

Two-sided Search in International Markets

J. Eaton^{a,b}, D. Jinkins^c, J. Tybout^{a,b}, D.Y. Xu^{d,b*}

^aPenn State, ^bNBER, ^cCopenhagen Business School, ^dDuke

(first draft: October 2016)

February 25, 2026

Abstract

We develop a dynamic model of international business-to-business transactions featuring two-sided endogenous search, bargaining, and many-to-many matching. Fitting our model to customs records on U.S. apparel imports, we quantify market features that shape the division of profits between suppliers and buyers and drive their distinct life cycles. In particular, suppliers are abundant relative to buyers, and large firms on each side of the market enjoy scale economies in search costs due to "visibility" effects. Using our estimated model we quantitatively investigate the IT revolution, the 2005 phaseout of the Agreement on Textiles and Clothing, and Trump's 2018 tariffs on Chinese apparel. Increasing the access of foreign exporters to the U.S. market can congest matching, dampening or even reversing the gains consumers enjoy from access to extra varieties. On the other hand, lower search costs can significantly improve consumer welfare, intensifying competition among both retailers and their upstream suppliers.

*We wish to acknowledge the fundamental contribution of our late coauthor Jonathan Eaton to this research project. His work and insight were invaluable to its development, and we miss him deeply. We also gratefully acknowledge support from the National Science Foundation (Grant SES-1426645) and the United States Census Bureau. We thank Costas Arkolakis, Andrew Bernard, Lukasz Drozd, Joris Pinkse, and participants in numerous seminars for many useful comments on earlier drafts. We also thank Pietra Rivoli and Julia Hughes (President, U.S. Fashion Industry Association) for very helpful discussions. Finally, we are grateful to Xuan Lin, Nicolas Urdaneta Andrade, Yingyan Zhao, and Zi Wang for excellent research assistance. 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, 5952-2017-0808, 11212-2024-0229, and 11397-2024-0514). 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.

1 Overview

International trade is built on relationships. To access foreign markets, firms must form partnerships with overseas buyers or suppliers, but these ties are often fragile and short-lived. Sustained participation requires continuous search and rematching, generating a constantly evolving network of supplier–buyer connections that governs the dynamics of 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 suppliers, domestic buyers (retailers), and domestic consumers. Heterogeneous suppliers and buyers search for each other, taking stock of their current situation and the aggregate search efforts of other agents. The matches that result determine which goods each buyer carries. Consumers then choose where to shop, purchasing the individual goods that buyers offer. For each active supplier-buyer relationship, ongoing forward-looking Nash bargaining determines wholesale and retail prices of the supplier’s good. Overall, the model connects the dynamic formation of trade relationships to evolving buyer offerings, prices, and consumer welfare.

1.1 Main messages

Fit to customs records on U.S. apparel imports, our model allows us to evaluate the impacts of several market shocks on network structure, trade, and welfare. First, we explore the effects of changing the number and mix of foreign apparel suppliers with access to the U.S. market. This experiment is calibrated to approximate the nearly simultaneous phaseout of a Chinese export license system and the Agreement on Textiles and Clothing (ATC) in 2005. We find that greater access of low-quality Chinese suppliers reduced overall U.S. welfare by generating congestion in the wholesale market, thereby inducing higher-quality suppliers to reduce their search efforts. The resulting reduction in apparel quality at the typical retailer (buyer) more than offset the standard love-of-variety benefit from adding products. Consumers were also hurt by the exit of some smaller buyers, which narrowed their shopping choices and shifted retail market shares toward “big box” stores.

In a second experiment, we simulate a decrease in search costs on both sides of the wholesale market. This exercise is meant to approximate advances in information and communication technology (ICT), and is calibrated to match the growth in buyer-supplier connections we

observe over the period 1998 – 2004. Here we find that a reduction in search costs increased consumer welfare significantly, but at the expense of the profits of both individual buyers and suppliers. The simple reason is that increases in search efficiency inspired entry on both sides of the market, spreading consumer spending more thinly across active firms.

Finally, in our third experiment, we simulate the short-run impact and longer-term effects of Trump’s 2018 (Section 301) tariffs on Chinese apparel. We find that initially consumer welfare fell several percentage points because the tariffs were passed through to consumers. But over time, as lower-quality Chinese exporters exited, higher-quality exporters in other countries intensified their search efforts and rapidly formed new matches with U.S. retailers. This endogenous reallocation reshaped the quality composition of retailer assortments and substantially offset the initial welfare loss.

Beyond serving as a quantitative laboratory, our structural framework recovers key latent objects that are not directly observable in customs data, including the distribution of match-specific rents and the search efforts those rents induce on both sides of the market. Identifying these objects is essential for understanding how changes in market access or policy reshape network formation. In our baseline estimation we find that buyers and suppliers spend comparable amounts on search in the aggregate. But since there are far more active suppliers than active buyers, buyers spend less on search per match. As a whole, buyers obtain larger shares of the joint retail surplus despite incurring lower search costs.

Our model also delivers a new perspective on firm dynamics. Fundamentally, firms’ life cycles emerge endogenously from search frictions and competitive pressure, rather than through exogenous productivity dynamics. As buyers and suppliers 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 accumulation of partners.

1.2 Relation to the literature

Our paper relates to a wide variety of earlier contributions. First, it connects to papers on firm-level export dynamics that feature customer accumulation processes (Albornoz et al., 2012; Drozd and Nosal, 2012; Eaton et al., 2025; Chaney, 2014; Carballo et al., 2018;

Piveteau, 2021; Rodrigue and Tan, 2019; Fitzgerald et al., 2024).¹ We extend this strand of the literature by studying customer accumulation at the supplier-buyer match level rather than the supplier-country level.² In doing so, we characterize many-to-many matching patterns between heterogeneous agents in a dynamic equilibrium, and we model the associated bargaining outcomes between each buyer-supplier pair.

Because our model predicts firms’ dynamic matching patterns, it also connects to the literature that links firms’ life-cycles 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; Boehm et al., 2024). Instead, as in Eaton et al. (2025), we generate these tails by incorporating “visibility effects” in our search cost function, thereby allowing firms with a large market presence to find new business partners with relative ease. This 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 relevant literature focuses on transnational firm-to-firm trading patterns and the question of who matches with whom (Rauch, 2001; Rauch and Trindade, 2002; Bernard et al., 2018, 2022; Benguria, 2021; Monarch, 2022; Sugita et al., 2023; Eaton et al., forthcoming).³ Our model assumes that matching is random, 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.

Finally, since our firms deal in clothing, our paper relates to a substantial literature on global apparel markets. In addition to the descriptive studies summarized in Section 2.3 below, three papers are particularly relevant. First, Cahal et al. (2023) analyze the sourcing strategies apparel importers pursue, distinguishing those that rely on spot markets for each order from those that pursue longer term relationships. We also analyze apparel importers’ strategies, but in order to do so in the context of our dynamic structural model, we rule out relational contracts. Second, Khandelwal et al. (2013) analyze the effects of China’s export

¹See Alessandria et al. (2021) for a recent review of the literature on firm-level export dynamics. Relevant contributions focusing on the accumulation of customers in a domestic context include Foster et al. (2016), Gourio and Rudanko (2014), and Boehm et al. (2024).

²Eaton et al. (2025) also use match level data to study export dynamics. However, theirs is a single agent model that does not involve importer search and bargaining.

³Also see Antras and Chor (2021) for a survey of the large recent literature on value chains and firm-to-firm production networks.

licensing regime under the Agreement on Textiles and Clothing (ATC), finding that the ATC phaseout helped relatively low-price firms that had been previously constrained. We also analyze the effects of the ATC phase-out. However, since we infer exporters' qualities from their sales volumes rather than their unit values, we arrive at different conclusions. Finally, Bai et al. (2017) analyze the effects of China's pre-2005 restrictions on direct exporting by small firms, finding that the phasing out of these restrictions improved small exporters' productivity. While we also analyze the dynamic effects of this policy reform, we rule out indirect exporting and we treat firms' quality as time-invariant in order to focus on customer accumulation.

2 Data and Stylized Facts

While nothing in our framework restricts its application to international trade more broadly, we study the network of U.S. apparel importers and their foreign suppliers for several reasons. First, since most importers are either wholesale or retail firms, we can keep the import side of the market relatively simple. In particular, the revenue function of such firms is nearly separable across categories of consumer goods, so 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 suppliers play 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 can be viewed through the lens of the model, extracting implications for profit distributions and welfare.

Before describing the details of our model, we review some aggregate patterns in U.S. apparel trade over the last several decades and some micro features of the associated buyer-supplier 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, 2025; Bernard et al., 2018), Chile (Blum et al., 2010), Mexico (Sugita et al., 2023), Norway (Bernard et al., 2018), and Ireland (Fitzgerald et al., 2024).⁴ Other dynamic network features have received less attention, particularly those concerning

⁴Bernard and Moxnes (2018) review many of the stylized facts in the firm-to-firm trade literature.

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 2011.⁵ Among other variables, each record includes a ten-digit Harmonized Schedule (HS) product code, shipment value, shipment quantity, date of transaction, 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-supplier pairs through time.

The name and address of a given exporter may be recorded differently for different shipments, and this noise in our identifier can lead to overstatement in the number of exporters and business relationships, as well as in the rates of relationship turnover (Kamal and Monarch, 2017; Krizan et al., 2020). However, as Kamal and Monarch (2017) note, suppliers are relatively likely to be accurately identified for textile and apparel products.⁶

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 2007.⁷ These trends reflect the emergence of China and other developing economies as exporters and the phasing out of the ATC in 2005.

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

⁵We 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 paper.

⁶Kamal and Monarch (2017) state: “it is clear from U.S. regulations that the [foreign manufacturer ID (MID)] is used to track compliance with U.S. restrictions for textile shipments. MID criteria for textiles are the most stringent, since non-textile products typically do not have the rule-of-origin restrictions that exist for textile and apparel products.”

⁷Domestic 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 suppliers, 1996-2011

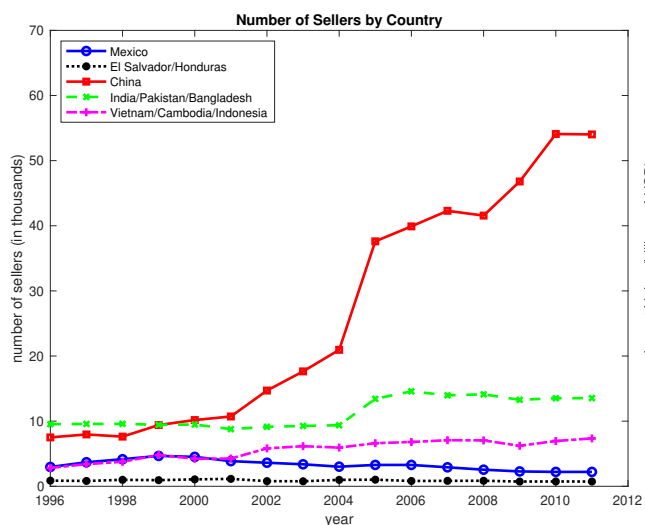


Figure 3: Number of suppliers by country, 1996-2011

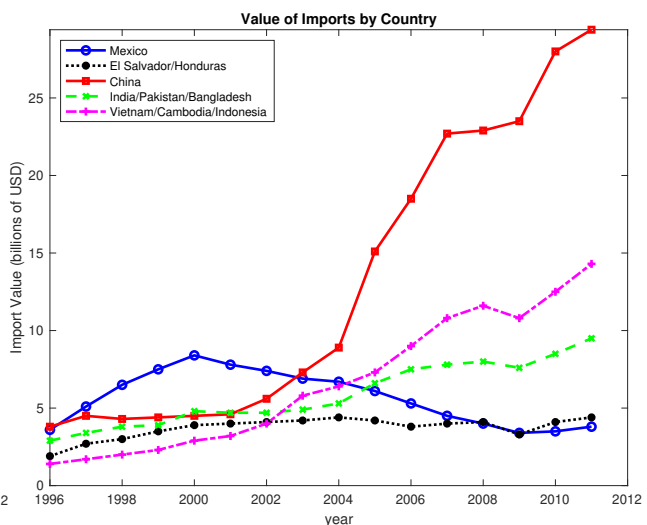


Figure 4: Value of imports by country, 1996-2011

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. These figures also imply that China and India shipped substantially less *per exporter* to the U.S. than other countries—a fact we will return to later. 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.⁸ 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.⁹ Such firms may sell directly to consumers as well, but they also sell to retailers.

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 retailer-supplier

⁸Plunkett-Analytics (2015) and Gereffi and Memedovic (2003) provide related classifications.

⁹VF Corp owns JanSport, The North Face, Timberland, Lee, Wrangler, and Nautica.

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 owns 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.¹⁰

For our analysis we approximate this structure by partitioning the players in the market into three mutually unaffiliated types: suppliers (manufacturers either direct or intermediated), buyers (both wholesale and retail importers), and consumers. We ignore wholesale and retail firms that own their production facilities. We also gloss over the distinction between firms with in-house sourcing departments and firms that use sourcing firms. Accordingly, we treat suppliers' sales to buyers that are intermediated by sourcing firms the same as suppliers' sales that are not intermediated, thereby ruling out a potential reason for double marginalization.

How much distance do these simplifications put between our model and the actual structure of the market? First, while a few branded importers 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.¹¹ McFarlan et al. (2012) report that, in 2004, 9 of the 10 largest apparel retailers in the U.S. ran their own

¹⁰Census classifies *establishments* according to their main activity, such as wholesaling or retailing. Customs records report the importing *firm*, which may own multiple establishments in different categories.

¹¹"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 many representatives of apparel manufacturers located in South Asia.




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

sourcing offices.¹² In a more recent report, the U.S. Fashion Industry Association surveyed 30 executives representing various segments of the apparel importing market (Lu, 2016).¹³ 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.¹⁴

¹²Kohl’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).

¹³Almost 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.)

¹⁴Online 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 have moved away from 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 (Cahal et al., 2023), but most buyer-supplier 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 0.77 we estimate from our data.¹⁵

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 suppliers deal with. Table 1 and Table 2 report annual transition rates for suppliers’ buyer counts and buyers’ supplier 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 suppliers and buyers. Second, the transition matrix also exhibits a more pronounced downward adjustment probability when the number of relationships is relatively small. For instance in Table 1, when a supplier has 1 – 9 buyers, it is likely to have fewer relationships next year. Similar patterns hold for suppliers 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.

2.3.3 Degree distributions

Consistent with the tendency to lose business partners in Tables 1 and 2, it is unusual for a firm to sustain large portfolios of foreign partners. Table 3 reports the frequency distributions

¹⁵We 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.

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

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, 1996-2011. Figures are cross-year averages of annual transition rates during the sample period. Buyer-supplier 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-supplier pairs that generate a single shipment are not considered to have matched.

Table 2: Year-to-year transition rates: suppliers 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, 1996-2011. Figures are cross-year averages of annual transition rates during the sample period. Buyer-supplier 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-supplier pairs that generate a single shipment are not considered to have matched.

of suppliers per buyer and buyers per supplier. Note that the vast majority of the firms have less than 10 partners in our data (86 percent of the buyers and 99 percent of the suppliers).¹⁶ Nonetheless, both degree distributions come close to obeying a power law, implying that a small fraction of firms attain very large sizes.¹⁷

Table 3: Firm distributions by partner counts, 2011

x	Share of buyers with at most x suppliers	Share of suppliers 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

2.3.4 Sales heterogeneity

In addition to cross-firm heterogeneity in partner counts, we observe substantial sales heterogeneity. Column 2 of Table 4 reports average log imports per supplier among buyers with a single supplier (row 1), two suppliers (row 2), and so on. All log imports per supplier are relative to the average log imports of buyers with a single supplier. So, for example, buyers with 2 suppliers spend 113 percent more per supplier than buyers with a single supplier. Given the supplier arrival patterns documented in Table 2, the association between purchases per supplier and number of suppliers will help us to identify the distribution of buyer types.

The remaining columns in Table 4 report the average within-buyer share of the m^{th} largest supplier among buyers with n suppliers, $m \in \{1, 2, \dots, n\}$. (For each row, columns 3-12 sum to one.) Apparel importers typically have a dominant supplier whose share in total imports drops only modestly as the number of suppliers increases. These dominant supplier within-buyer

¹⁶The supplier’s degree distribution is limited to its U.S. partners, so this pattern is not surprising.

¹⁷This power-law feature appears in data from other countries, including Colombia (Eaton et al., 2008, 2025; 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 supplier rises from 1.99 in 2000 to 2.84 in 2011, reflecting an increase in the number of suppliers per buyer.

revenue shares, and the market shares of the less important suppliers, will help us identify the distribution of supplier types, as we discuss in section 4.2 below.

Table 4: Buyers’ imports per supplier and within-buyer supplier shares

No. suppliers	mean log imports	Supplier shares in buyer purchases, ordered by supplier size									
		1 st	2 nd	3 rd	4 th	5 th	6 th	7 th	8 th	9 th	10 th
1	0.000	1.000	0	0	0	0	0	0	0	0	0
2	1.134	0.771	0.229	0	0	0	0	0	0	0	0
3	1.604	0.668	0.240	0.092	0	0	0	0	0	0	0
4	1.764	0.608	0.232	0.111	0.049	0	0	0	0	0	0
5	1.904	0.544	0.231	0.126	0.067	0.032	0	0	0	0	0
6	2.054	0.540	0.218	0.115	0.070	0.039	0.019	0	0	0	0
7	2.214	0.502	0.214	0.123	0.074	0.045	0.027	0.019	0	0	0
8	2.094	0.460	0.212	0.125	0.080	0.054	0.035	0.023	0.011	0	0
9	2.364	0.451	0.201	0.121	0.083	0.055	0.038	0.026	0.017	0.017	0
10	2.324	0.420	0.197	0.125	0.084	0.060	0.042	0.029	0.021	0.013	0.007

*Based on monthly U.S. customs records from 2011. Figures in column 2 are average log imports per supplier of buyers with n suppliers. They are expressed net of the mean log imports per supplier for buyers with a single supplier. Figures in columns 3-12, the n^{th} row give the average within-buyer share of the m^{th} largest supplier among buyers with n suppliers, $m \in \{1, 2, \dots, n\}$.

2.3.5 Search costs

The frequent dissolution of buyer-supplier 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.¹⁸ What form do these costs take?

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: “[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 at 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.¹⁹ Since each importer has its own standards regarding acceptable practice,

¹⁸WalMart’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.

¹⁹In 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

the industry norm is for each firm to perform an audit of each factory it deals with before placing any orders.²⁰

We know less about search efforts on the manufacturers’ side. As mentioned above, 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., 2023).

In summary, US apparel imports have risen substantially, with steady increases in the numbers of suppliers 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 suppliers in this market. Underlying these aggregate changes are the fluid creation and destruction of business relationships between heterogeneous importers and exporters. We now develop a model of two-sided search that accounts for these data features in an equilibrium framework.

3 A model of buyer-supplier networks

We model the interactions between three types of agents in continuous time: suppliers (exporting manufacturers), buyers (importers/retailers), and consumers.²¹ Figure 7 provides a schematic overview. In the retail market, a representative final consumer sources her flow purchases from an evolving set of heterogeneous buyers, each of which offers its own evolving collection of goods. The consumer values each buyer for the menu of products it currently sells and for the amenities it offers—e.g., a convenient location, pleasant ambience, or attentive service. In the wholesale market, heterogeneous suppliers export their products to discrete subsets of heterogeneous buyers, providing each with a variety that is custom-tailored to the buyer’s specifications.

The representative consumer cannot save, and goods are non-durable, so she simply max-

risks such as ‘labor disputes,’ ‘political unrest,’ and ‘lack of resources to manage supply chain risks.’ ” (Lu, 2016).

²⁰This observation is based on a telephone interview with the president of the U.S. Fashion Industry Association, December 14, 2016.

²¹In the previous section, we documented that there are several types of importers in the apparel sector. We abstract from importer type in order to simplify our model, but allow for heterogeneity in retailer amenities.

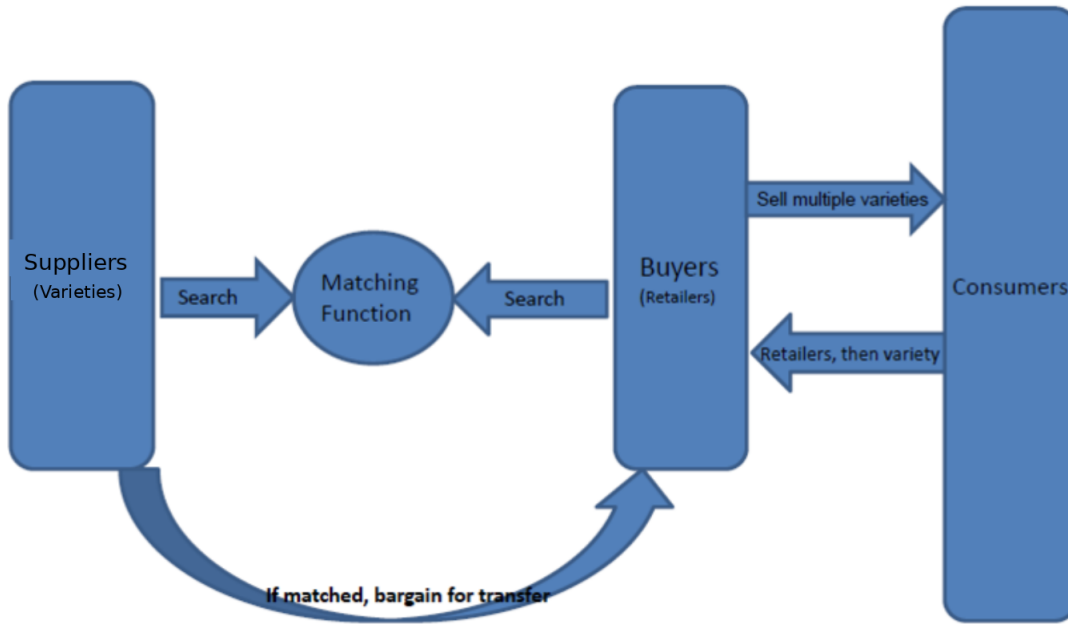


Figure 7: Model diagram

imizes her flow utility at each point in time. But buyers and suppliers are forward looking. They continuously make three types of choices. First, they choose the vector of retail prices for the products they currently offer. They base these choices on the final consumer’s demand elasticities and their product-specific marginal costs.²² Second, each buyer bargains multilaterally with her suppliers to determine the division of the surplus stream they jointly generate. In these negotiations, all agents recognize that the buyer’s portfolio of connections will evolve stochastically as it adds new suppliers and loses others. Finally, buyers and suppliers choose the amount they spend on searching for new business partners. These search efforts maximize firms’ present values, and govern the stochastic evolution of buyers’ and suppliers’ matching patterns.

More intense search increases the hazard of finding a new partner. This matching hazard also depends on wholesale market tightness. When the measure of buyers’ search for suppliers is high relative to the measure of suppliers’ search for buyers, matching hazards are low for buyers and high for suppliers, and vice versa. In addition, the ease with which agents find new partners depends on their previous successes. Agents who have already accumulated a large portfolio of partners find it easier to locate still more. This feature of our model, taken from

²²Accordingly, our model does not feature double marginalization.

Eaton et al. (2025), helps us capture the “fat-tailed” distributions of buyers across suppliers and suppliers across buyers discussed above.

3.1 The Retail Market

We now turn to model specifics, suppressing time subscripts where no ambiguity results. Consider the retail market first.

Preferences: At each point in time, the representative consumer allocates her budget across a continuum of buyers, indexed by y . Following Atkin et al. (2018) and Hottman et al. (2016), the utility she derives from purchasing from buyer y depends on both the amenities it provides and the discrete set of products it offers, J_y . Upstream suppliers provide a single unique product to each buyer they match with, so sets J_y and $J_{y'}$ are mutually exclusive for all $y \neq y'$. However, multiple buyers may source their goods from the same supplier(s).²³

The consumer’s preferences over products and buyers are given by the utility function:

$$U = \left[\int_y (\mu_y Q_y)^{(\eta-1)/\eta} \right]^{\eta/(\eta-1)},$$

where

$$Q_y = \left[\sum_{x \in J_y} (\xi_x q_{xy})^{(\alpha-1)/\alpha} \right]^{\alpha/(\alpha-1)}.$$

Here ξ_x measures the (time-invariant) appeal of any good produced by supplier x , μ_y measures the (time-invariant) appeal of buyer y , and q_{xy} is consumption of a good produced by supplier x and purchased from buyer y . It is straightforward to show that the associated flow demand for product x at buyer y is

$$q_{xy} = (\xi_x)^{\alpha-1} (\mu_y)^{\eta-1} \left(\frac{p_{xy}}{P_y} \right)^{-\alpha} \left(\frac{P_y}{P} \right)^{-\eta} \frac{E}{P}, \quad (1)$$

where p_{xy} is this good’s retail price, $P_y = \left[\sum_{x \in J_y} \left(\frac{p_{xy}}{\xi_x} \right)^{1-\alpha} \right]^{1/(1-\alpha)}$, $P = \left[\int_y \left(\frac{P_y}{\mu_y} \right)^{1-\eta} \right]^{1/(1-\eta)}$, and E is total consumer expenditure on apparel.²⁴

²³For instance, a single foreign supplier may manufacture T-shirts in distinct styles tailored to each buyer it serves.

²⁴This market-wide final demand function can also be interpreted to represent an aggregation of consumer-

Unit cost: When buyer y and supplier x collaborate to offer a product in the retail market, they jointly incur a supplier-specific, time-invariant unit cost, denoted c_x .²⁵ (We impose that the buyer’s share of the unit cost—covering expenses such as sales personnel, transportation, and retail operations—is proportional to the supplier’s marginal production cost, thereby eliminating the need for a y subscript on c_x .) Suppliers have constant returns and no capacity constraint, so c_x is independent of the number and volume of products supplier x produces.

Pricing and retail profits: Each buyer y bargains bilaterally and continually with its set of connected suppliers, J_y , over the retail prices it sets and the compensation it pays them.²⁶ The resulting payments cover each supplier’s production cost and, in addition, include a negotiated portion of the coalition’s total retail profits,

$$\pi_y^T = \sum_{x \in J_y} (p_{xy} - c_x) q_{xy}. \quad (2)$$

We discuss details of the bargaining game in Section 3.2.3 below. For now, it suffices to note that this game leads buyer y and its suppliers to agree to the set of prices that maximize π_y^T at each point in time. So the equilibrium price vector solves the system

$$q_{xy} + \sum_{x' \in J_y} \frac{\partial q_{x'y}}{\partial p_{xy}} (p_{x'y} - c_{x'}) = 0 \quad \forall x \in J_y, \quad (3)$$

with q_{xy} given by equation (1).

Given our nested CES utility function, these first-order conditions imply that the equilibrium mark-up for each product is simply (Appendix A.1):²⁷

$$\frac{p_{xy} - c_x}{p_{xy}} = \frac{1}{\eta}. \quad (4)$$

Thus, by equations (1) and (2), the total current retail profit—or flow surplus—shared between

specific nested-logit demand function, with each consumer purchasing a single unit of her most-preferred product (Verboven, 1996).

²⁵ c_x excludes search costs, which are sunk by the time prices are set for product x .

²⁶ There are no strategic interactions between buyers in the retail market because each buyer is measure 0 and each product is available from only one buyer. Hence the bargaining is one-to-many.

²⁷ The same result, along with some intuition, can be found in Hottman et al. (2016).

buyer y and its suppliers can be written as (Appendix A.1):

$$\pi_y^T = \frac{1}{\eta} \frac{E}{P^{1-\eta}} \left[\sum_{x \in J_y} \left(\frac{\eta}{\eta-1} \right)^{1-\alpha} \tilde{c}_x^{1-\alpha} \right]^{(1-\eta)/(1-\alpha)} \mu_y^{\eta-1}, \quad (5)$$

where $\tilde{c}_x = c_x/\xi_x$ is the quality-adjusted unit cost for buyer-supplier pair xy . Note that when $\alpha > \eta > 1$, this function exhibits diminishing expected returns with respect to y 's number of suppliers, $\|J_y\|$.

3.2 The Wholesale Market

Buyers source their merchandise from an upstream wholesale market populated by a continuum of suppliers. We now turn to the random matching process through which buyers and suppliers meet, and the resulting flow payoffs for each type of agent.

3.2.1 Total flow surplus with discrete types

To begin, we assume there are a finite number of types for both suppliers and buyers. Buyers are divided into I intrinsic types indexed by $i \in \{1, 2, \dots, I\}$, and suppliers are categorized into J intrinsic types indexed by $j \in \{1, 2, \dots, J\}$. For any type- j supplier, the quality-adjusted marginal cost of producing and retailing each of its goods is $\tilde{c}_j = c_j/\xi_j$. So while products can be vertically differentiated across suppliers, they are purely horizontally differentiated within each supplier's portfolio of offerings. Hereafter we will refer to suppliers with high \tilde{c} values as "low-quality" firms.

Let $\mathbf{s} = [s_1, s_2, \dots, s_J]$ be a vector of counts of the number of suppliers of each type currently matched with 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 with a particular supplier. Using this notation to restate equation (5), the total flow surplus accruing to a type- i buyer and its portfolio of suppliers \mathbf{s} is:

$$\pi_i^T(\mathbf{s}) = \frac{E}{\eta P^{1-\eta}} \left[\sum_{j=1}^J \left(\frac{\eta}{\eta-1} \right)^{1-\alpha} s_j \tilde{c}_j^{1-\alpha} \right]^{\frac{\eta-1}{\alpha-1}} \mu_i^{\eta-1} \quad (6)$$

3.2.2 Search

Since all buyers are known to the final consumer, there are no matching frictions in the retail market. But in the wholesale market, buyers and suppliers must both exert effort to find new business partners. In this section we describe the optimization problems solved by buyers (B) and suppliers (S) when choosing their effort levels.

Search costs All matches generate surplus in our baseline model, so agents on each side of the market always prefer to have more of them. To this end, each agent chooses her own search effort, σ , and finds a new business partner with hazard $\sigma\theta^A$, where θ^A is the buyer ($A = B$) or supplier ($A = S$) match hazard per unit of search effort.

These efforts are moderated by search costs. Specifically, agents with n^A partners who search at level σ must pay the flow cost

$$k^A(\sigma, n^A) = \frac{k_0\sigma^2}{(n^A + 1)^{\gamma^A}}, \quad A \in \{B, S\}, \quad (7)$$

where $n^B(\mathbf{s}) = \sum_{j=1}^J s_j$ and $n^S(\mathbf{b}) = \sum_{i=1}^I b_i$. This functional form generalizes Arkolakis (2010) by allowing search costs to rise or fall with market connections, n^A . It nests two possible cases. If $\gamma^A < 0$, firms with many partners find it relatively difficult to add still more, perhaps because their market is already “fished out,” as in Arkolakis (2010). If $\gamma^A > 0$, firms with many partners find it relatively easy to add still more, perhaps because their scale makes them highly visible to firms on the other side of the market. Regardless of the sign of γ^A , we will hereafter refer to the scale effect it captures as the “visibility effect.”

Match separations Once formed, relationships eventually terminate for one of two exogenous reasons.²⁸ First, with hazard rate δ , the buyer and supplier become incompatible. Second, the buyer or supplier may exit the wholesale market.²⁹ We denote the hazard rates of these exit events δ^B and δ^S , respectively.

²⁸In Appendix D.2 we discuss an alternative specification in which transitory match shocks generate endogenous separations.

²⁹When a buyer or supplier exits, it simultaneously loses all of its connections. This incrementally reduces the mass of active agents, but the total mass of *potential* agents remains constant by assumption, as will be seen in Section 3.3.1 below.

Buyers' search problem As we will discuss shortly, a bargaining game between buyers and their suppliers determines the division of the surplus (6) at each point in time. For a type- i buyer in state \mathbf{s} , this game yields flow payoff $\tau_{ji}(\mathbf{s})$ for each of its type- j suppliers, leaving it with flow payoff $\pi_i^T(\mathbf{s}) - \sum_{j=1}^J s_j \tau_{ji}(\mathbf{s})$. Accordingly, the buyer's value function $V_i^B(\mathbf{s})$ solves

$$\begin{aligned}
(\rho + \delta^B)V_i^B(\mathbf{s}) = & \pi_i^T(\mathbf{s}) - \sum_{j=1}^J s_j \tau_{ji}(\mathbf{s}) \\
& + \max_{\sigma_i^B} \left(\sigma_i^B \theta^B \sum_{j=1}^J v_j^S [V_i^B(\mathbf{s} + \mathbf{1}_j) - V_i^B(\mathbf{s})] - k^B(\sigma_i^B, n^B) \right) \\
& + (\delta + \delta^S) \sum_{j=1}^J s_j [V_i^B(\mathbf{s} - \mathbf{1}_j) - V_i^B(\mathbf{s})],
\end{aligned} \tag{8}$$

where $\mathbf{1}_j$ is a $J \times 1$ vector with j^{th} element 1 and 0's elsewhere, v_j^S is the probability that the next supplier the buyer meets will be type- j , and ρ is the discount rate. The optimal search policy for type- i buyers with a set of \mathbf{s} suppliers, $\sigma_i^B(\mathbf{s})$, therefore satisfies

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

The logic of equation (8) is straightforward. The buyer reaps its flow payoff, net of search costs, until the next event occurs. With hazard $s_j(\delta + \delta^S)$ this event is the exogenous termination of a type- j supplier relationship, and with hazard $\sigma_i^B \theta^B v_j^S$ it's a new match with a type- j supplier.

Suppliers' search problem Suppliers have constant marginal production costs, so the payoffs from each of their matches are separable, and their return to searching depends only on the expected value of an additional match. For a type- j supplier, denote this expected value V_j^S . Then if this type of supplier has n^S buyers, its optimal search intensity maximizes $\sigma_j \theta^S V_j^S - k^S(\sigma_j^S, n^S)$, and the associated first-order condition is:

$$\frac{\partial k^S(\sigma_j^S, n^S)}{\partial \sigma_j^S} = \theta^S V_j^S.$$

What determines V_j^S ? With our assumption of random matching, it can be written as:

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

where $v_i^B(\mathbf{s})$ is the probability that the next buyer the supplier meets will be type- i in state \mathbf{s} , and $V_{ji}^S(\mathbf{s})$ is the value it places on matches with this type of buyer:

$$\begin{aligned} (\rho + \delta + \delta^B + \delta^S) V_{ji}^S(\mathbf{s}) &= \tau_{ji}(\mathbf{s}) + \sigma_i^B \theta^B \sum_{k=1}^J v_k^S [V_{ji}^S(\mathbf{s} + \mathbf{1}_k) - V_{ji}^S(\mathbf{s})] \\ &\quad + (\delta + \delta^S) \sum_{k=1}^K (s_k - \mathbf{1}_{k=j}) [V_{ji}^S(\mathbf{s} - \mathbf{1}_k) - V_{ji}^S(\mathbf{s})]. \end{aligned} \quad (9)$$

Intuitively, a business relationship with a type- i buyer who has \mathbf{s} suppliers will terminate with exogenous hazard $(\delta + \delta^B + \delta^S)$, become a relationship with a type- i buyer who has $\mathbf{s} + \mathbf{1}_k$ suppliers with hazard $\sigma_i^B \theta^B v_k^S$, and become a relationship with a type- i buyer who has $\mathbf{s} - \mathbf{1}_k$ suppliers with hazard $(s_k - \mathbf{1}_{k=j}) \delta$. The factor $(s_k - \mathbf{1}_{k=j})$ adjusts for the risk of a supplier itself being dropped when exogenous separation or exit occurs.

3.2.3 Bargaining

Each buyer i and its suppliers bargain over the division of the total surplus they create, $V_i^T(\mathbf{s}) = V_i^B(\mathbf{s}) + \sum_j s_j V_{ji}^S(\mathbf{s})$. Under assumptions we shall discuss shortly, this game leads them to agree to the flow transfers $\tau_{ij}(\mathbf{s})$ that solve the following system of equations:

$$\tau_{ij}(\mathbf{s}) = \frac{1}{2} \left[\underbrace{\pi_i^T(\mathbf{s}) - \pi_i^T(\mathbf{s} - \mathbf{1}_j)}_{j\text{-type supplier's effect on surplus}} - \underbrace{\sum_{k|s_k > 0} (s_k - \mathbf{1}_{k=j}) (\tau_{ik}(\mathbf{s}) - \tau_{ik}(\mathbf{s} - \mathbf{1}_j))}_{j\text{-type supplier's effect on transfers to other suppliers}} \right], \quad (10)$$

$j = 1, \dots, J; \quad \mathbf{s} \in \mathbb{S}$

where \mathbb{S} is the set of all feasible supplier portfolios. Intuitively, equation (10) underscores the two ways a supplier benefits a buyer. Not only does a new supplier contribute to the total flow surplus, but it also reduces the transfers to all other suppliers by reducing their bargaining power. The supplier and buyer split this marginal surplus equally.

Although our model characterizes a many-to-many matching equilibrium, the bargaining game that determines this surplus division is one (buyer) to many (suppliers). The reason is that suppliers have constant marginal production costs and offer distinct products to each of their (measure zero) buyers. This makes the outcome of their bargaining with any one of their buyers independent of the outcome of their negotiations with the others. In contrast, a buyer’s bargaining sessions with her portfolio of suppliers are interdependent because the surplus function (6) exhibits diminishing returns with respect to the supplier count, except in the special case where $\alpha = \eta$ (to be discussed in Section D.2 below).

We obtain equation (10) under the following assumptions:³⁰

Assumption 1. *No commitment:* *Buyers and their suppliers cannot commit to long-term transfer schedules. They renegotiate whenever the buyer’s set of connected suppliers changes.*

Assumption 2. *Bilateral bargaining:* *Buyers use delegates to bargain on their behalf—one for each connected supplier—and the bilateral bargaining sessions these delegates conduct occur simultaneously, following the protocol described by Binmore et al. (1986).*

Assumption 3. *Private information:* *The specific offers and counter-offers made by each delegate and the supplier she bargains with are privately observed by the two parties involved. Bargaining breakdowns are, however, publicly observed.*

Assumption 4. *Passive beliefs:* *When a delegate or a supplier receives an out-of-equilibrium offer or an unexpected rejection, she does not revise her beliefs regarding the bilateral negotiations between other delegates and suppliers.*

Assumption 5. *Limited contract space:* *Agents cannot condition their contract offers on the buyer’s search effort or portfolio history.*

Together, these assumptions simplify the multilateral dynamic bargaining problem to a Nash-like problem in which each delegate-supplier pair maximizes its bilateral surplus, taking beliefs about the actions of other agents as given. Assumption 1 allows buyers and suppliers

³⁰Stole and Zwiebel (1996) motivate equation (10) using a distinct set of assumptions. We do not rely on their bargaining protocol because Brugemann et al. (2019) have shown it to be problematic. Specifically, since Stole and Zwiebel (1996) assume complete information and sequential bargaining, a supplier’s position in the bargaining queue affects her bilateral bargaining power. This leads to heterogeneous transfers for otherwise identical suppliers, contrary to equation (10). Brugemann et al. (2019) provide a set of assumptions that avoid this problem, but their version of equation (10) presumes homogeneous suppliers, so their “Rolodex game” protocol does not apply to our surplus splitting rule.

to adjust their agreements whenever the set of connected agents changes. Assumption 2 pins down the division of the marginal flow surplus generated by the two agents involved in any one of the bilateral games, holding fixed agents' beliefs about all of the other games. Assumptions 3 and 4 justify this treatment of beliefs and help ensure that a Bayes Nash equilibrium exists. Finally, Assumption 5 rules out tenure-dependent contracts. Otherwise, suppliers would have an incentive to try to discourage their buyers' search efforts, as in Lentz (2014).

Further details concerning the role of these assumptions appear in Appendix A.2. Here we briefly note that they lead to three lemmas. The first two—which follow the logic of de Fontenay and Gans (2014)—determine what bargaining outcomes would obtain if agents were to bargain over the flow surplus $\pi^T(\mathbf{s})$ at each point in time without regard to the future. Lemma A.1 establishes that the outcomes of these bargaining games over $\pi^T(\mathbf{s})$ are "bilaterally efficient," i.e., they maximize the marginal surplus generated by the supplier-buyer pair, taking the actions of the other bilateral games as given.³¹ Accordingly, they imply equation (3). Lemma A.2 establishes that, given bilateral efficiency, the bilateral surplus divisions are "fair," i.e., each supplier-buyer pair evenly splits the surplus it generates, as implied by equation (10). Finally, Lemma A.3 establishes that when a buyer and her suppliers negotiate over $V_i^T(\mathbf{s})$ rather than $\pi_i^T(\mathbf{s})$ —that is, when negotiations are forward-looking—the resulting transfers will *still* satisfy equation (10) at each point in time.

3.3 Market Dynamics

As suppliers and buyers match with one another, and as existing matches die, agents' portfolios of connections evolve. Let $M_i^B(\mathbf{s})$ be the measure of type i buyers in state \mathbf{s} , and let $M_j^S(\mathbf{b})$ be the measure of type j suppliers in state \mathbf{b} . We now develop the equations of motion for each of these objects.

3.3.1 Equations of Motion

Consider buyers first. And, for the moment, take the matching hazard per unit of search θ^B and conditional matching probabilities ν_j^S as given. Then, since relationships with a supplier end with exogenous hazard $\tilde{\delta} = \delta + \delta^S$, the equation of motion for the measure of type- i buyers

³¹The rule need not be *socially* optimal. In addition to the standard congestion externalities in search and matching models, buyers in our setting also have incentives to over-accumulate suppliers and thereby capture a larger share of total surplus.

with \mathbf{s} suppliers is:

$$\begin{aligned} \dot{M}_i^B(\mathbf{s}) &= \sum_j \left[\mathbf{1}_{s_j > 0} \sigma_i^B(\mathbf{s} - \mathbf{1}_j) \theta^B v_j^S M_i^B(\mathbf{s} - \mathbf{1}_j) + \tilde{\delta}(s_j + 1) M_i^B(\mathbf{s} + \mathbf{1}_j) \right] \\ &\quad - \left[\sigma_i^B(\mathbf{s}) \theta^B M_i^B(\mathbf{s}) + (\tilde{\delta} n^B(\mathbf{s}) + \delta^B) M_i^B(\mathbf{s}) \right], \\ &\quad \mathbf{s} \in \{\mathbb{S} | n^B(\mathbf{s}) > 0\}; \quad i = 1, \dots, I \end{aligned} \quad (11)$$

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 $\tilde{\delta}(s_j + 1)$. By analogous logic, the group loses existing members that either add a supplier (with hazard $\sigma_i^B(\mathbf{s}) \theta^B$), lose a supplier (with hazard $\tilde{\delta} n^B(\mathbf{s}) = \tilde{\delta} \sum_j s_j$), or exogenously exit (with hazard δ^B).

The measure of buyers of type i with $n^B = 0$ suppliers evolves according to:

$$\dot{M}_i^B(\mathbf{0}) = \delta^B \sum_{n^B(\mathbf{s}) \neq 0} M_i^B(\mathbf{s}) + \tilde{\delta} \sum_j M_i^B(\mathbf{1}_j) - \sigma_i^B(\mathbf{0}) \theta^B M_i^B(\mathbf{0}) \quad (12)$$

Increases in $M_i^B(0)$ amount to net exit by type- i buyers, since they do not generate sales or import goods. Buyers in this state nonetheless continue to search, and they appear as entrants when they successfully match.

Replacing \mathbb{S} with \mathbb{B} , B with S , \mathbf{s} with \mathbf{b} , and i with j in equations (11) and (12) gives the corresponding equations of motion for measures $M_j^S(\mathbf{b})$ of suppliers. Since the search behavior of a supplier only depends on the number of its partners, we can work with equations of motion which only depend on partner numbers n^S :

$$\begin{aligned} \dot{M}_j^S(n^S) &= [\sigma_j^S(n^S - 1) \theta^S M_j^S(n^S - 1) + (\delta + \delta^B)(n^S + 1) M_j^S(n^S + 1)] \\ &\quad - [\sigma_j^S(n^S) \theta^S M_j^S(n^S) + ((\delta + \delta^B)n^S + \delta^S) M_j^S(n^S)], \\ &\quad n^S = 1, 2, \dots, N^S; \quad j = 1, \dots, J \end{aligned} \quad (13)$$

Finally, the measure of supplier of type j with $n^S = 0$ suppliers evolve according to:

$$\dot{M}_j^S(0) = \delta^S \sum_{n^S \neq 0} M_j^S(n^S) + (\tilde{\delta} + \delta^B)M_j^S(1) - \sigma_j^S(0)\theta^S M_j^S(0) \quad (14)$$

We treat the measures of each intrinsic type as exogenous constants and impose the adding-up constraints:

$$M_i^B = \sum_{\mathbf{s} \in \mathbb{S}} M_i^B(\mathbf{s}) \quad (15)$$

$$M_j^S = \sum_{n^S} M_j^S(n^S), \quad (16)$$

Because the units of measurement for these objects are arbitrary, we impose $\sum M_i^B = 1$, and we treat the measure of potential suppliers per potential buyer, $M^S = \sum_j M_j^S$, as a parameter to be estimated. Note that at any point in time, some agents on each side of the wholesale market will not have any active matches. So *active* suppliers per *active* buyer, $\sum_j (M_j^S - M_j^S(\mathbf{0})) / \sum_i (M_i^B - M_i^B(\mathbf{0}))$, is an endogenous object, and it typically differs from the exogenous ratio of *potential* suppliers to *potential* buyers M^S/M^B .

To characterize the steady state of this system we set $\dot{M}_i^B(\mathbf{s}) = \dot{M}_j^S(n^S) = 0$ and solve the system of $I \cdot (\|\mathbb{S}\| + 1) + J \cdot (N^S + 1)$ equations, (11)-(16). Solving for transition dynamics is feasible but more involved. Appendix B.2 provides details.

3.3.2 Market Aggregates

It remains to develop expressions for the matching hazard per unit search measures, θ^B and θ^S , and the conditional matching probabilities, $v_i^B(\mathbf{s})$ and v_j^S . Since each type- i buyer in state \mathbf{s} searches with intensity $\sigma_i^B(\mathbf{s})$, the total search effort by these buyers is

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

and the aggregate search effort of *all* buyers is:

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

Analogously, each type- j supplier with n^S buyers searches with intensity $\sigma_j^S(n^S)$, so this group's total search effort is:

$$H_j^S(n^S) = \sigma_j^S(n^S)M_j^S(n^S),$$

and the overall search effort of *all* suppliers is:

$$H^S = \sum_{j=1}^J \sum_{n^S} H_j^S(n^S)$$

Together, H^B and H^S determine θ^B and θ^S . Following much of the labor search literature, we assume a matching function that is homogeneous of degree one in these two objects. Specifically, as in Petrongolo and Pissarides (2001), we write our measure of matches per unit time as ³²

$$X = f(H^S, H^B) = H^B \left[1 - \left(1 - \frac{1}{H^B}\right)^{H^S} \right] \approx H^B \left[1 - e^{-H^S/H^B} \right] \quad (17)$$

From the buyers' perspective, matching hazard per unit search is:

$$\theta^B = \frac{f(H^S, H^B)}{H^B} \quad (18)$$

The larger θ^B , the more matches take place per unit of buyer search. Likewise, matching hazard per unit search from the suppliers' perspective is:

$$\theta^S = \frac{f(H^S, H^B)}{H^S}. \quad (19)$$

$\frac{H^S}{H^B}$ is close to the conventional measure of market tightness $\frac{M^S}{M^B}$. Under our constant returns to scale (CRS) matching function, it maps 1 to 1 to the equilibrium hazards per unit search, θ^B and θ^S . Finally, 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 suppliers of type j is $v_j^S = \frac{\sum_{n^S} H_j^S(n^S)}{H^S}$.

3.3.3 Steady State Equilibrium

Given the mass of *potential* suppliers M_j^S for each type j , the mass of *potential* buyers M_i^B for each type i , and the total apparel import expenditure E , a steady state equilibrium consists

³²This matching function can be interpreted as randomly placing balls in urns where the aggregate search effort of the buyers play the role of urns and the suppliers the role of balls. $(1 - \frac{1}{H^B})^{H^S}$ is the probability of a specific urn not getting any ball, generating no match.

of: (i) retail prices p_j for each variety that buyers source from type j suppliers, (ii) bargaining transfers $\tau_{ji}(\mathbf{s})$ from any type i buyer with supplier portfolio \mathbf{s} to each of her type j suppliers, (iii) optimal search efforts $\sigma_i^B(\mathbf{s})$ and $\sigma_j^S(\mathbf{s})$, (iv) matching hazards per unit search θ^B and θ^S , and (v) distributions of buyers $M_i^B(\mathbf{s})$ and suppliers $M_j^S(\mathbf{s})$ over states $\mathbf{s} \in \mathbb{S}$ such that:

- (i) Consumers solve their purchase problem (equation 1), buyers charge optimal retail prices (equation 4), and the product market clears.
- (ii) Buyers and suppliers solve their respective value functions and optimal search problems (equations 8 and 9), and the negotiated transfers satisfy the bargaining condition (equation 10).
- (iii) The distribution of buyers and suppliers is stationary: $\dot{M}_i^B(\mathbf{s}) = 0$ and $\dot{M}_j^S(\mathbf{s}) = 0$ for all i, j , and \mathbf{s} (equations 11–14).
- (iv) Matching hazards θ^B and θ^S are consistent with the aggregation of individual search efforts through the matching functions (equations 18 and 19).

Appendix B summarizes the computational algorithm we use to solve for the steady state equilibrium and for transitional paths following an unexpected shock.

4 Fitting the model to data

We now describe our approach to estimating the structural parameters of our model.

4.1 Exogenously fixed parameters

We estimate most of the structural parameters in our model using the dataset summarized in Section 2.1 above, augmented by data on U.S. apparel buyers from the Census Bureau’s Annual Retail Trade Survey (ARTS). However, several parameters are not identified by these data, so we calibrate them using other sources. These include the within-store cross-product elasticity of substitution, α , the buyer death hazard, δ^B , and the supplier death hazard, δ^S . We set $\alpha = 4.35$ to match the elasticity of substitution across apparel varieties estimated by Hottman et al. (2016), we set $\delta^B = 0.07$ to match the exit rate among U.S. retail firms (Jarmin et al., 2009), and we set $\delta^S = 0.15$ to match the average exit rate among Chinese apparel producers during the period 2004–2006 (Zhu, 2014).

4.2 Model identification

To estimate the remaining parameters we proceed in two stages. First, since it is possible to identify the match separation hazard δ without solving our dynamic model, we do so in a preliminary step. Specifically, we estimate δ as the Poisson parameter that best fits the distribution of match death rates observed in the customs records, adjusted for the buyer death hazard and the supplier death hazard discussed above. This calculation yields $\delta = 0.774 - 0.07 - 0.15 = 0.554$.

Second, we estimate the remaining parameters jointly, exploiting the structure of our model for identification. These include the search cost scalar, κ_0 , the two visibility effect parameters, γ^B and γ^S , the cross-supplier elasticity of substitution, η , the measure of potential suppliers per potential buyer, M^S , and three parameters that characterize the agent type distributions. The first of these latter parameters, $\sigma_{\ln \mu}^2$, determines the distribution of buyer appeal under the assumption that $\ln \mu$ is (discretized) normal with mean 0. The other two parameters, Δ and ω , determine the distribution of quality-adjusted supplier costs, \tilde{c}_j , under the assumption that the log of this variable has zero mean and two mass points.³³ Specifically, Δ is the log difference between the two possible \tilde{c}_j realizations, and ω is the share of potential suppliers with the lower \tilde{c} value.

Our GMM estimator takes the standard form:

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

where $\Lambda = (k_0, \gamma^B, \gamma^S, M^S, \Delta, \omega, \sigma_{\ln \mu}^2, \eta)$, \bar{m} is a row vector of 209 targeted data moments, $m(\Lambda)$ is the corresponding model-based vector of moments, and W is a block-diagonal weighting matrix. The targeted vector \bar{m} includes the statistics that appear in Tables 1-4 and the ratio of variable costs to revenues among U.S. apparel retailers.³⁴ Accordingly, each of the first five blocks of W is the inverse of a covariance matrix for moments from a particular table

³³We limit the number of mass points to two in order to keep our state space manageable.

³⁴Some cells in Tables 1-4 are sparsely populated, and some cells can be inferred from others, so not all of the moments in these tables are used. To summarize the degree distributions, we target the cumulative 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 in Tables 1 and 2, including diagonals and up to 5 off-diagonal elements in each row. For the within-buyer supplier shares, we target all of the mean shares in columns 3-12 of table 4 except for the last non-zero element of each row. (The rows sum to one, so these elements are implied by the others.) For the relation between log imports per supplier and number of suppliers, we target all of the second column of Table 4 except the first row, which is one by construction.

in Section 2.3, and the last block is the inverted variance of the ratio of variable costs to total revenue among U.S. apparel retailers. Appendix C provides details.

Not surprisingly, the Andrews et al. (2017) sensitivity matrix tells us that the mapping from the moments in \bar{m} to the elements of Λ is many-to-many. However, certain moments are particularly helpful in identifying certain parameters. Consider first the elasticity of substitution across buyers, η . By equation (4), firms use the mark-up rule $q_{xy}p_{xy}\frac{\eta-1}{\eta} = q_{xy}c_x$. Summing over all matches and rearranging yields $\frac{\eta-1}{\eta} = \frac{C}{R}$, where $C = \sum_x \sum_y q_{xy}c_x$ is total variable costs incurred by all apparel buyers (exclusive of transfers to suppliers), and $R = \sum_x \sum_y q_{xy}p_{xy}$ is the total revenue they generate in the retail market.³⁵ Accordingly, $\frac{C}{R}$ essentially identifies η .

How do the other moments help with identification? As explained in Section 2.3, Tables 1 and 2 report transition probabilities for buyers per supplier and suppliers per buyer, respectively, while Table 3 summarizes the cross-buyer distribution of supplier counts (column 1) and the cross-supplier distribution of buyer counts (column 2). Because they are closely related, these tables help us to identify the same set of parameters.³⁶ Specifically, they are particularly useful for identifying κ_0 , γ^B , and γ^S , because these cost function parameters strongly influence the rates at which firms in different states add business connections. They also help identify the number of potential suppliers per potential buyer, M^S , because this affects the expected number of connections on each side of the market.

Table 4 differs from Tables 1-3 in that it describes revenues rather than firm counts. It includes two types of information. Column 2 reports payments per supplier for buyers with different numbers of suppliers. And the remaining columns report the average revenue shares of the primary supplier, the next most important supplier, etc., for buyers with two suppliers, three suppliers, and so on. Consider column 2 first. With random matching, the composition of supplier portfolios is similar across buyers.³⁷ So the distribution of payments per supplier

³⁵The costs of merchandise reported in the ARTS include total transfers to suppliers, τ_s . We therefore add these transfers to our simulated total variable cost measure in our moment condition, which becomes: $E\left[\frac{\eta-1}{\eta} + \frac{\tau_s}{E}\right] = \left(\frac{C_M + C_B}{R}\right)$, where R , C_M (buyers' merchandise purchases), and C_B (buyers' variable costs) are data.

³⁶In steady state, degree distributions can be inferred by compounding matrices of transition probabilities. We nonetheless include them as separate targets because they describe firms' partner counts rather than rates of transition between them.

³⁷Interestingly, despite random matching, the distribution of supplier types actually differs a bit between small and larger buyers. The reason is that buyers' search intensities depend upon the composition of their supplier portfolios. For example, buyers with a single supplier search harder for more suppliers when their current supplier happens to be low quality. This makes the incidence of high-quality suppliers relatively high among single supplier buyers.

across buyers of different sizes mainly reflects the distribution of buyer types, μ , and the column 2 moments help to identify $\sigma_{\ln \mu}^2$.

Finally, the remaining columns of Table 4 are helpful in identifying the supplier productivity difference Δ and the fraction of high-type suppliers ω . To understand why, suppose all suppliers were a single type. Then for any given buyer, the within-firm revenue share of each of its suppliers would be inversely proportional to the buyer’s supplier count. But with heterogeneous suppliers, different types serving the same firm receive different transfers, and the cross-supplier distribution of transfers reflects the cross-supplier distribution of quality-adjusted costs.

4.3 The estimates

Parameter values. Table 5 reports estimates of Λ . Note first that the buyer and supplier visibility parameters (γ^B and γ^S) are both positive and statistically significant, implying that agents on each side of the market find it easier to meet new potential business partners once they have established a market presence. However, visibility effects are stronger for buyers than for suppliers.³⁸ This partly reflects the structure of our model, which gives buyers, but not suppliers, diminishing returns to additional matches. Thus unlike buyers, suppliers do not need as much extra enticement to search from visibility-induced cost reductions. The estimates also reflect the fact that the suppliers-per-buyer distribution is skewed more to the right than the buyers-per-supplier distribution (Table 3). This means the model must explain the presence of some very large buyers, despite diminishing returns. Since large buyers must replace many expiring matches each period, they need relatively favorable search costs.

Turning to the distribution of potential buyers, the spread in buyer types, $\sigma_{\ln \mu}^2$ is large. Here too, this is helping to capture the “fat” tail of the supplier-per-buyer distribution in Table 3. It also reflects the fact that larger buyers tend to spend more per supplier, despite diminishing returns to adding suppliers (Table 4, column 1).

Next, the estimated ratio of potential suppliers to potential buyers is $M^S = 4.20$. Since there are 8.48 *active* suppliers per *active* buyer, this figure implies that searching with very low intensity is more common on the buyer side of the market. However, if one excludes buyers and suppliers with only one match, the ratio of active suppliers to active buyers falls

³⁸Consistent with this result, Eaton et al. (2025) found weak visibility effects for Colombian exporters using the same search cost function.

to 2.88. The prevalence of single match relationships indicates that participation in the U.S. market is incidental for many foreign apparel suppliers.

Finally, our ω estimate implies that only 3 percent of the potential exporter population is “high-quality.” And our estimate of Δ implies that these elite exporters are about 45 percent more appealing (per dollar spent) than their low-quality competitors. These figures are the model’s way of explaining the share in sales of the largest supplier, the next largest supplier, etc., at buyers of different sizes (Table 4). With only two types of suppliers the related parameters are heavily over-identified.³⁹

Table 5: Cost and Distributional Parameters*

	Estimate	Std. Error
k_0	0.009	0.003
γ^B	0.320	0.041
γ^S	0.230	0.046
ω	0.030	0.002
Δ	0.454	0.006
M^S	4.203	0.728
η	2.432	0.728
$\sigma_{\ln \mu}^2$	7.428	2.508
Objective function	10,461.73	

*GMM estimates of $\hat{\Lambda}$ based on equation 20. Construction of the weighting matrix is described in Appendix D. Aside from the retail sector expenditure share, the 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; then we re-scale the other elements of each row so that they sum to one. We also exclude elements describing jumps of more than 4 matches in either direction, which are very low probability events. For the degree distributions in Table 3, we use the measurements at 1, 2, 3, 4, 5, 10 and 15 partners. For the within-buyer supplier shares, we target all of the mean shares in columns 3-12 of Table 4 except for the last non-zero element of each row. (The rows sum to one, so these elements are implied by the others.) For the relation between log imports per supplier and number of suppliers, we target all of second column of Table 4 except the first row, which is zero by construction. The total number of targeted moments is 209.

Model fit. Figure 8 summarizes the model fit. The top two panels show that, overall, the model does a good job of explaining the 140 targeted elements of the transition matrices (Tables 1 and 2). Only one simulated moment is substantially off.⁴⁰ Hence it accurately

³⁹While adding supplier types would have improved the fit, it would have substantially increased the dimensionality of the computational problem.

⁴⁰Specifically, we substantially underestimate the fraction of suppliers with at least 10 buyers that will continue to have at least 10 buyers next year. This type of supplier is a very small fraction of all suppliers, so its transition probabilities don’t get weighted very heavily by our estimator.

captures the general tendency for firms to lose clients over time, as well as the tendency for buyers and suppliers that are new to the U.S. market to ramp up the number of their matches during their early years. However, the visibility effect in our search cost function is not sufficient to explain the search efforts of very large suppliers.

Our estimated transition matrices imply steady state distributions for buyers per supplier and suppliers per buyer. The middle panels of Figure 8 shows how well these distributions match up to their data-based counterparts in Table 3. Here too the model does quite well, with the small exception that the model slightly underestimates the frequency of buyers with large numbers of suppliers.

The lower left panel of Figure 8 summarizes the model’s ability to generate the pattern reported in the second column of Table 4. It shows that the model replicates the positive association between a buyer’s average log payments per supplier and its number of suppliers. However, it predicts a somewhat flatter relationship between these variables than we observe in the data. Most likely this is a consequence of the fact that computational constraints have forced us to only allow for two types of suppliers. Finally, the lower right panel of Figure 8 summarizes the model’s ability to generate the supplier share patterns reported in columns 3-12 of Table 4. Although we only allow for two intrinsic types of suppliers, the model does remarkably well in terms of its ability to predict suppliers’ shares in buyers’ total spending on merchandise at firms of different sizes.

There is one targeted data moment that does not appear in Figure 8: the ratio of variable costs to total revenues among apparel buyers. For completeness, we note here that the data-based value of this is 0.731, while the model-based simulated value is 0.730.

Alternative specifications. In Appendix D, we develop three alternative specifications to our baseline model and discuss their properties. First, we consider a mechanical model in which each buyer and each supplier has its own parametrically fixed search intensity. Although this version of the model has far more free parameters than our baseline specification, it fits the data only slightly better, and it fails to capture endogenous compositional effects on buyer and supplier portfolios in counterfactual experiments. Next we develop a version of the model in which match-specific fixed costs and match-level earning shocks induce endogenous separations, as in Mortensen and Pissarides (1994). Here we find that separations occur mainly among low-productivity firms which contribute little to market-wide aggregates. Finally, we add heterogeneous match death hazards to the model and show that this has little effect on

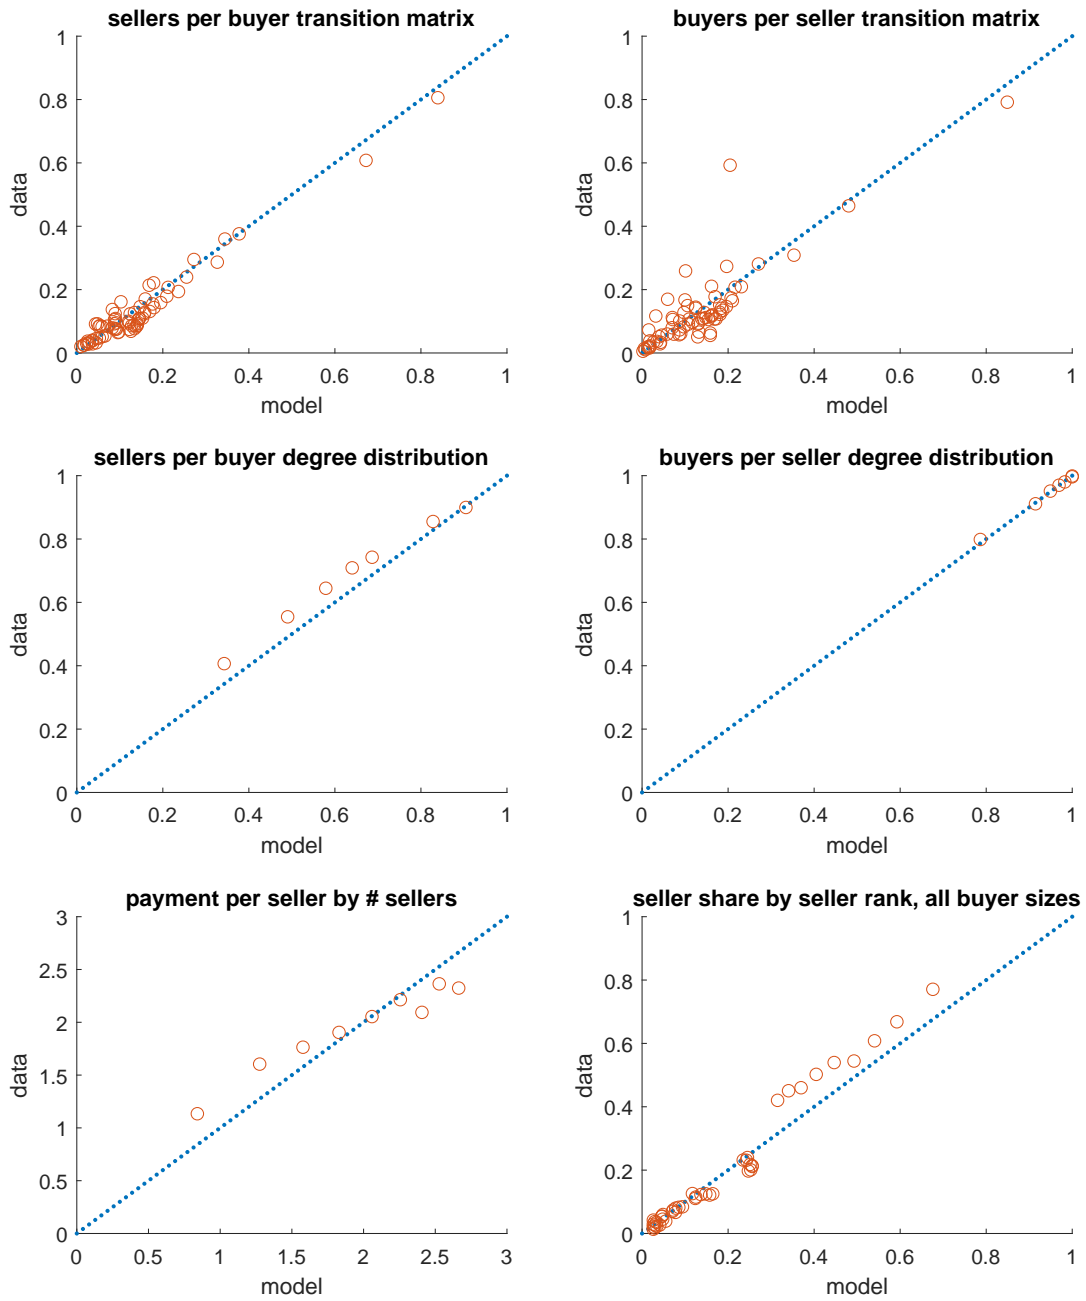


Figure 8: Data-based versus model-based moments

our estimates, mainly because there is only a weak relationship between buyer size and match death in the data.⁴¹

5 The baseline equilibrium

Using our estimated parameters, we now characterize the quantitative features of our baseline model. First, we summarize the profits (i.e., flow surplus) and assess the relative importance of search costs for various types of U.S. buyers and their foreign suppliers. We then examine the model’s implications for firm-level life-cycle dynamics.

5.1 Profits and search costs


Our model implies that suppliers and buyers together take a constant share $\frac{1}{\eta} = 0.411$ of final expenditures E as gross profits, using the rest to cover variable production and distribution costs (Section 3.2).⁴² But the division of these profits among individual buyers and suppliers is endogenously determined by their types and matching patterns. These patterns matter because an additional supplier not only contributes to a buyer’s retail profit, it also helps the buyer negotiate down transfers to her other suppliers. As a result, high- μ buyers—which match with more suppliers—are able to capture relatively large profit shares. And similarly, high-quality suppliers—which have a relatively large impact on sales—capture larger profit shares per dollar exported.

Figure 9a shows buyers’ average number of suppliers (dashed red line) and average gross profit share (solid blue line) as functions of the buyer appeal index, μ . Note that the least appealing buyers typically match with a single foreign supplier, rendering the negotiating effect irrelevant, and preventing them from capturing a larger profit share. In contrast, high- μ firms have many suppliers, none of whom is critical at the margin. This gives these buyers a strong negotiating position, allowing them to capture two-thirds of gross profits. Combined, these high- μ buyers account for nearly all buyer profits.

⁴¹ To highlight the effect of two-sided search on welfare, we also calibrate a standard sunk/fixed cost exporting model in the spirit of Alessandria et al. (2021). This comparison isolates the role of endogenous buyer–supplier matching and search congestion in shaping policy outcomes, demonstrating that in their absence, policy experiments can generate welfare responses that differ, even in sign, from our baseline search framework (see Appendix E for details).

⁴²To 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 was roughly \$100 billion in 2015.

We can similarly calculate the gross profits obtained by the two types of suppliers. Our parameter estimates imply that only 3 percent of *potential* suppliers are of high type. However, in the steady state equilibrium, the high type suppliers are disproportionately more likely to form at least one match with U.S. buyers. Consequently, approximately 12 percent of *active* suppliers are high type. Together, these high-quality suppliers obtain 0.11 of the total expenditure as profit, while low-quality suppliers receive only 0.03.⁴³



(a) Buyer Share of Gross Profit

(b) Buyer Search Cost as Share of Flow Profit

Figure 9: Buyer Profit and Search Cost Heterogeneity

We find similar heterogeneity in the search costs incurred by buyers. Our model implies that overall, they spend about 16 percent of their gross profits on efforts to match with suppliers. However, as Figure 9b shows, these costs—expressed as a share of firm-level profits—fall systematically with buyer appeal (solid blue line). For low- μ buyers, they represent around 45 percent of profits, but among the most appealing buyers, they account for less than 10 percent. There are two primary reasons for this. First, as mentioned above, low- μ buyers have relatively little bargaining power with their suppliers, so they have relatively low profit shares. Second, our estimate of γ^B implies that buyers with many suppliers enjoy significant visibility effects. In fact, to achieve any level of search intensity, σ^B , firms with 30 active matches spend 42 percent less than firms with a single match. Low appeal buyers have little incentive to exploit this scale effect, since the profits they earn from their matches are relatively modest.

On the supplier side, total search costs constitute 0.064 of total expenditure, representing 45 percent (0.064/0.14) of total supplier profit. Compared with the buyers, suppliers incur

⁴³In other words, high type suppliers earn close to 80 percent (0.11/0.14) of the aggregate supplier profit.

relatively higher search costs in proportion to their profit partly because our estimates indicate that the number of potential suppliers significantly exceeds the number of potential buyers.

5.2 Buyer and supplier connections over life cycle

The discussion thus far has focused on steady state patterns of profit sharing and search costs across heterogeneous buyers and suppliers. We now examine the model’s implications for the growth of buyer and supplier connections over their life cycles in the global apparel market.

We begin by analyzing the pattern observed in U.S. apparel trade data, as depicted in Figure 10. The blue dashed line illustrates that U.S. importers operating in the international market for 10 years accumulate, on average, nine more suppliers than those in their first year. It is encouraging that our model simulation, represented by the red solid line, effectively explains the observed increase. This increase results from the rapid accumulation of suppliers by high-type buyers in our model and their disproportionate representation in the older cohorts. In contrast, the average foreign supplier establishes a significantly lower number of U.S. busi-

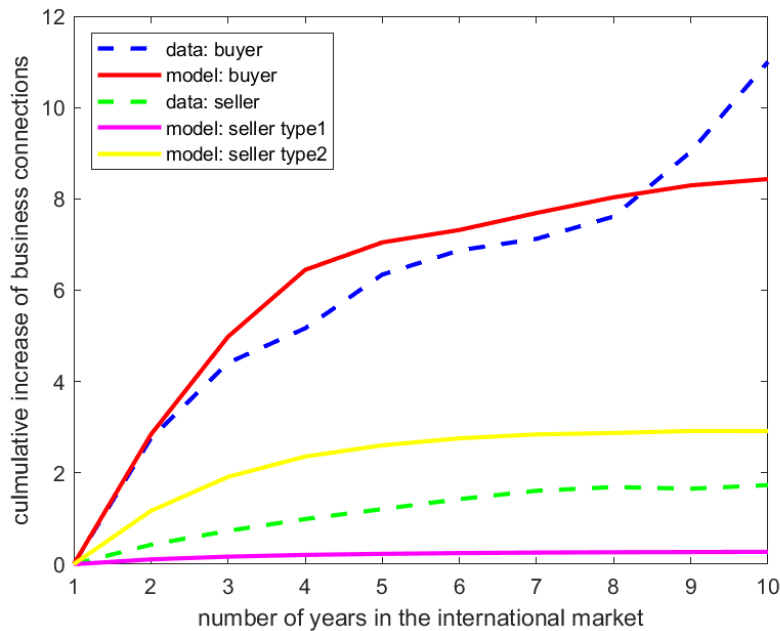


Figure 10: Buyers and Suppliers Accumulation of Connections

ness connections in both the model and the data. The green dashed line indicates that, on average, suppliers with 10 years of experience have only 1.5 more buyers than a new exporter. Our model generates qualitatively similar patterns: with the yellow solid line representing the

high-type supplier and the pink solid line representing the low-type supplier. However, since the low-type suppliers account for most of the suppliers in our model, we under-predict by about 50 percent the growth pattern in the data. In sum, our model can explain the rich dynamic life cycle pattern on both sides of the market without explicitly targeting these data moments in our estimation.

6 Experiments

We now use our model to conduct several experiments, focusing on two sets of questions. First, what combination of lower search costs (to proxy the spread of IT) and greater access to the U.S. apparel market (to proxy the ATC phaseout and Chinese reforms) allows the model to approximate the market developments we document in Section 2? And how did these shocks change the structure of the U.S. apparel market? Second, through the lens of our model, what were the effects of the 2018 Trump tariffs on the U.S. apparel market? Specifically, how did higher tariffs on low-quality apparel imports from China affect supplier-buyer matching patterns, profits, and consumer welfare? Also, how long did it take for the adjustment process to play out, and what did the transition path look like?

6.1 Interpreting market developments, 1996-2011

Our experiments concerning market developments are motivated by three patterns in Figures 1–4 (Section 2.2). First, the number of suppliers increased sharply by 28 percent in 2005. Second, exports per seller simultaneously dropped by about 15 percent. Third, despite essentially flat domestic apparel consumption, the total number of foreign suppliers serving the U.S. market rose steadily between 1996 and 2004, for a cumulative increase of about 25 percent.

The first two patterns reflected greater U.S. market participation by Chinese exporters, which are relatively small on average.⁴⁴ This sudden increase in Chinese exporters can be traced to the final phaseout of the ATC, and to China’s 2004 removal of regulations that had previously restricted small private firms from transacting directly with foreign buyers (Bai et al., 2017).⁴⁵ The third market development reflected a variety of factors, but for purposes of

⁴⁴Taken together, Figures 3 and 4 imply that over the period 1996–2011, average sales per exporter (in thousands of dollars) were: China, 84.2; India, 73.4; Indonesia, 103.1; Bangladesh, 105.1; Vietnam, 118.5; and Mexico, 172.2.

⁴⁵The ATC phaseout eliminated remaining quotas on U.S. imports from China and other low-income

our experiments we will interpret it to solely reflect falling search costs in the global apparel market. We will hereafter refer to the 2004-2005 reforms as "policy shocks" and to falling search costs as an "IT shock."

Because neither the number of active exporters nor sales per exporter shows a systematic trend after 2005, and because our model parameters are estimated using data from 2011, we interpret the baseline estimates reported in Table 5 as characterizing the 2006-2011 (post-policy shock, post-IT shock) equilibrium. And to isolate the effects of each type of shock, we compare this equilibrium with two others: the 2004 (post-IT shock, pre-policy shock) equilibrium, and the 1996 (pre-policy shock, pre-IT shock) equilibrium.

To construct the 2004 equilibrium, we choose the values of $(M^{S,2004}, \omega^{2004})$ that imply the observed growth in the number of active exporters and their average sales between 2004 and the baseline years (2006-2011), holding other parameters fixed. We find that the share of high-quality firms in potential exporters fell from $\omega^{2004} = 0.043$ to $\omega^{2006-2011} = 0.030$ as Chinese access to the U.S. market improved, while the measure of potential suppliers rose from $M^{S,2004} = 2.4$ to its post-shock estimated value of $M^{S,2006-2011} = 4.2$. To construct the 1996 equilibrium, we fix all parameters except κ_0 at their 2004 values. Then we increase κ_0 sufficiently to reduce the number of active exporters by the observed fraction, going backward from 2004 to 1996. We find that to explain the smaller number of active exporters in 1996 with search costs alone, we require $\kappa_0^{1996} = 1.65 \times \kappa_0^{2004}$.

Table 6 presents a comparison of the key steady state features in each equilibrium. For ease of interpretation, we express the magnitudes in columns (2) and (3) as percentage deviations from their corresponding values in the initial equilibrium (column 1). Consider first the pre-shock (1996) equilibrium in column (1), which features relatively few potential suppliers and relatively high search costs. The measures of active suppliers and buyers are expressed as ratios to our numeraire, namely, the population of potential buyers, M^B . Several features of this equilibrium are robust to the shocks we consider, and can be viewed as core implications of our estimated model. First, there are far more low-quality suppliers than high-quality suppliers, but the latter match with many more buyers, capture most of the supplier-side profits, and account for most of the supplier-side search effort. Second, there are many more active suppliers than active buyers. However, buyers collectively retain more profit than

countries. Although it was formally implemented in four stages beginning in 1995, it was heavily back-loaded: the final 50 percent of quota integration did not occur until January 1, 2005 (Kowalski and Molnar, 2009).

suppliers due to their monopsony power, as discussed in Section 5.1. Finally, a large fraction of the low- μ buyers have very low match counts and are unlikely to be connected to high-type suppliers.

Now consider the changes in key equilibrium outcomes when search costs fall (column 2). Consumer welfare improves significantly by 8 percent, driven by a 17 percent increase in the number of domestic retailers (buyers) and an 11 percent increase in the average number of suppliers per active buyer. (While not targeted, the proliferation of domestic retailers is consistent with the trend shown in Figure 6.) As a result of the relative increase in supply-side competition, domestic buyers were able to cut back on their search efforts and to extract a slightly higher share of industry profits.

Table 6: Counterfactual: Interpreting market developments

	Before policy and IT shocks: 1996 (low M^S , high κ_0)	Before policy shocks, after IT shock: 2004 (low M^S , baseline κ_0)	After policy and IT shocks: 2011 (baseline M^S and κ_0)
	(1)	(2)	(3)
1. measure, active low-quality suppliers	0.480	21.6%	57.2%
2. measure, active high-quality suppliers	0.087	4.2%	24.3%
3. measure, active buyers	0.190	16.7%	15.6%
4. total profit, low-quality suppliers	0.024	-0.3%	24.5%
5. total profit, high-quality suppliers	0.117	-1.4%	-7.8%
6. total profit, buyers	0.270	0.6%	1.2%
7. total search costs, low-quality suppliers	0.010	0.1%	24.2%
8. total search costs, high-quality suppliers	0.056	4.0%	-4.9%
9. total search costs, buyers	0.046	-14.8%	-4.8%
10. number of suppliers per buyer	1.695	11.3%	17.9%
11. high-quality suppliers per buyer	0.581	11.7%	-6.8%
12. consumer welfare	1.000	8.0%	7.1%

*Baseline figures reflect several normalizations. First, measures of active buyers and suppliers 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.

Finally, consider the simulated post-IT shock, post-policy shock equilibrium reported in Column 3. Here we keep κ_0 at its 2004 value, but to match the observed changes in both active exporters and sales per exporter thereafter, we increase the mass of potential suppliers (M^S) and we drop the share of high-quality firms among potential exporters (ω). Comparing column (3) with column (2), several notable patterns emerge. First, a congestion effect is evident. As the number of low-quality suppliers surged, high-quality suppliers found it increasingly difficult to match with buyers and therefore backed off their search efforts. So despite the substantial increase in the overall number of suppliers per buyer, the number of high-type suppliers per buyer actually decreased, falling below even the levels observed in the pre-IT shock,

pre-policy reform equilibrium (column 1).⁴⁶ Second, since high-type suppliers contributed significantly more value to buyers than low-types, their lower visibility weakened buyers’ search incentives and caused some of them with smaller markets (low- μ values) to exit. This reduced consumers’ shopping options and shifted market shares toward “big box” stores. Together, these adjustments actually reduced consumer welfare relative to the “pre-policy shock, post IT shock” level documented in column (2). These findings contrast with the predictions of the canonical model of export dynamics (Alessandria et al., 2021), which predicts a substantial welfare gain from the same experiment due to love-of-variety effects.⁴⁷ Appendix E provides details.

To summarize, the two major shocks we have considered both increased the number of suppliers actively serving the U.S. market. But they had different effects on consumer welfare and profits. Better search technologies dramatically increased consumer welfare by expanding the supply of all types of clothing and the number of retail outlets. But the removal of quota-based barriers to the U.S. market actually reduced welfare by triggering a proliferation of low-quality suppliers. This created market congestion, making it more difficult for buyers to find high-quality exporters, and thus reducing the average quality of their offerings.

6.2 Trump tariffs and trade disruptions

We next use our model to simulate the effects of Trump’s Section 301 tariffs on the U.S. apparel market in late 2019, both on impact and in the medium run.⁴⁸ Since Chinese apparel producers sold substantially lower value per exporter than producers in other countries, we will treat Chinese suppliers as low-quality firms. We will further simplify by treating the tariff bump as an unexpected permanent 15 percentage point shock.

Holding matching patterns fixed, some of the short-run effects of this policy shock are standard. By construction, our model implies that tariffs passed through to domestic U.S.

⁴⁶ If we had allowed buyers to choose a search intensity for each exporting country, and if the policy shocks in 2005 reduced the average quality of China’s potential exporters, they would have created an incentive for buyers to redirect their search toward other countries. This would have dampened the congestion effect. On the other hand, holding the Chinese quality mix fixed, an increase in the number of Chinese exporters would have made it easier to meet suppliers in that country, and encouraged buyers to search there. Figure 3 suggests this latter effect was dominant.

⁴⁷The Alessandria et al. (2021) model recognizes firm heterogeneity, idiosyncratic shocks, and sunk and fixed exporting costs. But it does not incorporate buyer heterogeneity or endogenous search on either side of the market.

⁴⁸For details of the tariff coverage and its impact on the apparel sector, see USITC (2023).

prices based on the purchase share from low-quality suppliers. Therefore, with total spending on apparel fixed, the Trump tariff hike moderately depresses aggregate joint profits (net of tariffs). However, individual agents are differentially impacted, depending upon their type and their portfolio of connections. Low-quality suppliers lose market shares relative to high-quality suppliers, and buyers with high initial concentrations of low-quality suppliers lose market share relative to others.⁴⁹

These initial impacts manifest exclusively in intensive margin adjustments to the profits and revenues of individual agents. But over time, firms on both sides of the market adjust their search intensities in response to the new payoff structure, matches die off, and the network structure evolves. Table 7 summarizes the changes in steady-state supplier composition, U.S. consumer welfare, and buyer profits following the tariff. As expected, imposing a tariff on low-quality suppliers substantially reduces their aggregate profits. This prompts a reduction in their search effort, leading to a 22 percent decline in the steady-state mass of active low-quality suppliers and a 7.6 percent reduction in the number of low-quality suppliers per buyer.

High-quality suppliers fare much better. Because they are not subject to tariffs, they benefit from both the reallocation of retail market shares and the reduction in wholesale market congestion. These developments cause them to intensify their search efforts, and lead to a 28.1 percent increase in the number of high-quality suppliers per buyer and an 18.3 percent increase in the total profit earned as a group.

The steady state welfare impact of the tariffs on consumers is quite limited. With search effort falling among low-quality suppliers and rising among high-quality suppliers, the portfolio of products offered by the typical buyer improves its appeal. This compositional effect on buyers' retail offerings largely offsets the reduction in apparel variety from the low-quality supplier exit and higher retail prices, as we will document shortly.

A more novel prediction of this paper concerns the speed of adjustment following the shock. Along the transition path, the share of total consumer spending on type j goods can

⁴⁹Consider a buyer with s_1 low-quality suppliers and s_2 high-quality suppliers. The share of this buyer's sales attributable to high-quality suppliers is $h_{2|i} = \frac{s_2 \tilde{c}_2^{1-\alpha}}{s_1 (\tilde{c}_1 (1+t_1))^{1-\alpha} + s_2 \tilde{c}_2^{1-\alpha}}$ and the buyer's price index is $p_i(\mathbf{s}) = \frac{\eta}{\eta-1} [s_1 ((1+t_1)\tilde{c}_1)^{1-\alpha} + s_2 \tilde{c}_2^{1-\alpha}]^{\frac{1}{1-\alpha}}$.

Table 7: Counterfactual: Trump Section 301 Tariff

	Long run responses to tariff
1. measure, active low-quality suppliers	-22.4%
2. measure, active high-quality suppliers	3.4%
3. measure, active buyers	1.6%
4. total profit, low-quality suppliers	-52.5%
5. total profit, high-quality suppliers	18.3%
6. total profit, buyers	-2.2%
7. number of suppliers per buyer	-7.6%
8. high-quality suppliers per buyer	28.1%
9. consumer welfare	-1.0%

be written as

$$R_j = \sum_{\mathbf{s}} s_j(\mathbf{s}) \underbrace{\left(\frac{\tilde{c}_j(1+t_j)}{\tilde{c}(\mathbf{s})} \right)^{1-\alpha}}_{\text{Intensive Margin within Buyer}} \underbrace{\frac{\sum_i M_{it}^B(\mathbf{s}) \mu_i^{\eta-1} \tilde{c}(\mathbf{s})^{1-\eta}}{\sum_{\mathbf{s}} \sum_i M_{it}^B(\mathbf{s}) \mu_i^{\eta-1} \tilde{c}(\mathbf{s})^{1-\eta}}}_{\text{Extensive Margin: Buyer's Supplier Composition}} \quad (21)$$

where $\tilde{c}(\mathbf{s}) = [\sum_k s_k(\mathbf{s}) ((1+t_k)\tilde{c}_k)^{1-\alpha}]^{\frac{1}{1-\alpha}}$ and $s_j(\mathbf{s})$ is the number of type- j suppliers connected to buyers in state \mathbf{s} . This decomposition makes clear that the fundamental source of dynamics in trade responses following an unexpected tariff shock is the evolution of U.S. buyers' supplier composition, $M_{it}^B(\mathbf{s})$. Importantly, this composition is endogenous, reflecting the interaction of buyer and supplier search and matching behavior. Figure 11a illustrates the transition path of high-quality and low-quality suppliers toward the new steady state. Since high-quality suppliers significantly increase their search effort immediately after the tariff is imposed, their mass rises rapidly, reaching the new steady-state level within approximately three years. In contrast, because buyer relationships with low-quality suppliers dissolve gradually, the transition takes about six years.

Are these trajectories broadly consistent with observed export responses to the Trump Section 301 tariffs in 2019? Using equation (21) and the transfer functions implied by equation (10), we can impute the revenue trajectories they imply for low-quality and high-quality suppliers.⁵⁰ Then, equating Chinese apparel suppliers with low-quality and equating the

⁵⁰ In our model, payments to type- j exporters have two components: compensation for variable production costs and surplus transfers. To construct the first component, note that the total quantity exported by type- j suppliers is $Q_j = (\frac{\eta}{\eta-1}c_j)^{-1}R_j$, so the total variable cost incurred by type- j suppliers is $c_jQ_j = (\frac{\eta-1}{\eta})R_j$. To construct the second component, express total transfers to type- j suppliers by buyers in state \mathbf{s} as $T_j(\mathbf{s}) \equiv \sum_i M_i^B(\mathbf{s})\tau_{ji}(\mathbf{s})$, where the $\tau_{ji}(\mathbf{s})$'s solve the system of equations (10). Aggregating $T_j(\mathbf{s})$ over buyer states and adding compensation for variable production costs, type- j exporters receive $IMP_j = \sum_{\mathbf{s}} s_j(\mathbf{s})T_j(\mathbf{s}) + (\frac{\eta-1}{\eta})R_j$.

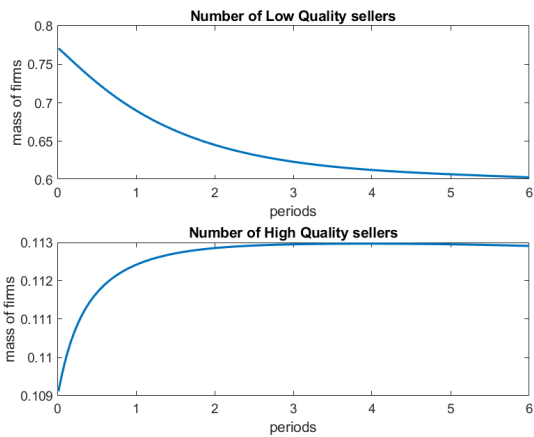
other major apparel suppliers with high-quality, we can compare the predicted evolution of China’s export revenues (normalized by predicted revenues of the other countries’ exporters) with its actual revenues (normalized by the actual revenues of other countries’ exporters).

To measure actual Chinese export revenues relative to others, we follow the empirical literature (e.g., Fajgelbaum et al. (2024)) and conduct an event-study analysis using major non-Chinese apparel exporters (Vietnam, Bangladesh, Indonesia, India, Mexico) as the control group. Specifically, we estimate U.S. imports of HS-10 product g from country j in year t using:

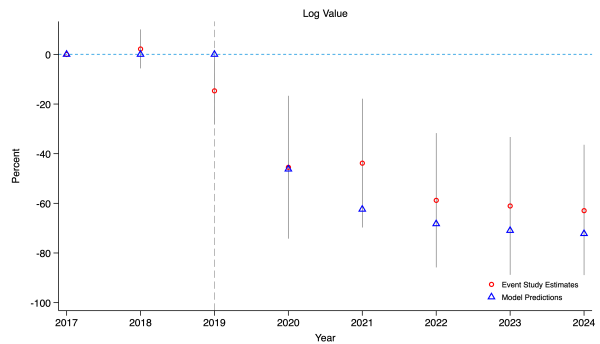
$$\ln y_{jgt} = \alpha_{jg} + \alpha_{gt} + \sum_{\tau=\underline{\tau}}^{\bar{\tau}} \beta_{\tau} \mathbf{1}[t - e = \tau] \times target_{jg} + \varepsilon_{jgt},$$

where e denotes the tariff implementation year (2019). The indicator $target_{jg}$ equals one if a given product is Chinese and subject to the Trump Section 301 apparel tariffs. The β_{τ} measure U.S. imports from China relative to imports from other countries, year by year, and net of the fixed effects.

Figure 11b reports these coefficient estimates (red circles) together with the model’s predicted evolution in the relative value of low-quality suppliers (blue triangles). Our model’s predictions align closely with the observed adjustment patterns. In particular, the sharp contraction in Chinese imports at the time of the tariff’s implementation (2020) is driven primarily by intensive-margin adjustment in the model, while the continued decline in subsequent years is fully accounted for by gradual extensive-margin reallocation in buyer–supplier relationships.



(a) Transition Dynamics of Suppliers



(b) Event Study: China vs. Other Exporters

Figure 11: Trump Tariff Effects on Supplier Dynamics

Given these responses, consumer welfare initially dropped 2.2 percent when the tariff was imposed. However, since high-quality suppliers swiftly gained market share, consumer welfare quickly rebounded and, in fact, slightly overshoot its steady state level. Thereafter, the gradual decline of low-quality suppliers brings the transition path toward a steady state welfare loss of about 1 percent. Overall, the limited loss of U.S. consumer welfare and the rapid responses of non-Chinese suppliers highlight that congestion in our two-sided search framework interacts with trade policy—particularly policies covering a *subset* of market participants.

7 Conclusion

We develop a dynamic model of international buyer-supplier matching in which agents on both sides of the market optimally choose how intensely to search. Fit to customs data on U.S. apparel imports, our framework captures key cross-sectional and time-series features of business-to-business relationships between foreign exporters and U.S. buyers. It also allows us to quantify search costs on each side of the market, and to simulate market responses to trade policy shocks and technological changes.

We find, first, that the aggregate costs of forming business relationships are borne almost equally by buyers and suppliers. But buyers' search costs are lower on a per-match basis, reflecting, *inter alia*, the fact that there are far more suppliers than buyers. Bargaining power and visibility effects also matter, making search costs per dollar of profit much smaller among buyers with many suppliers.

Second, buyers and suppliers adjust their search intensity over their life cycles, both because their market visibility changes and because buyers face diminishing returns to adding business partners. Different types of agents in the market mature differently. Conditional on survival, it takes about 10 years for a successful buyer to reach its long run size, while a supplier typically gets there in 5. Search frictions thus appear to constitute a major reason that aggregate trade flows react to shocks with long, unpredictable lags.

Third, reductions in search costs due to technological progress can induce substantial entry on both sides of the market, increasing consumer welfare while shifting rents away from buyers and suppliers. On the other hand, increasing the set of potential suppliers with access to the downstream market need not be welfare improving. In particular, our simulations suggest that the discrete jump in U.S. market access induced by the ATC phaseout in 2005 actually

reduced consumer welfare. The reason is that additional low-quality suppliers created market congestion, causing high-quality suppliers to reduce their search efforts. Although the number of available apparel varieties rose, this hurt consumers by reducing the quality of the offerings at the typical retailer (buyer), and by causing some smaller buyers to exit the market.

Finally, our simulations suggest that Trump’s 2018 tariffs on Chinese apparel imports reduced consumer welfare for the usual reason: they were passed through to consumers. But this effect was partly offset in the long run because the tariffs reduced search efforts among low-quality Chinese suppliers. In turn, this reduced market congestion and encouraged high-quality suppliers in other countries to increase their downstream market presence.

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,” *Annual Review of Economics*, 2021, 13, 253–280.
- Andrews, Isiah, Matthew Gentzkow, and Jesse Shapiro**, “Measuring the Sensitivity of Parameter Estimates to Estimation Moments,” *Quarterly Journal of Economics*, 2017, 132, 1553–1592.
- Antras, Pol and Davin Chor**, “Global Value Chains,” in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics, Volume 5*, 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.

- Bernard, Andrew, Adreas Moxnes, Emmanuel Dhyne, Glenn Magerman, and Kalina Manova**, “The Origins of Firm Heterogeneity: a Production Network Approach,” *Journal of Political Economy*, 2022, 130 (7), 1765–1804.
- 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.
- Binmore, K., A. Rubinstein, and A. Wolinsky**, “The Nash Bargaining Solution in Economic Modelling,” *Rand Journal of Economics*, 1986, 17, 176–188.
- Blum, Bernardo, Sebastian Claro, and Ignatius Horstmann**, “Facts and Figures on Intermediated Trade,” *American Economic Review*, 2010, 100 (2), 419–423.
- Boehm, Johannes, Ezra Oberfield, Ruairdh South, and Mazhar Waseem**, “The Network Origins of Firm Dynamics: Contracting Frictions and Dynamism with Long-Term Relationships,” Working Paper, Cornell University 2024.
- 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,” *Quarterly Journal of Economics*, 2023, 138 (4), 2391–2450.
- 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.
- de Fontenay, Catherine C. and Joshua S. Gans**, “Bilateral Bargaining with Externalities,” *The Journal of Industrial Economics*, 2014, 62 (4), 756–788.
- Dominguez, Juan Camilo, Jonathan Eaton, Marcela Eslava, and James Tybout**, “Search and Learning in Export Markets: Evidence from Interviews with Colombian Exporters,” *Review of International Economics*, 2023, 31 (3), 1093–1116.
- 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,” *Journal of International Economics*, 2025, 157.

- , – , **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,” *Econometrica*, forthcoming.
- Fajgelbaum, Pablo, Pinelopi Goldberg, Patrick Kennedy, Amit Khandelwal, and Daria Taglioni**, “The US-China Trade War and Global Reallocations,” *American Economic Review: Insights*, 2024, 6 (2), 295–312.
- Fitzgerald, Dorieann, Stephanie Haller, and Yaniv Yedid-Levi**, “How Exporters Grow,” *Review of Economic Studies*, 2024, 91 (4), 2276–2306.
- 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.
- 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.
- 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.
- Kowalski, Przemyslaw and Margit Molnar**, “Economic Impacts of the Phases-out in 2005 of Quantitative Restrictions Under the Agreement on Textiles and Clothing,” Trade policy Working Paper 90, OECD 2009.

- 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.
- Lentz, Rasmus**, “Optimal employment contracts with hidden search,” Technical Report, National Bureau of Economic Research 2014.
- 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.
- Monarch, Ryan**, “It’s Not You, It’s Me: Price, Quality, and Switching in U.S.-China Trade Relationships,” *Review of Economics and Statistics*, 2022, *104*, 909–928.
- Mortensen, Dale T. and Christopher Pissarides**, “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 1994, *61* (3), 397–415.
- 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), 23–63.
- 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.
- **and Vitor Trindade**, “Ethnic Chinese Networks in International Trade,” *The Review of Economics and Statistics*, 2002, *84* (1), 116–130.
- Rodrigue, Joel and Yong Tan**, “Price, Product Quality, and Exporter Dynamics: Evidence from China,” *International Economic Review*, 2019, *60* (4), 1911–1955.
- statista.com**, “U.S. online shop and mail-order sales of clothing and clothing accessories from 2003 to 2019,” 2022.
- 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*, 2023, *106* (5), 1544–1561.

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.

USITC, “Economic Impact of Section 232 and 301 Tariffs on U.S. Industries,” Technical Report 5405, United States International Trade Commission, Washington, DC 2023.

Verboven, Frank, “The nested logit model and representative consumer theory,” *Economics Letters*, 1996, 50 (1), 57–63.

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 Model Derivations

We provide more model details of the static profit, bargaining, and the transfer function in this section.

A.1 Demand, Pricing, and Total Profit

Following Atkeson and Burstein (2008) and Hottman et al. (2016), this section establishes that the optimal retail pricing rule is $p_{xy} = \frac{\eta}{\eta-1} c_x$, and the associated total flow surplus for buyer y and its suppliers is given by equation (5) in the main text.

Recall from Section 3.1 that consumer demand for good x from buyer y is

$$q_{xy} = (\xi_x)^{\alpha-1} (\mu_y)^{\eta-1} (p_{xy})^{-\alpha} (P_y)^{\alpha-\eta} (P)^{\eta-1} E, \quad (\text{A-1})$$

where

$$P_y = \left[\sum_{x \in J_y} \left(\frac{p_{xy}}{\xi_x} \right)^{1-\alpha} \right]^{1/(1-\alpha)} \quad \text{and} \quad P = \left[\int_y \left(\frac{P_y}{\mu_y} \right)^{(1-\eta)} \right]^{1/(1-\eta)}. \quad (\text{A-2})$$

Also, the total flow surplus generated by buyer y and its suppliers is

$$\pi_y^T = \sum_{x \in J_y} (p_{xy} - c_x) q_{xy}, \quad (\text{A-3})$$

and since negotiated prices are bilaterally efficient (see Section A.2 below), they satisfy the first-order conditions for profit maximization in the retail market:

$$q_{xy} + \sum_{x' \in J_y} \frac{\partial q_{x'y}}{\partial p_{xy}} (p_{x'y} - c_{x'}) = 0, \forall x \in J_y \quad (\text{A-4})$$

Dividing equation (A-4) by market-wide sales (E) and multiplying through by p_{xy} , we can

rewrite these first-order conditions as

$$h_{xy} + \frac{\partial q_{xy}}{\partial p_{xy}} \frac{p_{xy}}{q_{xy}} h_{xy} \left(\frac{p_{xy} - c_x}{p_{xy}} \right) + \sum_{x' \in J_y, x' \neq x} \frac{\partial q_{x'y}}{\partial p_{xy}} \frac{p_{xy}}{q_{x'y}} h_{x'y} \left(\frac{p_{x'y} - c_{x'}}{p_{x'y}} \right) = 0, \forall x \in J_y \quad (\text{A-5})$$

where $h_{xy} = \frac{p_{xy} q_{xy}}{E}$ is the share of buyer y 's product sourced from supplier x in market-wide sales. Further, recognizing that $\frac{\partial \ln P}{\partial \ln p_{xy}} = 0$ (because there is a continuum of buyers), equation (A-1) implies that elasticities appearing in equation (A-5) can be written as:

$$\frac{\partial \ln q_{xy}}{\partial \ln p_{xy}} = -\alpha + (\alpha - \eta) h_{x|y} \quad (\text{A-6})$$

$$\frac{\partial \ln q_{x'y}}{\partial \ln p_{xy}} = (\alpha - \eta) h_{x|y}, \forall x' \neq x \quad (\text{A-7})$$

where $h_{x|y} = \frac{\left(\frac{p_{xy}}{\xi_x}\right)^{1-\alpha}}{P_y^{1-\alpha}} = \frac{p_{xy} q_{xy}}{\sum_{x' \in J_y} p_{x'y} q_{x'y}}$ is the within-buyer- y revenue share of supplier x . Accordingly, we can restate equation (A-5) as:

$$h_{xy} + [-\alpha + (\alpha - \eta) h_{x|y}] h_{xy} \left(\frac{p_{xy} - c_x}{p_{xy}} \right) + \sum_{x' \in J_y, x' \neq x} [(\alpha - \eta) h_{x|y}] h_{x'y} \left(\frac{p_{x'y} - c_{x'}}{p_{x'y}} \right) = 0$$

Or, dividing through by h_{xy} and collecting terms:

$$1 - \alpha \left(\frac{p_{xy} - c_x}{p_{xy}} \right) + (\alpha - \eta) \sum_{x' \in J_y} \frac{h_{x'y}}{h_y} \left(\frac{p_{x'y} - c_{x'}}{p_{x'y}} \right) = 0$$

where $h_y = \frac{\left(\frac{P_y}{\mu_y}\right)^{1-\eta}}{P^{1-\eta}} = \frac{\sum_{x' \in J_y} p_{x'y} q_{x'y}}{E}$ is buyer y 's share in market-wide sales. Since the last term is the same for all $x \in J_y$, the product-specific Lerner index $l_y = \left(\frac{p_{xy} - c_{xy}}{p_{xy}} \right)$ is the same $\forall x \in J_y$. Further, since $\sum_{x' \in J_y} \frac{h_{x'y}}{h_y} = 1$, this equation reduces to $1 - \alpha l_y + (\alpha - \eta) l_y = 0$, or $l_y = \frac{1}{\eta}$. That is, buyers charge a constant markup $p_{xy} = \frac{\eta}{\eta-1} c_x$ for all of their products.

Finally, using this mark-up rule and equations (A-1)-(A-3), we obtain our expression for total flow surplus (equation 5):

$$\pi_y^T = \frac{1}{\eta} \sum_{x \in J_y} R_{xy} = \frac{1}{\eta} \frac{E}{P^{1-\eta}} \underbrace{\left[\sum_{x \in J_y} \left(\frac{\eta}{\eta-1} \right)^{1-\alpha} \tilde{c}_x^{1-\alpha} \right]^{\frac{1-\eta}{1-\alpha}}}_{\text{total sales}} \mu_y^{\eta-1}$$

where $\tilde{c}_x = \frac{c_x}{\xi_x}$

A.2 Bargaining and Transfer

This section describes the bargaining game that implies equation (10) in the text. We focus our discussion on the division of the surplus generated by any particular buyer and her set of suppliers, since there are no strategic interactions across buyers in our model, and suppliers view their relations with each buyer as independent of one another. (The latter follows from suppliers' constant marginal costs and cross-buyer product differentiation.)

We proceed by first showing that equations (3) and (10) describes the solution to a static game in which the flow surplus $\pi^T(\mathbf{s})$ is divided up between a buyer and her suppliers. To do so, we adopt the assumptions necessary to invoke Theorem 2a of de Fontenay and Gans (2014), hereafter dFG. Then we consider the dynamic bargaining problem in which buyers and their suppliers bargain over the present value of the surpluses that they jointly create, $V^T(\mathbf{s})$, as we assume in the text. Adding a modest restriction to the contract space, we show that equation (10) still applies at each point in time, even though membership in the buyer's connected set of suppliers evolves endogenously.

A.2.1 Notation

To summarize the static game, we adopt some of the notation of dFG. Let J be the set of suppliers connected to a particular buyer prior to the beginning of bargaining, and let $K \subseteq J$ be the set of agents in J that reach an agreement with this buyer. Together, these suppliers, indexed by $x \in K$, determine the composition of the buyer's vector of supplier types counts, \mathbf{s} . (Since it suffices to analyze a particular buyer in isolation, there is no need to carry along a buyer index.)

When the supplier set K is connected to the buyer, let $p_x(K)$ be the retail price of supplier x 's product delivered through the buyer to final consumers, and let $\tau_x(K)$ be the associated transfer from the buyer to supplier x . Also, collecting the set of prices for each of the suppliers' products in $\mathbf{p}(K) = \{p_x(K)\}_{x \in K}$, write the associated total flow surplus generated by the buyer and her portfolio of suppliers as $\pi^T(\mathbf{p}(K))$.⁵¹ Next, given these prices, write the flow

⁵¹ $\pi^T(\mathbf{p}(K))$ is given by equation (2) of the text evaluated at the price vector $\mathbf{p}(K)$ and the associated market-clearing quantities. Here, unlike in the text, there is no presumption that these prices maximize $\pi^T(\mathbf{p}(K))$.

payoff to the buyer, net of transfers to her set of suppliers (K), as:⁵²

$$\pi^B(\mathbf{p}(K)) = \pi^T(\mathbf{p}(K)) - \sum_{x \in K} \tau_x(K). \quad (\text{A-8})$$

Finally, prior to the bargaining sessions, let all agents share a common set of conjectures regarding the equilibrium set of suppliers to the buyer, $\hat{K} \subseteq J$, and their equilibrium actions and transfers:

$$[\hat{\mathbf{p}}(\hat{K}), \hat{\tau}(\hat{K})] = \left\{ \hat{p}_x(\hat{K}), \hat{\tau}_x(\hat{K}) \right\}_{x \in \hat{K}}.$$

A.2.2 The flow surplus bargaining equilibrium

To develop our characterization of bargaining over the current flow surplus, we invoke the following assumptions.

Assumption 1. No commitment: *A buyer and her suppliers cannot commit to long-term transfer schedules. They can renegotiate whenever the buyer’s set of connected suppliers changes.*

Assumption 2. Bilateral bargaining: *Buyers use delegates to bargain on their behalf—one for each connected supplier. And the bilateral bargaining sessions they conduct occur simultaneously, following the protocol described by Binmore et al. (1986). Specifically, for the session involving supplier x , one agent—the delegate or supplier x —is randomly selected to make an initial offer consisting of a retail price and a transfer for each possible network state, $\{p_x(K), \tau_x(K)\}_{K \subseteq J}$. The other agent can choose to accept or reject this offer. “Acceptance closes the bilateral bargaining between that pair. Rejection leads, with exogenous probability $1 - \sigma$, to a breakdown in negotiations between the pair” and the exclusion of x from K (dFG, p. 763). If the bargaining does not break down, the rejecting party proposes a new offer. The round of negotiations concludes when all suppliers in J have either reached an agreement with their delegated counterpart or have been excluded from her network.*

Assumption 3. Private information: *The specific offers and counter-offers made by each supplier and the delegate she bargains with are privately observed by the two parties involved. Bargaining breakdowns are, however, publicly observed.*

Assumption 4. Passive beliefs: *When a delegate or a supplier receives an offer of $p_x(K) \neq \hat{p}_x(K)$ or $\tau_x(K) \neq \hat{\tau}_x(K)$, or an unexpected rejection, she does not revise her beliefs regarding the bilateral negotiations between other delegates and suppliers.*

Together, these assumptions simplify the multilateral bargaining problem to a Nash-like problem in which each delegate-supplier pair maximizes its bilateral surplus, taking the actions

⁵²In the approximate notation of dFG, the flow payoffs in this equilibrium are $\Upsilon^B(K) \equiv \pi^B(K)$ for the buyer, and $\Upsilon_x^S(K) \equiv \tau_x(K)$ for supplier x .

of other agents as given. Further, by dFG Theorem 2a, as $\sigma \rightarrow 1$, these assumptions and the diagonal concavity of the flow surplus function $\pi^T(\mathbf{p}(K))$ imply there is a unique perfect Bayesian equilibrium in which $\hat{\mathbf{p}}(K)$ satisfies equation (3) of the main text, and the surplus splitting rule is implicitly defined by:

$$\tau_x(K) = \frac{1}{2} \left[\underbrace{\pi^T(\hat{\mathbf{p}}(K)) - \pi^T(\hat{\mathbf{p}}(K \setminus x))}_{x\text{'s contribution to total surplus}} - \underbrace{\sum_{k \in K \setminus x} (\tau_k(K) - \tau_k(K \setminus x))}_{\text{impact of } x \text{ on other supplier transfers}} \right], \quad x \in K. \quad (\text{A-9})$$

This is equation (10) in the main text expressed in the notation of Section A.2.1 above.

To demonstrate the role of Assumptions 1-4 in establishing equation (A-9) above, we now sketch the relevant derivations, relying on the logic of dFG. First, we explain why Binmore et al. (1986) bargaining (Assumption 2) and passive beliefs (Assumption 4) lead to bilateral efficiency. Then we show why, given bilateral efficiency, these same assumptions imply a "fair" division of the marginal surplus created when a buyer adds supplier x to her portfolio of suppliers.

Bilateral efficiency: The following Lemma establishes that the solution to our bargaining game is bilaterally efficient:

Lemma A.1. *Suppose agents hold passive beliefs and expect the equilibrium network state to be \hat{K} . Any perfect Bayesian equilibrium of the bargaining game over the flow surplus involves each $x \in \hat{K}$ choosing an action $\hat{p}_x(\hat{K})$ that is bilaterally efficient. That is, taking the outcomes of the other bargaining games as given, the equilibrium actions maximize the combined surplus of the buyer and her supplier, $x \in K$:*

$$\tilde{p}_x(K) = \arg \max_{p_x} \pi^B(p_x, \hat{\mathbf{p}}(K) \setminus \hat{p}_x(K)) + \tau_x(p_x, \hat{\mathbf{p}}(K) \setminus \hat{p}_x(K)) \quad (\text{A-10})$$

Proof. Consider the bilateral bargaining session involving supplier x . If the buyer's delegate is chosen to make the first (p_x, τ_x) offer, she will choose the smallest transfer acceptable to x ,

recognizing that with probability $1 - \sigma$, bargaining will break down before the next round:

$$\begin{aligned} & \max_{p_x, \tau_x} \pi^T(p_x, \hat{\mathbf{p}}(\hat{K}) \setminus \hat{p}_x(\hat{K})) - \tau_x - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) & (\text{A-11}) \\ \text{s.t. } & \tau_x \geq \sigma \hat{\tau}_x(\hat{K}) + (1 - \sigma) \cdot 0 \end{aligned}$$

Here, the delegate implicitly presumes that if supplier x were to reject its offer and make a counteroffer in the next round, it would make the equilibrium offer, $\hat{\tau}_x(\hat{K})$.

Similarly, if the supplier makes the first offer, he proposes the (p_x, τ_x) values that solve:

$$\begin{aligned} & \max_{p_x, \tau_x} \tau_x(p_x, \hat{\mathbf{p}}(\hat{K}) \setminus \hat{p}_x(\hat{K})) & (\text{A-12}) \\ \text{s.t. } & \pi^T(p_x, \hat{\mathbf{p}}(\hat{K}) \setminus \hat{p}_x(\hat{K})) - \tau_x - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) \geq \sigma \pi^B(\hat{K}) + (1 - \sigma) \pi^B(\hat{K} \setminus x) \end{aligned}$$

Here supplier x believes that if the buyer's delegate rejects an out-of-equilibrium offer, she will make the equilibrium counteroffer next period, earning $\pi^B(\hat{K})$. Also, the supplier believes that if her negotiations with the delegate are exogenously terminated, the buyer will receive the equilibrium payoff associated with the remaining network, $\pi^B(\hat{K} \setminus x)$.

The constraint for the delegate's problem (A-11) will bind, implying $\tau_x = \sigma \hat{\tau}_x(\hat{K})$. Using this result and equation (A-8) to restate (A-11) yields the buyer's optimal p_x offer:

$$\tilde{p}_x = \arg \max_{p_x} \pi^T(p_x, \hat{\mathbf{p}}(\hat{K}) \setminus \hat{p}_x(\hat{K})) - \sigma \hat{\tau}_x(\hat{K}) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) \quad (\text{A-13})$$

By Assumption 4, the delegate and supplier x take the outcomes of the other bilateral games as given. Further, the delegate assumes that supplier x will propose transfer $\hat{\tau}_x(\hat{K})$ if negotiations proceed to the next round. So when choosing p_x , the buyer's delegate ignores the last two terms in equation (A-13) and maximizes the total payoff to the buyer-supplier pair, given the outcomes of all other bilateral bargaining games. (By definition, this is bilateral efficiency.) Similarly, if supplier x were chosen to make the first offer, but a profitable deviation from x 's offer existed, the delegate would reject it and make the deviating offer. \square

Two final observations regarding retail price-setting merit mention. First, for any given

network of active agents, bilaterally efficient pricing implies that the set of equilibrium prices is unique for all $\eta > 1$. This is immediately apparent from fact that bilateral efficiency implies equation (A-4), which in turn implies the mark-up rule $p_x = \frac{\eta}{\eta-1}c_x$, as demonstrated in Section A.1 above. Second, since the marginal supplier always adds to total surplus in our model, i.e., $\pi^T(K) > \pi^T(K/x) \quad \forall K, x \in K$, each delegate-supplier pair wishes to avoid bargaining breakdowns and will reach agreement in the first round of its negotiations. It follows that all suppliers who are connected to the buyer will sell to her (i.e., $\hat{K} = J$), and since \hat{K} maps to the vector of connected buyer counts, \mathbf{s} , equation (3) of the text follows immediately from equation (A-13) above.

Fair surplus division: We next show that, when agents bargain at each point in time over the current flow surplus, the outcomes of the bilateral games described by Assumption 2 are “fair.” More precisely, presuming the probability of exogenous bargaining breakdowns $(1 - \sigma)$ is small, the buyer and each of her suppliers evenly split the flow surplus created by their coalition, net of payments to other suppliers. Accordingly, the equilibrium division of the total flow surplus is given by equation (10) of the text:

Lemma A.2. *As $\sigma \rightarrow 1$, there exists a perfect Bayesian equilibrium with equilibrium network state $\hat{K} \subseteq J$ such that equation (A-8) holds and agents’ payoffs satisfy $\pi^B(\hat{K}) - \pi^B(\hat{K} \setminus x) = \hat{\tau}_x(\hat{K}) \quad \forall x \in K$.*

Proof. Let the conjectured equilibrium $[\hat{\mathbf{p}}(K), \hat{\tau}(K)]$ be bilaterally efficient, with $\hat{\tau}_x(\hat{K})$ satisfying equation (A-9) evaluated at $\hat{K} \quad \forall x$. Also, for the bilateral game involving supplier x , let $\hat{R}_x^b(\hat{K})$ be the equilibrium payoff to x in network J when the buyer is chosen to make the first offer, and let $\hat{R}_x^x(\hat{K})$ be the equilibrium payoff when supplier x makes the first offer. Similarly, let $\hat{R}_b^x(\hat{K})$ be the equilibrium payoff to the buyer if supplier x is chosen to make the first offer, and let $\hat{R}_b^b(\hat{K})$ be the equilibrium payoff to the buyer if the buyer makes the first offer.

By Lemma A.1, the bilaterally efficient prices will be chosen in each bilateral game under our assumptions: $\mathbf{p}(\hat{K}) = \hat{\mathbf{p}}(\hat{K})$. But what will the equilibrium transfers be? Consider the game involving supplier x and suppose the buyer’s delegate makes the first offer. She will want to make supplier x ’s acceptance constraint bind, so she will choose a transfer that matches

the value to x of rejecting it:

$$\hat{R}_x^b(\hat{K}) = \sigma \hat{R}_x^x(\hat{K}) + (1 - \sigma) \hat{\tau}_x(\hat{K} \setminus x) \quad (\text{A-14})$$

(Although $\hat{\tau}_x(\hat{K} \setminus x) = 0$, we carry it along for clarity of exposition.) Likewise, if supplier x makes the first offer, he will choose a transfer just high enough to make the delegate indifferent between accepting and rejecting it:

$$\hat{R}_b^x(\hat{K}) = \sigma \hat{R}_b^b(\hat{K}) + (1 - \sigma) \pi^B(\hat{K} \setminus x) \quad (\text{A-15})$$

Finally, the two parties to the bargaining will exhaust the flow surplus created by their match, net of payments to other suppliers:

$$\hat{R}_b^b(\hat{K}) + \hat{R}_x^b(\hat{K}) = \pi^T(\hat{K}) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) \quad (\text{A-16})$$

$$\hat{R}_b^x(\hat{K}) + \hat{R}_x^x(\hat{K}) = \pi^T(\hat{K}) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) \quad (\text{A-17})$$

These relationships are sufficient to determine the equilibrium flow payoffs for the case when the buyer's delegate makes the first offer and the case where supplier x makes the first offer. Adding $\hat{R}_b^b(\hat{K})$ to both sides of (A-14) and $\hat{R}_x^x(\hat{K})$ to both sides of (A-15) yields:

$$\hat{R}_b^b(\hat{K}) + \hat{R}_x^b(\hat{K}) = \hat{R}_b^b(\hat{K}) + \sigma \hat{R}_x^x(\hat{K}) + (1 - \sigma) \hat{\tau}_x(\hat{K} \setminus x) \quad (\text{A-18})$$

$$\hat{R}_b^x(\hat{K}) + \hat{R}_x^x(\hat{K}) = \sigma \hat{R}_b^b(\hat{K}) + (1 - \sigma) \pi^B(\hat{K} \setminus x) + \hat{R}_x^x(\hat{K}) \quad (\text{A-19})$$

Since the right-hand sides of (A-18) and (A-19) are equal by (A-16) and (A-17), differencing them yields:

$$\begin{aligned} (\sigma - 1) \hat{R}_b^b(\hat{K}) + (1 - \sigma) \hat{R}_x^x(\hat{K}) &= (\sigma - 1) \left[\pi^B(\hat{K} \setminus x) - \hat{\tau}_x(\hat{K} \setminus x) \right] \\ \hat{R}_b^b(\hat{K}) - \hat{R}_x^x(\hat{K}) &= \left[\pi^B(\hat{K} \setminus x) - \hat{\tau}_x(\hat{K} \setminus x) \right] \end{aligned} \quad (\text{A-20})$$

Using this result to eliminate $\hat{R}_x^x(\hat{K})$ in (A-14) and exploiting (A-16), one obtains:

$$\begin{aligned}\hat{R}_x^b(\hat{K}) &= \sigma \left[\hat{R}_b^b(\hat{K}) - \left(\pi^B(\hat{K} \setminus x) - \hat{\tau}_x(\hat{K} \setminus x) \right) \right] + (1 - \sigma) \hat{\tau}_x(\hat{K} \setminus x) \\ &= \frac{\sigma}{1 + \sigma} \left(\pi^T(\hat{K}) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) \right) - \frac{\sigma}{1 + \sigma} \pi^B(\hat{K} \setminus x) + \frac{1}{1 + \sigma} \hat{\tau}_x(\hat{K} \setminus x) \quad (\text{A-21})\end{aligned}$$

Similarly, using (A-20) to eliminate $\hat{R}_b^b(\hat{K})$ in (A-15) and exploiting (A-17), one obtains

$$\begin{aligned}\hat{R}_b^x(\hat{K}) &= \sigma \left[\pi^B(\hat{K} \setminus x) - \hat{\tau}_x(\hat{K} \setminus x) + \hat{R}_x^x(\hat{K}) \right] + (1 - \sigma) \pi^B(\hat{K} \setminus x) \\ &= \frac{\sigma}{1 + \sigma} \left(\pi^T(\hat{K}) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) \right) + \frac{1}{(1 + \sigma)} \pi^B(\hat{K} \setminus x) - \frac{\sigma}{(1 + \sigma)} \hat{\tau}_x(\hat{K} \setminus x) \quad (\text{A-22})\end{aligned}$$

Substituting (A-21) into (A-16) yields an expression for $\hat{R}_b^b(\hat{K})$ in terms of payoffs:

$$\begin{aligned}\hat{R}_b^b(\hat{K}) &= \pi^T(\hat{K}) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) - \hat{R}_x^b(\hat{K}) \\ &= \frac{1}{1 + \sigma} \left(\pi^T(\hat{K}) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) \right) + \frac{\sigma}{1 + \sigma} \pi^B(\hat{K} \setminus x) - \frac{1}{1 + \sigma} \hat{\tau}_x(\hat{K} \setminus x) \quad (\text{A-23})\end{aligned}$$

Likewise, substituting (A-22) into (A-17) yields an expression for $\hat{R}_j^j(\hat{K})$ in terms of payoffs:

$$\begin{aligned}\hat{R}_x^x(\hat{K}) &= \pi^T(\hat{K}) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) - \hat{R}_b^x(\hat{K}) \\ &= \frac{1}{(1 + \sigma)} \left(\pi^T(\hat{K}) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) \right) - \frac{1}{(1 + \sigma)} \pi^B(\hat{K} \setminus x) + \frac{\sigma}{(1 + \sigma)} \hat{\tau}_x(\hat{K} \setminus x) \quad (\text{A-24})\end{aligned}$$

Comparing (A-21) with (A-24) and (A-22) with (A-23), one can confirm that as $\sigma \rightarrow 1$, $\hat{R}_x^b(\hat{K}) \rightarrow \hat{R}_x^x(\hat{K}) \rightarrow \hat{\tau}_x(\hat{K})$ and $\hat{R}_b^x(\hat{K}) \rightarrow \hat{R}_b^b(\hat{K}) \rightarrow \pi^B(\hat{K})$. Accordingly, as the probability of an exogenous breakdown in bilateral bargaining becomes small, it doesn't matter which agent makes the first offer. Further, each agent receives half of the surplus their match creates, net

of the changes in transfers to other suppliers it induces:

$$\hat{\tau}_x(\hat{K}) = \frac{1}{2} \left[\left(\pi^T(\hat{K}) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) \right) - \left(\pi^B(\hat{K} \setminus x) + \hat{\tau}_x(\hat{K} \setminus x) \right) \right] \quad (\text{A-25})$$

$$\pi^B(\hat{K}) - \pi^B(\hat{K} \setminus x) = \frac{1}{2} \left[\left(\pi^T(\hat{K}) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K}) \right) - \left(\pi^B(\hat{K} \setminus x) + \hat{\tau}_x(\hat{K} \setminus x) \right) \right] \quad (\text{A-26})$$

The fair allocation property therefore holds in the limit as $\sigma \rightarrow 1$.⁵³ □

As noted earlier, all suppliers connected to a buyer will reach an agreement with her, so the set $\hat{K} = J$ maps to the buyer's vector of supplier counts, \mathbf{s} . Using $\pi^B(\hat{K} \setminus x) = \pi^T(\hat{K} \setminus x) - \sum_{k \in \hat{K} \setminus x} \hat{\tau}_k(\hat{K} \setminus x)$ and $\hat{\tau}_x(\hat{K} \setminus x) = 0$, one can therefore confirm that equation (A-9) above, and its restatement as equation (10) of the text, are implied by equation (A-25).

A.2.3 The dynamic bargaining equilibrium

Given Assumptions 1-4, Lemma A.2 establishes that if suppliers and buyers bargain over the flow surplus $\pi^T(J)$ at each point in time, equation (A-9) provides a precise solution of the transfer. However, both suppliers and buyers in our model are forward-looking in their negotiations, so we need to establish that this static allocation rule applies in our dynamic setting. To do so, we need only distinguish suppliers' portfolios by their counts of supplier types, \mathbf{s} , so from this point forward we revert to the notation we use in the main text.

A particular dynamic concern arises from the endogeneity of search efforts. Recall that our total flow surplus function (6) exhibits diminishing returns with respect to the buyer's portfolio size, n^B . So adding a supplier reduces the marginal value of all existing matches to the buyer. This creates an incentive for suppliers to discourage buyers' search efforts through tenure-dependent contracts, as in Lentz (2014).⁵⁴ Allowing this type of contract would make our model impossibly complex, so we rule it out with the following assumption.

Assumption 5. *Limited contract space: suppliers cannot condition transfers on the buyer's search effort or portfolio history.*

⁵³dFG note that if each party is equally likely to move first, the fair allocation property also obtains in expectation for any $0 < \sigma < 1$.

⁵⁴We are grateful to a referee for pointing this out.

We can then show that under Assumptions 1-5, the system (A-9) characterizes the surplus splitting rule when agents negotiate over the expected present value of their match participation. Having established equation (10) of the text, we will hereafter revert to identifying the portfolio of suppliers using \mathbf{s} rather than K :

Lemma A.3. *Suppose a buyer in state \mathbf{s} bargains with her portfolio of suppliers over the total surplus they generate, $V^B(\mathbf{s}) + \sum_k s_k V_k^S(\mathbf{s})$. If Assumptions 1-5 hold, the stable bargaining outcome at each point in time is given by the solution to the system of equations (A-9), with $\pi^B(\mathbf{s}) = \pi^T(\mathbf{s}) - \sum_{j=1}^J s_j \tau_j(\mathbf{s})$.*

Proof. By Assumption 5, transfers can only depend upon agents' current states. Accordingly, by Assumptions 1-4 and the logic of Lemma A.2, the equilibrium surplus split must satisfy:

$$[V^B(\mathbf{s}) - V^B(\mathbf{s} - \mathbf{1}_j)] - V_j^S(\mathbf{s}) = 0 \quad \forall j.$$

Here the threat point of the buyer's negotiation, $V^B(\mathbf{s} - \mathbf{1}_j)$, is the equilibrium buyer value function when she renegotiates transfers with the remaining suppliers, choosing the corresponding optimal search intensity for state $\mathbf{s} - \mathbf{1}_j$. We now demonstrate that the flow surplus sharing rule (A-9) is implied by this expression.

To begin, difference the buyer's value function (8) with respect to the threat point, obtaining:

$$\begin{aligned} \tilde{\rho}(V^B(\mathbf{s}) - V^B(\mathbf{s} - \mathbf{1}_j)) &= [\pi^B(\mathbf{s}) - \pi^B(\mathbf{s} - \mathbf{1}_j)] - [k^B(\mathbf{s}) - k^B(\mathbf{s} - \mathbf{1}_j)] \\ + \sigma^B(\mathbf{s})\theta^B \sum_k \nu_k^S [V^B(\mathbf{s} + \mathbf{1}_k) - V^B(\mathbf{s})] &- \sigma^B(\mathbf{s} - \mathbf{1}_j)\theta^B \sum_k \nu_k^S [V^B(\mathbf{s} - \mathbf{1}_j + \mathbf{1}_k) - V^B(\mathbf{s} - \mathbf{1}_j)] \\ + \tilde{\delta} \sum_k s_k [V^B(\mathbf{s} - \mathbf{1}_k) - V^B(\mathbf{s})] &- \tilde{\delta} \sum_k (s_k - \mathbf{1}_{k=j}) [V^B(\mathbf{s} - \mathbf{1}_j - \mathbf{1}_k) - V^B(\mathbf{s} - \mathbf{1}_j)], \end{aligned} \tag{A-27}$$

where $\tilde{\rho} \equiv \rho + \delta^B$ and $\tilde{\delta} \equiv \delta + \delta^S$. Now simplify this equation in two steps. First, use a discrete approximation of the first order condition for the buyer's optimal search,

$$[k^B(\mathbf{s}) - k^B(\mathbf{s} - \mathbf{1}_j)] \approx (\sigma^B(\mathbf{s}) - \sigma^B(\mathbf{s} - \mathbf{1}_j)) \left(\theta^B \sum_k \nu_k^S [V^B(\mathbf{s} - \mathbf{1}_j + \mathbf{1}_k) - V^B(\mathbf{s} - \mathbf{1}_j)] \right),$$

to restate the relationship creation terms in (A-27):

$$\begin{aligned}
& -[k^B(\mathbf{s}) - k^B(\mathbf{s} - \mathbf{1}_j)] + \sigma^B(\mathbf{s}) \left(\theta^B \sum_k \nu_k^S [V^B(\mathbf{s} + \mathbf{1}_k) - V^B(\mathbf{s})] \right) \\
& -\sigma^B(\mathbf{s} - \mathbf{1}_j) \left(\theta^B \sum_k \nu_k^S [V^B(\mathbf{s} - \mathbf{1}_j + \mathbf{1}_k) - V^B(\mathbf{s} - \mathbf{1}_j)] \right) \\
& = \sigma^B(\mathbf{s}) \theta^B \left(\sum_k \nu_k^S [V^B(\mathbf{s} + \mathbf{1}_k) - V^B(\mathbf{s})] - \sum_k \nu_k^S [V^B(\mathbf{s} - \mathbf{1}_j + \mathbf{1}_k) - V^B(\mathbf{s} - \mathbf{1}_j)] \right)
\end{aligned}$$

Second, restate the relationship destruction terms in (A-27) as

$$\begin{aligned}
& \tilde{\delta} \sum_k s_k [V^B(\mathbf{s} - \mathbf{1}_k) - V^B(\mathbf{s})] - \tilde{\delta} \sum_k (s_k - \mathbf{1}_{k=j}) [V^B(\mathbf{s} - \mathbf{1}_j - \mathbf{1}_k) - V^B(\mathbf{s} - \mathbf{1}_j)] \\
& = -\tilde{\delta} [V^B(\mathbf{s}) - V^B(\mathbf{s} - \mathbf{1}_j)] \\
& \quad + \tilde{\delta} \sum_k (s_k - \mathbf{1}_{k=j}) \left([V^B(\mathbf{s} - \mathbf{1}_k) - V^B(\mathbf{s})] - [V^B(\mathbf{s} - \mathbf{1}_j - \mathbf{1}_k) - V^B(\mathbf{s} - \mathbf{1}_j)] \right)
\end{aligned}$$

Combining these expressions, rewrite (A-27) as:

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

Next, use the surplus sharing rule $(\tilde{\rho} + \tilde{\delta})[V^B(\mathbf{s}) - V^B(\mathbf{s} - \mathbf{1}_j)] - (\tilde{\rho} + \tilde{\delta})V_j^S(\mathbf{s}) = 0$ and the supplier value function (9) to restate the above expression as:

$$\begin{aligned}
& [\pi^B(\mathbf{s}) - \pi^B(\mathbf{s} - \mathbf{1}_j)] - \tau_j(\mathbf{s}) \\
& + \sigma^B(\mathbf{s}) \theta^B \sum_k \nu_k^S \left([V^B(\mathbf{s} + \mathbf{1}_k) - V^B(\mathbf{s})] - V_j^S(\mathbf{s} + \mathbf{1}_k) \right) \\
& - \sigma^B(\mathbf{s}) \theta^B \sum_k \nu_k^S \left([V^B(\mathbf{s} + \mathbf{1}_k - \mathbf{1}_j) - V^B(\mathbf{s} - \mathbf{1}_j)] - V_j^S(\mathbf{s}) \right) \\
& + \tilde{\delta} \sum_k (s_k - \mathbf{1}_{k=j}) \left([V^B(\mathbf{s} - \mathbf{1}_k) - V^B(\mathbf{s} - \mathbf{1}_k - \mathbf{1}_j)] - V_j^S(\mathbf{s} - \mathbf{1}_k) \right) \\
& - \tilde{\delta} \sum_k (s_k - \mathbf{1}_{k=j}) \left([V^B(\mathbf{s}) - V^B(\mathbf{s} - \mathbf{1}_j)] - V_j^S(\mathbf{s}) \right) = 0
\end{aligned}$$

Finally, use the surplus sharing rule to drop the destruction terms. Then re-arrange the remaining terms and exploit the surplus sharing rule once more to obtain:

$$\begin{aligned}
\tau_j(\mathbf{s}) &= \pi^B(\mathbf{s}) - \pi^B(\mathbf{s} - \mathbf{1}_j) + \sigma^B(\mathbf{s})\theta^B \sum_k \nu_k^S \left([V^B(\mathbf{s} + \mathbf{1}_k) - V^B(\mathbf{s} + \mathbf{1}_k - \mathbf{1}_j)] - V_j^S(\mathbf{s} + \mathbf{1}_k) \right) \\
&\quad - \sigma^B(\mathbf{s})\theta^B \sum_k \nu_k^S \left([V^B(\mathbf{s}) - V^B(\mathbf{s} - \mathbf{1}_j)] - V_j^S(\mathbf{s}) \right) \\
&= \pi^B(\mathbf{s}) - \pi^B(\mathbf{s} - \mathbf{1}_j)
\end{aligned}$$

Since $\pi^B(\mathbf{s}) = \pi^T(\mathbf{s}) - \sum_{j=1}^J s_j \tau_j(\mathbf{s})$, the above result implies the system of equations (A-9). □

B Computation of the Model

We provide more details on the computation of our model below.

B.1 Steady State

We iterate on the market slackness measures θ^B , θ^S , and P to find the fixed point that defines the steady state equilibrium. Conditional on these aggregate equilibrium objects, we start by computing the value functions for both buyers and suppliers.

Buyer Value Function and Intensity Matrix The value function of each type i buyer is defined on the state vectors (s_1, s_2) which indicates the number of type 1 and type 2 suppliers a buyer purchase from. We arrange the state vectors following the sequence $(s_1 = 0, s_2 = 0)$, $(s_1 = 1, s_2 = 0)$, ..., $(s_1 = N_1, s_2 = 0)$, ... $(s_1 = N_1, s_2 = N_2)$. To facilitate computation, we set the maximum number of suppliers to be $N_1 = 5$, $N_2 = 45$. Overall the dimension of the value function vector is $(N_1 + 1) \times (N_2 + 1)$.

We then start by constructing the continuous time intensity matrix Q^B . The intensity matrix has sparse structure, with row elements adding up to zero.

- The diagonal of Q^B is the total hazard of leaving current state (s_1, s_2) (row)
 - adding a supplier: $-\sigma^B(s_1, s_2)(\theta_1^B + \theta_2^B)$

- losing a supplier: $-\tilde{\delta}(s_1 + s_2)$
- exit the import market: $-\delta^B$
- The off-diagonal element gives the hazard of entering next state (column)
 - add a type 1 supplier: $\sigma^B(s_1, s_2)\theta_1^B$ to state $(s_1 + 1, s_2)$
 - add a type 2 supplier: $\sigma^B(s_1, s_2)\theta_2^B$ to state $(s_1, s_2 + 1)$
 - lose a type 1 supplier: $\tilde{\delta}s_1$ to state $(s_1 - 1, s_2)$
 - lose a type 2 supplier: $\tilde{\delta}s_2$ to state $(s_1, s_2 - 1)$
 - exit the import market: δ^B to state $(0, 0)$
- Boundary conditions
 - When $s_1 = N_1$, no element in $(s_1 + 1, s_2)$, the diagonal doesn't include $\sigma^B(s_1, s_2)\theta_1^B$. Similarly when $s_2 = N_2$, the diagonal doesn't include $\sigma^B(s_1, s_2)\theta_2^B$.
 - When $s_1 = 0$, no element in $(s_1 - 1, s_2)$, when $s_2 = 0$, no element in $(s_1, s_2 - 1)$, diagonal automatically adjusted.
- In practice, we compress the state (s_1, s_2) into a single index $I(s_1, s_2) = s_1 * (1 + N_1) + s_2$. We can then define the following states accordingly
 - $I(s_1 + 1, s_2) = I(s_1, s_2) + 1$ if $s_1 \neq N_1$
 - $I(s_1, s_2 + 1) = I(s_1, s_2) + (N_1 + 1)$ if $s_2 \neq N_2$
 - $I(s_1 - 1, s_2) = I(s_1, s_2) - 1$ if $s_1 \neq 0$
 - $I(s_1, s_2 - 1) = I(s_1, s_2) - (N_1 + 1)$ if $s_2 \neq 0$
 - Any elements in the vectors V^B and s^B can be indexed by $I(s_1, s_2)$. Meanwhile, elements in matrix Q^B can be indexed by $[I(s_1, s_2), I(s'_1, s'_2)]$ with $I(s'_1, s'_2)$ defined above.
- Given the intensity matrix Q^B , static payoff function $\pi^B(s_1, s_2; P)$, and market slackness θ^B , we can iteratively evaluate

$$\kappa'(\sigma^B(s_1, s_2)) = (s + 1)^{\gamma^B} \theta^B [\nu_1^S (V^B(s_1 + 1, s_2) - V^B(s_1, s_2)) + \nu_2^S (V^B(s_1, s_2 + 1) - V^B(s_1, s_2))] \quad \forall (s_1, s_2)$$

where the value function can be expressed in matrix form as

$$\vec{V}^B = [\rho \mathbf{I} - Q^B]^{-1} [\vec{\pi}^B - \kappa^B(\vec{\sigma})]$$

until the value function \vec{V}^B and intensity matrix Q^B converge.

Buyer State Distribution A particularly convenient feature of defining the intensity matrix Q^B is that we can now characterize the buyer steady state distribution, i.e. the measure of buyers in each state (s_1, s_2) .

$$\vec{M}^B = [\mathbf{I} + (Q^B)']^{-1} \mathbf{1}$$

The calculation above can be repeated for each type i of buyers, so we can obtain $\vec{\sigma}_i^B$, Q_i^B , and \vec{M}_i^B , $\forall i$.

Supplier Match-Specific Value Function We then compute the match-specific value of a supplier-buyer pair $V_{ji}(s_1, s_2)$. Note that, from the supplier perspective, they will need to additionally adjust for the hazard rate δ^S of her own exit. We outline the construction of continuous time intensity matrix T_{1i}^S for type 1 supplier's match below, T_{2i}^S can be constructed in a very similar fashion.

- The diagonal of T_{1i}^S is the total hazard of the buyer leaving current state (s_1, s_2) (row)
 - adding a supplier: $-\sigma_i^B(s_1, s_2)(\theta_1^B + \theta_2^B)$
 - losing a supplier: $-\tilde{\delta}(s_1 + s_2)$
 - buyer exit the import market: $-\delta^B$
- The off-diagonal element gives the hazard of entering next state (column)
 - add a type 1 supplier: $\sigma_i^B(s_1, s_2)\theta_1^B$ to state $(s_1 + 1, s_2)$
 - add a type 2 supplier: $\sigma_i^B(s_1, s_2)\theta_2^B$ to state $(s_1, s_2 + 1)$
 - lose a type 1 supplier, but not the supplier itself: $\tilde{\delta}(s_1 - 1)$ to state $(s_1 - 1, s_2)$
 - lose a type 2 supplier: $\tilde{\delta}s_2$ to state $(s_1, s_2 - 1)$
 - buyer exits the import market: δ^B to state $(0, 0)$
 - supplier exits the export market or relationship terminated: $\tilde{\delta}$ to state $(0, 0)$

- Unlike Q_i^B , state $(0, s_2), \forall s_2$ is absorbing since the match involving the type 1 supplier was dissolved. So we set the row vector representing these states in T_{1i}^S all equal to zero.
- Equipped with T_{ji}^S , we can compute the match-specific value function in matrix form

$$\vec{V}_{ji}^S = [\rho \mathbf{I} - T_{ji}^S]^{-1} \vec{\tau}_{ji}, j = 1, 2$$

Given random match and equal weights on the type distribution of buyers, we can calculate the expected value of a new match for supplier j

$$V_j^S = \sum_i (\vec{v}_i^B)' \vec{V}_{ji}$$

where

$$\vec{v}_i^B = \frac{(\vec{\sigma}_i^B \odot \vec{M}_i^B)}{\sum_i (\vec{\sigma}_i^B \odot \vec{M}_i^B)' \mathbf{1}}$$

We can then use the FOC of supplier search to find the optimal policy $\sigma_j^S(n)$, where n is the number of buyers the supplier is already matched with

$$\frac{\partial k^S(\sigma_j^S, n)}{\partial \sigma_j^S} = \theta^S V_j^S.$$

Supplier Intensity Matrix The intensity matrix of supplier is simpler, since the state variable is n (number of buyers), where $n = 0, 1, \dots, N$. For symmetry, we assume $N = 50$. For each type $j = 1, 2$ supplier, we compute the following procedures.

- The diagonal of Q_j^S is the total hazard of leaving current state n (row)
 - adding a buyer: $-\sigma_j^S(n)\theta_s$
 - losing a buyer: $-(\delta + \delta^B)n$
 - exit the export market: $-\delta^S$
- The off-diagonal element gives the hazard of entering next state (column)
 - add a buyer: $\sigma_j^S(n)\theta_s$
 - lose a buyer: $(\delta + \delta^B)n$
 - exit the export market: δ^S to state 0

- Boundary conditions
 - When $n = 0$, no chance of losing a buyer
 - When $n = N$, no chance of adding a buyer

Supplier State Distribution Similar to the buyer side, we can then use Q_j^S to compute the steady state supplier state distribution as

$$\vec{M}_j^S = [\mathbf{I} + (Q_j^S)']^{-1} \mathbf{1}$$

Finally, we can use $\vec{\sigma}^B$, $\vec{\sigma}^S$, \vec{M}^B , \vec{M}^S to update the market slackness θ^B and θ^S as well as the aggregate price index P .

Existence and Uniqueness Neither uniqueness nor existence is guaranteed for the dynamic equilibrium. However, at the candidate parameter vectors evaluated by our estimator of the baseline model (Table 5), we have found that solutions almost always exist and fall in a plausible range. Not surprisingly, if we make the visibility parameters (γ^B and γ^S) very large, the scale effects become strong enough that solutions no longer obtain.

B.2 Transition Dynamics

We now describe the procedures of solving the transition dynamics from an initial steady state to a new steady state with a permanent economic environment change (i.e., the expiration of ATC). The basic idea involves the backward recursion of the value function (formerly defined as the Hamilton-Jacobi-Equation) and the forward iteration of the mass of firms at each state (formerly defined as the Kolmogorov forward equation). Again for clarity of notation, we abstract from buyer type i in the description below.

- Take the value functions $\vec{V}^B(T)$ and $\vec{V}^S(T)$ and its associated policy functions $\vec{\sigma}^B(T)$ from the new terminal steady state. Similarly take industry state distribution $\vec{M}^B(0)$ and $\vec{M}^S(0)$ from the initial steady state.
- Now start with a guess of the sequence of $\{\theta^B(t), \theta^S(t), P(t)\}_{t=0}^T$. In practice, we choose small step of Δt to discretize the time interval.

- We can construct the associated intensity matrix $Q^B(T)$ and $T_j^S(T)$ with $\sigma^B(T)$, $\tilde{\delta}$ and $\theta^B(T)$. We can also construct the buyer flow payoff function $\bar{\pi}^B(T)$ with $P(T)$.
- Update buyer value function with *backward* recursion and finite difference

$$\begin{aligned}\rho \vec{V}^B(t - \Delta t) &= \bar{\pi}^B(t) - c(\vec{\sigma}^B(t)) + Q^B(t) \vec{V}^B(t - \Delta t) + \frac{1}{\Delta t} (\vec{V}^B(t) - \vec{V}^B(t - \Delta t)) \\ \vec{V}^B(t - \Delta t) &= [(\rho + \frac{1}{\Delta t}) * \mathbf{I} - Q^B(t)]^{-1} [\bar{\pi}_t^B - c(\vec{\sigma}^B(t)) + \frac{1}{\Delta t} \vec{V}^B(t)]\end{aligned}$$

Given $\vec{V}^B(t - \Delta t)$, we use the optimal search FOC to update $\vec{\sigma}^B(t - \Delta t)$ and the associated $Q^B(t - \Delta t)$ each step. Using the similar procedure outlined in the previous subsection on steady state, we can construct $T_j^S(t - \Delta t)$ with the updated $\vec{\sigma}^B(t - \Delta t)$ and use it to compute

$$\vec{V}_j^S(t - \Delta t) = [(\rho + \frac{1}{\Delta t}) * \mathbf{I} - T_j^S(t)]^{-1} [\vec{\tau}^S(t) + \frac{1}{\Delta t} \vec{V}_j^S(t)]$$

- Next start iterate *forward* from $\vec{M}^B(0)$ to $\{\vec{M}^B(t)\}, t > 0$. Again we use finite difference approximation

$$\vec{M}^B(t + \Delta t) = \vec{M}^B(t) + (Q^B(t))' \vec{M}^B(t) \Delta t$$

Combining $\vec{M}^B(t)$, $\vec{V}_j^S(t)$, and $\theta^S(t)$, we can solve for search intensity $\vec{\sigma}^S(t)$. Update implied supplier intensity matrix $Q^S(t)$ with $\vec{\sigma}^S(t)$, and iterate supplier distribution forward with

$$\vec{M}^S(t + \Delta t) = \vec{M}^S(t) + Q^S(t)' \vec{M}^S(t) \Delta t$$

- Use the $\vec{M}^B(t)$, $\vec{M}^S(t)$, $\vec{\sigma}^B(t)$, and $\vec{\sigma}^S(t)$ to update the guess of $\theta^B(t)$, $\theta^S(t)$, $P(t)$, $\forall t$. Repeat the above until the sequence of these aggregate equilibrium objects converges.
- Once these objects converge, we will also need to check the following
 - whether the implied $\vec{M}^B(T)$ and $\vec{M}^S(T)$ are close enough to the terminal new steady state
 - if not, increase T and restart everything

C The moment weighting matrix

Our GMM estimator (equation 20) is based on a vector of sample moments with six components: $\bar{m}' = [\bar{m}'_1, \bar{m}'_2, \bar{m}'_3, \bar{m}'_4, \bar{m}'_5, \bar{m}'_6]$. Here \bar{m}_1 contains the buyers per supplier (*BPS*) transition probabilities from Table 1, \bar{m}_2 contains the suppliers 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, \bar{m}'_5 contains the intra-firm supplier share statistics reported in Table 4, and \bar{m}_6 is the ratio of total variable costs to total revenue among U.S. apparel retailers. In this appendix we describe our construction of the associated block-diagonal weighting matrix, component by component:

$$W = \text{diag}[\text{cov}(\bar{m}_1), \text{cov}(\bar{m}_2), \text{cov}(\bar{m}_3), \text{cov}(\bar{m}_4), \text{cov}(\bar{m}_5), \text{var}(\bar{m}_6)]^{-1}$$

C.1 Transition matrices

Consider first \bar{m}_1 and \bar{m}_2 . Elements of these vectors are estimates of probabilities that suppliers (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:⁵⁵

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

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

To derive covariance matrices for these vectors, it is convenient to reshape them as matrices. For the *BPS* transition probabilities, we write:

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

⁵⁵Here i and j refer to possible values for buyer and supplier counts. They should not be confused with the i and j subscripts that appear in the text, which refer to buyer and supplier types.

where K is the (topcoded) maximum number of buyers attainable by an individual supplier. Replacing BPS superscripts with SPB superscripts gives our notation for the suppliers 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 supplier.

Next, restate 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 as⁵⁶

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

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

⁵⁶Suppose 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 the larger probability is always the one subtracted from 1.

recovered. Specifically, since we can write $\left(\hat{\Pi}_i^{BPS}\right)' = \mathbf{A} \cdot \left(\hat{\mathbf{F}}_i^{BPS}\right)'$ where

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

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

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

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

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

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

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

However, for two reasons, we exclude some elements of $\hat{\Pi}^{BPS}$ from \bar{m}_1 . First, each row of $\hat{\Pi}^{BPS}$ sums to one, so $\hat{\Pi}^{BPS}$ contains some redundant information and Ψ^{BPS} is singular.

Second, $\hat{\mathbf{\Pi}}^{BPS}$ contains many zeros in cells more than 2 positions from 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:

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

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.⁵⁷ 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:

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

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 $cov(\bar{m}_2)$.

⁵⁷Examples: $\tilde{\Psi}_2^{BPS} = [\pi_{1|2}^{BPS}, \pi_{2|2}^{BPS}, \pi_{3|2}^{BPS}, \pi_{4|2}^{BPS}]$ and $\tilde{\Psi}_{K-1}^{BPS} = [\pi_{K-3|K-1}^{BPS}, \pi_{K-2|K-1}^{BPS}, \pi_{K-1|K-1}^{BPS}, \pi_{K|K-1}^{BPS}]$.

C.2 Degree distributions

In addition to elements of the transition matrices, we target the degree distributions presented in Table 3.⁵⁸ These are collected in \bar{m}_3 and \bar{m}_4 , which correspond to the supplier and buyer degree distributions, respectively. As with the transition probabilities, our expressions apply equally to buyers and suppliers, so we limit our exposition to the buyers per supplier 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 suppliers 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:

$$cov(\hat{\mathbf{G}}^{BPS}) = \frac{1}{N_S} \begin{bmatrix} G_1^{BPS}(1 - G_1^{BPS}) & G_1^{BPS}(1 - G_2^{BPS}) & \dots & G_1^{BPS}(1 - G_k^{BPS}) \\ G_1^{BPS}(1 - G_2^{BPS}) & G_2^{BPS}(1 - G_2^{BPS}) & \dots & G_2^{BPS}(1 - G_k^{BPS}) \\ \vdots & \vdots & \ddots & \vdots \\ G_1^{BPS}(1 - G_k^{BPS}) & G_2^{BPS}(1 - G_k^{BPS}) & \dots & G_k^{BPS}(1 - G_k^{BPS}) \end{bmatrix} \quad (\text{A-28})$$

where the total number of observations on suppliers is N_S . Hence, $cov(\bar{m}_3) = cov(\hat{\mathbf{G}}^{BPS})$, $cov(\bar{m}_4) = cov(\hat{\mathbf{G}}^{SPB})$, and we can approximate both objects by replacing the cumulative probabilities that underlie them with their sample analogs.

C.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 supplier that we observe among buyers with two suppliers, three suppliers, and four suppliers 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 $cov(\bar{m}_5)$.

⁵⁸The 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.

D Alternative specifications

In this appendix, we describe and calibrate three alternative models and compare their performance with our baseline model. The models are a mechanical model that takes search intensities as exogenous, a model including match productivity shocks and endogenous match death, and a model with heterogeneous death hazards across low- and high-quality seller types. We describe each in turn below.

D.1 A Mechanical Model

D.1.1 Overview

The baseline model characterizes firms’ search efforts as optimal forward-looking behavior, given their portfolio of clients and market-wide conditions. Presuming that the model is well-specified, this allows us to quantify search costs, and to perform counterfactual experiments that are not subject to the “Lucas critique.” But one might reasonably ask whether our formulation captures observed matching patterns and transition dynamics better than a simple mechanical model, in which firms do not adjust their search intensities as their client portfolios and market conditions evolve.

To address this question, we develop such a mechanical search model and fit it to our data. It is similar to our baseline model in most dimensions, but it differs in two respects. First, and most importantly, each type of buyer and supplier is freely endowed with a fixed level of search intensity.⁵⁹ Second, since sales are irrelevant to search in this model, it does not characterize match-specific sales or transfers.

We treat each search intensity as a free parameter so, given that we allow 30 buyer types and two supplier types (as in the baseline model), the mechanical model involves 32 search intensities as well as the share of “high” type suppliers and the number of potential suppliers relative to potential buyers. The mechanical model thus has 26 more estimated parameters than the eight-parameter baseline model. Nonetheless, it fits the data only 0.4 percent better than the baseline model.⁶⁰

⁵⁹One could imagine an even richer parameter space, with one search intensity estimate for each type of firm in each state. Such a saturated model would trivially replicate the transition matrices and degree distributions we target without adding much insight.

⁶⁰Because the mechanical model doesn’t predict match sales, it is fit without those moments. For comparability, we calculate the fit of our benchmark model after removing the statistics in Table 4 from the baseline fit metric. We have not re-fit the baseline model using the restricted moment vector; presumably it would fit

D.1.2 Model structure

Let there be I intrinsic buyer types, each with its own intrinsic visibilities, $(\sigma_1^B, \sigma_2^B, \dots, \sigma_I^B)$. Analogously, let there be J intrinsic supplier types with visibility levels $(\sigma_1^S, \sigma_2^S, \dots, \sigma_J^S)$. Further let $(n_1^B, n_2^B, \dots, n_I^B)$ and $(n_1^S, n_2^S, \dots, n_J^S)$ be the associated exogenous measure of each type of potential buyer and supplier. Accordingly, the probability that a randomly chosen match is between a type- i buyer and a type- j supplier is $v_i^B v_j^S$ where

$$v_i^B = \frac{\sigma_i^B n_i^B}{\sum_{i'} \sigma_{i'}^B n_{i'}^B} \text{ and } v_j^S = \frac{\sigma_j^S n_j^S}{\sum_{j'} \sigma_{j'}^S n_{j'}^S}.$$

The arguments of the equilibrium measures, $\theta^B = \frac{m(H^S, H^B)}{H^B}$ and $\theta^S = \frac{m(H^S, H^B)}{H^S}$, are $H^B = \sum_{i'} \sigma_{i'}^B n_{i'}^B$ and $H^S = \sum_{j'} \sigma_{j'}^S n_{j'}^S$. In contrast to the baseline behavioral model, aggregate search effort and matching hazards are exogenous objects because agents' connections do not affect their search intensities.

Next, following the notation of the baseline model, let the current measures of buyers in a particular state \mathbf{s} be $[M_1^B(\mathbf{s}), M_2^B(\mathbf{s}), \dots, M_I^B(\mathbf{s})]$, where \mathbf{s} is a $J \times 1$ vector counting buyers' numbers of connections with each type of supplier. And let the measures of suppliers in a particular state \mathbf{b} , by type, be $[M_1^S(\mathbf{b}), M_2^S(\mathbf{b}), \dots, M_J^S(\mathbf{b})]$ where \mathbf{b} is an $I \times 1$ vector counting the number of connections with buyers of each type.

All relationships end with exogenous hazard δ , and buyers and suppliers die with exogenous hazards δ^B and δ^S . Letting $\tilde{\delta} = \delta + \delta^S$, the equation of motion for type- i buyers with \mathbf{s} suppliers is:

$$\begin{aligned} \dot{M}_i^B(\mathbf{s}) &= \sum_j \left[\sigma_i^B \theta^B v_j^S M_i^B(\mathbf{s} - \mathbf{1}_j) + \tilde{\delta}(s_j + 1) M_i^B(\mathbf{s} + \mathbf{1}_j) \right] \\ &\quad - \left[\sigma_i^B \theta^B + \tilde{\delta} n^B(\mathbf{s}) + \delta^B \right] M_i^B(\mathbf{s}). \end{aligned} \tag{A-29}$$

$\mathbf{s} \in \mathbb{S}; i = 1, \dots, I$

The set of type- i buyers in state \mathbf{s} 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 $\sigma_i^B \theta^B$. 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

the restricted moment vector even better if we had done so.

of exogenous attrition, which occurs with hazard $\tilde{\delta}(s_j + 1)$. By analogous logic, the group loses existing members that either add a supplier (with hazard $\sigma_i^B \theta^B$) or lose one (with hazard $\tilde{\delta} n^B(\mathbf{s})$). The group also loses an existing member if it is hit by the death shock δ^B .

The measure of buyers of type i with $s = 0$ suppliers evolves according to:

$$\dot{M}_i^B(\mathbf{0}) = \delta^B \sum_{\mathbf{s} \neq \mathbf{0}} M_i^B(\mathbf{s}) + \tilde{\delta} \sum_j M_i^B(\mathbf{1}_j) - \sigma_i^B \theta^B M_i^B(\mathbf{0}) \quad i = 1, \dots, I \quad (\text{A-30})$$

Replacing \mathbb{S} with \mathbb{B} , \mathbf{s} with \mathbf{b} , and i with j in equations (A-29) and (A-30) gives the corresponding equations of motion for measures $M_j^S(\mathbf{b})$ of suppliers.

D.1.3 Estimation

While we will estimate a large number of search intensity parameters, for comparability we use the number and distribution of types from the baseline model. We assume that buyer types are of uniform measure and sum to one: $\sum_{\mathbf{s} \in \mathbb{S}} M_i^B(\mathbf{s}) = 1/30$ for all i . As in the baseline, we assume that there are two supplier types, and estimate the relative measure of suppliers and the fraction of high-type suppliers.

Degree distributions Setting $\dot{M}_j^S(\mathbf{b}) = 0$ for all $\mathbf{b} \in \mathbb{S}$; $i = 1, \dots, I$ and $\dot{M}_i^B(\mathbf{s}) = 0$ for all $\mathbf{s} \in \mathbb{B}$; $j = 1, \dots, J$, we obtain a system of equations that can be solved for the steady state levels of $M_j^S(\mathbf{b})$ and $M_i^B(\mathbf{s})$, as in the behavioral version of the model. With this solution in hand, we can calculate the fraction of type- i buyers with c suppliers and the fraction of type- j suppliers with c clients as:

$$f_i^B(c) = \sum_{\mathbf{s} | \mathbf{s}\boldsymbol{\iota} = c} M_i^B(\mathbf{s}), \quad f_j^S(c) = \sum_{\mathbf{b} | \mathbf{b}\boldsymbol{\iota} = c} M_j^S(\mathbf{b})$$

where $\boldsymbol{\iota}$ is a $J \times 1$ (or $I \times 1$) vector of ones in the buyer (or supplier) expression. Aggregating across buyer types yields the measure of buyers and suppliers with c partners:

$$C^B(c) = \sum_i n_i^B f_i^B(c), \quad C^S(c) = \sum_j n_j^S f_j^S(c)$$

These expressions directly imply the degree distributions for each side of the market.

Transitions Regardless of a buyer's type and current state, the hazard with which it will lose one of its c clients is $(\delta + \delta^S)c$, and the hazard with which it will lose all of its clients is δ^B . Likewise, the hazard with which a type- i buyer gains a client is always $\sigma_i^B \theta^B$. Accordingly, pooling all buyer types, the hazards of jumping from c clients down to $c - 1$ or 0, or up to $c + 1$ clients are state-independent:

$$\begin{aligned}\lambda^B(c - 1|c) &= (\delta + \delta^S)c \text{ and} \\ \lambda^B(0|c) &= \delta^B \text{ and} \\ \lambda^B(c + 1|c) &= \sum_i \sigma_i^B \theta^B \left(\frac{f_i^B(c) n_i^B}{\sum_{i'} f_{i'}^B(c) n_{i'}^B} \right),\end{aligned}$$

And analogous expressions hold on the supplier side of the market:

$$\begin{aligned}\lambda^S(c - 1|c) &= (\delta + \delta^B)c \text{ and} \\ \lambda^S(0|c) &= \delta^S \text{ and} \\ \lambda^S(c + 1|c) &= \sum_j \sigma_j^S \theta^S \left(\frac{f_j^S(c) n_j^S}{\sum_{j'} f_{j'}^S(c) n_{j'}^S} \right),\end{aligned}$$

These hazards can be fed through an intensity matrix to calculate analogs to observed transition rates and degree distributions for buyers and suppliers, unconditional on their types. For each supplier type and each buyer type, the common death rates δ , δ^S , δ^B and one type-specific parameter (σ_i^B or σ_j^S) govern the transitions and thus also the steady state number of connections. These transitions also help identify the relative fraction of the two supplier types and the relative share of potential suppliers to buyers.

Results We estimate 30 buyer search intensity parameters σ_i^B , and two supplier search intensity hazards σ_j^S . Search intensity results are presented in Figure 12. Firm types display widely varying search behavior, with the most intensively searching buyer type searching more than one hundred times harder than the least intensively searching buyer type. Most buyer types search at a low intensity, while just over a third search at higher intensities, with two types searching at very high intensity. In addition, we estimate the exogenous share of the high-type supplier (the one that searches more intensely) to be to be 3.8%, and the exogenous relative measure of suppliers as 15.5.

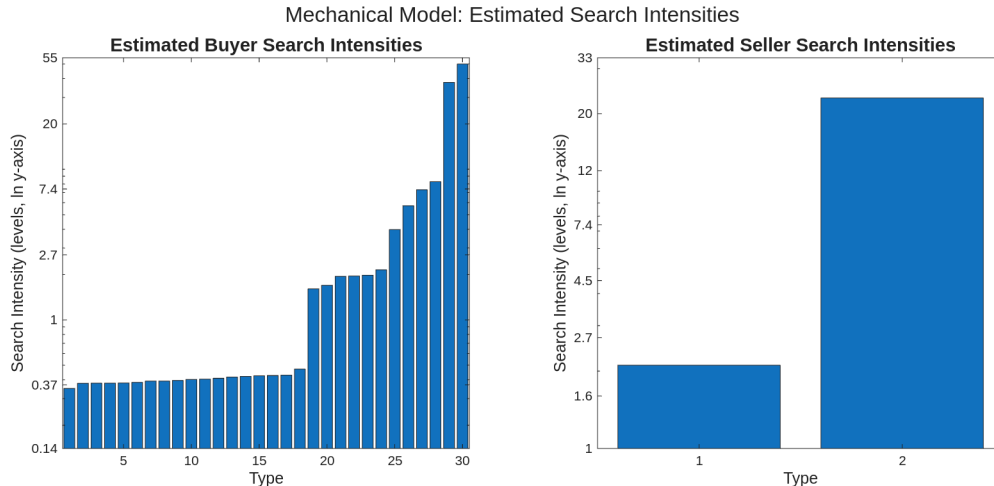


Figure 12: Mechanical model estimated search intensities

We cannot directly compare the fit metric of the mechanical model with that of the baseline. First, the mechanical model is fit on a smaller number of moments. Second, the models are not nested. The mechanical model has more parameters (34 vs 8), but it does not characterize buyer-supplier transfers, and it does not allow for changes in search intensity with changes in the client portfolio.⁶¹

The objective function of the fitted mechanical model is 6,721.8. The objective function of the baseline model, only calculated on the moments used to fit the mechanical model, is 6,746.7. The baseline model fit is only 0.4% higher than that of the mechanical model, even though it has far fewer parameters, and even though we have not re-estimated its parameters using the mechanical model’s truncated moment vector. The reason the baseline is able to perform so well is that firms endogenously change search intensity as they accumulate clients and this pattern is fully consistent with the mechanical model’s findings.

D.1.4 Mechanical model counterfactuals

We generate the effect of the 2004 easing of apparel imports from China in our mechanical model to compare with the baseline model. As in Section 7.1, the policy shock increases potential suppliers of both types, with a larger increase for low-quality suppliers.⁶² We summarize

⁶¹Allowing for changes in search intensities with client number would have saturated the moments we are fitting the models on.

⁶²Concretely, we first recover from the baseline transition calibration the type-specific 2004-to-2011 scaling factors for potential low- and high-quality suppliers. We then apply those same scaling factors to the mechanical model’s own estimated 2011 potential low- and high-quality supplier masses, and infer the implied 2004

the results of this experiment in Table 8.⁶³ Relative to the baseline model, the mechanical

Table 8: Effect of 2004 policy shock, baseline model and mechanical model

	Baseline % Change	Mechanical % Change
Measure, active low-quality suppliers	29.28	19.44
Measure, active high-quality suppliers	19.29	19.60
Measure, active buyers	-0.94	0.05
Number of suppliers per buyer	5.93	4.34
high-quality suppliers per buyer	-16.56	-12.35

model delivers much more entry of low-type suppliers, slightly more entry of high-type suppliers, and a small increase in buyers relative to a decrease in the baseline. Because buyers' search intensities are fixed in the mechanical model, active buyer mass is roughly unchanged, whereas it declines slightly in the baseline model when buyers optimally reduce search in response to a worse supplier mix. It is a similar story for low-type suppliers. In the mechanical model they do not adjust search behavior, while in the baseline model they search less intensely in response to increased competition for matches.

D.2 Match-specific shocks and fixed costs

D.2.1 Overview

One limitation of our benchmark formulation is that it does not allow for match-specific fixed costs and earnings shocks. Thus matches only die for exogenous reasons, and in stationary equilibria, match-specific revenues evolve solely because buyers' portfolios evolve.

It is possible to add these features to the model, but to keep it computationally tractable we need also to impose that the cross-buyer elasticity of substitution (η) is the same as the cross-product, within-buyer elasticity of substitution (α). This shuts down inter-dependencies between the marginal surpluses contributed by each supplier in the gross surplus function (5). We must also shut down visibility effects on both sides of the market; otherwise, each

values of M^S and ω . This preserves comparability across models without forcing the mechanical model to match the baseline's absolute 2004 supplier mass.

⁶³The "Baseline % Change" entries report policy-shock effects only. For each outcome x , we compute $100 \times \left[\frac{1 + \Delta x_{2011|1996}/100}{1 + \Delta x_{2004|1996}/100} - 1 \right]$, where $\Delta x_{2004|1996}$ and $\Delta x_{2011|1996}$ are the baseline model entries in Table 6, columns (2) and (3).

match-specific shock would be a state variable.⁶⁴

We describe this model in detail below. In order to identify the fixed costs of maintaining a match, we add a new target moment: the slope coefficient from a regression of match death probabilities on the log number of supplier matches. The model fits slightly worse than our benchmark model, primarily because with $\alpha = \eta$ it cannot explain why buyers' revenues exhibit diminishing returns to additional clients (Refer to equation 5). Taking the estimates at face value, they imply that idiosyncratic match earnings shocks are not the main cause of match deaths. The deaths these shocks induce primarily occur among low-type buyers, which contribute little to market-wide aggregates.⁶⁵

D.2.2 Model Outline

To keep this model tractable, we make two simplifying assumptions. First, we impose that the elasticity of substitution across retailers is the same as the within-retailer elasticity of substitution across products ($\eta = \alpha$). This makes each supplier's contribution to a buyer's gross surplus independent of the contributions of all other suppliers:

$$\pi_i^T(\mathbf{s}) = \frac{E}{\eta P^{1-\eta}} \left[\sum_{j=1}^J \left(\frac{\eta}{\eta-1} \right)^{1-\eta} s_j \tilde{c}_{ij}^{1-\eta} \right] \mu_i^{\eta-1} = \sum_j s_j \pi_{ij}^T, \quad (\text{A-31})$$

where $\pi_{ij}^T = \frac{E}{\eta P^{1-\eta}} \left(\frac{\eta}{\eta-1} \right)^{1-\eta} \tilde{c}_{ij}^{1-\eta} \mu_i^{\eta-1}$ is the match-specific surplus. Accordingly, buyers bargain with each of their suppliers without regard to the outcome of their bargaining with other suppliers.

Second, we shut down visibility effects on both sides of the market ($\gamma^B = \gamma^S = 0$). Without this assumption, agents would need to keep track of the partners they are matched with, and the current state of each of their matches, since these would affect their expectations regarding their future visibility. In consequence, the dimension of the state space would become unmanageably large.

⁶⁴ The resulting formulation resembles the Mortensen and Pissarides (1994) model, relabeling employers as buyers and workers as suppliers. But unlike that model, it allows for endogenous search efforts and it does not incorporate macro shocks.

⁶⁵ At the best fit estimates, the most productive buyer types, representing 97.5 percent of visibility, always accept low-productivity supplier matches. Our estimates thus imply that only a small share of matches may separate due to match earnings shocks. Eaton et al. (2025) also find that fixed match continuation costs are unimportant.

Match values To incorporate match-specific shocks and fixed costs, we follow Mortensen and Pissarides (1994). Let ε be a multiplicative shock to gross match profits, π_{ij}^T , with arrival hazard λ and, conditional on arrival, distribution function $G(\varepsilon)$. Further, assume the ε realizations are *i.i.d.* across matches and time. Finally, assume that suppliers must pay a flow fixed cost, F , to maintain a match.

With these additional model features, the flow value to a type- i buyer of a match with a type- j supplier in state ε becomes:

$$(\rho + \delta + \delta^B + \delta^S)V_{ji}^B(\varepsilon) = \pi_{ij}^T\varepsilon - \tau_{ji}(\varepsilon) + \lambda \int [\max\{V_{ji}^B(x), 0\} - V_{ji}^B(\varepsilon)] dG(x), \quad (\text{A-32})$$

where $\tau_{ji}(\varepsilon)$ is the portion of the flow match surplus the buyer transfers to the supplier. Similarly, the flow value of a match for a type- j supplier matched with a type- i buyer in state ε is:

$$(\rho + \delta + \delta^B + \delta^S)V_{ji}^S(\varepsilon) = \tau_{ji}(\varepsilon) - F + \lambda \int [\max\{V_{ji}^S(x), 0\} - V_{ji}^S(\varepsilon)] dG(x) \quad (\text{A-33})$$

Bargaining Since match surpluses are independent of each other, the bargaining game in this version of the model is simpler than in the baseline model. Match by match, the Nash bargaining solves:

$$\max_{\tau_{ij}(\varepsilon)} [V_{ji}^S(\varepsilon) - 0] [V_{ji}^B(\varepsilon) - 0] \quad (\text{A-34})$$

Thus the usual sharing rule obtains:

$$V_{ji}^S(\varepsilon) = V_{ji}^B(\varepsilon) \quad (\text{A-35})$$

Substituting equations (A-32) and (A-33) into (A-35), we obtain the optimal transfer:

$$\tau_{ji}(\varepsilon) = \frac{\pi_{ij}^T\varepsilon + F}{2} \quad (\text{A-36})$$

So the supplier is partly compensated for bearing the fixed costs of sustaining the match.

Shock cutoffs The total value of the surplus, $S_{ij}(\varepsilon) = V_{ji}^S(\varepsilon) + V_{ji}^B(\varepsilon)$, increases monotonically in ε . And since each party to the match receives a fraction of its total value, there

is a shock cutoff value, ε_{ij}^* , below which the buyer and the supplier terminate the match by mutual agreement. Using equations (A-36) and (A-32), one can express the total value of the match to the buyer at $\varepsilon = \varepsilon^*$ as:⁶⁶

$$0 = (\pi_{ij}^T \varepsilon_{ij}^* - F) + \lambda \int_{\varepsilon_{ij}^*}^{\infty} \frac{\pi_{ij}^T (x - \varepsilon_{ij}^*)}{\rho + \delta + \delta^B + \delta^S + \lambda} dG(x) \quad (\text{A-37})$$

Given $G(x)$, this equation implicitly determines ε^* . For example, if the shock is log-normally distributed with mean 0 and variance σ^2 , the cutoff value solves:

$$0 = (\pi_{ij}^T \varepsilon_{ij}^* - F) + \lambda \pi_{ij}^T \frac{(\mathbb{E}[\varepsilon | \varepsilon > \varepsilon_{ij}^*] - \varepsilon_{ij}^*) \left(1 - \Phi\left(\frac{\ln \varepsilon_{ij}^*}{\sigma}\right)\right)}{\rho + \delta + \delta^B + \delta^S + \lambda}, \quad (\text{A-38})$$

where $\Phi(\cdot)$ is the standard normal distribution function, and the conditional expectation of ε is:

$$\mathbb{E}[\varepsilon | \varepsilon > \varepsilon_{ij}^*] = e^{\frac{\sigma^2}{2}} \frac{\Phi\left(\sigma - \frac{\ln \varepsilon_{ij}^*}{\sigma}\right)}{1 - \Phi\left(\frac{\ln \varepsilon_{ij}^*}{\sigma}\right)}. \quad (\text{A-39})$$

Match surplus Each new match begins with a random draw from $G(\varepsilon)$, and prior to its formation, buyers and suppliers cannot foresee the associated ε realization. So their search intensities depend upon expected match values, $\mathbb{E}[S_{ij}(\varepsilon)]$. Given $G(\varepsilon)$ and ε_{ij}^* , one can calculate these objects as functions of the non-stochastic components of flow surpluses, π_{ij}^T .

To see this, note that the flow value equation,

$$(\rho + \delta + \delta^B + \delta^S) S_{ij}(\varepsilon) = \pi_{ij}^T \varepsilon - F + \lambda \int [\max\{S_{ij}(x), 0\} - S_{ij}(\varepsilon)] dG(x),$$

can be re-stated as

$$(\rho + \delta + \delta^B + \delta^S + \lambda) S_{ij}(\varepsilon) = \pi_{ij}^T \varepsilon - F + \lambda (1 - G(\varepsilon_{ij}^*)) \mathbb{E}[S_{ij}(\varepsilon) | \varepsilon > \varepsilon_{ij}^*]$$

Taking conditional expectations over $\varepsilon_{ij} > \varepsilon_{ij}^*$, defining $\phi_{ij} = G(\varepsilon_{ij}^*)$, and collecting terms

⁶⁶First evaluate equation (A-32) at $\varepsilon = \varepsilon_{ij}^*$, use $V_{ji}^B(\varepsilon_{ij}^*) = 0$, and substitute out $\tau_{ji}(\varepsilon^*)$ using equation (A-36). Then, subtract the resulting equation from equation (A-32) and solve for the buyer value function. Finally, weight this function by $dG(x)$ and integrate the result over $\varepsilon \geq \varepsilon^*$ to get the buyer's match continuation value. Using this continuation value, equation (A-37) obtains.

yields:

$$\mathbb{E} [S_{ij}(\varepsilon)|\varepsilon > \varepsilon_{ij}^*] = \frac{1}{(\rho + \delta + \delta^B + \delta^S + \lambda\phi_{ij})} (\pi_{ij}^T \mathbb{E}[\varepsilon|\varepsilon > \varepsilon_{ij}^*] - F),$$

$$\text{implying } \mathbb{E} [S_{ij}(\varepsilon)] = \frac{1-\phi_{ij}}{(\rho+\delta+\delta^B+\delta^S+\lambda\phi_{ij})} \cdot (\pi_{ij}^T \mathbb{E}[\varepsilon|\varepsilon > \varepsilon_{ij}^*] - F).$$

Optimal search With constant returns ($\alpha = \eta$) and without visibility effects ($\gamma^B = \gamma^S = 0$), buyers' and suppliers' search efforts do not depend upon their current set of connections. They simply equate the marginal cost of search to the expected payoff from forming a match, adjusted for matching hazard per unit search:

$$\frac{d\kappa^B(\sigma_i^B)}{d\sigma_i^B} = \theta^B \frac{1}{2} \sum_{j=1}^J v_j^S \mathbb{E} [S_{ij}(\varepsilon)] \quad (\text{A-40})$$

$$\frac{d\kappa^S(\sigma_j^S)}{d\sigma_j^S} = \theta^S \frac{1}{2} \sum_{i=1}^I v_i^B \mathbb{E} [S_{ij}(\varepsilon)] \quad (\text{A-41})$$

These first-order conditions imply policy functions for σ_j^S and σ_i^B that vary only with exogenous buyer and supplier types (\tilde{c}_i, μ_i) and matching hazards (θ^S, θ^B) . Hence this version of the model is closely related to the mechanical model discussed in Appendix D.1 above. And since the solution algorithm is essentially the same, we do not repeat it here.

D.2.3 Market equilibrium

This subsection characterizes the market's steady state equilibrium. We make the same normalization as the baseline model, in which $\sum M_i^B = 1$, and the measure of potential suppliers per potential buyer, $M^S = \sum_j M_j^S$, is a parameter to be estimated. Then, the only missing pieces are the market aggregates and market slackness. With the difference that buyer search is independent of its state, the overall visibility of buyers is the same as in the baseline model:

$$H^B = \sum_{i=1}^I \sum_{\mathbf{s} \in \mathbb{S}} H_i^B(\mathbf{s}) = \sum_{i=1}^I \sum_{\mathbf{s} \in \mathbb{S}} \sigma_i^B M_i^B(\mathbf{s}) \quad (\text{A-42})$$

Then, given the normalization of the buyer measure, compute the overall visibility of suppliers as:

$$H^S = \sum_{j=1}^J \sigma_j^S v_j^S N^S \quad (\text{A-43})$$

Finally, the measure of matches per unit time and the market slackness are given by the same expressions as in the baseline model.

D.2.4 Computation of the Model

This subsection describes how to compute the transition/intensity matrices. For the buyers we can still characterize an intensity matrix. For the suppliers we compute the “expected” intensity matrix given that supplier transitions now depend on the type(s) of buyers they are matched with. Therefore, the relevant intensity matrix for suppliers is one in which the entries compute the average hazard when suppliers have a certain number of partners.

Buyer intensity matrix Adjusting the buyer intensity matrix to account for the endogenous match destruction is straightforward. We now need to keep track of the possibility of matches being destroyed from the match-specific shock. Also, note that the optimal buyer search is independent of the state and only depends on buyer type. Define the buyer intensity matrix, Q^B , with the same dimensions as in the baseline model.

- The diagonal of Q^B is the total hazard of leaving current state (s_1, s_2) (row)
 - adding a supplier: $-\sigma^B(\theta_1^B + \theta_2^B)$
 - losing a supplier: $-\tilde{\delta}(s_1 + s_2) - \lambda\phi_{i1}s_1 - \lambda\phi_{i2}s_2$
 - exit the import market: $-\delta^B$
- The off-diagonal element gives the hazard of entering the next state (column)
 - add a type 1 supplier: $\sigma^B\theta_1^B$
 - add a type 2 supplier: $\sigma^B\theta_2^B$
 - lose a type 1 supplier: $\tilde{\delta}(s_1 + s_2 + 1) + \lambda\phi_{i1}(s_1 + 1)$
 - lose a type 2 supplier: $\tilde{\delta}(s_1 + s_2 + 1) + \lambda\phi_{i2}(s_2 + 1)$
 - exit the import market: δ^B to state $(0, 0)$
- Boundary conditions
 - When $s_1 = N_1$, no element in $(s_1 + 1, s_2)$, the diagonal doesn’t include $\sigma^B\theta_1^B$. Similarly when $s_2 = N_2$, the diagonal doesn’t include $\sigma^B\theta_2^B$.

- When $s_1 = 0$, no element in $(s_1 - 1, s_2)$, when $s_2 = 0$, no element in $(s_1, s_2 - 1)$, diagonal automatically adjusted.

Supplier Intensity Matrix Unlike the baseline model, supplier transitions depend on the type(s) of the buyer they are matched with. Using only the buyer’s intensity matrix, we can get the mass of steady state matches between any two types i and j . By extension, we can compute the unconditional probability of any match involving supplier type j of being with a buyer type i , which we call $P_j(i)$. This is not generically the same as the probability of meeting a particular buyer type ν_{ij} , because different types of relationship now end with different hazards. Consider a typical row m in the intensity matrix. The cell $(m, m + 1)$ is simply $\sigma_j \theta_S$, since search behavior is independent of state. The cell $(m, m - 1)$ is given by the expression:

$$m(\delta + \delta_B) + \sum_{k=1}^m \phi_{ij(k)}$$

The first term is the exogenous match death hazard, and the second is the endogenous death hazard. We would like to compute the aggregate transition matrix for suppliers, so we would like to know the average hazard for suppliers with m partners. Since all matching is random, the expectation of the last expression with respect to the unconditional probability of buyer types is:

$$m(\delta + \delta_B + \mathbb{E}_{P_j}[\phi_{ij}])$$

The rest of the intensity matrix is standard. Once the aggregate intensity matrix is complete, we can derive both the transition matrix for buyers per supplier and the steady state distribution of buyers per supplier and other moments.

D.2.5 Model performance

This subsection discusses the performance of the match shock model. Parameters are fit to all moments from the baseline model, as well as one additional moment. In order to identify fixed costs, we add the slope coefficient from a regression of the probability of match death on the log number of buyers of the supplier associated with the match.⁶⁷ Since we shut down visibility effects, we estimate two fewer parameters. In order to recover a bit of flexibility, we allow search costs to differ between buyers and suppliers. We also add two parameters related

⁶⁷The data regression also includes the log number of suppliers of the buyer associated with the match.

to match shocks and match destruction: the standard deviation of shocks, which are assumed to be distributed log-normal with mean zero, and a fixed cost normalized as a fraction of the price index. We follow Eaton et al. (2025) and set the arrival rate of match shocks to four.⁶⁸ We thus have nine parameters to estimate, one more than in the baseline.

Table 9: Parameter Estimates and Fit for Baseline, Match Shock, and No Dispersion Models

Parameter	Model		
	Baseline	Match Shock	No Dispersion
Buyer search cost scalar	0.009 (0.003)	0.0090 (0.0014)	0.0071 (0.0014)
Supplier search cost scalar	same as buyer	0.0020 (0.2254)	0.0019 (0.1833)
Buyer visibility parameter	0.320 (0.041)	— —	— —
Supplier visibility parameter	0.230 (0.046)	— —	— —
Share of high-type suppliers	0.030 (0.002)	0.0341 (0.0003)	0.0341 (0.0003)
High-type supplier cost advantage	0.454 (0.006)	0.0559 (0.0008)	0.0551 (0.0008)
Buyer type dispersion	7.428 (2.508)	6.5224 (0.2542)	5.9207 (0.2372)
Supplier to buyer ratio	4.203 (0.728)	6.7579 (0.2114)	8.0639 (0.3999)
Within-store elasticity	calibrated	1.8771 (0.0034)	1.8687 (0.0035)
Cross-store elasticity	2.432 (0.728)	same as within —	same as within —
Match shock scale	—	0.0414 (0.5724)	0.0000 —
Fixed cost (shr. of median match surp.)	—	0.0015 (0.0006)	0.0018 (0.0005)
Match death moment (data: -0.0119)	—	-0.003	0.000
Objective function	11,346.79	11,522.2	11,892.5
Objective function (baseline moments)	10,461.73	10,983.1	11,007.4

Table 9 presents the estimated parameters and fit compared with the baseline. Figure 13 presents the detailed fit of the match shock model against the data. There are a number

⁶⁸Since firms are risk neutral, we cannot separately identify the match shock arrival and scale without adding moments from the evolution of match sales.

of differences in the parameter estimates. In the match shock model, supplier search cost is much lower than buyer search cost, cross- and within-store elasticities are estimated much lower, and there are more potential suppliers per buyer. The model is using these parameters in an attempt to generate the large buyers which the baseline model does via the visibility effects. The fixed costs are not large in absolute value. They are only 1.5 percent of median match surplus. They do, however, induce selection. The lowest 46.7 percent of buyer types are always rejected by low-type suppliers, and the next 6.7 percent are only accepted if the match shock is high enough. The 46.7 percent of buyer types with the highest market appeal accept low-type suppliers regardless of their state. As these top buyer types generate 97.5% of visibility, and the buyers subject to endogenous separation generate 0.9% of visibility, the share of observed matches which are subject to endogenous separation is below 1%.

The overall fit of the match shock model is slightly worse compared with the baseline model (11,522.2 vs 11,346.79).⁶⁹ In the lower left panel of Figure 13, the match shock model understates the increase in payments per supplier as buyers grow large. This is due to the lack of decreasing returns in the match shock model. As expected, the match death probability regression coefficient is negative due to endogenous match death. The magnitude is smaller in absolute value than what we see in the data.

Because our match shock model requires constant returns to scale on both sides of the market, it does not nest our baseline model. In order to gauge the extent to which endogenous match destruction affects model performance, we instead shut down match shocks by setting the match shock dispersion to (only slightly above) zero. The results of this “no dispersion” experiment are reported in column three of Table 9. One effect of shutting down match shocks is that the regression of match death probability on number of supplier partners becomes zero. This confirms our intuition, as the only margin which generates differential match death across supplier types is that low-type suppliers sometimes endogenously drop low-state buyers. It also means that the no dispersion model has no hope of fitting the match death regression coefficient. The overall fit of the no dispersion model is unsurprisingly worse than the match shock model, since it is a restriction on a parameter. There are only a couple of substantive differences in the parameter estimates between the match shock and no dispersion model. The no dispersion model has higher fixed costs, more suppliers per buyer, a lower buyer

⁶⁹For comparison, we present the fit of the baseline model with the extra moment we added in estimating the match shock model. Adding this model does not change the baseline best fit parameters. As death hazards are exogenous and there is no endogenous match separation, the baseline will deliver a regression coefficient of zero for all sets of parameters.

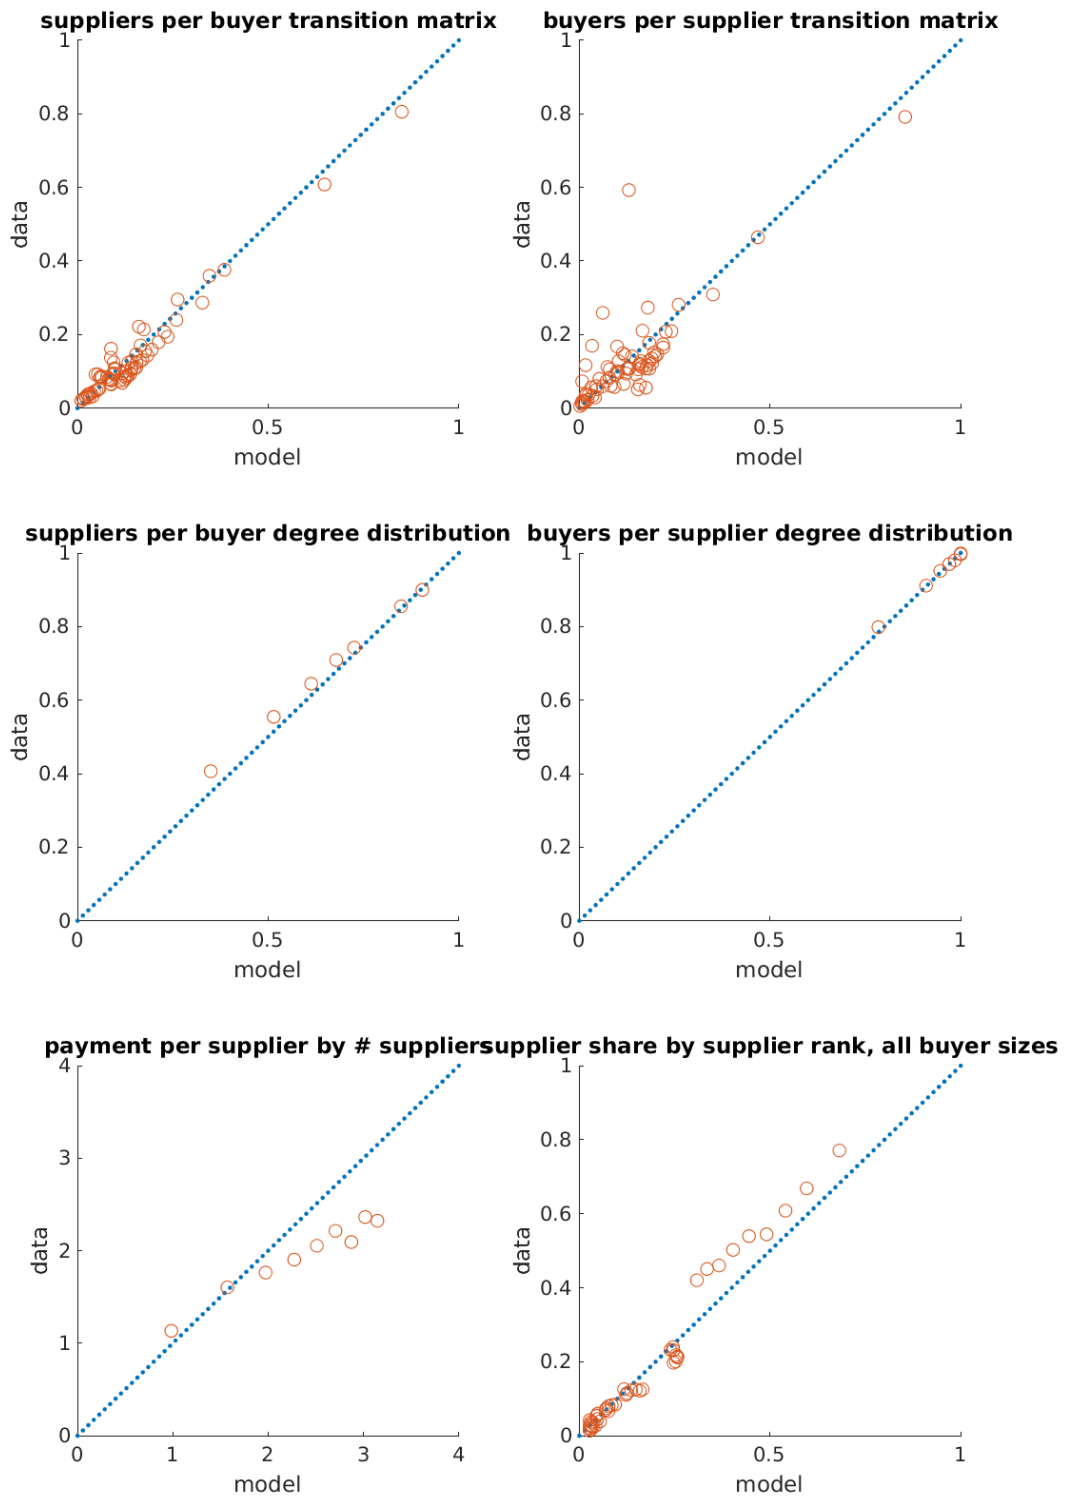


Figure 13: Match Shock Model: Data-based versus model-based moments

type dispersion, and lower buyer search cost scalar. We do not include the moment plot figures corresponding to Table 13 for the no dispersion model, because the figures are virtually identical to those from the match shock model.

D.3 Heterogeneous match death hazards

D.3.1 Overview

An alternative approach to allowing for heterogeneous match death rates does not require that we impose $\alpha = \eta$ or that we shut down visibility effects. However, it is more *ad hoc* in the sense that match death hazards are presumed to depend upon intrinsic supplier characteristics and not upon market forces. We describe this model in detail below. In short, we allow the two supplier types to have different match death hazards. In order to identify the additional parameter, we again use the slope coefficient from a regression of match death probabilities on the log number of supplier matches.

Because of the weak (but negative) relationship between match death hazard and number of supplier matches in the data, the estimated best fit match death hazard of high-type suppliers is only slightly lower than that of low-type suppliers (0.523 vs 0.563). The rest of the fitted parameters are similar to those of the baseline model. Moreover, the fit of the heterogeneous death hazard model on only the baseline moments is slightly worse than our baseline fit. For these reasons and for parsimony, we maintain the benchmark assumption of a uniform death hazard across supplier types.

D.3.2 Value functions

The value functions for this model are the same as those in the baseline model, up to the flow value of the buyer. With heterogeneity in supplier death hazards:

$$\begin{aligned}
(\rho + \delta^B)V_i^B(\mathbf{s}) = & \pi_i^B(\mathbf{s}) - \sum_{j=1}^J s_j \tau_{ji}(s) - k_s^B(\sigma_i^B, n^B) + \sigma_i^B(\mathbf{s}) \theta^B \sum_{j=1}^J v_j^S [V_i^B(\mathbf{s} + \mathbf{1}_j) - V_i^B(\mathbf{s})] \\
& + \sum_{j=1}^J (\delta_j + \delta^S) s_j [V_i^B(\mathbf{s} - \mathbf{1}_j) - V_i^B(\mathbf{s})]
\end{aligned} \tag{A-44}$$

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 supplier the buyer meets will be type- j . The only difference with

the baseline model is that the match death hazard δ_j has a j subscript, and is inside of the summation across supplier types in the last term. The necessary condition for the buyer's optimal search is:

$$\frac{\partial k^B(\sigma_i^B, n^B)}{\partial \sigma_i^B} = \theta^B \sum_{j=1}^J v_j^S [V_i^B(\mathbf{s} + \mathbf{1}_j) - V_i^B(\mathbf{s})]. \quad (\text{A-45})$$

We make a similar modification to the value function for a supplier of type j matched with a buyer of type i :

$$(\rho + \delta_j + \delta^B + \delta^S)V_{ji}^S(\mathbf{s}) = \tau_{ji}(\mathbf{s}) + \sigma_i^B(\mathbf{s})\theta^B \sum_{k=1}^J v_k^S [V_{ji}^S(\mathbf{s} + \mathbf{1}_k) - V_{ji}^S(\mathbf{s})] \quad (\text{A-46})$$

$$+ \sum_{k=1}^J (\delta_k + \delta^S)(s_k - \mathbf{1}_{k=j}) [V_{ji}^S(\mathbf{s} - \mathbf{1}_k) - V_{ji}^S(\mathbf{s})] \quad (\text{A-47})$$

As search is random, the expected value of a supplier's next relationship is:

$$V_j^S = \sum_i \sum_{\mathbf{s} \in \mathbb{S}} v_i^B(\mathbf{s}) V_{ji}^S(\mathbf{s}). \quad (\text{A-48})$$

And the necessary condition for optimal search is:

$$\frac{\partial k^S(\sigma^S, n^S)}{\partial \sigma^S} = \theta^S V_j^S. \quad (\text{A-49})$$

D.3.3 Laws of Motion

Change in the measure of buyers of type i with profile of $\mathbf{s} \neq \mathbf{0}$ is:

$$\begin{aligned} \dot{M}_i^B(\mathbf{s}) = & \sum_j [\sigma_i^B(\mathbf{s} - \mathbf{1}_j)\theta^B v_j^S M_i^B(\mathbf{s} - \mathbf{1}_j) + (\delta_j + \delta^S)(s_j + 1)M_i^B(\mathbf{s} + \mathbf{1}_j)] \\ & - \left[\sigma_i^B(\mathbf{s})\theta^B M_i^B(\mathbf{s}) + \left(\delta^B + \sum_j (\delta_j + \delta^S)s_j \right) M_i^B(\mathbf{s}) \right], \mathbf{s} \in \mathbb{S}; i = 1, \dots, I. \end{aligned} \quad (\text{A-50})$$

And the same for $\mathbf{s} = \mathbf{0}$:

$$\dot{M}_i^B(\mathbf{0}) = \delta^B \sum_{n^B(\mathbf{s}) \neq 0} M_i^B(\mathbf{s}) + \sum_j (\delta_j + \delta^S) M_i^B(\mathbf{1}_j) - \sigma_i^B(\mathbf{0})\theta^B M_i^B(\mathbf{0}) \quad i = 1, \dots, I. \quad (\text{A-51})$$

On the supplier's side:

$$\begin{aligned} \dot{M}_j^S(\mathbf{b}) &= \sum_k [\sigma_j^S(\mathbf{b} - \mathbf{1}_k) \theta^S v_k^B M_k^S(\mathbf{b} - \mathbf{1}_k) + (\delta_j + \delta^B) (s_k + 1) M_j^S(\mathbf{b} + \mathbf{1}_k)] \\ &\quad - \left[\sigma_j^S(\mathbf{b}) \theta^S M_j^S(\mathbf{b}) + \left(\delta^S + (\delta_j + \delta^B) \sum_k s_k \right) M_j^S(\mathbf{b}) \right], \mathbf{b} \in \mathbb{B}; j = 1, \dots, J. \end{aligned} \quad (\text{A-52})$$

And for $\mathbf{b} = \mathbf{0}$:

$$\dot{M}_j^S(\mathbf{0}) = \delta^S \sum_{n^S(\mathbf{b}) \neq 0} M_j^S(\mathbf{b}) + \sum_k (\delta_j + \delta^B) M_j^S(\mathbf{1}_k) - \sigma_j^S(\mathbf{0}) \theta^S M_j^S(\mathbf{0}) \quad j = 1, \dots, J. \quad (\text{A-53})$$

Aggregate mass of buyers and suppliers, search effort, and matching hazard calculations are as in the baseline model.

Identification In the baseline, we use the average match death probability observed in the data to back out δ . Since there are two supplier types, each with a δ_j , we need an additional moment to identify the additional parameter. In our data, we regress a match death dummy (D_{ijt+1}) on the log of the supplier's buyer count (N_{jt}^B), the log of the buyer's supplier count (N_{it}^S), and annual time effects (τ_t):

$$D_{ijt+1} = \beta_{NB} \ln N_{it}^B + \beta_{NS} \ln N_{jt}^S + \tau_t + \epsilon_{ijt}. \quad (\text{A-54})$$

The estimated value of β_{NB} is our additional target. This helps identify the two δ_j 's because suppliers with more matches are likely to be of the lower marginal-cost type.

The model average hazard of a match ending is:

$$\bar{\delta} = \frac{\sum_{\mathbf{s}} (\delta_1 |\mathbf{s}_1| + \delta_2 |\mathbf{s}_2|) M^b(\mathbf{s})}{\sum_{\mathbf{s}} |\mathbf{s}| M^b(\mathbf{s})} + \delta_B + \delta_S \quad (\text{A-55})$$

That is, it is the average hazard of a match death shock added to the hazard that the buyer dies and the hazard that the supplier dies. Given the match death hazard of one supplier type, we set the match death hazard of the other type to perfectly match the observed average death hazard (see Table 11 below).

We recover the regression coefficient in our model by regressing a match death dummy (D_{ijt+1}) on the log of the supplier's buyer count (N_{it}^B). We do not include the buyer's supplier

count or annual time effects because these have no effect on death hazards in our model, by construction.

D.3.4 Results

Table 10 compares results from our baseline model, also reported in Table 5, with results from estimating the heterogeneous match death hazard model.

Table 10: Comparison of Heterogeneous Death Hazard Model with Baseline*

	Baseline	Hetero Match Death Hazard
k_0	0.009 (0.003)	0.008 (0.003)
γ^B	0.320 (0.041)	0.311 (0.051)
γ^S	0.230 (0.046)	0.177 (0.068)
ω	0.030 (0.002)	0.031 (0.001)
Δ	0.454 (0.006)	0.457 (0.006)
M^S	4.203 (0.728)	4.496 (1.012)
η	2.432 (0.728)	2.435 (0.177)
$\sigma_{\ln \mu}^2$	7.428 (2.508)	7.121 (3.178)
δ_1	—	0.563 (0.010)
δ_2	—	0.523
Objective function	11,346.79	10,867.17
Baseline moments only	10,461.73	10,672.62

*GMM estimates of $\hat{\Lambda}$ based on equation (7). Moments targeted are the same as in Table 2, with the addition of the coefficient from regressing a match-death dummy on the supplier's number of matches. The objective function value for the baseline is computed including the extra moment, which adds 885.06 to the objective reported in Table 2.

Table 10 shows that the heterogeneous death hazard model improves on the baseline fit by around four percent. The fits of both models include the contribution of the additional

Table 11: Model vs Data-based Moments*

	Data	Model-based Estimates
Average match death probability	0.774	0.774
Buyer count coefficient ($\hat{\beta}_{N_B}$)	-0.0119	-0.0087

*Data vs. estimated model values of the average annual match-death probability and the coefficient on the supplier’s buyer count in equation A-54.

moment, the regression coefficient presented in Table 11. It is not surprising that the heterogeneous death hazard model fits somewhat better, since not only does it include an additional degree of freedom, but the baseline model will generate an exact zero for the newly included moment at any parameter vector. The bottom row of Table 10 reports that the baseline model fit is somewhat better than the heterogeneous death hazard model fit on the original moments.

The estimated parameters across the two models are similar. The death hazards, which we allow to vary across supplier types, do not differ much. High-type suppliers have a match death hazard only 0.04 less than that of low-type suppliers. Since the estimated parameters are similar and we do not focus on changes in match exit rates in our policy counterfactuals, we choose to use a common death hazard in our baseline model.

E Fixed/sunk cost exporting model

In this appendix we explore how our model’s implications differ from those of the canonical model of firm-level export dynamics described by Alessandria et al. (2021). To this end we develop and calibrate a stripped-down version of that model (hereafter our AAR model), retaining its key features: heterogeneous firms, fixed and sunk exporting costs, and serially correlated iceberg costs. Then we use it to redo the experiments reported in Section 6.1 of the text.

E.1 Our AAR model

Time is discrete. Each potential exporter has a permanent productivity type $z \in \{z_L, z_H\}$. Type shares are denoted by ω_L and $\omega_H = 1 - \omega_L$. Let $h \in \{0, 1\}$ denote whether the firm exported in the previous period. Firms that consider exporting draw an additive normal shock

when paying the fixed cost,

$$f(h) = \bar{f} + \mathbf{1}\{h = 0\}\Delta f + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, 1),$$

Each exporter also draws a two-state iceberg trade-cost shifter $\xi \in \{\xi_L, \xi_H\}$ that evolves according to

$$\Pr(\xi_{t+1} = \xi_t) = \rho_\xi, \quad \Pr(\xi_{t+1} \neq \xi_t) = 1 - \rho_\xi, \quad (\text{A-56})$$

with $\xi_H > \xi_L$. New exporters always start with $\xi = \xi_H$.

The foreign representative consumer has Cobb–Douglas preferences over a home aggregate C_H^* and an import aggregate C_M^* with import expenditure $E^* = (1 - \alpha)Y^*$. Within the import aggregate, preferences are CES over varieties with elasticity $\sigma > 1$. Hence prices follow the standard markup rule

$$p(z, \xi) = \mu\tau\xi\frac{w}{z}, \quad \mu \equiv \frac{\sigma}{\sigma - 1}, \quad (\text{A-57})$$

where w is the domestic factor price and $\tau \geq 1$ captures common iceberg costs. Revenue given demand $r = pq$ satisfies

$$r(z, \xi) = E^* \left(\frac{p(z, \xi)}{P^*} \right)^{1-\sigma}, \quad (\text{A-58})$$

with P^* the import price index. A firm exports if its profits from exporting are greater than its fixed costs, $\frac{1}{\sigma}r(z, \xi) \geq f(h)$. Aggregating across active exporters gives

$$(P^*)^{1-\sigma} = \sum_{z \in \{z_L, z_H\}} \sum_{\xi \in \{\xi_L, \xi_H\}} N(z, \xi) p(z, \xi)^{1-\sigma}, \quad (\text{A-59})$$

where $N(z, \xi)$ is the measure of active exporters in state (z, ξ) . We solve for equilibrium by iterating on the denominator in Eq. (A-59) until the implied P^* is consistent with the masses generated by firms' optimal export decisions.

E.2 Calibration

Parameters taken from EJTX Some of the parameters in AAR have analogs in EJTX. We set these equal to their values in our best-fit EJTX calibration. Specifically, we set the high-type share to $\omega_H = 0.03$ and the productivity gap $z_H/z_L = e^{0.454} \approx 1.575$. We set the elasticity of substitution across firms to $\sigma = 3.25$, the geometric mean of the EJTX within-

store and cross-store elasticities, to provide a single-nest proxy for the nested CES structure. We set the discount factor to $\beta = 0.951$ to correspond to the annual discount rate used in EJTX. We normalize cost of production $w = 1$ and the total incumbent mass $M_L + M_H = 1$. Finally, we set total spending on imported goods E^* to unity.

Simulated targets As we do not have data on the relevant population of foreign firms, we cannot directly observe their export market participation patterns. We therefore calibrate remaining parameters $\theta = (\bar{f}, \Delta_f, \xi_H, \rho_\xi)$ to fit moments generated by a simulation of EJTX. These are chosen to be informative about exporter dynamics:

1. the conditional entry rate, i.e., the probability that a non-exporter begins exporting within the year;
2. the probability of surviving to age two conditional on surviving to age 1;
3. the ratio of average export sales at ages 2 and 1 among firms that survive through age 2 (\bar{r}_2/\bar{r}_1); and
4. the analogous ratio at ages 3 and 1 among firms that survive through age 3 (\bar{r}_3/\bar{r}_1).

We fit our AAR model using an identity weighting matrix.

E.3 Estimation Results

Table 12 reports the parameter estimates and the corresponding EJTX targets. The entry rate

Table 12: MSM results (EJTX moments vs. our AAR model)

Parameter/moment	Estimate	EJTX target	our AAR model
\bar{f}	2.1507	—	—
Δ_f	0.3071	—	—
ξ_H	1.3529	—	—
ρ_ξ	0.5552	—	—
Entry rate	—	0.1093	0.1093
Conditional survival	—	0.8209	0.8209
Within-firm sales ratio \bar{r}_2/\bar{r}_1	—	1.4332	1.4332
Within-firm sales ratio \bar{r}_3/\bar{r}_1	—	1.5090	1.5090

primarily disciplines the gap between entry and continuation costs Δ_f , while the conditional

survival probability anchors the level of \bar{f} . The two growth ratios identify the trade-cost process (ξ_H, ρ_ξ) by pinning down how quickly survivor sales expand after entry. The estimated trade-cost process is moderately persistent and features a significant difference in iceberg costs across types.

While both models explain the moments in Table 12, they rely on different mechanisms to do so. In particular, to explain post-entry sales growth, the EJTX model uses convex search costs and reputation effects. But the AAR model interprets this growth as low-type suppliers drawing the low-trade-cost state after paying the sunk entry cost. Also, to match foreign market entry and survival rates, the two models invoke different types of foreign market participation costs. In EJTX they are supplier search costs; but in our AAR model they are the fixed cost paid by entrants and continuing exporters. (The costs also differ in magnitude: aggregate search costs make up 45.0% of aggregate supplier profits in the EJTX model, while aggregate sunk and fixed costs are 62.4% of total exporter profit in our AAR model.)

It is instructive to compare the fixed costs we estimate with those reported by Alessandria et al. (2021). We find the continuation-to-entry cost ratio to be $f_1/f_0 = \bar{f}/(\bar{f} + \Delta_f) = 0.875$, while Alessandria et al. (2021) report $f_1/f_0 = 0.263$. One plausible interpretation is that apparel exporters, which are subject to frequently changing fashion trends, exhibit more churning than the typical American exporter studied in Alessandria et al. (2021).⁷⁰

E.4 Experiment, Revisited (Section 6.1)

Using the policy experiment described in Section 6.1, we now juxtapose the predictions of our AAR model with those of our EJTX (baseline) model. The experiment, recall, is to increase the mass of potential exporters from $M^{S,\text{pre}} = 2.4$ to $M^{S,\text{post}} = 4.2$ while simultaneously reducing the fraction of potential suppliers that are high quality from $\omega_H^{\text{pre}} = 0.043$ to $\omega_H^{\text{post}} = 0.030$. All other parameters are fixed at their estimated baseline values from our AAR model calibration.

The results are summarized in Table 13. The supplier-entry margins are directionally similar across models, but they differ in several important respects. In our AAR model, high-productivity exporters are disproportionately close to the margin of exporting, so they

⁷⁰Selection implies that both the average entry cost paid by firms that actually enter and the average fixed costs paid by firms that continue exporting are both lower than the unconditional average of those costs. The conditional entry cost $\mathbb{E}[f_0 + \varepsilon \mid \text{entry}] = 0.747$, which is 22% of average entrant revenue. Among firms that continue exporting, the average fixed cost is $\mathbb{E}[f_1 + \varepsilon \mid \text{continue}] = 1.571$, or 18% of average continuing revenue.

Table 13: Section 6.1 experiment: EJTX vs. our AAR model (% changes)

Metric	EJTX*	our AAR model
Measure active low-quality suppliers	+29.3	+13.5
Measure active high-quality suppliers	+19.3	+21.6
Measure active buyers	-0.9	—
High-quality suppliers per buyer	-16.6	—
Consumer welfare ($1/P$)	-0.8	+5.5

*Figures in this column are based on differences between figures in columns 2 and 3 of Table 6 in the text.

disproportionately expand into foreign markets as M^S rises, even though ω_H falls. But in the EJTX model, low-quality suppliers expand relative to high-quality suppliers, while high-quality suppliers per buyer fall substantially. These adjustments reflect, inter alia, the discouraging effects of market congestion on the search efforts of buyers and high-quality suppliers. (Refer to Section 6.1 for details.)

The welfare implications of the two models are correspondingly different. In EJTX, the reduction in buyer counts and the reductions in the average quality of buyers' retail offerings combine to reduce consumer welfare nearly 1 percent. But in our AAR model there are no buyer-side adjustments—indeed, there are no buyers—and endogenous adjustments in search efforts are absent. Accordingly, the expanded supplier mass simply lowers the CES price index through love-of-variety effects, generating substantial welfare gains.