calculate that $var(\ln \mu_i) = 0.73$ and $var(\ln \xi_j) = 0.24$.³⁴ Finally, we discretize these distributions, allowing for $I = 30$ buyer types and $J = 2$ seller types. For buyers we impose a discretized normal distribution and use the method suggested by Kennan (2006). For sellers we use a non-parametric distribution with parameter $\omega = \Pr(\ln \xi = \ln \xi_2)$, where $\xi_2 > \xi_1$.³⁵ These choices reflect the relatively low variance in seller types and the fact that the dimensionality of our numerical problem scales more quickly with J .

4.2 Search and network parameters

It remains to estimate the match separation hazard δ , the search cost parameters (k_0, γ^B, γ^S), the discount rate plus firm death hazard ρ , the ratio of potential sellers to potential buyers $N_s = \sum_j M_j^S / \sum_i M_i^B$, and the seller distribution parameter ω . To identify δ , we use the Poisson parameter value that best fits the match death hazard observed in the customs records, adjusted for the buyer death hazard of 0.07 and the seller death hazard of 0.15. This calculation yields $\delta = 0.80 - 0.07 - 0.15 = 0.58$.³⁶

For the remaining parameters, we use the method of simulated moments, basing identification on the patterns documented in section 2 above. For the degree distributions, we target the empirical distributions described in Table 3, evaluated at partner counts of 1, 2, 3, 4, 5, 10 and 15. For the transition matrices, we directly target the conditional probabilities reported

³⁴The correlation between the buyer and seller fixed effects is -0.031, suggesting assortativity does not play a large role in matching patterns. Together, the buyer and seller effects account for 75 percent of the variation in the match effects. Hence our payoff function misses some of the variation in revenues. An explanation for this residual noise in match-specific payoffs is idiosyncratic variation in marginal costs unrelated to product or store appeal, c_{jb} . We leave it out of our model to simplify calculations.

³⁵To infer the distribution of $\ln \xi$ we impose $E(\ln \xi_j) = 0$ and fix $(\ln \xi_1, \ln \xi_2)$ to be the 25th and 75th percentile of the estimated distribution of seller fixed effects.

³⁶The buyer death hazard is the exit rate among U.S. retail firms (Jarmin et al., 2009). The seller death hazard is the average exit rate among Chinese apparel producers during the period 2004-2006 (Zhu, 2014), which we take to be representative of the emerging market exporters who supplied the U.S. market during the sample period.

in Tables 1 and 2, including diagonals and up to 5 off-diagonal elements in each row. For the share of the largest seller, we target the elements in the right-hand column of Table 4, omitting the first row. Finally, for the number of potential suppliers per potential buyer, we target the observed ratio of active sellers to active buyers, i.e., the data-based counterpart to $\left[\sum_j (M_j^S - M_j^S(\mathbf{0})) \right] / \left[\sum_i (M_i^B - M_i^B(\mathbf{0})) \right]$.

Our estimator takes the standard GMM form:

$$\hat{\Lambda} = \arg \min (m(\Lambda) - \bar{m}) \cdot W \cdot (m(\Lambda) - \bar{m})' \quad (20)$$

where \bar{m} is the vector of 158 targeted sample moments, $m(\Lambda)$ is the corresponding model-based vector of moments, W is the efficient weighting matrix, and $\Lambda = (k_0, \gamma^B, \gamma^S, \omega, N_s)$. Because the targeted degree distributions and transition matrices can be expressed as collections of cumulative probabilities, it's straightforward to construct closed-form expressions for W . Appendix section D provides details.

Parameter estimates. Table 6 reports the 5 estimated parameters. Note first that the network parameters γ^B and γ^S are positive and highly significant, with $\gamma^B < \gamma^S$. So it appears that visibility effects are positively related to the number of existing connections, and are especially strong for sellers. In large part, identification of these parameters comes from the “fat-tail” of the empirical degree distributions in Table 3. Matches frequently expire, and the larger a firm’s portfolio of business partners, the greater are its replacement needs. Without the advantage of higher visibility for highly connected firms, it would be unprofitable for them to maintain their large portfolios of partners.³⁷

To capture the difference between the buyer and the seller degree distributions, our model also needs to generate 2.9 active sellers per active buyer. The difference in network parameters helps do this ($\gamma^B < \gamma^S$) but more is needed, so the estimated ratio of potential sellers to potential buyers is high ($N^S = 7.97$).

Finally, the fraction of potential sellers that are high-appeal (ω) captures the relationship between the share in sales of the largest seller and the total number of sellers. With only

³⁷Eaton et al. (2021) found strong visibility effects through similar identification.

Table 6: Cost and Distributional Parameters*

	estimate (std. error)
k_0	0.011 (0.001)
γ^B	0.459 (0.022)
γ^S	0.691 (0.008)
ω	0.045 (0.000)
N^S	7.967 (0.278)
objective function	23,132.5

*Efficient GMM Estimates of $\hat{\Lambda}$ based on equation 20. Moments used for identification are reported in Tables 1, 2, 3, and 4. For Tables 1 and 2, we exclude the first column to limit the impact of noisy importer identifiers on our estimates. We also exclude elements describing jumps of more than 4 matches in either direction, which are very low probability events. This leaves 70 targeted elements from each table. For Table 4 we use all elements except the first row, which is trivially equal to 1. For Table 3, we use the measurements at 1, 2, 3, 4, 5, 10 and 15 partners. The total number of targeted moments is thus 158.

two types of sellers the related moments are heavily over-identified.³⁸ We find that around 5 percent of the potential exporter population is “high-appeal.”

Model fit. Figure 8 summarizes the model fit. Among the 140 elements of the transition matrices we target, only one is substantially off.³⁹ (Refer to the first quadrant.) Hence the model accurately captures the general tendency for firms to lose clients over time, as well as the tendency for buyers and sellers that are new to the U.S. market to ramp up the number of their matches during their early years. By accurately predicting these features of transitions, we can replicate the distribution of partner counts for buyers and sellers. (Refer to the second and third quadrant.) The only moments our model systematically misses are the fractions of sales due to the largest supplier among firms of different size, which we understate by 0.07 to 0.10, depending upon the number of partners. (Refer to the fourth quadrant.)

³⁸While adding seller types would have improved the fit, it would have substantially increased the dimensionality of the computational problem.

³⁹We underestimate the probability of retaining at least 10 buyers among sellers who have already reached this level, a very small group in our sample.

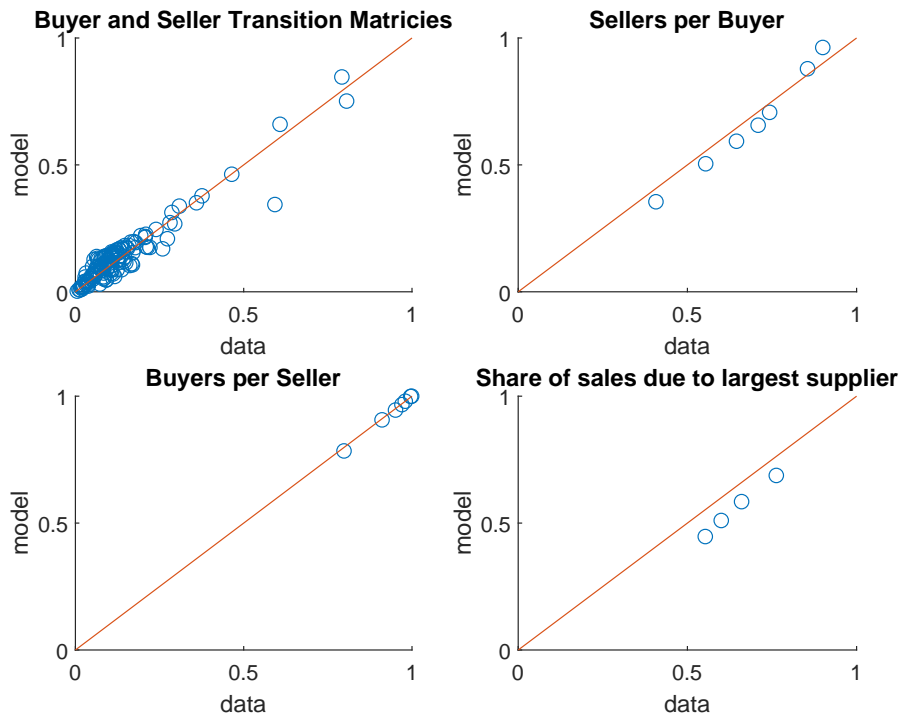


Figure 8: Data-based versus model-based moments

5 Putting the model to work

We now turn to some of the model’s key implications.

5.1 Profits, search costs, and maturation

We start by using our baseline model estimation to quantify the search costs borne by importers (buyers) and exporters (sellers), and their evolution with market tenure. We report these results in column 1 of Table 7.

Recall from Section 3.2 that preference and bargaining parameters pin down several aggregates. In particular, given our parameterization, suppliers and buyers together receive a share $1/\eta = 0.272$ of total final expenditure E as profits gross of search costs, with the remaining industry revenues covering variable production and distribution costs. Out of this total share, sellers get $(\beta/\eta) \cdot ((\eta - 1)/(\alpha - 1)) = 0.109$ and buyers the remaining 0.163.⁴⁰ (Refer to row

⁴⁰To get a crude sense for the dollar value of these aggregates, refer back to Figure 1, which shows that the f.o.b. value of apparel imports amounted to roughly \$100 billion.

6 of Table 7.)

How do these gross profits compare with the search costs that firms incur to earn them? In the base year, we find that buyers spend a share 0.043 on search (column 1, row 9), leaving 0.120 as net profits. For their part, sellers devote 0.044 to search (column 1, rows 7-8), leaving 0.065 as net profits. So buyers and sellers each bear about half of the total cost of linking with each other. However, there are far fewer buyers than sellers (column 1, rows 1-3), so matches per buyer exceed matches per seller, and search costs per *match* are lower for the average buyer than for the average seller. Partly for this reason, net profits per buyer are substantially higher than net profits per seller (column 1, row 13-15).

What determines variation in search costs across different types of agents? For a given side of the market (buyer or seller), it mainly reflects cross-firm variation in the return to search. The reason is that the long run marginal cost schedule for all firms depend solely on search effort.⁴¹ But the marginal return to search for sellers, $\theta^S V_j^S$, increases with seller quality, ξ_j , and similarly, for any given buyer state \mathbf{s} , the marginal return to search for buyers, $\theta^B \sum_{j=1}^J v_j^S [V_i^B(\mathbf{s} + \ell_j) - V_i^B(\mathbf{s})]$ increases with μ_i . Hence, for example, active low- ξ sellers spend only a combined $0.006 \cdot E$ on search, while high- ξ sellers spend $0.038 \cdot E$, even though active low- ξ sellers are more than 4 times as common as high- ξ sellers (column 1, rows 1-2).

Several other factors determine search spending among buyers versus among sellers. First, the market tightness indices θ^S and θ^B scale the returns to search and thus directly influence search efforts. Second, the visibility effects, γ^B and γ^S , affect the slope of the search cost function, and happen to favor sellers over buyers ($\gamma^B < \gamma^S$). These model features, in combination with the distributions of potential buyer types and potential seller types, determine the distributions of clients across firms represented in Table 3, as well as the search efforts reported in Table 7. Finally, a retailer's marginal revenue declines as it adds product varieties. But, for an exporter, the expected revenue from a match with new retailer is independent of the number of its existing customers.

The discussion thus far has focused on the cross-section of search costs across heterogeneous

⁴¹From equation (9), the marginal cost of search is $\frac{\partial k^\ell(\sigma^A, n^A)}{\partial \sigma^A} = \frac{k_0 \sigma^A}{(n^A + 1)^{\gamma^A}}$. But for any particular firm, the long run average number of connections, $E(n^A)$, depends only on its search intensity, σ^A . So $E\left[\frac{\partial k^\ell(\sigma^A, n^A)}{\partial \sigma^A}\right]$ is, roughly speaking, a function of σ^A alone.

buyers and sellers. Interestingly, holding agent type fixed, spending on matching also varies with market tenure. One reason is that the cost of any given level of search falls with the size of a firm’s existing match portfolio (n^B or n^S), reflecting market visibility effects. (Refer to equation 9.) In itself this effect discourages buyers and sellers from searching when their portfolios are small, and it allows large firms to replace their expiring business relationships at lower cost. Second, as noted above, buyers face diminishing returns to additional matches. In combination with our convex search cost function, this reinforces their incentive to search relatively intensively when their client portfolios are small.

These two model features allow us to replicate accurately life cycle dynamics for surviving buyers and sellers as reflected in the transition matrices in Tables 1 and 2 (refer back to Figure 8). Both buyers and sellers eventually settle into their long run portfolio sizes, at which point they simply search to replace their randomly expiring matches. But the average buyer takes more than 25 years to do so, after which it’s matched with about 7 suppliers, while the average seller takes half as long and thereafter maintains connections with only 2 buyers. An implication of these results is that the adjustment to a trade shock or search technology shock may take many years to play out.

5.2 Interpreting market developments

Equipped with our baseline model and estimated parameters, we use our framework to conduct two types of experiments: a reduction in search costs and changes in the population of potential suppliers. We calibrate the search cost reduction to the expansion in the number of agents participating in the apparel market between 1998 and 2004, as depicted in Figures 2 and 6.⁴² We calibrate changes in the population of potential exporters to the sudden expansion in the number of active suppliers from 2004-2005. This year coincided with the phaseout of the ATC and the beginning of the boom in China’s exports to the United States (refer to Figure 2).

Given that we have estimated the baseline model using data from 2011, and that we treat 2005 – 2011 as the *post*-shock equilibrium, our experiments amount to simulating *pre*-shock counterfactuals, then characterizing movements from the *pre*- to the *post*-shock periods. For

⁴²We used the years 1998 – 2004 because 2005 and thereafter include the ATC phaseout, and the financial crisis of 2007-08, as well as its lingering effects. An implicit assumption is that search cost remains stable after 2004 and comparable in magnitude to those in our baseline equilibrium period, 2005 – 2011.

Table 7: Counterfactual experiments

	Baseline	Search cost reduction (lower k_0)	ATC phaseout (higher N_s)	Larger share, low- ξ exporters (lower ω)
1. measure, active low- ξ sellers	0.906	1.342	1.205	1.639
2. measure, active high- ξ sellers	0.214	1.172	1.364	0.462
3. measure, active buyers	0.344	1.246	0.988	0.886
4. total surplus, low- ξ sellers	0.021	1.008	1.211	2.218
5. total surplus, high- ξ sellers	0.088	0.998	0.959	0.883
6. total surplus, buyers	0.163	1.000	1.000	1.000
7. total search costs, low- ξ sellers	0.006	1.019	1.196	2.250
8. total search costs, high- ξ sellers	0.038	1.095	0.916	0.978
9. total search costs, buyers	0.043	0.856	1.053	1.279
10. search cost per low- ξ seller	0.007	0.759	0.993	1.373
11. search cost per high- ξ seller	0.175	0.934	0.672	2.116
12. search cost per buyer	0.126	0.687	1.065	1.443
13. net profits per low- ξ seller	0.016	0.748	1.011	1.345
14. net profits per high- ξ seller	0.233	0.799	0.729	1.779
15. net profits per buyer	0.348	0.855	0.994	1.045
16. share, low- ξ seller	0.809	1.028	0.978	1.487
17. share, low- ξ seller matches	0.627	0.967	1.093	1.547
18. buyer/seller ratio	0.307	0.954	0.802	0.804
19. θ_B : buyer match rate factor	0.959	1.001	1.032	0.989
20. consumer welfare	1.000	1.061	1.002	0.991

*Figures in columns 2-4 are ratios of post-shock to pre-shock values. Baseline figures reflect several normalizations. First, measures of active buyers and sellers are expressed as shares of the population of potential buyers. Second, surpluses, profits, and search costs are expressed as shares of total consumer expenditures. Finally, baseline consumer welfare is normalized to unity.

the search cost experiment (column 2, Table 7), this means increasing k_0 sufficiently that the measure of active market participants increases from the *pre*-shock equilibrium to the baseline by 25 percent, approximately the observed cumulative growth in active agents between 1998 and 2004. For experiments involving changes in the population of potential exporters, this means adjusting N_s (column 3, Table 7) or ω (column 4, Table 7) sufficiently that there is a 25 percent increase in the ratio of active sellers to active buyers from the *pre*-shock equilibrium to the baseline. This is roughly the amount by which the ratio jumped in 2005, in line with the ATC phaseout.

5.2.1 Shocks to search costs

We design our search cost counterfactual to approximate the effects of the IT revolution on communications and market visibility. We capture such a technological shift by reducing the search cost scalar for buyers and sellers, k_0 .⁴³ Of course, this growth also reflected increases in market size and other developments which we didn't control for.

The second column of Table 7 reports the results. As expected, the measures of active buyers and active sellers both increase as search costs fall. But the measure of sellers increases more than the measure of buyers, so the ratio of buyers to sellers falls by 4.6 percent. This reduction in buyers per seller is due entirely to the rapid expansion of low-appeal sellers, even though the fraction ω of high-appeal sellers in the *potential* seller population remains unchanged. The larger adjustment by low-appeal firms reflects the fact that they tend to be closer to the entry-exit margin, while high-appeal sellers are more likely to be searching actively already, and thus to adjust only on the number of matches margin. In fact, the adjustment by high-type sellers is so large that the portion of matches accounted for by low-appeal sellers actually falls more than 3 percent.

Combined, the shift in match composition toward high-appeal sellers and the increase in the total number of active agents generate a 6.1 percent improvement in consumer welfare. But the increase in buyers and sellers means that consumer spending is spread more thinly across both types of agents. Thus net profits fall 14.5 percent for the average buyer, 20 percent for the average high-quality seller, and 15 percent for the average low-quality seller.

5.2.2 Shocks to the population of potential suppliers

The phaseout of the ATC on January 1, 2005 eliminated the need for a firm from China (or other low income country) to have a quota license to export apparel to the United States. At roughly the same time, China phased out regulations prohibiting small firms from dealing directly with foreign buyers. Khandelwal et al. (2013) and Bai et al. (2017) discuss how these

⁴³One could alternatively have targeted the measure of active participants by adjusting the the network parameters γ^B and γ^S . However, generating a 25 percent increase in market participants would require that γ^B and γ^S drop from 0.46 and 0.69 to -0.92 and -1.38, respectively, implying visibility effects become strongly dominated by congestion. It would also dramatically shift the sellers-per-buyer distribution by favoring small exporters over large ones. We view these features of such an experiment as sufficiently implausible to rule out network effects as the primary cause of observed growth in the population of active sellers.

reforms both increased the population of potential exporters and changed its composition.⁴⁴

Our last two experiments explore some possible effects of these shocks. The first asks how an increase in the measure of potential sellers per potential buyer (N_s) affects the equilibrium, holding constant the fraction of potential sellers that are high quality (ω). The second asks about raising the fraction of the potential seller population that are low quality ($1-\omega$), holding N_s constant.

Changes in potential sellers per potential buyer. Column 3 of Table 7 reports the effects of increasing N_s , the number of potential sellers per potential buyer. This shock has little effect on the measure of active buyers, and the increase in active sellers per buyer is not completely offset by adjustments in search intensities, so the market tightness index (θ^B) improves for buyers, rising about 3 percent. This congestion effect reduces the likelihood of matching per unit of search effort for any type of seller.

Despite a substantial increase in active sellers there is almost no effect on welfare. The main reason is that high-appeal sellers reduce their search effort dramatically, with no corresponding drop in search efforts per low-appeal seller. As a consequence, for the typical retailer the ratio of low- to high-appeal sellers rises by about 9 percent. The decline in quality offsets the increase in variety in overall retailer appeal.

A higher N^S actually raises profit net of search costs by about 1 percent for a typical low-appeal seller. The reason is that, with high-appeal sellers backing off searching, the expected gross profit per match rises for low-appeal firms. In contrast, the reduction in the number of matches reduces profit net of search costs for high-appeal sellers by about 27 percent.

Changes in share of high-quality sellers: The crowding out effect generated by increasing N^S is of course stronger when the fraction of low-appeal firms in the population of potential firms rises, as might have happened if China's phaseout of minimum size requirements had been the only policy reform. Our final experiment considers the effect of reducing the fraction of high-quality suppliers in the potential seller population (ω) from about 12 percent to the baseline 4 percent. This drop is sufficient to generate the same observed 25

⁴⁴Khandelwal et al. (2013) argue that rent-seeking and political favoritism meant the licenses did not go to the most productive firms. Bai et al. (2017) note that, while small firms were allowed to export through intermediaries before 2004, direct access to foreign markets improved their relative productivities.

percent increase in sellers per buyer during 2005, holding N^S constant.⁴⁵

The last column of Table 7 reports the results. Going forward in time (from the *pre*-shock equilibrium to the baseline), such a shock increases the measure of active low-appeal suppliers by 64 percent, while reducing the measure of active high-appeal suppliers by 54 percent, with the share of matches involving a low-quality seller rising by 55 percent. The measure of active buyers falls by about 11 percent (Overall, the set of active agents rises about 4 percent.)

The net effect of these adjustments is a decline in consumer welfare of 0.9 percent with an increase in search costs among all buyers and sellers.⁴⁶ Does anyone benefit from the reduction in seller quality? Buyers who remain active enjoy greater per-agent net profits, as do all low-appeal sellers and non-exiting high-appeal sellers. Interestingly, it's the larger (bigger μ) buyers who stick around, expanding relative to their competitors, and shifting their offerings toward low-appeal goods as they do so. Thus a surge in access of low appeal suppliers to the U.S. market may help explain the rising market share of "big box" chain retailers in the U.S. over the past 30 years (Jarmin et al., 2009).

6 Conclusion

We have developed a dynamic model of international buyer-seller matching in which agents on both sides of the market optimally choose how intensely to search. Quantifying the model with customs data on U.S. apparel imports, the framework captures key cross-sectional and time-series features of business-to-business relationships between foreign exporters and U.S. buyers. We've used the framework to examine how these relationships respond to several types of shocks.

We find, first, that the aggregate costs of forming business relationships are borne almost equally by buyers and sellers. But buyers' search costs are lower on a per-match basis, reflecting, *inter alia*, the fact that there are far more sellers than buyers, making it relatively easy

⁴⁵This experiment isn't meant to capture the actual change in the mix of suppliers, since *inter alia*, ATC quotas were phased out at roughly the same time. But it's nonetheless instructive to explore how changes in the composition of the potential seller population affect the market equilibrium.

⁴⁶These results suggest a case for trade restrictions that limit market access to low-appeal producers. As discussed by Hosios (1990), welfare effects in search models depend upon the particular congestion and market thickness externalities created by search activities on both sides of the market, which in turn depend on the bargaining share β and the matching function.

for buyers to find suppliers.

Second, buyers and sellers adjust their search intensity over their life cycles, both because their market visibility changes and because buyers face diminishing returns to adding business partners. There are big differences in how different types of agents in the market mature. It takes about 25 years for a successful buyer to reach its long run size, while a seller typically gets there in 10. Thus search frictions thus appear to constitute a major reason that aggregate trade flows react to shocks with long, unpredictable lags.

Third, heterogeneity plays a key role in determining the search behavior of individual agents and, as in the labor literature, the externalities generated by this behavior mean that market equilibria can be inefficient. For example, adding low-appeal suppliers to the market can reduce welfare because the positive effects of the extra varieties they offer are more than offset by reductions in the search efforts of existing high-quality sellers, which in turn reduces the quality of the bundles available to final consumers.

Finally, counterfactual experiments show that, despite bringing many new suppliers to the U.S. market, the China shock may have had only a small effect on consumer welfare in the apparel sector if the quality mix of the *potential* suppliers remained unaffected. On the other hand, if the China shock improved the relative access of low quality sellers to the U.S. market, it may have reduced consumer welfare through congestion effects while encouraging “big box” retailing. Finally, reductions in search frictions due to the IT revolution may well have generated large welfare gains while shifting the size distribution of buyers toward larger firms.

References

- Albornoz, Facundo, Hector F Calvo Pardo, Gregory Corcos, and Emanuel Ornelas,** “Sequential Exporting,” *Journal of International Economics*, 2012, *88* (1), 17–31.
- Alessandria, George, Costas Arkolakis, and Kim Ruhl,** “Firm Dynamics and Trade,” Working Paper 27934, NBER 2020.
- Allen, Treb,** “Information Frictions in Trade,” *Econometrica*, 2014, *86*(6), 2041–2083.
- Alvarez, Vanessa, Michele Fioretti, Ayumu Ken Kikkawa, and Monica Morlacco,** “Two-Sided Market Power in Firm-to-Firm Trade,” Working Paper, UBC Sauder 2020.
- Antras, Pol and Davin Chor,** “Global Value Chains,” Working Paper 28549, NBER 2021.
- Arkolakis, Costas,** “Market Penetration Costs and the New Consumers Margin in International Trade,” *Journal of Political Economy*, 2010, *118* (6), 1151–1199.
- , “A Unified Theory of Firm Selection and Growth,” *Quarterly Journal of Economics*, 2016, *131* (1), 89–155.
- Atkeson, A. and A. Burstein,** “Pricing-to-Market, Trade Costs, and International Relative Prices,” *American Economic Review*, 2008, *98* (5), 1998–2031.
- Atkin, David, Ben Faber, and Marco Gonzalez-Navarro,** “Retail Globalization and Household Welfare: Evidence from Mexico,” *Journal of Political Economy*, 2018, *126* (1), 1–73.
- Bai, Xue, Kala Krishna, and Hong Ma,** “How You Export Matters,” *Journal of International Economics*, 2017, *104*, 122–137.
- Benguria, Felipe,** “The Matching and Sorting of Exporting and Importing Firms: Theory and Evidence,” *Journal of International Economics*, 2021, *131*.
- Berman, Nicolas, Vincent Rebeyrol, and Vincent Vicard,** “Demand Learning and Firm Dynamics: Evidence from Exporters,” *Review of Economics and Statistics*, 2019, *101* (1), 91–106.

- Bernard, Andrew, Adreas Moxnes, Emmanuel Dhyne, Glenn Magerman, and Kalina Manova**, “The Origins of Firm Heterogeneity: a Production Network Approach,” *Journal of Political Economy*, forthcoming.
- **and Andreas Moxnes**, “Networks and Trade,” *Annual Review of Economics*, 2018, *10*, 65–85.
- , – , **and Karen Helene Ulltveit-Moe**, “Two-Sided Heterogeneity and Trade,” *Review of Economics and Statistics*, 2018, *100* (3), 424–439.
- Blum, Bernardo, Sebastian Claro, and Ignatius Horstmann**, “Facts and Figures on Intermediated Trade,” *American Economic Review*, 2010, *100* (2), 419–423.
- Brugemann, B., P. Gautier, and G. Menzio**, “Intra-firm Bargaining and Shapley Values,” *Review of Economic Studies*, 2019, *86* (2), 564–592.
- Cahal, Julia, Rocco Macchiavello, and Guillermo Nogueras**, “International Buyers’ Sourcing and Suppliers’ Markups in Bangladeshi Garmets,” Working Paper F-37119-PAK-2, International Growth Centre 2021.
- Carballo, Jeronimo, Gianmacro Ottaviano, and Christian volumepe Martincus**, “The Buyer Margins of Firms’ Exports,” *Journal of International Economics*, 2018, *112*, 33–49.
- Cavallo, Alberto**, “Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers,” *American Economic Review*, 2017, *107* (1), 283–303.
- Chaney, Thomas**, “The Network Structure of International Trade,” *American Economic Review*, 2014, *104* (11), 3600–3034.
- Dominguez, Juan Camilo, Jonathan Eaton, Marcela Eslava, and James Tybout**, “Search and Learning in Export Markets: Case Studies for Colombia,” Working Paper, Pennsylvania State University 2021.
- Drozd, Lukasz and Jaromir Nosal**, “Understanding International Prices: Customers as Capital,” *American Economic Review*, 2012, *102* (1), 364–395.

- Eaton, Jonathan, Marcela Eslava, David Jinkins, C.J. Krizan, and James Tybout**, “A Search and Learning Model of Export Dynamics,” Working Paper 29100, NBER 2021.
- , – , **Maurice Kugler, and James Tybout**, “The Margins of Entry into Export Markets: Evidence from Colombia,” in Elhanan Helpman, Dalia Marin, and Thierry Verdier, eds., *The Organization of Firms in a Global Economy*, Cambridge, MA: Harvard University Press, 2008.
- , **Samuel Kortum, and Francis Kramarz**, “Firm-to-Firm Trade: Imports, Exports and the Labor Market,” Working Paper, Pennsylvania State University 2018.
- Fitzgerald, Dorieann, Stephanie Haller, and Yaniv Yedid-Levi**, “How Exporters Grow,” Working Paper, Federal Reserve Bank of Minneapolis 2019.
- Foster, Lucia, John Haltiwanger, and Chad Syverson**, “The Slow Growth of New Plants: Learning About Demand?,” *Economica*, 2016, 83, 91–129.
- Gereffi, Gary and Olga Memedovic**, “The Global Apparel Value Chain: What Prospects for Upgrading by Developing Countries,” Technical Report, United Nations Industrial Development Organization 2003.
- Gourio, F. and L. Rudanko**, “Customer Capital,” *Review of Economic Studies*, 2014, 81 (3), 1102–1136.
- Gumpert, Anna, Haishi Li, Andreas Moxnes, Natalia Ramondo, and Felix Tintelnot**, “The Life-Cycle Dynamics of Exporters and Multinational Firms,” *Journal of International Economics*, 2020, 126.
- Ha-Brookshire, Jung and Barbara Dyer**, “Apparel Import Intermediaries: The Impact of a Hyper-Dynamic Environment on U.S. Apparel Firms,” *Clothing and Textiles Research Journal*, 2008, 26 (1), 66–90.
- Hosios, Arthur J.**, “On the Efficiency of Matching and Related Models of Search and Unemployment,” *Review of Economic Studies*, 1990, 57 (4), 279–298.

- Hottman, C., S. Redding, and D. Weinstein**, “Quantifying the Sources of Firm Heterogeneity,” *Quarterly Journal of Economics*, 2016, *131* (3), 1291–1364.
- Jarmin, Ronald, Shawn Klimek, and Javier Miranda**, “The Role of Retail Chains: National, Regional and Industry Results,” in Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts, eds., *Producer Dynamics: New Evidence from Micro Data*, University of Chicago Press for the NBER, 2009, pp. 237–262.
- Kamal, Fariha and Ryan Monarch**, “Identifying Foreign Suppliers in U.S. Import Data,” *Review of International Economics*, 2017, *26* (1), 117–139.
- Kennan, John**, “A Note on Discrete Approximations of Continuous Distributions,” Working Paper, University of Wisconsin Madison 2006.
- Khandelwal, Amit, Peter Schott, and Shang-Jin Wei**, “Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters,” *American Economic Review*, 2013, *103* (6), 2169–2195.
- Krizan, C.J., James Tybout, Zi Wang, and Yingyan Zhou**, “Are Customs Records Consistent Across Countries? Evidence from the U.S. and Colombia,” Working Paper 20-11, Center for Economic Studies, Bureau of the Census 2020.
- Li, Shengyu**, “A Structural Model of Productivity, Uncertain Demand, and Export Dynamics,” *Journal of International Economics*, 2018, *115*, 1–15.
- Lu, Sheng**, “2016 Fashion Industry Benchmarking Study,” Technical Report, United States Fashion Industry Association 2016.
- Luttmer, Erzo**, “Selection, Growth, and the Size Distribution of Firms,” *Quarterly Journal of Economics*, 2007, *122* (3), 1103–1144.
- , “On the Mechanics of Firm Growth,” *Review of Economic Studies*, 2011, *78* (3), 1042–1068.
- McFarlan, F., Michael Shih-Ta Chen, and Keith Chi-Ho Wong**, “Li and Fung 2012,” Case Study 9-312-102, Harvard Business School 2012.

- Petrongolo, Barbara and Christopher Pissarides**, “Looking into the Black Box: a Survey of the Matching Function,” *Journal of Economic Literature*, 2001, 39 (2), 390–431.
- Piveteau, Paul**, “An Empirical Dynamic Model of Trade With Customer Accumulation,” *American Economic Journal: Macroeconomics*, 2021, 13 (4).
- Plunkett-Analytics**, “Apparel and Clothing Brands, Designers, Importers and Distributors Industry (U.S.),” Technical Report, Plunkett Research 2015.
- Rauch, James E.**, “Business and Social Networks in International Trade,” *Journal of Economic Literature*, 2001, 39 (4), 1177–1203.
- Rauch, James E and Joel Watson**, “Starting Small in an Unfamiliar Environment,” *International Journal of Industrial Organization*, 2003, 21 (7), 1021–1042.
- Rauch, James E. and Vitor Trindade**, “Ethnic Chinese Networks in International Trade,” *The Review of Economics and Statistics*, 2002, 84 (1), 116–130.
- statista.com**, “U.S. online shop and mail-order sales of clothing and clothing accessories from 2003 to 2019,” 2022.
- Steinwender, Claudia**, “Real Effects of Information Frictions: When the States and the Kingdom became United,” *American Economic Review*, 2018, 108 (3), 657–696.
- Stole, L.A. and J. Zwiebel**, “Intra-Firm Bargaining Under Non-Binding Contracts,” *Review of Economic Studies*, 1996, 63, 375–410.
- Sugita, Yoichi, Kensuke Teshima, and Enrique Seira**, “Assortative Matching of Exporters and Importers,” *Review of Economics and Statistics*, forthcoming.
- Taplin, Ian**, “Global Commodity Chains and Fast Fashion: How the Apparel Industry Continues to Re-Invent Itself,” *Competition and Change*, 2014, 18 (3), 246–64.
- Terry, Lisa**, “Snapshot: Sewing up the Apparel Supply Chain,” *Inbound Logistics*, 2008.
- Timoshenko, Olga A.**, “Learning versus Sunk Cost Explanations of Export Persistence,” *European Economic Review*, 2015, 79, 113–128.

Zhu, Shengjun, “Go Up, Go West, and Go Out: Changing Industrial Policy and Firm Strategies in China’s Apparel Industry,” Doctoral Dissertation, University of North Carolina 2014.

Appendices

A Consequences of Imperfect Matching

This appendix explores the biases in network statistics that arise when customs records do not accurately report the true exporter for incoming shipments. These reporting errors generally occur because the importer does not report the name and address of the exporter in a sufficiently standardized way, and thus makes a given supplier appear to be multiple suppliers. Because of this problem, the number of firms exporting to the U.S. from foreign country A is generally far fewer according to the customs records of country A than it is according to the U.S. customs records (Kamal and Monarch, 2018; Krizan et al. 2020). In contrast, the number of importing firms in the U.S. and the number of shipments they receive are measured accurately because these records are based on “employer identification numbers” and/or “alpha” codes rather than names and addresses.

A.1 The environment

To keep our analysis tractable, we need to make some assumptions about the nature of recording errors. First, we will treat them as i.i.d. across shipments and for any individual shipment occurring with probability π_e . Second, we will assume that each time a reporting error occurs, it creates a “phantom” exporter in the customs records that does not actually exist, so that each phantom exporter appears to send a single shipment to an actual importer. Third, we’ll treat all business relationships as generating the same number of shipments per period, h . Finally, we’ll assume the true (unobserved) distribution of sellers per buyer (SPB) across buyers and the true (unobserved) distribution of buyers per seller (BPS) across sellers are both Pareto. These we will denote $f_B^*(c; \theta_B)$ and $f_S^*(c; \theta_S)$, where c ’s are partner counts and θ ’s are shape parameters. For numerical examples, we will choose parameterizations $(\pi_e, h, \theta_B, \theta_S)$ that are crudely consistent with observed outcomes in the data.

A.2 The sellers per buyer distribution

Let $y \in \mathbb{B}$ index buyers and let $x \in \mathbb{S}_y$ index true sellers shipping to y . Suppose we naively interpret all h shipments $x \rightarrow y$ as accurately recorded, including those attributed to phantom exporters. Then the apparent number of relationships created by this single $x \rightarrow y$ association, denoted n , will have probability distribution:

$$\Pi_{n|h} = \begin{cases} (\pi_e)^h + h(1 - \pi_e)(\pi_e)^{h-1} & n = h \\ \binom{h}{n-1} (1 - \pi_e)^{h-(n-1)} (\pi_e)^{n-1} & 1 \leq n < h \end{cases}$$

Intuitively, the most matches that could possibly be attributed to a buyer is the number of shipments the buyer receives, h . This can occur when all h shipments from x are attributed to phantom exporters (with probability $(\pi_e)^h$), or when one shipment is correctly attributed to x and the remaining $n - 1$ are attributed to phantoms (with probability $h(1 - \pi_e)(\pi_e)^{h-1}$). For values of n less than h , it must be that $n - 1$ shipments are assigned to phantom exporters and the remaining $h - (n - 1)$ shipments are correctly attributed to x . The probabilities of these outcomes are given by the binomial distribution.

A.2.1 The apparent distribution

Using these probabilities, we can construct the mean *apparent* client count when buyer y has $c = \|\mathbb{S}_y\|$ true clients. Calling n_x the total number of apparent sellers generated by the h shipments from x to y , the probability that y 's total client count appears to be m when it is actually c is

$$\Pi_{m|c}^0 = \Pr \left[m = \sum_{x \in \mathbb{S}_y} n_x \right] = \sum_{\mathbf{n} \ni m = \sum_{x \in \mathbb{S}_y} n_x} \Pr[\mathbf{n}]$$

where $\mathbf{n} = \{n_x\}_{x \in \mathbb{S}_y}$ is the vector of apparent client counts associated with each of the $\|\mathbb{S}_y\|$ sellers. Given this distribution, the observed distribution of buyers across sellers-per-buyer counts is:

$$f_B^0(m) = \sum_{c=1}^{c^{max}} \Pi_{m|c}^0 \cdot f_B^*(c)$$

To give a sense for orders of magnitude, Table A.1 below compares the true and apparent

frequency counts for a case where buyers can have a maximum of $c^{max} = 15$ true suppliers.⁴⁷ Referring back to Section 1 of the paper, we choose $N_B = \|\mathbb{B}\| = 10,000$ and $\theta = 1.1$ as plausible figures. Also, we choose $\pi_e = 0.2$ and $h = 3$, implying that the maximum number of observed clients is $h \cdot c^{max} = 45$, and the ratio of actual to apparent matches is $N_B/N_B^0 = 1/\sum_{m=1}^{m^{max}} m f_B^0(m) = 0.63$. This latter figure roughly replicates the ratio of identified foreign exporters in the LFTTD to sellers identified by trading partner customs records as shipping to the U.S.⁴⁸

Table A.1: Sellers per buyer, actual vs. apparent

sellers	$N_B f_B^*(c)$	$N_B^0 f_B^0(m)$	$\frac{N_B f_B^*(m) - N_B^0 f_B^0(c_s)}{N_B f_B^*(m)}$	sellers	$N_B f_B^*(c)$	$N_B^0 f_B^0(m)$	$\frac{N_B f_B^*(m) - N_B^0 f_B^0(c_s)}{N_B f_B^*(m)}$
1	5600.1	4564.7	0.18	14	42.088	109.39	-1.59
2	1762.1	4158.9	-1.36	15	36.581	95.216	-1.60
3	850.43	2212	-1.6	16	0	83.592	-
4	497.25	1175.7	-1.36	17	0	73.936	-
5	324.8	822.83	-1.53	18	0	65.776	-
6	228.11	575.67	-1.52	19	0	58.662	-
7	168.63	428.76	-1.54	20	0	52.1	-
8	129.52	331.47	-1.55	21	0	45.578	-
9	102.46	263.12	-1.56	22	0	38.71	-
10	82.993	213.88	-1.57	23	0	31.418	-
11	68.532	177.09	-1.58	24	0	24.013	-
12	57.506	148.93	-1.58	25	0	17.073	-
13	48.912	126.92	-1.59				

Several messages emerge. First and most obviously, since the average match generates $\sum_{n=1}^h n \cdot \Pi_{n|h} = 1.59$ apparent matches, the match count substantially overstates the true number of matches. Second, the $f_B^0(\cdot)$ distribution is substantially flatter than the true distribution, $f_B^*(\cdot)$. The reason is that firms with true relationship count c are spread across the range of apparent relationships from c to hc , and the amount of spreading depends upon c . So if one takes the data at face value, one is likely to seriously mischaracterize the SPB distribution, and a simple re-scaling of the data will not correct the problem.

⁴⁷Beyond this maximum, the combinatorial possibilities grow very rapidly, making computation impractical.

⁴⁸Using customs records from a large sample of countries, and controlling for HS2 sector, Kamal and Monarch (2017) find that the number of foreign firms exporting to the U.S. according to their home country customs records is 73 percent of the number of such firms in the sample countries according to U.S. customs records. They also find large variation in this discrepancy across source countries. Without controlling for HS2, Krizan et al. (2020) find that the number of Colombian firms exporting to the U.S. according to Colombian customs records is only 54 percent of the number of such firms according to U.S. customs records

A.2.2 Dropping single shipment matches

Since all phantom exporters generate a single shipment (by assumption), one simple remedy is to exclude all buyer-seller matches that generate a single shipment. But some actual relationships appear to generate single shipments, so if we drop all sellers that we observe sending a single shipment, our resulting seller counts will understate the true counts (in expectation) by the expected number of instances in which a true match is missed.

How severe will the resulting bias be? When single observed-shipment matches are dropped, we need at least 2 observed shipments to count a relationship. The probability that any true $x \rightarrow y$ match is properly recorded is thus

$$\begin{aligned}\pi_{rec} &= 1 - \sum_{m=0}^1 \binom{h}{m} (\pi_e)^{h-m} (1 - \pi_e)^m \\ &= 1 - (\pi_e)^h - h (\pi_e)^{h-1} (1 - \pi_e)\end{aligned}\tag{A-1}$$

And the probability of observing m connections when c actually exist is:

$$\Pi_{m|c}^1 = \binom{c}{m} (1 - \pi_{rec})^{c-m} (\pi_{rec})^m\tag{A-2}$$

So when singletons are excluded, the fraction of buyers that appear to have m suppliers is:

$$f_B^1(m) = \sum_{c=1}^{c^{max}} \Pi_{m|c}^1 \cdot f_B^*(c)\tag{A-3}$$

How does $N_B^1 f_B^1(m)$ compare to the true distribution, $N_B f_B^*(c_S)$? Once again using $\pi_e = 0.2$, $h = 3$, $\theta_B = 1.1$, and $N_B = 10,000$, we calculate the true and apparent (sans singleton) buyer frequency counts in Table A.2 below. Here, for compatibility with Table A.1, we set $c^{max} = 45$.⁴⁹

Clearly, dropping the single-shipment relationships isn't a perfect solution, but it brings

⁴⁹The "true" frequency counts in Table A.2, $N_B f_B^*(c)$, differ slightly from those in Table A.1 because computational constraints forced us to truncate the Pareto distribution for true clients in Table A.1 at $c^{max} = 15$ and to re-scale this distribution so that it summed to unity.

Table A.2: Sellers per buyer, actual vs. apparent without singletons

sellers	$N_B f_B^*(c)$	$N_B^1 f_B^1(m)$	$\frac{N_B f_B^*(m) - N_B^1 f_B^1(c_S)}{N_B f_B^*(m)}$	sellers	$N_B f_B^*(c)$	$N_B^1 f_B^1(m)$	$\frac{N_B f_B^*(m) - N_B^1 f_B^1(c_S)}{N_B f_B^*(m)}$
1	5426.1	5206.1	0.041	14	40.78	36.667	0.101
2	1707.4	1605.4	0.060	15	35.445	31.843	0.102
3	824.02	765.52	0.071	16	31.08	27.9	0.102
4	481.8	444.04	0.078	17	27.464	24.637	0.103
5	314.71	288.4	0.084	18	24.437	21.908	0.103
6	221.02	201.69	0.087	19	21.878	19.603	0.104
7	163.39	148.61	0.090	20	19.695	17.638	0.104
8	125.49	113.84	0.093	21	17.82	15.951	0.105
9	99.277	89.863	0.095	22	16.196	14.492	0.105
10	80.414	72.659	0.096	23	14.782	13.222	0.106
11	66.403	59.908	0.098	24	13.543	12.109	0.106
12	55.719	50.205	0.099	25	12.451	11.129	0.106
13	47.392	42.654	0.100				

us much closer to the true distribution. Further, the discrepancy between the true and actual distribution is nearly a constant factor of proportionality and therefore nearly innocuous. The deviation from proportionality is largest (but still modest) among those firms with a single supplier, since they are relatively likely to be dropped.

A.3 The buyers per seller distribution

What about buyers per seller? Before excluding singletons, the apparent *BPS* distribution is:

$$f_S^0(m) = \begin{cases} \frac{(1-\pi_{rec}) \cdot h \sum_c c \cdot f_S^*(c) + \sum_c \Pi_{1|c}^1 \cdot f_S^*(c)}{\Lambda} & m = 1 \\ \frac{\sum_c \Pi_{m|c}^1 \cdot f_S^*(c)}{\Lambda} & m > 1 \end{cases} .$$

where $\Lambda = (1 - \pi_{rec}) \cdot h \sum_c c \cdot f_S^*(c) + 1$ is the ratio of apparent to true exporters, or equivalently, the expected number of phantom exporters per true exporter, plus 1.⁵⁰

The phantom exporters are all contained in the single buyer category, so they can be eliminated by dropping all single-shipment relationships. But doing so means dropping those true relationships that generated a single detected shipment, just as with the $f_B^1(m)$ distribution discussed above. The resulting probability distribution resembles equation (A-3), the only difference coming from the distinction between true buyer and seller degree distributions, $f_S^*(c)$

⁵⁰At the parameter values introduced above and $\theta_S = 2.5$ (as estimated), $\Lambda = 0.104 \cdot 3 \cdot \frac{2.5}{1.5} + 1 = 1.52$.

and $f_B^*(c)$:

$$f_S^1(m) = \sum_c \Pi_{m|c}^1 \cdot f_S^*(c).$$

Since we have already generated this distribution for the case of buyers in Table A.2, and we are treating the true degree distributions as Pareto for both buyers and sellers, we do not repeat the exercise here.

B Demand and Pricing

Following Atkeson and Burstein (2008) and Hottman et al. (2016), this appendix establishes equation 2. Using standard CES results, we begin by characterizing prices and market shares for a particular retailer y offering a particular subset of product varieties in the group, $x \in J_y$:

$$C_y = \left(\sum_{x \in J_y} (\xi_x C_{xy})^{\frac{\alpha-1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}}, \quad C = \left(\int_y (\mu_y C_y)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \quad (\text{A-4})$$

$$P_y = \left[\sum_{x \in J_y} \left(\frac{P_{xy}}{\xi_x} \right)^{1-\alpha} \right]^{1/(1-\alpha)}, \quad P = \left[\int_y \left(\frac{P_y}{\mu_y} \right)^{(1-\eta)} \right]^{1/(1-\eta)} \quad (\text{A-5})$$

$$h_y = \frac{\left(\frac{P_y}{\mu_y} \right)^{1-\eta}}{P^{1-\eta}}, \quad h_{x|y} = \frac{\left(\frac{P_{xy}}{\xi_x} \right)^{1-\alpha}}{P_y^{1-\alpha}}$$

These expressions imply the revenue generated by retail sales of product x at store y is:

$$\begin{aligned} R_{xy} &= P_{xy} q_{xy} \\ &= h_{x|y} h_y E \\ &= \mu_y^{\eta-1} \xi_x^{\alpha-1} P_{xy}^{1-\alpha} P_y^{\alpha-\eta} P^{\eta-1} E \end{aligned} \quad (\text{A-6})$$

Since we assume a continuum of buyers, $\frac{\partial \ln P}{\partial \ln P_y} = 0$. Also,

$$\begin{aligned} \frac{\partial \ln P_y}{\partial \ln P_{xy}} &= \frac{\partial P_y}{\partial P_{xy}} \frac{P_{xy}}{P_y} \\ &= 1/(1-\alpha) \left[\sum_{x' \in J_y} \left(\frac{P_{x'y}}{\xi_{x'}} \right)^{1-\alpha} \right]^{1/(1-\alpha)-1-1/(1-\alpha)} \left[(1-\alpha) \left(\frac{P_{xy}}{\xi_x} \right)^{1-\alpha} \right] \\ &= \left[\sum_{x' \in J_y} \left(\frac{P_{x'y}}{\xi_{x'}} \right)^{1-\alpha} \right]^{-1} \left[\left(\frac{P_{xy}}{\xi_x} \right)^{1-\alpha} \right] = h_{x|y} \end{aligned}$$

Bertrand-Nash pricing therefore implies:

$$\begin{aligned} \frac{\partial \ln R_{xy}}{\partial \ln P_{xy}} &= (1-\alpha) + h_{x|y}(\alpha - \eta) \\ \frac{\partial \ln R_{xy}}{\partial \ln P_{x'y}} &= h_{x'|y}(\alpha - \eta) \quad \forall x' \neq x \end{aligned}$$

Plugging these expressions into the first-order conditions for pricing,

$$q_{xy} + \sum_{x' \in J_y} \frac{\partial q_{x'y}}{\partial P_{xy}} (P_{x'y} - c_{x'y}) = 0 \quad \forall x \in J_y,$$

we obtain:

$$\begin{aligned} \frac{q_{xy}}{E} + \frac{\partial q_{xy}}{\partial P_{xy}} \frac{P_{xy}}{E} \left(\frac{P_{xy} - c_{xy}}{P_{xy}} \right) + \sum_{x' \in J_y, x' \neq x} \frac{\partial q_{x'y}}{\partial P_{xy}} \frac{P_{x'y}}{E} \left(\frac{P_{x'y} - c_{x'y}}{P_{x'y}} \right) &= 0 \\ \frac{q_{xy}}{E} + \frac{\partial q_{xy}}{\partial P_{xy}} \frac{1}{q_{xy}} \left(\frac{P_{xy} q_{xy}}{E} \right) \left(\frac{P_{xy} - c_{xy}}{P_{xy}} \right) + \sum_{x' \in J_y, x' \neq x} \frac{\partial q_{x'y}}{\partial P_{xy}} \frac{1}{q_{x'y}} \left(\frac{q_{x'y} P_{x'y}}{E} \right) \left(\frac{P_{x'y} - c_{x'y}}{P_{x'y}} \right) &= 0 \\ \frac{P_{xy} q_{xy}}{E} + \frac{\partial q_{xy}}{\partial P_{xy}} \frac{P_{xy}}{q_{xy}} \left(\frac{P_{xy} q_{xy}}{E} \right) \left(\frac{P_{xy} - c_{xy}}{P_{xy}} \right) + \sum_{x' \in J_y, x' \neq x} \frac{\partial q_{x'y}}{\partial P_{xy}} \frac{P_{xy}}{q_{x'y}} \left(\frac{q_{x'y} P_{x'y}}{E} \right) \left(\frac{P_{x'y} - c_{x'y}}{P_{x'y}} \right) &= 0 \\ h_{xy} + \frac{\partial q_{xy}}{\partial P_{xy}} \frac{P_{xy}}{q_{xy}} h_{xy} \left(\frac{P_{xy} - c_{xy}}{P_{xy}} \right) + \sum_{x' \in J_y, x' \neq x} \frac{\partial q_{x'y}}{\partial P_{xy}} \frac{P_{xy}}{q_{x'y}} h_{x'y} \left(\frac{P_{x'y} - c_{x'y}}{P_{x'y}} \right) &= 0 \\ h_{xy} + (-\alpha + (\alpha - \eta) h_{x|y}) h_{xy} \left(\frac{P_{xy} - c_{xy}}{P_{xy}} \right) + \sum_{x' \in J_y, x' \neq x} ((\alpha - \eta) h_{x|y}) h_{x'y} \left(\frac{P_{x'y} - c_{x'y}}{P_{x'y}} \right) &= 0 \\ 1 - \alpha \left(\frac{P_{xy} - c_{xy}}{P_{xy}} \right) + (\alpha - \eta) \sum_{x' \in J_y} \frac{h_{x'y}}{h_y} \left(\frac{P_{x'y} - c_{x'y}}{P_{x'y}} \right) &= 0 \end{aligned}$$

where $h_{xy} = h_{x|y}h_y = \frac{P_{xy}q_{xy}}{E}$. Since this relationship holds for all $x \in J_y$, and the last term is common to all products, the mark-up for each product must be the same. Call it $m = \frac{P_{xy} - c_{xy}}{P_{xy}}$ and reduce this equation to $1 - \alpha m + (\alpha - \eta)m = 0$, or

$$m = \frac{1}{\eta}.$$

C Bargaining and Transfer

This appendix establishes that Nash bargaining over the total value of match surpluses leads to a simple sharing rule for flow profits. It then shows how the flow profits accruing to buyers and sellers can be expressed in terms of their current states and marginal costs.

C.1 Division of match surplus

Differencing the buyer's value function and suppressing buyer type i , we have:

$$\begin{aligned} \rho(V^B(\mathbf{s}) - V^B(\mathbf{s} - \ell_j)) &= [\pi^B(\mathbf{s}) - \pi^B(\mathbf{s} - \ell_j)] - [k^B(\mathbf{s}) - k^B(\mathbf{s} - \ell_j)] \\ &+ \sigma^B(\mathbf{s})\theta^B \sum_k \nu_k^S [V^B(\mathbf{s} + \ell_k) - V^B(\mathbf{s})] - \sigma^B(\mathbf{s} - \ell_j)\theta^B \sum_k \nu_k^S [V^B(\mathbf{s} - \ell_j + \ell_k) - V^B(\mathbf{s} - \ell_j)] \\ &+ \delta \sum_k s_k [V^B(\mathbf{s} - \ell_k) - V^B(\mathbf{s})] - \delta \sum_k (s_k - \mathbf{1}_{k=j}) [V^B(\mathbf{s} - \ell_j - \ell_k) - V^B(\mathbf{s} - \ell_j)] \end{aligned}$$

Now we simplify this equation in two steps. First apply a discrete approximation of the first order condition at $(\mathbf{s} - \ell_j)$ for the buyer search:

$$[k^B(\mathbf{s}) - k^B(\mathbf{s} - \ell_j)] \approx (\sigma^B(\mathbf{s}) - \sigma^B(\mathbf{s} - \ell_j)) \left(\theta^B \sum_k \nu_k^S [V^B(\mathbf{s} - \ell_j + \ell_k) - V^B(\mathbf{s} - \ell_j)] \right)$$

Using the above, we can simplify

$$\begin{aligned}
& -[k^B(\mathbf{s}) - k^B(\mathbf{s} - \ell_j)] + \sigma^B(\mathbf{s})(\theta^B \sum_k \nu_k^S [V^B(\mathbf{s} + \ell_k) - V^B(\mathbf{s})]) \\
& -\sigma^B(\mathbf{s} - \ell_j)(\theta^B \sum_k \nu_k^S [V^B(\mathbf{s} - \ell_j + \ell_k) - V^B(\mathbf{s} - \ell_j)]) \\
& = \sigma^B(\mathbf{s}) \left(\theta^B \sum_k \nu_k^S [V^B(\mathbf{s} + \ell_k) - V^B(\mathbf{s})] - \theta^B (V^B(\mathbf{s} - \ell_j + \ell_k) - V^B(\mathbf{s} - \ell_j)) \right)
\end{aligned}$$

Second, we can also simplify the destruction side using

$$\begin{aligned}
& \delta \sum_k s_k [V^B(\mathbf{s} - \ell_k) - V^B(\mathbf{s})] - \delta \sum_k (s_k - \mathbf{1}_{k=j}) [V^B(\mathbf{s} - \ell_j - \ell_k) - V^B(\mathbf{s} - \ell_j)] \\
& = -\delta [V^B(\mathbf{s}) - V^B(\mathbf{s} - \ell_j)] + \delta \sum_k (s_k - \mathbf{1}_{k=j}) ([V^B(\mathbf{s} - \ell_k) - V^B(\mathbf{s})] - [V^B(\mathbf{s} - \ell_j - \ell_k) - V^B(\mathbf{s} - \ell_j)])
\end{aligned}$$

To summarize, the above gives us:

$$\begin{aligned}
& (\rho + \delta)[V^B(\mathbf{s}) - V^B(\mathbf{s} - \ell_j)] \tag{A-7} \\
& = [\pi^B(\mathbf{s}) - \pi^B(\mathbf{s} - \ell_j)] + \sigma^B(\mathbf{s})\theta^B \sum_k \nu_k^S ([V^B(\mathbf{s} + \ell_k) - V^B(\mathbf{s})] - [V^B(\mathbf{s} + \ell_k - \ell_j) - V^B(\mathbf{s} - \ell_j)]) \\
& \quad + \delta \sum_k (s_k - \mathbf{1}_{k=j}) ([V^B(\mathbf{s} - \ell_k) - V^B(\mathbf{s} - \ell_k - \ell_j)] - [V^B(\mathbf{s}) - V^B(\mathbf{s} - \ell_j)])
\end{aligned}$$

By the definition of the type j seller's value function, we have

$$\begin{aligned}
(\rho + \delta)V_j^S(\mathbf{s}) & = \tau^j(\mathbf{s}) + \sigma^B(\mathbf{s})\theta^B \sum_k \nu_k^S [V_j^S(\mathbf{s} + \ell_k) - V_j^S(\mathbf{s})] \tag{A-8} \\
& \quad + \delta \sum_k (s_k - \mathbf{1}_{k=j}) [V_j^S(\mathbf{s} - \ell_k) - V_j^S(\mathbf{s})]
\end{aligned}$$

Finally, using equations (A-7), (A-8), and the surplus sharing rule, we have

$$\begin{aligned}
& \beta(\rho + \delta)[V^B(\mathbf{s}) - V^B(\mathbf{s} - \ell_j)] - (1 - \beta)(\rho + \delta)V_j^S(\mathbf{s}) \\
= & \beta[\pi^B(\mathbf{s}) - \pi^B(\mathbf{s} - \ell_j)] - (1 - \beta)\tau^j(\mathbf{s}) \\
& + \sigma^B(s)\theta^B \sum_k \nu_k^S (\beta[V^B(\mathbf{s} + \ell_k) - V^B(\mathbf{s})] - (1 - \beta)V_j^S(\mathbf{s} + \ell_k)) \\
& - \sigma^B(s)\theta^B \sum_k \nu_k^S (\beta[V^B(\mathbf{s} + \ell_k - \ell_j) - V^B(\mathbf{s} - \ell_j)] - (1 - \beta)V_j^S(\mathbf{s})) \\
& + \delta \sum_k (s_k - \mathbf{1}_{k=j}) (\beta[V^B(\mathbf{s} - \ell_k) - V^B(\mathbf{s} - \ell_k - \ell_j)] - (1 - \beta)V_j^S(\mathbf{s} - \ell_k)) \\
& - \delta \sum_k (s_k - \mathbf{1}_{k=j}) (\beta[V^B(\mathbf{s}) - V^B(\mathbf{s} - \ell_j)] - (1 - \beta)V_j^S(\mathbf{s})) \equiv 0
\end{aligned}$$

Using the surplus sharing rule $\beta(V^B(\mathbf{s}) - V^B(\mathbf{s} - \ell_j)) - (1 - \beta)V_j^S(\mathbf{s}) = 0, \forall j, \mathbf{s}$ and cancelling terms and re-arranging, the flow transfer to a type- j seller by a buyer in state \mathbf{s} is share β of the total flow surplus generated by their match⁵¹:

$$\tau^j(\mathbf{s}) = \beta[\pi^B(\mathbf{s}) - \pi^B(\mathbf{s} - \ell_j) + \tau^j(\mathbf{s})] \quad (\text{A-9})$$

The *total* flow surplus created by the marginal match between a type- j seller by a type- i buyer in state \mathbf{s} must equal the sum of the flow surpluses reaped by the buyer and the seller:

$$\pi_i^B(\mathbf{s}) - \pi_i^B(\mathbf{s} - \ell_j) + \tau^j(\mathbf{s}) = \pi_i^T(\mathbf{s}) - \pi_i^T(\mathbf{s} - \ell_j)$$

So we can re-state (A-9) as:

$$\tau_{ji}(\mathbf{s}) = \beta [\pi_i^T(\mathbf{s}) - \pi_i^T(\mathbf{s} - \ell_j)] \quad (\text{A-10})$$

C.2 Payoffs, states and costs

We now express buyers' and sellers' surpluses in terms of their states and marginal costs. Consider first the gross surplus generated by a type- i buyer in state \mathbf{s} jointly with its portfolio

⁵¹To derive this equation, we also need to assume that $\beta(V^B(\mathbf{s} + \ell_k - \ell_j) - V^B(\mathbf{s} - \ell_j)) - (1 - \beta)V_j^S(\mathbf{s}) \approx 0$. This holds strictly when sellers are homogeneous. In the case of heterogeneous sellers, we will need a type j 's seller's valuation of a relationship $V_j^S(s)$ to be very close to a case when another seller of type j of the buyer is replaced with type k , i.e. $V_j^S(s - \ell_j + \ell_k)$.

of suppliers. The exact retail price index for this buyer's offerings is

$$P_i(\mathbf{s}) = \left[\sum_{j=1}^J s_j \left(\frac{\eta}{\eta-1} \right)^{1-\alpha} \tilde{c}_j^{1-\alpha} \right]^{\frac{1}{1-\alpha}},$$

so the share of this buyer in total retail sales is

$$h_i(\mathbf{s}) = \frac{\left(\frac{P_i(\mathbf{s})}{\mu_i} \right)^{1-\eta}}{P^{1-\eta}},$$

where P is the exact market-wide price deflator. And since the mark-up factor for all buyers is $\frac{1}{\eta}$ (see Appendix ??), the gross surplus before transfers at a type- i buyer in state \mathbf{s} is:

$$\pi_i^T(\mathbf{s}) = \frac{h_i(\mathbf{s})E}{\eta} \tag{A-11}$$

$$= \left(\frac{\eta}{\eta-1} \right)^{1-\eta} \frac{E}{\eta P^{1-\eta}} \left[\sum_{j=1}^J s_j \tilde{c}_j^{1-\alpha} \right]^{\frac{1-\eta}{1-\alpha}} \mu_i^{\eta-1} \tag{A-12}$$

How much of this surplus do buyers transfer to their suppliers? The increment to gross surplus when a type- i buyer in state \mathbf{s} adds a type- j seller is:

$$\frac{\partial \pi_i^T(\mathbf{s})}{\partial s_j} = \frac{\eta-1}{\alpha-1} \left(\frac{\eta}{\eta-1} \right)^{1-\eta} \tilde{c}_j^{1-\alpha} \frac{E}{\eta P^{1-\eta}} \left[\sum_{j'=1}^J s_{j'} \tilde{c}_{j'}^{1-\alpha} \right]^{\frac{1-\eta}{1-\alpha}-1} \mu_i^{\eta-1}$$

And by equation A-10, each type- j seller gets some exogenous share β of the marginal seller's contribution to this surplus. So, summing up the buyer's transfers to all its suppliers yields

$$\begin{aligned}
\tau_i(\mathbf{s}) &= \sum_{j'=1}^J s_{j'} \beta \frac{\partial \pi_i^T(\mathbf{s})}{\partial s_{j'}} = \beta \frac{\eta-1}{\alpha-1} \left(\frac{\eta}{\eta-1} \right)^{1-\eta} \sum_{j'=1}^J s_{j'} \left(\tilde{c}_{j'}^{1-\alpha} \frac{E}{P^{1-\eta}} \left[\sum_{j=1}^J s_j \tilde{c}_j^{1-\alpha} \right]^{\frac{1-\eta}{1-\alpha}-1} \mu_i^{\eta-1} \right) \\
&= \beta \frac{\eta-1}{\alpha-1} \left(\frac{\eta}{\eta-1} \right)^{1-\eta} \frac{E}{P^{1-\eta}} \sum_{j'=1}^J s_{j'} \tilde{c}_{j'}^{1-\alpha} \left[\sum_{j=1}^J s_j \tilde{c}_j^{1-\alpha} \right]^{\frac{1-\eta}{1-\alpha}-1} \mu_i^{\eta-1} \\
&= \beta \frac{\eta-1}{\alpha-1} \left(\frac{\eta}{\eta-1} \right)^{1-\eta} \frac{E}{P^{1-\eta}} \left(\sum_{j=1}^J s_j \tilde{c}_j^{1-\alpha} \right)^{\frac{1-\eta}{1-\alpha}} \mu_i^{\eta-1}
\end{aligned}$$

Transfers as a share of surplus are thus the same for all buyers:

$$\frac{\tau_i(\mathbf{s})}{\pi_i^T(\mathbf{s})} = \frac{\beta \frac{\eta-1}{\alpha-1} \left(\frac{\eta}{\eta-1} \right)^{1-\eta} \frac{E}{P^{1-\eta}} \left(\sum_{j=1}^J s_j \tilde{c}_j^{1-\alpha} \right)^{\frac{1-\eta}{1-\alpha}} \mu_i^{\eta-1}}{\left(\frac{\eta}{\eta-1} \right)^{1-\eta} \frac{E}{\eta P^{1-\eta}} \left[\sum_{j=1}^J s_j \tilde{c}_j^{1-\alpha} \right]^{\frac{1-\eta}{1-\alpha}} \mu_i^{\eta-1}} = \beta \frac{\eta-1}{\alpha-1}, \quad (\text{A-13})$$

and buyers' profits after transfers are:

$$\pi_i^B(\mathbf{s}) = \pi_i^T(\mathbf{s}) - \tau_i(\mathbf{s}) = \pi_i^T(\mathbf{s}) \left(1 - \beta \frac{\eta-1}{\alpha-1} \right).$$

Finally, it is straightforward to construct industry-wide aggregates for transfers and net buyer profits. By equations A-11 and A-13, a type- i buyer in state \mathbf{s} transfers a total of $\tau_i(\mathbf{s}) = \beta \frac{\eta-1}{\alpha-1} \frac{h_i(\mathbf{s})E}{\eta}$ to its suppliers. Summing these transfers across all buyers, industry-wide transfers to sellers are:

$$\sum_{i=1}^I \sum_{\mathbf{s} \in \mathbb{S}} \beta \frac{\eta-1}{\alpha-1} M_i^B(\mathbf{s}) \frac{h_i(\mathbf{s})E}{\eta} = \frac{\beta}{\eta} \frac{\eta-1}{\alpha-1} E,$$

and similarly, industry-wide buyer profits after transfers are:

$$\sum_{i=1}^I \sum_{\mathbf{s} \in \mathbb{S}} \left[1 - \beta \frac{\eta-1}{\alpha-1} \right] M_i^B(\mathbf{s}) \frac{h_i(\mathbf{s})E}{\eta} = \left[1 - \beta \frac{\eta-1}{\alpha-1} \right] \frac{E}{\eta}.$$

Note that as sellers' bargaining power $\beta \rightarrow 0$, sellers' transfers cease to exceed their production costs. Their search efforts then go to zero and the market disappears. However, as $\beta \rightarrow 1$,

buyers continue to earn rents from their infra-marginal matches, earning a total surplus of $\frac{E}{\eta} \left[\frac{\alpha-\eta}{\alpha-1} \right]$.

D The moment weighting matrix

Our GMM estimator (equation 20) is based on a vector of sample moments with five components: $\bar{m}' = [\bar{m}'_1, \bar{m}'_2, \bar{m}'_3, \bar{m}'_4, \bar{m}'_5]$. Here \bar{m}_1 contains the buyers per seller (*BPS*) transition probabilities from Table 1, \bar{m}_2 contains the sellers per buyer (*SPB*) transition probabilities from Table 2, \bar{m}_3 contains the *BPS* degree distribution from Table 3, \bar{m}_4 contains the *SPB* degree distribution from the same Table, and \bar{m}_5 contains the intra-firm supplier share statistics reported in Table 4. In this appendix we describe our construction of the associated block-diagonal weighting matrix, component by component:

$$W = \text{diag} [cov(\bar{m}_1), cov(\bar{m}_2), cov(\bar{m}_3), cov(\bar{m}_4), cov(\bar{m}_5)]^{-1}$$

D.1 Transition matrices

Consider first \bar{m}_1 and \bar{m}_2 . Elements of these vectors are estimates of probabilities that sellers (described by \bar{m}_1) or buyers (described by \bar{m}_2) with i partners in period t will have j partners in period $t + 1$. They are sample analogs to the population transition probabilities:⁵²

$$\begin{aligned} \pi_{j|i}^{BPS} &= P(n_{t+1}^B = j | n_t^B = i). \\ \pi_{j|i}^{SPB} &= P(n_{t+1}^S = j | n_t^S = i). \end{aligned}$$

where n_t^B is number of buyers in year t , and n_t^S is number of sellers in year t .

To derive covariance matrices for these vectors, it is convenient to reshape them as matrices.

⁵²Here i and j refer to possible values for buyer and seller counts. They should not be confused with the i and j subscripts that appear in the text, which refer to buyer and seller types.

For the *BPS* transition probabilities, we write:

$$\mathbf{\Pi}^{BPS} = \begin{bmatrix} \pi_{1|1}^{BPS} & \pi_{2|1}^{BPS} & \cdots & \pi_{K-1|1}^{BPS} & \pi_{K|1}^{BPS} \\ \pi_{1|2}^{BPS} & \pi_{2|2}^{BPS} & \cdots & \pi_{K-1|2}^{BPS} & \pi_{K|2}^{BPS} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \pi_{1|K-1}^{BPS} & \pi_{2|K-1}^{BPS} & \cdots & \hat{\pi}_{K-1|K-1}^{BPS} & \hat{\pi}_{K|K-1}^{BPS} \\ \pi_{1|K}^{BPS} & \pi_{2|K}^{BPS} & \cdots & \pi_{K-1|K}^{BPS} & \pi_{K|K}^{BPS} \end{bmatrix}$$

where K is the (topcoded) maximum number of buyers attainable by an individual seller. Replacing *BPS* superscripts with *SPB* superscripts gives our notation for the sellers per buyer transition matrix. There is no difference in our derivation of the covariance matrices for the two sets of moments, so hereafter we will focus on buyers per seller.

Next, re-state each row of $\mathbf{\Pi}^{BPS}$ in terms of cumulative probabilities:

$$\mathbf{F}^{BPS} = \begin{bmatrix} \pi_{1|1}^{BPS} & \pi_{1|1}^{BPS} + \pi_{2|1}^{BPS} & \cdots & \pi_{1|1}^{BPS} + \pi_{2|1}^{BPS} + \cdots + \pi_{K-1|1}^{BPS} & 1 \\ \pi_{1|2}^{BPS} & \pi_{1|2}^{BPS} + \pi_{2|2}^{BPS} & \cdots & \pi_{1|2}^{BPS} + \pi_{2|2}^{BPS} + \cdots + \pi_{K-1|2}^{BPS} & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \pi_{1|K-1}^{BPS} & \pi_{1|K-1}^{BPS} + \pi_{2|K-1}^{BPS} & \cdots & \pi_{1|K-1}^{BPS} + \pi_{2|K-1}^{BPS} + \cdots + \pi_{K-1|K-1}^{BPS} & 1 \\ \pi_{1|K}^{BPS} & \pi_{1|K}^{BPS} + \pi_{2|K}^{BPS} & \cdots & \pi_{1|K}^{BPS} + \pi_{2|K}^{BPS} + \cdots + \pi_{K-1|K}^{BPS} & 1 \end{bmatrix},$$

and call the i^{th} row of this matrix $\mathbf{F}_i^{BPS} = (F_{1|i}^{BPS}, F_{2|i}^{BPS}, \dots, F_{K|i}^{BPS})$, where $F_{\ell|i}^{BPS} = \pi_{1|i}^{BPS} + \pi_{2|i}^{BPS} + \cdots + \pi_{\ell|i}^{BPS}$. Then the covariance matrix for the i^{th} row of the sample analog to this matrix, call it $\hat{\mathbf{F}}_i^{BPS}$, can be written as⁵³

⁵³Suppose we wish to calculate the covariance between two sample-based cumulative probabilities, $\hat{F}_q = \hat{F}(x_q)$ and $\hat{F}_m = \hat{F}(x_m)$. These are calculated at chosen cutoffs x_m and $x_q \geq x_m$ using a sample of n draws from the distribution $F(X)$. Then

$$\begin{aligned} cov(\hat{F}_q, \hat{F}_m) &= E(\hat{F}_q \hat{F}_m) - E(\hat{F}_q)E(\hat{F}_m) = E\left[\frac{\sum_i I_{\{X_i \leq x_q\}} \sum_j I_{\{X_j \leq x_m\}}}{n^2}\right] - F_q F_m \\ &= E\left[\frac{\sum_i \sum_{i \neq j} I_{\{X_i \leq x_q\}} I_{\{X_j < x_m\}}}{n^2}\right] + E\left[\frac{\sum_i I_{\{X_i \leq x_m\}}}{n^2}\right] - F_q F_m \\ &= \left[\frac{n(n-1)F_q F_m}{n^2}\right] + \left[\frac{nF_m}{n^2}\right] - F_q F_m = \frac{F_m(1-F_q)}{n} \end{aligned}$$

Note that larger probability is always the one subtracted from 1.

$$cov(\hat{\mathbf{F}}_i^{BPS}) = \frac{1}{N_i^S} \begin{bmatrix} F_{1|i}^{BPS}(1 - F_{1|i}^{BPS}) & F_{1|i}^{BPS}(1 - F_{2|i}^{BPS}) & \dots & F_{1|i}^{BPS}(1 - F_{K|i}^{BPS}) \\ F_{1|i}^{BPS}(1 - F_{2|i}^{BPS}) & F_{2|i}^{BPS}(1 - F_{2|i}^{BPS}) & \dots & F_{2|i}^{BPS}(1 - F_{K|i}^{BPS}) \\ \vdots & \vdots & \ddots & \vdots \\ F_{1|i}^{BPS}(1 - F_{K|i}^{BPS}) & F_{2|i}^{BPS}(1 - F_{K|i}^{BPS}) & \dots & F_{K|i}^{BPS}(1 - F_{K|i}^{BPS}) \end{bmatrix},$$

where N_i^S is the number of observations on firms with i clients that we use to construct $\hat{\mathbf{\Pi}}_i^{BPS}$.

Finally, since $\hat{\mathbf{F}}_i^{BPS}$ is a linear transformation of $\hat{\mathbf{\Pi}}_i^{BPS}$, the variances of $\hat{\mathbf{\Pi}}_i^{BPS}$ can be easily recovered. Specifically, since we can write $(\hat{\mathbf{\Pi}}_i^{BPS})' = \mathbf{A} \cdot (\hat{\mathbf{F}}_i^{BPS})'$ where

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 \\ -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & 0 & -1 & 1 \end{bmatrix},$$

the covariance matrix for the (transposed) i^{th} row of $\hat{\mathbf{\Pi}}^{BPS}$ is

$$\Psi_i^{BPS} = cov \left[(\hat{\mathbf{\Pi}}_i^{BPS})' \right] = \mathbf{A} \cdot cov \left[(\hat{\mathbf{F}}_i^{BPS})' \right] \cdot \mathbf{A}'.$$

If we were to use all elements of the sample transition matrix $\hat{\mathbf{\Pi}}^{BPS}$ as targets, the associated sample moment vector would be:

$$vec \left[(\hat{\mathbf{\Pi}}^{BPS})' \right] = \begin{pmatrix} (\hat{\mathbf{\Pi}}_1^{BPS})' \\ (\hat{\mathbf{\Pi}}_2^{BPS})' \\ \vdots \\ (\hat{\mathbf{\Pi}}_K^{BPS})' \end{pmatrix}_{K^2 \times 1}$$

And treating $\hat{\mathbf{F}}_i^{BPS}$ and $\hat{\mathbf{\Pi}}_i^{BPS}$ as independent of $\hat{\mathbf{F}}_j^{BPS}$ and $\hat{\mathbf{\Pi}}_j^{BPS}$, $j \neq i$, the covariance for this vector would be:

$$\text{cov} \left(\text{vec} \left[\left(\hat{\mathbf{\Pi}}^{BPS} \right)' \right] \right) \stackrel{\text{def}}{=} \Psi^{BPS} = \begin{bmatrix} \Psi_1^{BPS} & 0 & 0 & \cdots & 0 \\ 0 & \Psi_2^{BPS} & 0 & \cdots & \vdots \\ 0 & 0 & \Psi_3^{BPS} & 0 & \vdots \\ \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & \cdots & \cdots & 0 & \Psi_K^{BPS} \end{bmatrix}_{K^2 \times K^2}$$

However, for two reasons, we exclude some elements of $\hat{\mathbf{\Pi}}^{BPS}$ from \bar{m}_1 . First, each row of $\hat{\mathbf{\Pi}}^{BPS}$ sums to one, so $\hat{\mathbf{\Pi}}^{BPS}$ contains some redundant information and Ψ^{BPS} is singular. Second, $\hat{\mathbf{\Pi}}^{BPS}$ contains many zeros in cells more than 2 positions from on the diagonal, since few firms dramatically change their client counts from period to period. We therefore keep only 2 elements on each side of the diagonal, as well as the diagonal itself.

Specifically, for any given initial state, $n_t^B = i$, $2 < i < K - 1$, we include $\hat{\pi}_{i-2|i}^{BPS}$, $\hat{\pi}_{i-1|i}^{BPS}$, $\hat{\pi}_{i|i}^{BPS}$, $\hat{\pi}_{i+1|i}^{BPS}$, and $\hat{\pi}_{i+2|i}^{BPS}$ in \bar{m}_1 . So out of the i^{th} matrix, Ψ_i^{BPS} , we use the 5×5 submatrix:

$$\begin{aligned} \tilde{\Psi}_i^{BPS} &= \text{cov}[\pi_{i-2|i}^{BPS}, \pi_{i-1|i}^{BPS}, \pi_{i|i}^{BPS}, \pi_{i+1|i}^{BPS}, \pi_{i+2|i}^{BPS}] \\ &= \begin{bmatrix} \Psi_{i,(i-2,i-2)}^{BPS} & \Psi_{i,(i-1,i-2)}^{BPS} & \Psi_{i,(i,i-2)}^{BPS} & \Psi_{i,(i+1,i-2)}^{BPS} & \Psi_{i,(i+2,i-2)}^{BPS} \\ \Psi_{i,(i-2,i-1)}^{BPS} & \Psi_{i,(i-1,i-1)}^{BPS} & \Psi_{i,(i,i-1)}^{BPS} & \Psi_{i,(i+1,i-1)}^{BPS} & \Psi_{i,(i+2,i-1)}^{BPS} \\ \Psi_{i,(i-2,i)}^{BPS} & \Psi_{i,(i-1,i)}^{BPS} & \Psi_{i,(i,i)}^{BPS} & \Psi_{i,(i+1,i)}^{BPS} & \Psi_{i,(i+2,i)}^{BPS} \\ \Psi_{i,(i-2,i+1)}^{BPS} & \Psi_{i,(i-1,i+1)}^{BPS} & \Psi_{i,(i,i+1)}^{BPS} & \Psi_{i,(i+1,i+1)}^{BPS} & \Psi_{i,(i+2,i+1)}^{BPS} \\ \Psi_{i,(i-2,i+2)}^{BPS} & \Psi_{i,(i-1,i+2)}^{BPS} & \Psi_{i,(i,i+2)}^{BPS} & \Psi_{i,(i+1,i+2)}^{BPS} & \Psi_{i,(i+2,i+2)}^{BPS} \end{bmatrix} \end{aligned}$$

For the boundary cases $i \leq 2$ and $i \geq K - 1$, we truncate the vector $[\pi_{i-2|i}^{BPS}, \pi_{i-1|i}^{BPS}, \pi_{i|i}^{BPS}, \pi_{i+1|i}^{BPS}, \pi_{i+2|i}^{BPS}]$ as needed and adjust the associated covariance matrix $\tilde{\Psi}_i^{BPS}$ accordingly.⁵⁴ In total, we end up with 70 moments in \bar{m}_1 (and, of course, the same number of moments in \bar{m}_2).

Collecting these moments and assuming the elements of $\hat{\mathbf{\Pi}}^{BPS}$ are not correlated across rows, the block-diagonal covariance matrix for all of the BPS transition probabilities of interest is:

⁵⁴For example, $\tilde{\Psi}_2^{BPS} = [\pi_{1|2}^{BPS}, \pi_{2|2}^{BPS}, \pi_{3|2}^{BPS}, \pi_{4|2}^{BPS}]$ and $\tilde{\Psi}_{k-1}^{BPS} = [\pi_{k-3|k-1}^{BPS}, \pi_{k-2|k-1}^{BPS}, \pi_{k-1|k-1}^{BPS}, \pi_{k|k-1}^{BPS}]$.

$$cov(\bar{m}_1) = \begin{bmatrix} \tilde{\Psi}_1^{BPS} & 0 & 0 & \cdots & 0 \\ 0 & \ddots & 0 & \cdots & \vdots \\ 0 & 0 & \tilde{\Psi}_i^{BPS} & 0 & \vdots \\ \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & \cdots & \cdots & 0 & \tilde{\Psi}_K^{BPS} \end{bmatrix}_{(3K-2) \times (3K-2)}$$

Replacing population transition probabilities $\mathbf{\Pi}^{BPS}$ with sample transition probabilities $\hat{\mathbf{\Pi}}^{BPS}$, we obtain our estimator for $cov(\bar{m}_1)$. And by the same logic, replacing all *BPS* superscripts with *SPB* superscripts gives our estimate of $cov(\bar{m}_2)$.

D.2 Degree distributions

In addition to elements of the transition matrices, we target the degree distributions presented in Table 3.⁵⁵ These are collected in \bar{m}_3 and \bar{m}_4 which correspond to the seller and buyer degree distributions, respectively. As with the transition probabilities, our expressions apply equally to buyers and sellers, so we limit our exposition to the buyers per seller degree distribution.

Define the cumulative cutoffs of the cumulative distribution of firms to be (c_1, c_2, \dots, c_k) . (These correspond to the leftmost column of Table 3.) Call the observed fraction of sellers with at most c_1 partners \hat{G}_1^{BPS} , the fraction with at most c_2 partners \hat{G}_2^{BPS} and so on. Then relying on results used above, the covariance matrix for the vector $\hat{\mathbf{G}}^{BPS} = (\hat{G}_1^{BPS}, \hat{G}_2^{BPS}, \dots, \hat{G}_k^{BPS})$ takes the form:

$$cov(\hat{\mathbf{G}}^{BPS}) = \frac{1}{N_S} \begin{bmatrix} G_1^{BPS}(1 - G_1^{BPS}) & G_1^{BPS}(1 - G_2^{BPS}) & \cdots & G_1^{BPS}(1 - G_k^{BPS}) \\ G_1^{BPS}(1 - G_2^{BPS}) & G_2^{BPS}(1 - G_2^{BPS}) & \cdots & G_2^{BPS}(1 - G_k^{BPS}) \\ \vdots & \vdots & \ddots & \vdots \\ G_1^{BPS}(1 - G_k^{BPS}) & G_2^{BPS}(1 - G_k^{BPS}) & \cdots & G_k^{BPS}(1 - G_k^{BPS}) \end{bmatrix} \quad (\text{A-14})$$

where the total number of observations on sellers is N_S . Hence, $cov(\bar{m}_3) = cov(\hat{\mathbf{G}}^{BPS})$,

⁵⁵The degree distributions in Table 3 are related to the transition matrices in Tables 1 and 2. In fact, in a stationary equilibrium, if one had access to the complete transition matrices (without topcoding), one could construct them from the transition matrices. We do not attempt to account for the correlation between our transition matrix moments and degree distribution moments in our weighting matrix.

$cov(\bar{m}_4) = cov(\hat{\mathbf{G}}^{SPB})$, and we can approximate both objects by replacing the cumulative probabilities that underly them with their sample analogs.

D.3 Within-buyer shares

Finally, \bar{m}_5 contains the within-buyer market shares reported in Table 4. The variances of the shares of the top seller that we observe among buyers with two sellers, three sellers, and four sellers can reasonably be treated as independent of one another, since they are constructed from different sets of firms. Accordingly, we use their sample variances directly to construct $cov(\bar{m}_5)$.