Feb 2019 No.403

International Buyers' Sourcing and Suppliers' Markups in Bangladeshi Garments

Julia Cajal Grossi, Rocco Macchiavello and Guillermo Noguera

WORKING PAPER SERIES

Centre for Competitive Advantage in the Global Economy

Department of Economics

International Buyers' Sourcing and Suppliers' Markups in Bangladeshi Garments*

Julia Cajal Grossi† Rocco Macchiavello‡ Guillermo Noguera§

January 21, 2019

Abstract

Large international buyers play a key role in global value chains. We exploit detailed transaction-level data on the usage of material inputs to study how Bangladeshi garment suppliers' markups vary across international buyers. We find substantial dispersion in markups across export orders of a given seller for the same product. Buyer effects explain a significant share of this variation, while destination effects do not. Buyers adopting relational sourcing strategies pay higher markups than non-relational buyers. This pattern holds within seller-product-year combinations, is robust to controlling for the buyer's size, traded volumes, and quality, and, together with larger volumes, implies higher profits for suppliers dealing with relational buyers.

Keywords: Markups, Sourcing Strategies, Global Buyers, Buyer-Driven Value Chains.
JEL Codes: L11, L14, D23, F63

*We are particularly indebted to Chris Woodruff with whom this project was started. We thank Jan De Loecker, Amit Khandelwal, Kalina Manova, Nina Pavcnik, Tim Schimdt-Einselhor, Catherine Thomas and Eric Verhoogen for helpful discussions. We also thank seminar participants at Barnard, Düsseldorf, London School of Economics, Namur, Nottingham, Oslo, Oxford, Warwick, Yale, and conference participants at the 2019 ASSA Meeting, Bank of Italy, Kellogg-Northwestern Conference on Markets and Organizations in Emerging Economies, NYU Trade Jamboree, PEDL Inaugural Conference, CEPR/WB Conference on Global Values Chains, Trade and Development, GEP Imports and Global Value Chains conference, 2018 SAET conference, and CAGE Research day for helpful discussions and comments. We gratefully acknowledge financial support from the Centre for Competitive Advantage in the Global Economy (CAGE).

†Graduate Institute, Geneva. E-mail: julia.cajal@graduateinstitute.ch.
‡London School of Economics and CEPR. E-mail: r.macchiavello@lse.ac.uk.
§Yale University and CAGE. E-mail: guillermo.noguera@yale.edu.

1 Introduction

There is broad consensus that large international buyers play a key role in global value chains (Gereffi, 1999; Antràs, 2015). Yet, it is still being debated what influence these large buyers have on their suppliers, particularly in developing countries. On the one hand, there is a concern that in some industries, suppliers face increasingly powerful buyers who squeeze suppliers' margins by engaging in highly competitive sourcing strategies.[1] On the other hand, global buyers increasingly rely on collaborative, or relational, sourcing practices in which suppliers receive higher markups and share in the gains from trade.[2] The debate thus revolves around the role played by buyers adopting radically different approaches to sourcing. While these polar sourcing strategies are well documented and understood theoretically (see, e.g., Taylor and Wiggins, 1997), empirical evidence on how they relate to suppliers' markups is lacking. Such evidence would improve our understanding of the distributional consequences of trade (Goldberg and Pavcnik, 2007) and our ability to formulate adequate policy responses (Lederman et al., 2010).[3]

Despite the importance of the question, data limitations have hampered empirical progress. Shedding light on the debate requires knowing how markups vary across different buyers, holding constant exporter and product heterogeneity. While recent work has made progress in measuring markup differences across manufacturers (see, e.g., De Loecker and Warzynski, 2012; Atkin et al., 2015) as well as within manufacturers across products (De Loecker et al., 2016), the lack of detailed data has so far prevented the analysis of buyer-seller-specific markups.

In this paper, we study how markups charged by Bangladeshi ready-made-garment suppliers vary across international buyers adopting different sourcing strategies. Besides its intrinsic interest,[4] the context presents two features that facilitate our analysis. First, we observe the price and quantity of a variable input (fabric) used in each export transaction.

[1] See Competition Commission (2000) and FTC (2001) for reports on buyer power in different sectors, and Oxfam (2002) and Gereffi (1999), among others, for the concerns that this raises for developing countries.

[2] Antràs (2015) examines the contractual imperfections that render relational sourcing necessary, Baker et al. (2002) model the rents needed to sustain cooperation in long-term relationships between firms, and Macchiavello and Morjaria (2015) provide direct evidence of such rents in export markets. Beyond static gains in markups, relational sourcing may also foster knowledge transfers and supplier's upgrading; see, e.g., Dwyer et al. (1987), Egan and Mody (1992), Kalwani and Narayandas (1995), and Atkin et al. (2017).

[3] More broadly, estimating markups is key to understanding how firms respond to trade liberalizations (see, e.g., De Loecker et al., 2016), how large are the welfare gains from trade (see, e.g., Atkin and Donaldson, 2015), and how firms transmit international shocks to local markets (see, e.g., Campa and Goldberg, 2005; Nakamura and Zerom, 2010).

[4] The garment industry has played a critical role in the early phases of export-oriented industrialization, most recently in East Asia (see, e.g., Dickerson, 1999; Gereffi, 1999). Bangladesh is the world's second largest exporter of garments (after China) and the industry, which accounts for over 80% of the country's exports and an estimated 12% of its GDP, employs over four million workers, mostly women.

This allows us to recover marginal costs and markups at the level of the export order. Second, we are able to link each export order, together with its estimated marginal cost and markup, to the international buyer the order is produced for. These two features allow us to explore how marginal costs and markups correlate with buyers' characteristics, holding constant factors at the seller-product-time level. From this exploration, our main result is that a given seller exporting a given product receives higher markups and makes higher profits when supplying a buyer that adopts relational sourcing (e.g., The Gap) than when dealing with a buyer that does not (e.g., J.C. Penney).

We begin by highlighting salient features of the production process for woven garments. These features serve as a guide for the multi-product firm-level cost minimization framework that we take to the data. Woven garments are manufactured order-by-order through a common sequence of steps, organized into an inspection and cutting stage and a sewing and finishing stage. The main advantage of our data is that we observe the so-called *buy-to-ship ratio* at the order level, which measures the relationship between the volume of fabric input and the volume of garment output in a given order. This ratio is a key performance indicator in the industry, capturing fabric efficiency at the inspection and cutting (*buy-to-cut*) and sewing and finishing (*cut-to-ship*) stages. In our data, order-level buy-to-ship ratios are highly dispersed, and we discuss novel evidence documenting how this dispersion may arise from differences in the order-level production processes, and/or from variable input choices that reflect substitution between fabric and other inputs.

Our theoretical framework considers a seller producing multiple export orders and choosing both fully flexible (e.g., fabric) and quasi-fixed (e.g., labor) factors of production, subject to order-specific productivity shifters, capacity constraints, and factors' cost shares that are seller-product-specific. Building on Hall (1986) and De Loecker et al. (2016), the firm's cost minimization problem implies that order-specific markups depend on (i) the share of expenditures on fabric in the order's revenues and (ii) the output elasticity with respect to fabric. The unique feature of our data is that we directly observe (i) for each export order. In addition, unlike the existing literature on markup estimation, we identify (ii) by deriving a structural equation that relates the order-level buy-to-ship ratio to fabric price and order size. We propose an instrumental variable (IV) approach for estimating this equation to recover the elasticity of output to fabric, which we then use to construct the order-level markups and costs. Our IV strategy borrows from the empirical literature on networks and entails constructing a neighbor-of-neighbor instrument (see, for example, Bramoullé et al., 2009; De Giorgi et al., 2010). We instrument the size of an order between a given supplier and buyer using the orders that buyers not directly connected to this supplier place with other suppliers which also sell to this buyer. We verify that our estimates conform with both industry knowledge and previous findings in the literature.[5]

Our analysis starts by presenting two novel facts. First, we show that there is substantial dispersion in the markups that a given seller charges to different buyers for the same product in a given year. Second, we show that buyer effects account for a large proportion of this dispersion. Destination effects, in contrast, explain only a small share of the observed variation. That is, within destination markets, buyers are highly heterogeneous in terms of the markups that they pay.

Motivated by these facts, we examine how markups vary systematically with buyers' characteristics. Our analysis is shaped by the aforementioned debate in the literature, whereby a buyer's market power should be negatively correlated with supplier markups, while a buyer's relational sourcing should be positively associated with them. We thus focus on two dimensions of cross-sectional heterogeneity across international buyers: *size* and *sourcing strategy*. As documented in the global value chains literature in general (e.g., De Toni and Nassimbeni, 2000; Taylor and Wiggins, 1997; McMillan, 1990), and in apparels in particular (e.g., Gereffi, 1999), buyers' sourcing practices range from *spot* procurement at one extreme to *relational* sourcing at the other extreme.[6] Following Heise et al. (2017), we capture the degree to which a buyer's sourcing strategy is relational by normalizing its number of suppliers by the overall number of shipments received by the buyer. Our baseline specification focuses on within seller-product-year differences in markups across buyers, conditioning on the destination market and the overall size of the buyer.[7] We find that relational buyers pay

[5]The structural equation is derived assuming a Cobb-Douglas production function at the order level and is estimated allowing the fabric elasticity to vary across products. It is important to note that the elasticity of output to fabric and, therefore, these assumptions, are needed only to recover the level of markups and to explore patterns *across* sellers which we use to confirm previous findings in the literature, namely that larger sellers enjoy higher (lower) markups (marginal costs), particularly so in their core products (see, e.g., De Loecker et al., 2016). For our main results, which focus on differences in markups across buyers *within* seller-product-time combinations, the output elasticity to fabric is allowed to flexibly vary at the seller-product-year level.

[6]These two models are sometimes referred to as 'adversarial' or 'American-style' sourcing in contrast to 'collaborative' or 'Japanese-style' sourcing (see, e.g., Kawasaki and McMillan, 1987; Richardson, 1993; Helper and Saki, 1997). In the reminder of the paper, we will refer to the two strategies as *spot* and *relational* sourcing respectively. Due to complementarities in sourcing practices across sourcing countries and suppliers (Milgrom and Roberts, 1990; Antràs et al., 2017), the choice of procurement system is to a large extent a buyer-level strategic decision. Our analysis thus mainly focuses on cross-buyer variation in sourcing strategies, rather than on differences across buyer-seller relationships. We document that even buyers operating in the same market source the same garments using different procurement strategies. We take these differences as given without attempting an explanation for the buyers' adoption of different sourcing strategies.

[7]Appendix B explores differences in markups across buyers without conditioning on seller-product-year fixed effects. We give less prominence to these correlations as they confound the role of buyer characteristics we are most interested in with sorting patterns between buyers and sellers. For example, within product-years, larger buyers are associated with lower marginal costs and higher markups. This correlation, however, could simply reflect that larger buyers source from larger sellers that have lower costs and charge higher markups.

higher prices and markups than spot buyers, while they are not associated with higher or lower marginal costs.

We investigate the robustness of this result to several potential confounders. First, buyers adopting relational sourcing practices tend to be larger. While our baseline specification conditions on the overall size of the buyer, it ignores the fact that relational buyers also tend to trade larger volumes with fewer suppliers. We thus decompose the size of the buyer into overall size and the volume traded in the relationship with the supplier. The result of this exercise is that relational buyers are associated with higher marginal costs conditional on the volume traded in the relationship, which itself is negatively (positively) correlated with marginal costs (markups).

A second potential confounder is product quality, which could explain the finding that relational buyers pay higher markups on garments that also have higher marginal costs.[8] That is, for any given supplier and narrow product category, relational buyers may source higher quality garments than non-relational ones. The key driver of output quality in garments is fabric (see, e.g., Medina, 2017), consistent with a large recent literature linking the quality of output to the quality of inputs used in production (see, e.g., Verhoogen, 2008; Kugler and Verhoogen, 2012; Bastos, Silva, and Verhoogen, 2018). We thus exploit the fact that we observe the price, variety, and country of origin of the fabric used in each export order to control for physical product quality. We also control for other aspects of quality along the buyer or the relationship dimensions, such as the seasonality in sourcing and specialization patterns. Even after controlling for these dimensions of product quality, we find that relational buyers afford higher marginal costs and pay higher markups than their non-relational counterparts.

Our results have important implications for suppliers' performance. Relational buyers source significantly larger volumes from their suppliers, relative to non-relational buyers of comparable size. Combined with the higher markups that they pay, these larger volumes imply that sellers' overall profits are larger when their portfolio of buyers tilts in favor of trade partners adopting relational sourcing. We confirm this implication of our results by aggregating our specifications at the relationship and seller levels. We conclude that the different sourcing practices adopted by large international buyers are, in fact, an important margin affecting exporters' overall performance.

Our preferred interpretation of the results is that the higher marginal costs and markups paid by relational buyers reflect the basic contractual differences between the two alternative sourcing strategies, as modeled in Taylor and Wiggins (1997). Under spot procurement,

[8] Destination-specific characteristics that correlate with prices (see, e.g., Manova and Zhang, 2012) are controlled for by destination fixed effects in our baseline specification.

buyers pay for each order without any explicit or implicit agreement regarding future trade. Under relational sourcing, buyers use the ongoing relationship to incentivize suppliers to undertake costly non-contractible actions, such as to guarantee reliable and on-time delivery. These incentives stem from the threat of terminating the relationship following contractual non-performance. For this threat to effectively provide incentives, relational buyers must pay rents to the suppliers. This interpretation is consistent with the evidence that, conditional on the volume traded in the relationship and the physical quality of the product, relational buyers' orders are associated with higher marginal costs (arising from sellers' costly non-contractible actions) and with higher markups (offered as incentives).

Other interpretations appear less plausible. An alternative reading of our results is that relational buyers select into trading relationships that have better match-specific capabilities. This hypothesis is *prima facie* inconsistent with the evidence that orders supplied to relational buyers entail higher marginal costs. Furthermore, we explore specifications that control for buyer-seller fixed effects and document that relationship dynamics differ across the two procurement strategies in ways consistent with relational buyers offering stronger implicit incentives. Another possibility is that higher markups simply reflect stronger competition among relational buyers to source from a limited number of suppliers with adequate capabilities. Besides it being at odds with information collected from industry participants during field interviews, this interpretation is not supported by the evidence that we present. We would expect the best suppliers to take advantage of competition forces and expand the volume they supply to relational buyers, on which they make higher profits. Instead, suppliers have a mix of both relational and non-relational buyers.

In sum, our results provide what is, to the best of our knowledge, the first direct evidence of the equilibrium relationship between markups and marginal costs on the one hand and a buyer's sourcing practice on the other, a relationship that has received significant attention in the literature. More broadly, we show that trading with relational buyers is associated with higher suppliers' profits, and we find no evidence that larger buyers, potentially commanding greater market power, squeeze suppliers' margins by paying lower markups.

Related Literature. This paper contributes to different strands of literature. First, the paper builds on the literature on heterogeneity in markups across firms.[9] On the empirical side of this literature, our paper is closest to supply-side approaches such as De Loecker and Warzynski (2012).[10] More recently, De Loecker et al. (2016) (see also Voigtlaender and

[9]See, e.g., Atkeson and Burstein (2008), Melitz and Ottaviano (2008), and Edmond et al. (2015) for theoretical frameworks featuring variable markups.

[10]The supply-side approach, which relies on cost minimization, requires data on prices and volumes of output and inputs to avoid otherwise necessary assumptions on market structure and demand functions

Garcia-Marin, forthcoming) extend this framework to estimate within-seller differences in markups across products. Our approach is similar to theirs in that we use the first order conditions from the firm's cost minimization problem to recover markups, without imposing a specific demand structure or a particular pricing rule. Unlike these papers, by directly observing input utilization for each export order, we can circumvent the corrections that would be necessary due to the unobservability of input-output allocations within multi-product firms. In the same vein, we observe order-specific input prices in our data, which eliminates the difficulties of dealing with industry-level prices to recover heterogeneous costs and markups. We complement De Loecker et al. (2016) by estimating variable markups at a more disaggregated level, focusing on heterogeneity in the markups charged to different buyers by a given seller in a product category. Our study of a specific industry and the level of detail in our data also brings us close in spirit to Atkin et al. (2015), who directly measure markups through survey questions in the Sialkot soccer ball cluster in Pakistan. In an Appendix, we confirm the main across-firms and across-products within-firm results in De Loecker et al. (2016) and Atkin et al. (2015) and also document similar, but novel, across-buyers patterns.

Second, this paper contributes to the literature studying the role of buyers in global sourcing.[11] Starting with Appelbaum and Gereffi (1994) and Gereffi (1999), a large, mostly qualitative and descriptive literature, has studied sourcing practices of global buyers and their impact on producers in developing countries. This literature integrates concepts from management and supply-chain studies (see, e.g., Kawasaki and McMillan, 1987; Richardson, 1993; Helper and Saki, 1997) into the analysis of global value chains. Taylor and Wiggins (1997) offer a theoretical model based upon which Heise et al. (2017) derive an observable metric of relational sourcing. We adapt this measure to our context and, taking advantage of our data, directly test the prediction in Taylor and Wiggins (1997) that relational buyers are associated with higher marginal costs and markups compared to non-relational ones.[12] A related, recent body of literature also documents the importance of relational practices when contracts between firms are hard to enforce (see, e.g., Macchiavello, 2010; Antràs and Foley, 2015; Macchiavello and Morjaria, 2015; Blouin and Macchiavello, forthcoming). Like these papers, we focus on a specific industry and exploit detailed data and contextual knowledge to investigate the channels through which relationships matter. Those papers point at the importance of future rents in deterring opportunism at the *relationship* level. The large number of buyers and sellers in our sample allows us to investigate the role of across

(see, e.g., Berry et al., 1995; Goldberg, 1995).

[11] See Antràs et al. (2017) for a state-of-the-art treatment of a buyer's sourcing problem.

[12] Startz (2018) estimates search and contracting frictions faced by Nigerian consumer goods importers with a framework à la Taylor and Wiggins (1997).

buyers variation in sourcing strategies. By showing that buyers adopting relational sourcing practices pay higher markups we thus follow a distinct, but complementary, approach.

Finally, by focusing on the buyer's margin, we also contribute to a fast growing literature on buyer-seller relationships in international trade (see, e.g., Bernard et al., forthcoming; Bernard et al., 2018b; Heise, 2018; Monarch, 2018; Eaton et al., 2016; Monarch and Schmidt-Eisenlohr, 2017). While this literature does not focus on differences in markups, more recent models (e.g., Carballo et al., 2018; Heise et al., 2017; Bernard et al., 2018a; Kikkawa et al., 2018) generate predictions on cost or markup heterogeneity across buyers or relationships. These predictions are indirectly tested on observable variables, such as prices, revenues or relationships' durations. To the best of our knowledge, however, there is no empirical evidence on how markups and marginal costs differ across this 'new' important margin of trade. This is because relationship-level markups cannot be recovered from commonly available data on buyer-seller trade alone without imposing significantly more structure on the price setting mechanism, or making strong assumptions about product quality or input allocation in multi-product firms. By contrast, our paper avoids these assumptions by resorting to detailed data on input usage at the export order level. This allows us to recover the markups charged in different orders and study how those vary across buyers with different characteristics.

The paper proceeds as follows. Section 2 describes key aspects of garment production in Bangladesh and the data. Section 3 introduces the theoretical framework and estimation approach and describes estimated markups and marginal costs. Section 4 presents our main results on marginal costs and markups heterogeneity and on its relationship with buyers' sourcing strategies. Section 5 concludes.

2 Industry and Data

We study the ready-made-garment industry in Bangladesh. Garment production is one of the oldest and largest export industries in the world, and one that has played a critical role in the early phases of export-oriented industrialization in several countries, most recently in East Asia (see, e.g., Dickerson, 1999 and Gereffi, 1999). In Bangladesh, this industry started developing in the 1980s. After decades of rapid growth, Bangladesh is now the world's second largest exporter of garments (after China). The Bangladeshi garment industry currently comprises more than 3,000 factories employing over four million workers, mostly women, and it accounts for over 80% of the country's exports and an estimated 12% of its GDP.[13]

[13]The garment industry in general, and in Bangladesh in particular, has been the object of extensive research. Among others, see Gereffi (1999) on organizational modes of export and upgrading, and Heath

Below we describe our data and sample and how we construct our main variables of study. We also discuss the garment production process and how its main characteristics inform our framework in the next section.

2.1 Data and Sample

Our main source of data consists of all the transaction-level export and import records from Bangladesh over the period 2005-2012. These records include information on the product transacted, its value, and its volume (in kilos), the date of the transaction, as well as unique identifiers for the Bangladeshi exporting/importing firms. For export transactions, the records also include the names of the international buyers. The main novelty of the data is that we can match material inputs usage to output at the transaction level.

We focus on woven garments. Two features of the Bangladeshi woven garment sector enable us to link material inputs use to output at the export order level. First, in contrast to other major garment exporters such as China, India, and Pakistan, Bangladesh lacks a domestic woven textile industry. Woven products exported by Bangladeshi firms are thus produced using imported fabric (e.g., cotton fabric) exclusively, as there are no suitable domestic substitutes. This implies that total fabric use by these woven garment exporters can be measured directly from customs records of fabric imports.

The second main feature of this sector is that export-oriented garment manufacturers participate in a customs bonded warehouse regime, which allows them to import material inputs for production of export orders duty free. To participate in this regime, exporters must indicate the export order for which the imported fabric will be used. Specifically, after receiving an order from an international buyer—a contract describing terms like payment, order size, and delivery dates, as well as technical specifications—the manufacturer submits a utilization declaration (UD) to the Bangladesh Garment Manufacturers and Exporters Association. If approved, a unique UD identifier is then assigned to all export and import transactions belonging to this contract. Importantly, this utilization declaration system operates at the level of the export order. This is in contrast to similar regimes (e.g., export processing zones) that operate at the firm level in other sectors in Bangladesh and in other countries.[14]

and Mobarak (2015) and Macchiavello et al. (2015) on women's labor force participation, education, and promotion to managerial roles in Bangladesh.

[14]The first of the two features described here does not apply to knitted garments, which account for about half of Bangladeshi garment exports. The production of knitted garments differs from that of woven garments in that it requires an additional production step from yarn to fabric. This step is sometimes performed in Bangladesh by either vertically integrated or independent units. As we do not have access to domestic data, we are unable to match imported yarn material to output of knitted garments in a systematic manner.

Given these two features of the woven garment sector, we are able to identify the material inputs that correspond to each export order, which is typically of a single garment product. In particular, each export and import customs record in our data contains a UD identifier alongside standard information as described above. We aggregate transaction-level records at the order (i.e., UD) level, producing, for each order, a single entry that contains information on: the buyer's identity and destination country, garment product code, value and volume of garment exported, seller's identity, fabric product code, value and volume of fabric imported, and country of origin of fabric.[15] To illustrate, a hypothetical observation in our dataset would look as follows: based on UD 2/124/46/902, *Nice Apparel Co. Ltd.* imported 400 kg of unbleached woven fabric (containing 85% or more by weight of cotton, in 3-thread or 4-thread twill, including cross twill, weighing not more than 200g/m^2, i.e., HS520813) at \$6 per kg from China on 01/20/2008 and subsequently exported to *Walmart Inc.* 450 kg of men's or boys' woven cotton shirts (HS620520) on 03/01/2018 at \$10 per kg.

We are able to match export and import transactions by UD identifier for 33,490 woven garment orders between 2005 and 2012. In the empirical analysis, we focus on the 17 six-digit HS codes in the two largest woven apparels: shirts and trousers. To ensure that we compute meaningful statistics on within-exporter markup dispersion, and to implement the IV estimation strategy that we describe in the next section, we further restrict our analysis to orders associated with the 500 largest exporters, which account for 87% of the relevant sample.[16]

Table 1 provides descriptive statistics for our final sample. On average, orders last around four months from the date of the first import of fabric to the date of the last export of apparel (Panel A). Sellers are active for an average of 6.63 years in the sample. The average seller exports 3.25 different woven products (six-digit HS codes) and trades with 5.93 buyers on average in a given year (Panel B). Buyers are slightly more diversified both in terms of products and sellers: the average buyer sources 4.25 products from 21.94 sellers on average in a given year (Panel C). Our sample contains trades between 5,633 buyer-seller pairs. Many transactions occur within repeated buyer-seller relationships; the average buyer-seller relationship lasts almost two years and involves around 3.39 orders per year (Panel D).

[15] In some cases, multiple types of fabric and additional material inputs (e.g., accessories like zippers and buttons) are imported in the same order. Our baseline analysis uses the total amount of fabric imported but abstracts from other material inputs as these may also be sourced domestically. We explicitly consider these additional nuances when exploring robustness of our main findings to differences in product quality.

[16] Table A1 in the Appendix compares transactions in the UD system, with those outside of the UD system. More details on the sample construction can be found in Appendix A.1.

2.2 Production Stages and Performance Metrics

The production process of garments comprises two sequential stages: (1) inspection and cutting, and (2) sewing and finishing.

Inspection and cutting. In the first stage, manufacturers inspect the quality and quantity of the purchased fabric, plan fabric utilization, and then proceed to cutting. To cut the fabric, manufacturers make markers, which are thin sheets of paper with diagrams indicating the pattern pieces to be cut for the specific style and sizes of an order. Manufacturers then spread the fabric on the cutting tables, cut the patterns, and finally ticket and bundle the pieces.

Each of these tasks has a direct impact on fabric efficiency and relative labor, capital, and fabric input use. For example, manufacturers may use fabric inspection machines to check for fabric and print defects and shading. The markers for cutting can be made either by hand or by using software that automatically arranges the pattern pieces to reduce fabric waste. Spreading can also be done by hand (using a spreading table with roll racks, tracks, clamps, lifters, and end cutters) or by automatic spreading machines. Finally, cutting can be performed using manual, semi-automatic, or automatic systems, employing a variety of portable cutters (rotary or straight knives) or stationary cutters (band, die, laser, etc.), and either manually handling the fabric or holding it in place using a vacuum to avoid distortions and misalignments in the spread.

Sewing and finishing. In the second stage, cut fabric is sent to the sewing department, where cut fabric pieces are sewn together along production lines. Depending on the type of garments, fabric, and machines, production lines typically employ between 30 and 70 sewing operators and one or more line supervisors. The quality of the machines, the number and quality of sewing operators, and the effort of supervisors all affect the efficiency with which garments are sewn together.

From the point of view of fabric use, there can be fabric losses at the sewing stage due to quality defects such as stains or garments sewn incorrectly. These losses can be reduced by including additional quality-control workers alongside the sewing lines. Factories can organize one or more inspection points along and at the end of the production line, or simply inspect quality in the finishing section, when the garment is pressed or ironed, finished, and packed.

The buy-to-ship ratio. Manufacturers typically compute a performance indicator for each of the two stages of production. The *buy-to-cut ratio* measures performance at the inspection and cutting stage. This is the ratio of purchased fabric to cut fabric, with lower values representing lower wastage and thus higher fabric efficiency in transforming purchased fabric into cut fabric bundles. The *cut-to-ship ratio* measures performance at the sewing and finishing stage. This is the ratio of cut fabric to shipped garments, with lower values representing lower wastage arising from defects and thus higher fabric efficiency at this stage.

The buy-to-ship ratio is a commonly used key performance indicator that captures fabric efficiency over the two stages. This is the ratio of purchased fabric to shipped garments, namely the product of the buy-to-cut and the cut-to-ship ratios. Lower values of the buy-to-ship ratio represent lower wastage and thus higher efficiency over the two stages of production. As an illustration, for our hypothetical data observation on *Nice Apparel Co. Ltd.* in Section 2.1, the buy-to-ship ratio corresponding to that order would be $400/450 \approx 0.89$. In our data, we find that the average order-level buy-to-ship ratio is 0.87; see Table 1.[17]

Our description of the two stages of production above provides insight into the possible sources of variation in the buy-to-ship ratio. To understand these further and guide our estimation framework, we next highlight important characteristics of the garment production process and discuss how to model them.

2.3 Modeling Garment Production

There are three key characteristics of the garment production process that will discipline our modeling assumptions in the next section:

Order-level production. Due to the buyer-driven nature of the garment value chain and to the utilization declaration system, manufacturers make production decisions based on the export orders they receive. Subject to productive capacity constraints, manufacturers allocate production orders to production lines often shifting operators (and, more rarely, machines and supervisors) across lines to optimize the use of resources. Export orders are thus the natural level of aggregation at which fabric utilization choices are made.

Dispersion in buy-to-ship ratios. We expect significant variation in buy-to-ship ratios. This is confirmed in our data. Table 1 reports a coefficient of variation of the order-level buy-to-ship ratio of 33%. As explained in the previous subsection, this observed dispersion

[17] The buy-to-ship ratio is computed using net export volumes (kilos) that include accessories and packaging (garments are folded in plastic envelopes and then stored in carton boxes). This explains why the ratio in our data is typically below one. We have access to detailed data on the type of packaging and number of packages as received by the processing customs office, and we verify that our results are unchanged when our baseline estimation of the buy-to-ship ratio equation is modified to condition for these standard packaging characteristics.

may result from both differences in efficiency in the inspection and cutting and sewing and finishing stages of production as well as from the possibility of substituting fabric with other inputs, at least to a certain extent. In the first case, the dispersion in the buy-to-ship ratio reflects dispersion in the buy-to-cut and cut-to-ship ratios. We next describe further evidence from within-firm studies suggesting that these sources of variation are indeed at play in the data.

The engineering study of Tanvir and Mahmood (2014) examines 30 Bangladeshi factories producing single jersey standard shirts. This study finds that fabric wastage is on average 8%; that is, out of 100 kilos of fabric that enter a factory, on average only 92 leave the factory in the form of garments. This metric varies significantly across factories, ranging between 1.6% and 19.2%. The authors find that most of this dispersion originates in the inspection and cutting stage, namely in differences in the buy-to-cut ratios (see Appendix Table A2).

Using data on reject rates and other defects from Macchiavello et al. (2015), we find that there is also variation in the sewing and finishing stage, namely in the cut-to-ship ratios. We examine a sub-sample of 6,000 orders that are included in the daily production records from the 51 factories (1,344 sewing lines) studied in Macchiavello et al. (2015). The reject rates at the final inspection point on the sewing line vary from 0% to 5% across these orders. This figure however is only a lower bound for the actual dispersion in the cut-to-ship ratios. One reason is that, while rejections lead to a complete waste of the garment's fabric, there is also partial fabric waste at this stage due to defects. A piece of garment that passes the end-of-line quality control may have required fabric-wasting corrections or alterations at intermediate points in the sewing process. A second reason is that data on end-of-line inspection points are available for relatively better managed factories, which tend to have inspection points alongside the sewing lines. Other factories instead only inspect quality in the finishing section, and given related observations in Tanvir and Mahmood (2014), we may expect these factories to exhibit even higher dispersion in wasted fabric.

Input substitutability. To reduce production costs, garment manufacturers have flexibility to substitute, to a certain extent, between fabric and other inputs. This can also contribute to the variation in buy-to-ship ratio observed in the data. An increase in the price of fabric incentivizes manufacturers to adopt fabric-saving practices, whereas an increase in the wage of sewing line operators incentivizes them to cut worker hours. Table A3 in the Appendix provides supportive evidence on the presence of this input substitutability. We relate the (logged) amount of fabric imported at the order level to two exogenous sources of variation in input prices (while controlling for order size, firm-product fixed effects, and a time trend). First, we study the effects of changes in the international price of cotton, noting that cotton is the most common material found in fabrics used for garment production in Bangladesh. Appendix Table A3 shows that, as anticipated, higher cotton prices translate into lower import volumes of fabric to produce orders of a given size. Second, we consider the effects of a policy that significantly increased the minimum wage in Bangladesh in November 2010. Also as expected, we find that the increase in labor costs translated into higher import volumes of fabric. While these correlations should be interpreted with caution, the results are in line with accounts of the industry.

In sum, we find that a model of garment production should accommodate three important characteristics: (1) a production process that operates at the order level, (2) fabric efficiency that may differ across orders, and (3) a technology that allows for substitution across inputs.

Our framework in the next section incorporates these characteristics into a Cobb-Douglas technology that transforms material fabric inputs and labor into garments. We address (1) by specifying this production function (output, inputs, and parameters) at the order level. We address (2) by allowing for a productivity shock at the order level. Finally, we address (3) as the Cobb-Douglas technology allows for substitution between inputs, and we limit this substitution by including a capacity constraint that accounts for the possibility of fixed or quasi-fixed factors of production.

Naturally, there are different formulations that could be used to address (1)-(3). The Cobb-Douglas formulation has at least two other features that we find appealing. The first one is that a Cobb-Douglas function emerges from the aggregation of production stages when the stages' technology either is also Cobb-Douglas or is Leontief with Pareto distributed technical coefficients (see, e.g., Houthakker, 1955 and Jones, 2005). This is particularly fitting for the garment sector that we study, where, as we have documented, the production process involves a number of sequential stages. In fact, as a microfoundation, if this process is partitioned into sufficiently small activities, then the technology for each such activity could be represented as a Leontief technology with fixed proportions, and as such our Cobb-Douglas function may be suitable to model the aggregate process. This aggregate process is the one that we observe in our data.[18]

The second convenient feature of the Cobb-Douglas function relates to our estimation strategy. As we explain in the next section, to estimate markups and marginal costs, we need to first estimate output elasticities. The Cobb-Douglas production functional form assumes constant output elasticities and thus allows us to perform this estimation even though we only observe fabric use. For example, as we do not observe capital and labor used on each

[18]More broadly, a more flexible production function, such as a translog, can be approximated to a first-order with a Cobb-Douglas function.

order, it would not be possible to pursue our estimation strategy under a more flexible production function like the translog. It is however important to stress that we need the elasticity of output to fabric only to compute the levels of markups and to explore patterns across firms. Our main results, which focus on exploring difference in markups across buyers within seller-product-time combinations, *do not* rely on the exact measurement of the output elasticities and are consistent with very flexible production functions in which the output elasticity varies at the seller-product-year level.

3 Estimation of Markups and Marginal Costs

This section describes our framework, estimation strategy, and estimation results. We begin by presenting in Section 3.1 a parsimonious model of garment production. This model captures the main aspects of garment production described in Section 2 and, building on Hall (1988) and De Loecker et al. (2016), allows us to construct sufficient statistics for order-level markups and marginal costs. In Section 3.2, we derive a structural equation that relates the order-level buy-to-ship ratio to fabric price and order size. We propose an IV approach to estimate this equation and recover the output elasticity of fabric used to construct order-level markups and marginal costs. Finally, in Section 3.3, we present our estimation results and characterize the distribution of prices, markups, and marginal costs.

3.1 Theoretical Framework

Setup. We model trade between buyers indexed by b and sellers indexed by s. Figure 1 describes the timing of events for any period t. First, buyers b and sellers s form links. Second, each seller s chooses fixed factors (production capacity), denoted by $\overline{L}_{st}$. Third, each buyer's demand is realized and buyers place product orders. We impose no restrictions on the mechanism via which orders are allocated to sellers. Finally, each seller s produces the orders it received and delivers them to the respective buyers. We index products by j and orders by o, and we denote the set of orders placed to seller s in period t (by all buyers and in all products) by O_{st}. Note that order o is seller-buyer-product-time specific (i.e., $sbjt$ specific); we omit these indices to ease the exposition. Each order specifies a volume Q_o, a unit output price P_o, and the material input (fabric) to be used in production, as well as designs and other contract terms as explained in Section 2.

To produce an order o, a seller combines labor L_o and material fabric inputs M_o with order-specific productivity ω_o. These are combined using the following order-level Cobb-

Douglas technology:

$$Q_o = L_o^{\beta_{sjt}} M_o^{\theta_{sjt}} e^{\omega_o}, \tag{1}$$

where β_{sjt} and θ_{sjt} are output elasticities with respect to labor and fabric, respectively, and the framework allows for these to vary flexibly across seller-product-time combinations. We allow for the possibility that, at the time of choosing the fabric inputs M_o, the seller may face a capacity constraint in labor given by the chosen capacity $\overline{L}_{st}$. Specifically, seller s chooses how much labor to use in each order $o \in O_{st}$ subject to the capacity constraint

$$\overline{L}_{st} = \sum_{o \in O_{st}} L_o, \tag{2}$$

where observe that summing over orders $o \in O_{st}$ is equivalent to summing over buyers, products, and orders for seller s in period t.

Denote the wages and fabric prices corresponding to order o by W_o and P_o^M respectively. We make two assumptions. First, we assume that wages can vary by product, time period, and seller, but not across orders or buyers for the same product-time-seller combination (i.e., we assume $W_o = W_{sjt}$). While this is still a restriction, note that this assumption significantly relaxes assumptions commonly made in the literature. Second, we assume that fabric prices are invariant to the size of the order. Importantly, this is only a restriction at the order level. Unlike most of the literature, we allow fabric prices to vary at the seller-product-time level, and thus to be correlated with determinants like upstream market power which may vary at such level.[19]

Seller's problem. Seller s in period t chooses $\{L_o, M_o\}_{o \in O_{st}}$ to minimize costs, subject to order-specific technology constraint (1) and capacity constraint (2), and taking order characteristics and prices as given. The Lagrangian for the seller's problem is

$$\mathcal{L}_{st} = \sum_o \Big(W_{sjt} L_o + P_o^M M_o \Big) + \sum_o \lambda_o \Big(Q_o - L_o^{\beta_{sjt}} M_o^{\theta_{sjt}} e^{\omega_o} \Big) + \lambda_{st}^L \Big(\overline{L}_{st} - \sum_o L_o \Big),$$

where λ_o and λ_{st}^L are the Lagrange multipliers associated with the technology constraints (1) and the capacity constraint (2), respectively. By standard logic, the order-specific multipliers λ_o represent the increase in total cost associated with producing one additional unit of output in order o. That is, λ_o represents the marginal cost for order o.

[19] It is possible to extend the model to multiple production factors without altering the structural equation derived below. That is, sellers in our model could be allowed to choose different bundles of operators, supervisors, and machines across different products, provided that the prices of these inputs vary at the seller-product-time level only.

The multiplier λ_{st}^{L} reflects the value of relaxing the capacity constraint for the seller. Having an extra unit of labor to be allocated across orders would allow the seller to reduce fabric input use and thus costs. Seller s can only adjust labor across orders in period t subject to not exceeding $\overline{L}_{st}$. Naturally, our analysis also applies if labor can instead be adjusted freely; in this case the multiplier λ_{st}^{L} is equal to zero.

The order-specific first order conditions with respect to labor L_o and fabric M_o can be written as

$$L_o = \frac{\beta_{sjt}}{\widetilde{W}_{sjt}} Q_o \lambda_o \tag{3}$$

and

$$M_o = \frac{\theta_{sjt}}{P_o^M} Q_o \lambda_o, \tag{4}$$

with $\widetilde{W}_{sjt} \equiv W_{sjt} + \lambda_{st}^{L}$. Note that orders $o \in O_{st}$ are interrelated only via the capacity constraint, and this is captured by the corresponding Lagrange multiplier λ_{st}^{L} which appears in the first order condition with respect to labor.

If the output elasticities θ_{sjt} and β_{sjt} in (1) were observable, we could directly calculate each order's marginal cost λ_o using the first order conditions above since Q_o and, critically, M_o and P_o^M are directly observed in our data for each order. Moreover, with a measure for the marginal cost, we could also compute each order's markup factor μ_o as the ratio between the order price P_o and the marginal cost λ_o:

$$\mu_o \equiv \frac{P_o}{\lambda_o}.$$

However, the input-output elasticity θ_{sjt} is not observable. That is, the data allow to directly measure deviations of markups and marginal costs from their seller-product-period average, but do not allow to directly compute the levels of markups and marginal costs. To do so, we need to estimate θ_{sjt}. We turn to this next.

3.2 Estimation

We combine equations (1), (3), and (4) to solve for M_o/Q_o. Taking logs, we obtain a structural equation that relates an order-specific buy-to-ship ratio to the order's size, the price of fabric used to produce the order, and two additional terms:

$$\ln \frac{M_o}{Q_o} = \frac{1 - \beta_{sjt} - \theta_{sjt}}{\beta_{sjt} + \theta_{sjt}} \ln Q_o - \frac{\beta_{sjt}}{\beta_{sjt} + \theta_{sjt}} \ln P_o^M + \frac{\beta_{sjt}}{\beta_{sjt} + \theta_{sjt}} \ln \left(\frac{\theta_{sjt} \widetilde{W}_{sjt}}{\beta_{sjt}} \right) - \frac{1}{\beta_{sjt} + \theta_{sjt}} \omega_o. \tag{5}$$

In principle, the framework allows for flexible production function parameters θ_{sjt} and β_{sjt}. In practice, in estimating (5) we are constrained by the amount of variation in the data and we obtain more precise estimates when these parameters only vary across product groups. We first consider a specification in which the production function parameters are restricted to be the same across products since a specification in which they are allowed to vary across the two broad apparel categories (shirts and trousers) reveals nearly identical estimates. In what follows we ease exposition by setting $\theta_{sjt} = \theta$ and $\beta_{sjt} = \beta$, but note that the estimation relaxes this constraint to have $\theta_{sjt} = \theta_\iota$ and $\beta_{sjt} = \beta_\iota$, $\iota \in \{shirts, trousers\}$. Let $\gamma_1 \equiv \frac{1-\beta-\theta}{\beta+\theta}$, $\gamma_2 \equiv -\frac{\beta}{\beta+\theta}$, $\delta_{sjt} \equiv -\gamma_2 \ln(\theta \widetilde{W}_{sjt}/\beta)$, and $\varepsilon_o \equiv -\omega_o/(\beta+\theta) + \nu_o$, where ν_o is an econometric error. Allowing for this error and simplifying terms in (5) using this notation yields the following estimating equation:

$$\ln \frac{M_o}{Q_o} = \gamma_1 \; q_o + \gamma_2 \; p_o^M + \delta_{sjt} + \varepsilon_o, \tag{6}$$

where lowercase letters denote logged variables.

The dependent variable in (6) is the buy-to-ship ratio at the order level, which is directly observed in our data. The first two explanatory variables on the right-hand side are also observable; these are the order size q_o and the price of fabric p_o^M. Instead, the third explanatory variable, δ_{sjt}, is not observable in our data. This term is a function of the wage W_{sjt}, which is common across orders for a given seller-product-time combination, and the Lagrange multiplier λ_{st}^L associated with the capacity constraint, which varies at the seller-time level. In particular, the Lagrange multiplier λ_{st}^L, is a sufficient statistic capturing interdependence in input choices across orders arising from prices and capacity constraints. Our main departure from the existing literature is that we can flexibly control for δ_{sjt} by including seller-product-time (i.e., sjt) fixed effects: while we lack information on labor and capital, our order-level data allows us to circumvent this challenge by exploiting the structural equation of order-level buy-to-ship ratios. The sjt fixed effects then control for both the interdependence across orders as well as for unobservable factors and productivity shocks affecting buy-to-ship ratios and common across orders at the sjt level. Finally, the fourth explanatory variable on the right-hand side of (6) includes an order-specific productivity shock, ω_o, which is not observable.

Estimating equation (6) allows us to construct all our variables of interest. Specifically, once we obtain the estimated coefficients $\widehat{\gamma}_1$ and $\widehat{\gamma}_2$, we compute the estimated elasticities $\widehat{\theta} = (1+\widehat{\gamma}_2)/(1+\widehat{\gamma}_1)$ and $\widehat{\beta} = -\widehat{\gamma}_2/(1+\widehat{\gamma}_1)$. We then combine $\widehat{\theta}$ with observable prices and quantities to obtain estimated marginal costs and markups at the order level, $\widehat{\lambda}_o = P_o^M M_o/(\widehat{\theta} Q_o)$ and $\widehat{\mu}_o = P_o/\widehat{\lambda}_o$. We next discuss the approach that we use to estimate equation (6).

Instrumental variable approach. The estimation of equation (6) poses two challenges. First, since quantities q_o are obtained from customs records, measurement error is likely present in our data. Measurement error would bias our estimate of $\gamma_1 = \frac{1-\beta-\theta}{\beta+\theta}$ towards zero, thus yielding $\beta+\theta = 1$ even when the production technology does not exhibit constant returns to scale. Second, we derived equation (6) under the assumption that productivity and the shadow price of labor are captured by a seller-product-year-specific shifter of the buy-to-ship ratio. Systematic deviations of productivity or the underlying production constraints that are correlated with volumes would bias our estimate of γ_1 and, with it, that of θ. In particular, misspecified productivity would overstate the scale coefficient and bias downwards our estimates of marginal costs. To address these two challenges, we instrument the size of the order q_o in our estimation.

Our IV strategy leverages the observed network of trade partnerships. The idea is that buyers cannot adjust their orders in response to shocks that are realized after orders have been allocated to sellers. Buyers take into account any information they have on demand and seller-product-year characteristics when placing their orders, but they cannot respond to ex-post production shocks (e.g., unexpected issues on the sewing line) that occur after orders have been assigned and production decisions have been made. This assumption does not appear to be too restrictive in light of the actual timing of events in the negotiation, production and delivery of a typical order.

Specifically, consider an example in which buyer b places an order with seller s, where we denote the order size by q_{sb}. Suppose that b also sources from another seller, s', who in turn sells to another buyer, b'. Importantly, in this example, b' is not a trade partner of s. We thus use the volume traded between s' and b', call it $q_{s'b'}$, as an instrument for q_{sb}. The argument is as follows. If b' receives a positive demand shock in its domestic market at the time of allocating orders, then it will order a large volume $q_{s'b'}$ from seller s'. Under capacity constraints, this means that seller s' will not be able to accept large volumes from buyer b, who, as a result, will tend to allocate a larger volume to seller s. To understand the exclusion restriction, note that since orders are allocated before production shocks occur, $q_{s'b'}$ is not a function of ω_o (or, in our example, ω_{sb}), the order-specific shocks that s faces in the production of the order for buyer b.

More generally, take an order o of size q_o placed by buyer b with seller s in quarter τ. We identify the sellers other than s who trade with b, and we use as an instrument for q_o the volume that these sellers trade in quarter τ with buyers other than b who are not trading with s. That is, for any firm (buyer or seller) i, denote by $\mathcal{N}_i$ the set of i's trade partners in quarter τ, and let $\mathcal{N}_i \backslash \{k\}$ be this set excluding partner k. Then the instrument for q_0 is the log of:

$$z_{sb\tau} = \frac{1}{\#\{\mathcal{N}_b \backslash \{s\}\}} \sum_{m \in \mathcal{N}_b \backslash \{s\}} \frac{1}{\#\{\mathcal{N}_m \backslash \mathcal{N}_s\}} \sum_{n \in \mathcal{N}_m \backslash \mathcal{N}_s} Q_{mn\tau},$$

where $\#\{\cdot\}$ is the cardinality of the set in the argument.

Two remarks on our empirical strategy. First, note that while the instrumented regressor q_o is an order-level variable, the instrument is at the seller-buyer-quarter-level. This higher level of aggregation is needed due to sparsity in our data but has almost no impact on our estimation in practice since our sample is dominated by buyer-seller-quarter triplets with unique orders. Second, we do not instrument for fabric price P_o^M. Note that fixed effects δ_{sjt} control for several sources of input price endogeneity, e.g., supplier's market power in the fabric's market. A concern more specific to our context, however, is that the international buyer directly negotiates the fabric price with the foreign fabric supplier. To overcome this concern we estimate a version of our IV specification including buyer fixed effects and find nearly identical results.

3.3 Estimation Results

The estimation strategy described above delivers estimates of the production function coefficients, which we then use to compute marginal costs and markups. These are recovered at the disaggregated level of the order and then combined as a weighted average for each seller-(buyer)-product-year combination.[20]

Production function estimates. Our estimates of the output elasticities, $\widehat{\beta}$ and $\widehat{\theta}$, are presented in Table 2. Panel A reports the results from a specification that restricts the output elasticities to be the same across product groups. The left-hand side of the panel shows the OLS results and the right-hand side shows the results from the IV estimation.[21] The latter yields estimated elasticities of 0.62 and 0.36 for fabric and labor respectively, adding up to essentially constant returns to scale. These estimates are lower than those

[20]It is important to note at this point that the estimation procedure described in Section 3.2 is necessary for recovering marginal costs and markups in levels. The analysis in Section 4 studies the outcomes of interest, markups and marginal costs, removing all seller-product-year-specific variation. In practice, this implies that we study *deviations* of the outcomes from seller-product-year means. In our setup, there is no variation in the elasticity of material inputs to output, θ, beyond those triplets (in fact, beyond the product-group level). With this, the estimation of that elasticity is relevant for the discussion that follows in this Section (3.3), but carries no bearing on the results we present in Section 4. As noted above, the main results in Section 4 are consistent with output elasticities to fabric varying at the seller-product-year level.

[21]The estimated coefficient for the first stage of the IV regression (not reported in Table 2) is 0.096, with a standard error of 0.012. The F-test statistic value is 175.6, indicating a strong IV.

obtained from the OLS specification, implying that the effects of quantities on marginal costs are overstated when the order size is not instrumented for. The IV estimates that we obtain are remarkably consistent with industry reports and costing sheets, which show that fabric represents roughly two thirds of variable unit costs in garment production.

The buy-to-ship estimating equation, (6), accommodates fabric prices that are correlated with unobservable characteristics of the seller-product-year combination or the destination of the export order. We study whether the *deviations* of those input prices from their fixed effect means are driven by buyer-specific attributes. Such correlation may arise if buyers of garment have market power upstream or if order-level quality choices are endogenous to the buyer. We thus expand our baseline IV specification to include buyer fixed effects. This exercise reveals an estimate of γ_2, the coefficient on the price of fabric in the buy-to-ship equation, of -0.42, close to the -0.37 of the baseline IV. Our parameter of interest, the elasticity of output with respect to fabric, is estimated to be 0.59, well within the confidence interval of our preferred specification.[22]

Panel B of Table 2 reports the results from a specification that allows the output elasticities to vary across two product groups, shirts and trousers. In the case of fabric, we find that the elasticities for shirts and trousers are not statistically significantly different from each other. In the case of labor, the elasticity is lower for shirts than for trousers. According to conversations we had with firm managers in Bangladesh, this difference reflects that, relative to trousers, shirts are simpler products, often produced in sewing lines using fewer workers per unit of fabric.

Marginal costs and markup estimates. Table 3 presents our estimates of the order-level marginal costs and markups, $\widehat{\lambda}_o$ and $\widehat{\mu}_o$. The table show that, on average, the price per kilo of garment paid by buyers is \$13.64. This average price is composed by \$3.26 of markup and \$10.38 of marginal cost, where the latter is in turn composed by \$7.58 of fabric and \$2.80 of labor and other costs. The implied average markup factor is 1.42. This estimate is in line with the findings of De Loecker et al. (2016), who report mean and median (seller-product) markup factors of 1.57 and 1.33 for the textiles and apparel sector in India.[23] Table 3 shows

[22]As an additional robustness exercise, we augment estimating equation (6) to include quarter fixed effects (results not reported for brevity). These would absorb any variation in buy-to-ship ratios common to all orders occurring at the same point in time. For example, seasonal demand patterns that correlate with order-level fabric prices would affect the estimation of the parameter of interest. The IV estimation of this augmented specification renders an elasticity of output to materials of 0.61, very close to our baseline estimate.

[23]Our estimates are also in line with annual reports available from sellers. For instance, Generation Next Fashions Ltd. and Beximco, both large Bangladeshi manufacturers of garments, report gross profit margins of 33 and 45% respectively in 2012. These margins are highly correlated with firm-wide measures of markups and are in the same range as the markups reported in Table 3.

that both markups and marginal costs exhibit significant dispersion. At one end, a small share of the orders are sold at a loss for the seller, while at the other end, the most profitable orders command a markup factor greater than 2. We find that order-level markup values are more dispersed than order-level marginal costs: the interquartile ratio is 5.33 for markups and 1.76 for marginal costs.

Firm-level patterns. While our main focus is on exploring within-seller variation in markups (charged to different buyers for the same product in the same year), it is useful to consider more aggregate patterns that can be compared with the findings of the literature. To this end, in Appendix B we aggregate order-level outcomes at the seller-product-year level. We find that at this level, (1) markups are more dispersed than marginal costs, as in Atkin et al. (2015); (2) exported quantities are negatively correlated with marginal costs and positively correlated with markups, in line with the results of De Loecker et al. (2016) for India and Atkin et al. (2015) for the soccer ball sector in Sialkot, Pakistan; and (3) core products of multi-product firms exhibit lower marginal costs and higher markups than other products of these firms, consistent with the core product hypothesis discussed in Mayer et al. (2014). We relegate the details to Appendix B.

4 Markups, Costs, and International Buyers

This section starts by exploring the extent to which a given seller obtains different markups from different buyers. It then studies whether such differences are systematically correlated with key characteristics of the buyers. Our focus of inquiry is motivated by the observation that the apparel industry is the prototypical 'buyer-driven' global value chain (see, e.g., Gereffi, 1999 and Gereffi et al., 2005). In these chains, on the one hand, large buyers may be able to exert market power in their sourcing countries and squeeze sellers' margins. On the other hand, buyers may adopt collaborative sourcing strategies that increase the profits of their suppliers and facilitate upgrading. The extent to which one or the other force dominates determines the benefits that developing countries' suppliers extract from trading with large international buyers. Our analysis then focuses on the relationship between sellers' performances and the *sourcing strategies* adopted by their international buyers.

After documenting the extent of markup dispersion within seller-product-year combinations, and after describing the relevant dimensions of heterogeneity across buyers, this section presents our main result: buyers adopting relational sourcing strategies pay higher markups to their suppliers. This statement reflects the conditional correlation between buyers' sourcing practices and sellers' markups, when such conditioning addresses the causes of dispersion commonly studied in the literature. In particular, we exploit variation only within seller-product-year triplets, therefore absorbing differences in markups that arise from a seller's specific productivity, scale, competence or quality. We further condition on the country of destination and size of the buyer, as well as the volumes traded in a specific relationship. The section concludes discussing a battery of robustness checks and alternative interpretations to our main findings.

4.1 Dispersion in Markups and International Buyers

Here we document two novel facts: *first*, there is large dispersion in markups within seller-product-year, and *second*, buyer effects account for a large share of the observed dispersion. Before we document these facts, we note that there is adequate variation in the data. As noted above, the average seller exports 3.25 products and trades with 5.93 buyers in a typical year. Table 1 also provides descriptive statistics of buyers and buyer-seller pairs. Buyers are even more diversified than sellers both in terms of products (4.25) and partners (21.94) in a given year (Panel C). Our final sample contains trade interactions in 5,633 buyer-seller pairs.

Figure 2 shows that the within-seller dispersion in markups (across buyers) is similar in magnitude to the dispersion across sellers. More specifically, the figure aggregates order-level markup factors for each seller-buyer-product-year combination. After residualizing these markups against product-year fixed effects, we construct the simple average, 25$^{\text{th}}$ and 75$^{\text{th}}$ percentile residual markup for each seller. The horizontal axis arranges sellers ascendingly in percentiles according to their average markup. Across the full range of sellers, the within-seller interquartile range is everywhere wide. Moreover, the average interquartile range collecting within-seller dispersion in markups is similar in magnitude to the interquartile range observed across sellers.

Table 4 shows the share of the observed dispersion in markups that is accounted for by buyer-specific characteristics. We estimate

$$y_{sbjy} = \delta_b + \delta_{sjy} + \varepsilon_{sbjy}$$

where the δ_b term collects buyer fixed effects. The specification absorbs seller-product-year variation in the term δ_{sjy}, which features as a baseline control for the analysis in the reminder of this section. We report results on different outcomes, y_{sbjy}, all in logs: traded volumes (q_{sbjy}), weighted averages of garment prices (p_{sbjy}), fabric prices (p^f_{sbjy}), marginal costs (mc_{sbjy}), and markup factors (μ_{sbjy}). The table shows the percentage of the variation in the corresponding outcome (along rows) accounted for by buyer and by seller-product-year fixed effects (along columns). Buyer effects capture about one third of the total explained variation in markups and marginal costs.[24] It is important to note that the prominence of this buyer effect is not down to the country of destination: an alternative specification that replaces the buyer fixed effects with destination fixed effects shows that the country explains between 5% and 6.7% of the observed variation in the outcomes.[25]

In sum, this variance decomposition exercise reveals that buyer-specific effects account for a very substantial share of the within-seller dispersion in markups and costs.[26] We now turn to a characterization of the international buyers in our sample, to uncover the dimensions of the buyer margin that induce such dispersion in markups and costs.

4.2 Buyer Characteristics

International buyers sourcing garments in Bangladesh are an extremely heterogeneous group ranging from dedicated apparel brands (e.g., The Gap), to non-specialized mass retailers (e.g., Walmart) and upscale branded marketers (e.g. Tommy Hilfiger). Although we observe more than 1,500 active buyers in the sample (see Panel C of Table 1 for summary statistics) the distribution of buyers' size is highly skewed with the largest 100 buyers accounting for 66% of the traded volumes. Table 5 provides a detailed look at the 25 largest buyers in the product groups of interest. The table ranks buyers according to their (upstream) market shares in Bangladesh, reported in Column (1). H&M, Walmart, and VF Corporation - a multi-brand American apparel company - lead the board with market shares of 5.2%, 4.9% and 4.2% respectively, more than 500 times larger than the median buyer in the sample.

Column (2) shows that buyers greatly differ with respect to the number of suppliers they source from. This is true also for buyers of similar size. For example, while Levi Strauss & Co. and J.C. Penney have similar market shares (2.26% and 2.00% respectively) in a typical year the former only sources from 7.5 suppliers whilst the latter does so from more than 25 sellers. This difference reflects radically different approaches to sourcing from these two companies. During an interview conducted by one of the authors with a sourcing director

[24] These are obtained as $17.9\%/(17.9\% + 34.8\%)$ and $20.2\%/(20.2\% + 44.8\%)$ for markups and marginal costs, respectively.

[25] The sample includes 55 different destinations. The figures corresponds to $2.5\%/(2.5\% + 34.8\%)$ for markups and $2.4\%/(2.4\% + 44.8\%)$ for marginal costs.

[26] Table B2 in the Appendix reports results from a specification at the seller-product-year level in which a dummy for the seller's main buyer is included. The main buyer effect accounts for approximately 75% (30%) of the variation in markups explained by seller (seller-product) fixed effects. Note that while the structure of the decomposition is as in the matched employee-employer literature (see, e.g., Card et al. (2018) for a recent example), the interpretation is different since, unlike workers, sellers and buyers trade with multiple partners in any given product-year combination (see Table 1 for summary statistics). Under this light, the variability in the average markup of the seller-product-year that is accounted for by one buyer only is very large.

for Levi Strauss & Co., it was reported that the company's origin as a manufacturer created a corporate culture centered around production capabilities. When the company started outsourcing production to foreign suppliers, it retained that focus by creating very strong partnerships with a limited number of suppliers. In exchange for loyalty and compliance, Levi Strauss & Co. transfers production capabilities to core suppliers (e.g., introducing new fabric material, assisting with planning and industrial engineering). In contrast, J.C. Penney has traditionally been closer to a strategy of "squeezing cost out of the supply chain" (see, *Sourcing Journal*, January 11$^{\text{th}}$, 2013) and during our sample's years, "decimated [their] sourcing department and trampled on trusted relationships established in foreign countries" under the leadership of Ron Johnson (see, e.g., *Forbes*, April 25$^{\text{th}}$, 2014).

The difference between Levi Strauss & Co. and J.C. Penney reflects a broader distinction between two polar sourcing models in the apparel industry: *spot* interactions at one end and *relational* sourcing at the other (see, e.g., De Toni and Nassimbeni, 2000; Taylor and Wiggins, 1997; McMillan, 1990).[27] Under *spot* procurement, suppliers are selected based on short-run cost minimization criteria exclusively: buyers source from multiple suppliers, with whom trade relationships tend to be short lived and ended by out-bids from cheaper suppliers. Procurement orders tend to be large and either one-off or sporadic. In contrast, under *relational* sourcing buyers concentrate orders on a small number of suppliers on which they rely for the on-time delivery of shipments of consistent quality. Under this model, buyers tend to engage with the supplier's production practices to foster the sellers' capabilities, develop adequate customization and guarantee synchrony with the buyer's just-in-time requirements. Price premia and the longer horizon of the relationship provide incentives for specific investments and mitigate the risks of opportunistic behavior.[28]

Column (3) in Table 5 explores differences in sourcing strategies across buyers. We capture a buyer's sourcing strategy through a measure of relational procurement along the lines of Heise et al. (2017). We normalize the number of sellers the buyer trades with, by the number of shipments the buyer receives in each product-year. We then construct a weighted

[27]Due to complementarities (Milgrom and Roberts, 1990, Antràs et al., 2017), the choice of procurement system is to a large extent a buyer-level decision. Even buyers operating in the same segment of the same destination market tend to source identical garments adopting different procurement strategies. We take these differences as given, without attempting an explanation for the buyers' idiosyncratic choices of sourcing strategies. Note also that vertical integration between international buyers and manufacturers is virtually inexistent in the industry we study.

[28]These two models are sometimes referred to in the literature as 'adversarial' or 'American-style' sourcing in contrast to 'collaborative' or 'Japanese-style' sourcing (see, e.g., Kawasaki and McMillan, 1987; Richardson, 1993; and Helper and Saki, 1997). The characterization is well documented in the analysis of governance structures in global value chains (Gereffi, 1999). More recently, formalizations of the buyer-supplier contracting problem, as in Taylor and Wiggins (1997), have been incorporated into standard trade models (Heise et al., 2017; Defever et al., 2016; Startz, 2018).

average for each buyer across all its product-year combinations. This metric reflects that buyers reliant on spot sourcing tend to spread out their shipments across multiple suppliers, while relational buyers would concentrate them in a set of core suppliers. Column (3) in the table reports the buyer's ranking with respect to this metric of relational sourcing, with 1 collecting the most relational buyer in the data. The resulting ranking maps closely to well known industry practices. For example, Levi Strauss & Co. ranks 3$^{\text{rd}}$, close to other large buyers known for their relational behavior, such as The Gap and H&M, ranked 1$^{\text{st}}$ and 5$^{\text{th}}$ respectively. At the other extreme, large German discount retailers Kik Textilen and JCK (G. Gueldenpfennig GmbH) clearly pursue a sourcing strategy based on spot procurement.[29]

Finally, Columns (4) and (5) aggregate order-level marginal costs and markups at the buyer level. Among the top 25 buyers, average marginal costs vary substantially around the industry median ($9.22), from $7.33 and $7.65 per kilo for discount retailers (Asda and Kik Textilen respectively) to $11.86 per kilo for H&M. Among the same set of buyers, markup factors also vary substantially around the industry median (1.38), ranging from 1.28 and 1.29 for JCK and J.C. Penney, to 1.70 and above for brands such as The Gap and C&A.

Appendix B explores differences across buyers more systematically. We document several novel facts. Figure B3 shows that buyers' average markups are even more dispersed than sellers' average markups. Figure B4 shows that markups (marginal costs) are positively (negatively) related to buyer's size. Finally, Table B4 provides evidence of core-product effects for buyers as well: within a buyer-year combination, core products for the buyer have lower (higher) marginal costs (markups). While these facts are analogous to those documented for sellers in the previous section, they are difficult to interpret. For example, the correlations of markups and costs with the buyer's size could simply reflect sorting patterns between sellers and buyers (e.g., only large sellers with lower costs have sufficient capacity to supply larger buyers). We thus conduct our main analysis with specifications in which we study differences in markups across buyers, conditional on seller-product-year effects.

[29] Appendix Table B3 describes the main characteristics of relational buyers. Relational buyers tend to be located in larger, more distant, richer markets (Panel A). On average, they import higher volumes and, conditional on these volumes, have fewer suppliers, with their largest and median seller being larger (Panel B). Also, conditioning on the buyer's overall volume of trade, more relational buyers tend to be more specialized (import a lower number of products). In addition, they place fewer orders, but receive significantly more shipments in any given year. Comparing across different export orders, for a given seller-product-year combination, those placed by relational buyers are larger in overall size relative to orders placed by spot buyers of comparable size. They also consist of more, smaller shipments, on average (Panel C).

4.3 Buyers' Sourcing Strategy and Markups

4.3.1 Main Results

We now investigate how a buyer's sourcing strategy correlates with the outcomes of interest. We estimate the following baseline specification

$$y_{sbjy} = \delta_{sjy} + \delta_d + \boldsymbol{\beta}' X_b + \varepsilon_{sbjy}$$

where the notation is as before: s stands for seller, b for buyer, j for product (six-digit HS code), d for destination and y for year. We focus on (log) prices, marginal costs and markup factors as dependent variables, i.e., $y_{sbjy} \in \{p_{sbjy}, mc_{sbjy}, \mu_{sbjy}\}$. The fixed effects δ_{sjy} absorb seller-product-year specific variation and allows us to study differences across buyers within sellers. The term X_b represents a vector of characteristics at the buyer level. The main explanatory variable in our analysis is one such characteristic: $Relational_b$, a dummy taking value of one if the buyer belongs to the top 10 percent of the distribution of the continuous relational characteristic described above. In all specifications we control for the buyer's size, defined as the overall volume the buyer imports across all woven products, trade partners and years in the data.[30] We condition on destination fixed effects, δ_d, to absorb differences explained by characteristics common to all buyers in a given country. We focus exclusively on buyer's cross-sectional variation rather than buyer-time or buyer-product variation, which would more likely reflect factors directly affecting the outcomes of interest. We would like to use measures of the buyer's overall size and procurement practices across its potentially many sourcing countries, to reduce concerns that a buyer's choice of size and sourcing strategy in Bangladesh are driven by factors correlating with the outcomes of interest. This information is, however, not available.[31] To assuage this concern we investigate robustness checks in which the buyer's sourcing characteristic is computed leaving-out observations for the products in the sample (HS codes corresponding to trousers and shirts). Those specifications, presented and discussed in Appendix C.2.2, yield nearly identical results.

Results are presented in Table 6. Columns (1) to (3) show that a buyer's relational sourcing strategy is associated with higher prices and markups and not statistically different

[30] In practice, we exclude the incumbent seller in the aggregation of volumes to construct the buyer size metric. We denote this variable q_{b-s}, which stands for the (log) volume that the buyer trades in the main woven categories, across all years, with all its partners except for s.

[31] For the buyer's size we have merged our data with external sources (Orbis) but those have yielded a relatively poor match. Constructing the measure of sourcing strategies for the buyers in the sample in other origin markets would require transaction level-data with buyer names and supplier's identifiers across several countries which are not available.

marginal costs. As the baseline specification conditions on the buyer's overall size and destination fixed effects, the result implies that relational buyers pay higher prices and markups relative to spot buyers of comparable size, competing in the same destination market.[32] As for the buyer's size, Appendix Table B5 shows that the overall volume imported by the buyer positively correlates with markups. The estimated relationship is however very small and entirely driven by the volume traded in the specific product between the buyer and the supplier. These correlations are consistent with either increasing returns to scale at the relationship-product level or with selection mechanisms (e.g., buyers source larger volumes from sellers with whom they have product-specific good matches). Regardless of the mechanism, we note that the results are not consistent with large buyers being able to squeeze the margins paid to their suppliers, conditional on selecting into these relationships.

We note that the relational sourcing dummy is mechanically correlated with sourcing volumes (see Table B3). Since larger traded volumes correlate with higher markups via lower marginal costs (see Table B5), Columns (4) to (6) in Table 6 add the bilateral trade volume as further control. When bilateral trade volumes are conditioned upon, we find that relational buyers pay 6.5% higher prices and 2.7% higher markups. Conditional on buyer's size and bilateral trade volume, then, a relational sourcing practice is also associated with *higher* marginal costs.

4.3.2 (Controlling for) Product Quality

A possible explanation for the finding that relational buyers pay higher markups for export orders with higher marginal costs is that they source, from the same supplier and within the same narrow product category, higher quality products (Atkin et al., 2015). A recent literature has linked the quality of output to the quality of inputs used in production (see, e.g., Verhoogen, 2008;Kugler and Verhoogen, 2012; Bastos et al., 2018). Section 3.3 showed that the price of the (main) input, fabric, constituted a large share of the marginal cost. We exploit the fact that we observe the price, variety and origin country of imported fabric on each export transaction to control for physical product quality. Columns (7) to (9) of Table 6 explore the robustness of the finding to the inclusion of fabric price as a control. This augmented specification sees a reduction by about one quarter to one third in the magnitude of the estimated coefficient on prices (Column (7)) and marginal costs (Column (8)). Even after controlling for input quality, relational buyers pay 2.8% higher markups. Consistent with the order-level specifications in Section 3.2, a 1% higher fabric price translates into a 0.6% higher marginal costs and into a 0.5% higher price, suggesting a relatively high

[32] Unreported exercises show qualitatively equivalent results in specifications that do not control for buyer's size and/or for destination fixed effects.

but incomplete price pass-through. The same 1% increase in fabric prices leads to a small decrease in markup. The pass-through is thus only marginally different across buyers with different sourcing strategies.[33]

Table 7 explores differences in quality with a battery of further robustness checks. Columns (1) to (3) augment the preferred specification in Table 6 to condition on product-*season*-year effects. Conditional on the baseline set of fixed effects (seller-product-year and destination) we also include dummies that control for the possibility that the buyer and seller trade the majority of their volume at a particular point in the fashion cycle. For concreteness, a category in this set of fixed effects would be *Men's shirts made of cotton produced in the Summer-Spring season of 2008*. Columns (4) to (6) include a large set of dummies collecting the type and origin of the fabric most used in the orders traded by the buyer and seller in the product category and year combination. That is, we control for the fabric variety, narrowly defined (e.g., *Woven fabrics of cotton, containing 85% or more by weight of cotton, plain weave, weighing not more than 100 g/m2 imported from China*). Finally, columns (7) to (9) control for the complexity of the traded garments, proxied with the number of different varieties used as inputs, where, again, a variety is defined by a type and origin of fabric. Across all these specifications, the baseline results on relational sourcing remain both qualitatively and quantitatively unchanged.

Besides the product's physical quality, the orders sourced by relational buyers from a given supplier and within a narrow product category might differ in other dimensions. Table C1 investigates two dimensions we observe in the data: specialization patterns and seasonality. With regards to seasonality, we document patterns consistent with buyers having some suppliers from whom they source all year long and 'sporadic' suppliers that supplement the core supply during specific seasons. Those core suppliers tend to receive (relatively) higher markups. Furthermore, we can show that the result on higher markups paid by relational buyers is not driven by patterns of product specialization nor by the differential product scope of relationships. In sum, our results suggests that product quality alone (quite broadly defined) does not explain why buyers adopting relational sourcing practices pay higher markups.

[33]Table C2 explores the robustness of the baseline results using alternative measures of a buyer's sourcing strategy. We consider various alternative measures: baseline, different cut-offs to define the dummy, the discrete and the continuous measure in Heise et al. (2017) (as close as we can replicate it), and a variant of our baseline measure constructed leaving out all products in the sample of interest (the 17 HS6 codes comprising shirts and trousers). Estimates for the relationship between the relational characteristic and markups vary between 0.013 and 0.029.

4.4 Mechanisms and Discussion

Our preferred interpretation for the main results is that the higher markups paid by relational buyers reflect the basic contractual difference between the two alternative procurement systems as modeled in, e.g., Taylor and Wiggins (1997). Under the relational system, the buyer might require suppliers to perform along other dimensions of differentiation, e.g., reliability and on-time delivery (see, e.g., Garvin, 1987). These dimensions require that the seller undertakes costly non-contractible investments. These higher costs are consistent with our finding that, conditional on the volume traded in the relationship and the physical quality of the product, relational buyers are associated with *higher* marginal costs.[34] In order for the supplier to engage in these undertakings, it must be the case that relationship termination following contractual non-performance is costly. The cost of contractual non-performance arises from the fact that the buyer pays rents to the supplier, in the form of higher markups.

Other explanations appear less plausible. For instance, it could be that relational buyers select into trading relationships that have better match-specific capabilities. This hypothesis is *prima facie* inconsistent with the evidence that orders supplied to relational buyers have higher marginal costs. Furthermore, the unfolding of relationships is consistent with the presence of dynamic incentives in the form of a promise of future trade. This is particularly so in the case of relational buyers, who then not only pay higher markups but also develop, over time, larger relationships. This makes dealing with such buyers, profitable. We note that approximately 12% of seller-buyer-year combinations exhibit negative profits (i.e., markups factors just-below-one). It turns out that 70% of cases with negative profits occur when the buyer-seller pair trade for the first time. That is, profits in a relationship tend to increase over time and relational buyers establish longer-lasting trade partnerships.[35]

Appendix Section C.2.1 explores differential within-relationship dynamics across buyer types. We consider specifications at the buyer-seller-year level that control for buyer-seller fixed effects and interact the relationship's age with the relational buyer dummy. Buyer-seller fixed-effects control for (time invariant) match-specific capabilities that could drive our main result. Conditional on relationship survival, we find that traded volumes grow in all relationships, but particularly so in those involving relational buyers. This steeper relationship growth reflects the stronger dynamic incentives that relational buyers can put in place to induce compliance of their suppliers.

[34]Note that buyers with relational sourcing practices might also require suppliers to have higher pay for workers, superior industrial relations and better environmental standard. While those factors could also be associated with higher costs, they vary at the seller-year level and are thus absorbed by the fixed effects.

[35]In the relevant sample the mean relationship age is 2.44 years for relational buyers and 1.35 years for non-relational ones.

Although at odds with our conversations with industry practitioners, an alternative explanation for our main finding is that higher markups simply reflect stronger competition between relational buyers to source from a limited number of suppliers with adequate capabilities. Given that relationships with relational buyers are more profitable, if this was the case we would expect the best suppliers to bid slightly lower markups and expand the volume they supply to relational buyers.

To see why, Table 8 explores the extent to which sellers earn higher profits when supplying relational buyers. We aggregate outcomes at the seller-buyer-year level (i.e., combine orders in a buyer-seller pair across different products categories) and estimate the following specification:

$$y_{sby} = \delta_{sy} + \delta_d + \beta_1 Relational_b + \boldsymbol{\gamma}' \boldsymbol{X}_{sby} + \varepsilon_{sby}$$

where the term δ_{sy} collects all seller-year specific variation, δ_d absorbs country level effects and $Relational_b$ corresponds to the baseline definition for relational buyers. The specification also controls for the buyer's overall size, traded volumes and, again, the price of fabric. The dependent variables are log of prices, costs and markups, p_{sby}, mc_{sby} and μ_{sby}, as well as volumes and profits, q_{sby} and π_{sby}.[36]

Columns (1) to (3) of Table 8 confirm the findings of Table 6: relational buyers pay higher prices resulting in markup factors 6.3% higher. To interpret the magnitude of the effect consider the average markup factor (1.43) and marginal costs (\$10.22) in the sample. For the average trade, then, relational buyers pay an additional \$0.82 per kilo of garment over an average (median) markup value of \$3.87 (\$3.38) per kilo. This represents a sizable increase in markups of approximately 21% (24%), amounting to \$103,280 (\$32,838) a year for a relationship of average (median) size.[37]

Column (4), however, shows that in addition to paying higher markups, relational buyers also trade 46% larger volumes. The combination of the markups and volumes effects implies that profits are about 75% higher with relational buyers (column (5)). Exporters thus make significantly higher profits when supplying relational buyers. Note that this relationship is identified using within seller-year variation only, and having conditioned on the country of destination and other relevant characteristics, including the buyer's size. The positive relationship between firm's profits and supplying relational buyers carries through at the

[36] Profits are constructed using observed traded volumes and estimated (weighted average) markup values. Naturally, the control for traded volumes, q_{sby}, is not included when the outcome is q_{sby} or π_{sby}.

[37] Note that the estimated coefficient on markups here is significantly larger the one obtained in the more disaggregated specification in Table 6. The difference between the two specifications arises from the fact that aggregating at the buyer-seller-year level gives relatively more weight to 'core products' for which marginal costs are lower (see, Tables B4 and B1) and particularly so for relational buyers, who tend to be more specialized, conditional on size (see Appendix C for details).

seller level: the higher the participation of relational buyers in the seller's overall volumes, the higher the seller's profits.[38]

These results are thus inconsistent with the hypothesis that higher markups simply reflect a thinner supply of suitable partners for relational buyers. Among those exporters that do supply relational buyers, most have a portfolio of partners comprising both relational and spot buyers. Given that the seller's overall profits are increasing in the share exported to relational buyers, sellers have strong incentives to lower markups and gain market share with relational buyers. In contrast, in our preferred interpretation, buyers would not accept to pay lower markups if those reflect rents that need to be paid to incentivize non-contractible investments: as in the standard efficiency wage model with unemployment, lower markups would fail to discipline the supplier.

In sum, while we cannot perfectly separate the 'treatment' effect from the 'selection' effect associated with supplying relational buyers, the evidence here provides sufficient ground to advance a preferred interpretation: the higher profits earned by a given exporter when supplying a relational buyer correspond to premia paid to incentivize the supplier to undertake non-contractible actions. Our results provide what is, to the best of our knowledge, the first direct evidence for the (equilibrium) relationship between markups and buyers' sourcing strategies, discussed in a vast qualitative literature and formalized in several models.

5 Conclusion

As production fragmentation deepens, the opportunities for manufacturers in developing countries to join global value chains and serve large international buyers multiply. The incidence that these buyers may have on suppliers' markups and marginal costs is unclear from an empirical standpoint. On the one hand, there is a concern that international buyers may exert significant market power, extracting most of the gains from trade from their suppliers. On the other hand, it has also been suggested, mostly based on anecdotal evidence, that international buyers adopting relational sourcing strategies may pay high markups and share the gains from trade with local suppliers. Testing these different views empirically has been difficult so far, as the markups in these trade transactions are unobservable.

This paper has contributed to this discussion with a methodological approach for recovering markups and marginal costs at the order level. We leveraged a unique dataset on

[38] Appendix Section C.2.2 explores specifications at the seller-year level. We find that the share of products exported to relational buyers positively correlates with prices, markups, volumes and thus profits at the seller level. The specifications control for the average size of the buyers the seller exports to, the (average) price of fabric used by the exporter and, when appropriate, for the seller's export volume. Taken together, this is evidence that when sellers trade more with relational buyers, they perform better.

transactions between international buyers and their suppliers of ready-made garments in Bangladesh, including records on the usage of material inputs for each export order. Using this data and an estimation approach grounded on detailed knowledge of the context, we were able to study how order-level markups and marginal costs vary with buyers' characteristics.

Our analysis reveals that exporters charge significantly different markups for the same product, at the same point in time, to different buyers. Buyers' sourcing strategies play a key role in generating these differences. While some buyers place their orders with the cheapest supplier available (based on some competitive bidding mechanism), other buyers engage in relational sourcing practices with their suppliers. We showed that whether a buyer adopts relational sourcing or not is associated with significant differences in markups and marginal costs within seller-product-year triplets, conditioning on destination and buyer size.

Our results constitute novel evidence of the importance of buyers' sourcing strategies. We studied other potential drivers, such as quality and competition for suppliers, and found that they do not explain the patterns in marginal costs and markups. Overall, our analysis supports the view that relational buyers offer relational rents, in the form of higher markups, to incentivize suppliers to undertake costly non-contractible actions. In turn, trading with international buyers that adopt relational sourcing practices may help producers in developing countries reap the benefits of participating in global value chains.

References

Antràs, P. (2015): *Global Production: Firms, Contracts, and Trade Structure*, Princeton University Press, 1st ed.

Antràs, P. and C. F. Foley (2015): "Poultry in Motion: A Study of International Trade Finance Practices," *Journal of Political Economy*, 123, 853–901.

Antràs, P., T. C. Fort, and F. Tintelnot (2017): "The Margins of Global Sourcing: Theory and Evidence from U.S. Firms," *American Economic Review*, 107, 2514–64.

Appelbaum, R. and G. Gereffi (1994): "Power and Profits in the Apparel Commodity Chain," in *Global Production: The Apparel Industry in the Pacific Rim*, ed. by E. Bonacich, L. Cheng, N. Chinchilla, N. Hamilton, and P. Ong, 43–62.

Atkeson, A. and A. Burstein (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 98, 1998–2031.

Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen (2015): "Markup and Cost Dispersion across Firms: Direct Evidence from Producer Surveys in Pakistan," *American Economic Review*, 105, 537–544.

Atkin, D. and D. Donaldson (2015): "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs," *NBER Working Paper 21439*.

Atkin, D., A. K. Khandelwal, and A. Osman (2017): "Exporting and Firm Performance: Evidence from a Randomized Experiment," *The Quarterly Journal of Economics*, 132, 551–615.

Baker, G., R. Gibbons, and K. J. Murphy (2002): "Relational Contracts and the Theory of the Firm," *The Quarterly Journal of Economics*, 117, 39–84.

Bastos, P., J. Silva, and E. Verhoogen (2018): "Export Destinations and Input Prices," *American Economic Review*, 108, 353–392.

Bernard, A. B., E. Dhyne, G. Magerman, K. Manova, and A. Moxnes (2018a): "The Origins of Firm Heterogeneity: A Production Network Approach," *mimeo*.

Bernard, A. B., A. Moxnes, and Y. U. Saito (forthcoming): "Production Networks, Geography and Firm Performance," *Journal of Political Economy*.

Bernard, A. B., A. Moxnes, and K. H. Ulltveit-Moe (2018b): "Two-Sided Heterogeneity and Trade," *Review of Economics and Statistics*, 100, 424–439.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

Blouin, A. and R. Macchiavello (forthcoming): "Strategic Default in the International Coffee Market," *The Quarterly Journal of Economics*.

Bramoullé, Y., H. Djebbari, and B. Fortin (2009): "Identification of Peer Effects Through Social Networks," *Journal of Econometrics*, 150, 41–55.

Campa, J. M. and L. S. Goldberg (2005): "Exchange Rate Pass-Through Into Import Prices," *Review of Economics and Statistics*, 87, 679–690.

Carballo, J., G. I. P. Ottaviano, and C. Volpe Martincus (2018): "The Buyer Margins of Firms' Exports," *Journal of International Economics*, 112, 33–49.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018): "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 36, S13–S70.

Competition Commission (2000): *Supermarkets: Competition Inquiries into the Groceries Market.*

De Giorgi, G., M. Pellizzari, and S. Redaelli (2010): "Identification of Social Interactions through Partially Overlapping Peer Groups," *American Economic Journal: Applied Economics*, 2, 241–275.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016): "Prices, Markups, and Trade Reform," *Econometrica*, 84, 445–510.

De Loecker, J. and F. Warzynski (2012): "Markups and Firm-Level Export Status," *American Economic Review*, 102, 2437–2471.

De Toni, A. and G. Nassimbeni (2000): "Just-in-Time Purchasing: an Empirical Study of Operational Practices, Supplier Development and Performance," *Omega-International Journal of Management Science*, 28, 631–651.

Defever, F., C. Fischer, and J. Suedekum (2016): "Relational Contracts and Supplier Turnover in the Global Economy," *Journal of International Economics*, 103, 147–165.

Dickerson, K. G. (1999): *Textiles and Apparel in the Global Economy*, Prentice Hall.

Dwyer, R., P. H. Schurr, and S. Oh (1987): "Developing Buyer-Seller Relationships," *The Journal of Marketing*, 51, 11–27.

Eaton, J., D. Jinkins, J. Tybout, and D. Xu (2016): "Two-sided Search in International Markets," *mimeo.*

Edmond, C., V. Midrigan, and D. Y. Xu (2015): "Competition, Markups, and the Gains from International Trade," *American Economic Review*, 105, 3183–3221.

Egan, M. L. and A. Mody (1992): "Buyer-Seller Links in Export Development," *World Development*, 20, 321–334.

FTC (2001): *Report on the Federal Trade Commission Workshop on Slotting Allowances and Other Marketing Practices in the Grocery Industry.*

Garvin, D. A. (1987): "Competing on the Eight Dimensions of Quality," *Harvard Business Review*, 65.

Gereffi, G. (1999): "International Trade and Industrial Upgrading in the Apparel Commodity Chain," *Journal of International Economics*, 48, 37–70.

Gereffi, G., J. Humphrey, and T. Sturgeon (2005): "The Governance of Global Value Chains," *Review of International Political Economy*, 12, 78–104.

Goldberg, P. (1995): "Product Differentiation and Oligopoly in International Markets: the Case of the U.S. Automobile Industry," *Econometrica*, 63, 891–951.

Goldberg, P. K. and N. Pavcnik (2007): "Distributional Effects of Globalization in Developing Countries," *Journal of Economic Literature*, 45, 39–82.

Hall, R. E. (1986): "Market Structure and Macroeconomic Fluctuations," *Brookings Papers on Economic Activity*, 1986, 285–338.

——— (1988): "The Relation between Price and Marginal Cost in U.S. Industry," *Journal of Political Economy*, 96, 921–947.

Heath, R. and A. M. Mobarak (2015): "Manufacturing Growth and the Lives of Bangladeshi Women," *Journal of Development Economics*, 115, 1–15.

Heise, S. (2018): "Firm-to-Firm Relationships and Price Rigidity: Theory and Evidence," *mimeo*.

Heise, S., J. R. Pierce, G. Schaur, and P. K. Schott (2017): "Trade Policy and the Structure of Supply Chains," *mimeo*.

Helper, S. and M. Saki (1997): "Supplier Relations in Japan and the United States: Are They Converging?" *MIT-Japan Program*.

Houthakker, H. S. (1955): "The Pareto Distribution and the Cobb-Douglas Production Function in Activity Analysis," *Review of Economic Studies*, 23, 27–31.

Jones, C. I. (2005): "The Shape of Production Functions and the Direction of Technical Change," *The Quarterly Journal of Economics*, 120, 517–549.

Kalwani, M. U. and N. Narayandas (1995): "Long-Term Manufacturer-Supplier Relationships: Do They Pay off for Supplier Firms?" *Journal of Marketing*, 59, 1–16.

Kawasaki, S. and J. McMillan (1987): "The Design of Contracts: Evidence From Japanese Subcontracting," *Journal of the Japanese and International Economies*, 1, 327–349.

Kikkawa, A. K., G. Magerman, and E. Dhyne (2018): "Imperfect Competition in Firm-to-Firm Trade," *mimeo*.

Kugler, M. and E. Verhoogen (2012): "Prices, Plant Size, and Product Quality," *Review of Economic Studies*, 79, 307–339.

Lederman, D., M. Olarreaga, and L. Payton (2010): "Export Promotion Agencies: Do They Work?" *Journal of Development Economics*, 91, 257–265.

Macchiavello, R. (2010): "Development Uncorked: Reputation Acquisition in the New Market for Chilean Wines in the UK," *mimeo.*

Macchiavello, R., A. Menzel, A. Rabbani, and C. Woodruff (2015): "Challenges of Change: An Experiment Training Women to Manage in the Bangladeshi Garment Sector," *mimeo.*

Macchiavello, R. and A. Morjaria (2015): "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports," *American Economic Review*, 105, 2911–45.

Manova, K. and Z. Zhang (2012): "Export Prices across Firms and Destinations," *The Quarterly Journal of Economics*, 127, 379–436.

Mayer, T., M. J. Melitz, and G. I. P. Ottaviano (2014): "Market Size, Competition, and the Product Mix of Exporters," *American Economic Review*, 104, 495–536.

McMillan, J. (1990): "Managing Suppliers: Incentive Systems in Japanese and U.S. Industry," *California Management Review*, 32, 38–55.

Medina, P. (2017): "Import Competition, Quality Upgrading and Exporting: Evidence from the Peruvian Apparel Industry," *mimeo.*

Melitz, M. J. and G. I. P. Ottaviano (2008): "Market Size, Trade, and Productivity," *Review of Economic Studies*, 75, 295–316.

Milgrom, P. and J. Roberts (1990): "The economics of modern manufacturing: Technology, strategy, and organization," *The American Economic Review*, 511–528.

Monarch, R. (2018): ""It's Not You, It's Me": Prices, Quality and Switching in U.S.-China Trade Relationships," *mimeo.*

Monarch, R. and T. Schmidt-Eisenlohr (2017): "Learning and the Value of Trade Relationships," *International Finance Discussion Papers 1218.*

Nakamura, E. and D. Zerom (2010): "Accounting for Incomplete Pass-Through," *Review of Economic Studies*, 77, 1192–1230.

Oxfam (2002): *Mugged: Poverty in Your Coffee Cup.*

Richardson, J. (1993): "Parallel Sourcing and Supplier Performance in the Japanese Automobile Industry," *Strategic Management Journal*, 14, 339–350.

Startz, M. (2018): "The Value of Face-to-Face: Search and Contracting Problems in Nigerian Trade," *mimeo.*

Tanvir, S. I. and T. Mahmood (2014): "Solid Waste for Knit Fabric: Quantification and Ratio Analysis," *Journal of Environmental and Earth Science*, 4, 68–80.

Taylor, C. R. and S. N. Wiggins (1997): "Competition or Compensation: Supplier Incentives Under the American and Japanese Subcontracting Systems," *American Economic Review*, 87, 598–618.

Verhoogen, E. A. (2008): "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector," *The Quarterly Journal of Economics*, 123, 489–530.

Voigtlaender, N. and A. Garcia-Marin (forthcoming): "Exporting and Plant-Level Efficiency Gains: It's in the Measure," *Journal of Political Economy*.

Tables and Figures

Table 1: Summary Statistics

	Obs.	Mean	Std. Dev.	P10	P25	P50	P75	P90
				Panel A: Orders				
$Buy-to-Ship_o$	22,445	0.87	0.29	0.51	0.67	0.86	1.04	1.22
$Length_o$ $(months)$	22,445	4.22	3.24	1.47	2.16	3.3	5.2	8
				Panel B: Sellers				
$Count^o_{sy}$	3,146	14.70	13.14	3	6	11	19	30
$Count^o_{sjy}$	6,506	6.30	7.68	1	2	4	8	15
$Count^j_{sy}$	3,146	3.25	1.87	1	2	3	4	6
$Share^j_{sy}$	3,146	58.03	34.80	6.42	25.21	63.03	92.63	100
$Count^b_s$	500	21.10	17.26	4	9	17	28.50	43
$Count^b_{sy}$	3,146	5.93	4.93	1	2	5	8	12
$Count^b_{sjy}$	6,506	3.02	2.99	1	1	2	4	6
$Share^b_{sy}$	3,146	43.83	36.94	1.72	8.33	34.37	82.10	100
$Length_s$ $(years)$	500	6.63	1.55	4.08	5.58	7.58	7.75	7.75
				Panel C: Buyers				
$Count^o_{by}$	4,453	13.38	29.72	1	2	5	13	27
$Count^o_{bjy}$	7,739	5.86	11.70	1	1	3	6	12
$Count^j_{by}$	4,453	4.25	3.82	1	2	3	5	9
$Share^j_{by}$	4,453	59.82	35.55	6.29	27.44	64.38	100	100
$Count^s_b$	2,671	46.14	43.26	9	18	31	62	107
$Count^s_{by}$	7,511	21.94	20.44	4	7	14	29	58
$Count^s_{bjy}$	11,252	8.94	9.06	1	3	6	12	21
$Share^s_{by}$	4,453	48.94	37.93	0	12.15	42.29	92.62	100
$Length_b$ $(years)$	1,576	5.48	2.43	1.58	3.58	6.46	7.67	7.75
				Panel D: Relationships				
$Count^o_{sby}$	10,368	3.39	4.60	1	1	2	4	7
$Count^o_{sbjy}$	12,556	2.55	3.16	1	1	1	3	5
$Count^j_{sby}$	10,368	1.45	0.84	1	1	1	2	2
$Length_{sb}$ $(years)$	5,633	1.86	2.03	0.08	0.25	1.17	2.75	5.08

Super- and sub-scripts are as follows: o corresponds to orders, b to buyers, s to sellers, j to HS6 product categories, y to years. $Count^x_y$ is the number of x per y. For example, $Count^o_{sjy}$ is the number of orders per seller-product-year combination. $Length_o$ is the number of months between the first import shipment and the last export shipment of the order. $Length_{sb}$, $Length_b$, and $Length_s$ are the number of years the buyer-seller pair, buyer, and seller is observed trading in the dataset, respectively. A value of 7.75 in these variables implies censoring. That is, more than 25% of the sellers under study and more than 10% of international buyers are active in all years of our panel. $Share^x_y$ is the share of x in y expressed in percentage terms. For example, for $Share^s_{by}$, the average seller's share in buyer's trade in a year is 48.94%.

Table 2: Input-Output Elasticities and Returns to Scale

Panel A: Elasticities Common Across Products	OLS				IV			
	Coefficient	St. Error	95% Low Bound	95% Up Bound	Coefficient	St. Error	95% Low Bound	95% Up Bound
Materials: θ	0.625	0.017	0.592	0.658	0.619	0.016	0.587	0.651
Labor: β	0.443	0.018	0.409	0.478	0.356	0.030	0.296	0.415
RTS: $\theta + \beta$	1.069	0.004	1.061	1.077	0.975	0.029	0.919	1.031
- Test $RTS = 1$ (χ^2)			293.90				0.78	
- First Stage Underid (LM)			-				219.29	
- First Stage Weak (F)			-				175.55	

Panel B: Product-Specific Elasticities	OLS				IV			
	Coefficient	St. Error	95% Low Bound	95% Up Bound	Coefficient	St. Error	95% Low Bound	95% Up Bound
Materials: θ^T	0.582	0.021	0.540	0.623	0.581	0.020	0.541	0.620
Labor: β^T	0.482	0.022	0.439	0.525	0.396	0.063	0.271	0.520
RTS: $\theta^T + \beta^T$	1.065	0.005	1.055	1.074	0.977	0.061	0.856	1.097
Materials: θ^S	0.683	0.026	0.632	0.734	0.671	0.026	0.619	0.722
Labor: β^S	0.390	0.027	0.337	0.443	0.303	0.033	0.238	0.368
RTS: $\theta^S + \beta^S$	1.073	0.006	1.061	1.085	0.974	0.029	0.917	1.031
- Test $RTS^T = 1$ (χ^2)			168.05				0.14	
- Test $RTS^S = 1$ (χ^2)			143.53				0.78	
- Test $\theta^T = \theta^S$ (χ^2)			9.09				7.35	
- Test $\beta^T = \beta^S$ (χ^2)			7.14				1.82	
- First Stage Underid (LM)			-				92.89	
- First Stage Weak (F)			-				36.89	

Elasticities and related statistics are computed using the Delta Method, combining the estimates of equation (6) in the main text. The baseline specification regresses the log buy-to-ship ratio on seller-product-year fixed effects, the price of fabric and the size of the order (Panel A). The augmented specification allows the coefficients on the price of fabric and the size of order to vary across product groups - trousers and shirts (Panel B). Both specifications are estimated via OLS (left columns) and also instrumenting the size of the order as described in the main text (right columns). All regressions run on a total of 16,561 orders, 56% of which are trousers, with the remaining corresponding to shirts. The standard errors in the underlying regressions are bootstrapped and clustered by seller.

Table 3: Order-level Summary Statistics

	But-to-Ship Ratio (kg/kg)	GarmentPrice ($/kg)	Fabric Price ($/kg)	Marginal Cost ($/kg)	Markup Factor (Units of MC)	Markup Value ($/kg)
	(1)	(2)	(3)	(4)	(5)	(6)
Mean	0.87	13.64	7.58	10.38	1.42	3.26
Median	0.86	13.05	7.27	9.66	1.30	2.87
10^{th} Percentile	0.51	8.60	4.65	5.65	0.96	-0.46
25^{th} Percentile	0.67	10.41	5.65	7.28	1.09	0.98
75^{th} Percentile	1.04	16.32	9.17	12.78	1.64	5.23
90^{th} Percentile	1.22	19.77	11.06	16.21	2.09	7.61
St. Deviation	0.29	4.21	2.41	4.18	0.45	3.23
Coeff. Variation	0.33	0.31	0.32	0.40	0.32	0.99
$90^{th}/10^{th}$ Ratio	2.38	2.30	2.38	2.87	2.17	-16.46
$75^{th}/25^{th}$ Ratio	1.56	1.57	1.62	1.76	1.51	5.33
Number of Orders				22,445		

All statistics are computed over all orders for which a markup was computed. Columns (1) to (3) are directly observed in the data, while columns (4) to (6) are constructed using the elasticities recovered as described in the text. The markup factor is defined as Price/Marginal Cost while the markup value is (Markup Factor - 1) × Marginal Cost.

Table 4: Variance Decomposition of Prices, Marginal Costs, and Markups

Seller-Buyer-Product-Year Outcomes		
	Fixed Effect	
Outcomes:	Buyer (δ_b) (1)	Seller-Product-Year (δ_{sjy}) (2)
Quantities (q_{sbjy})	38.0%	32.0%
Garment Price (p_{sbjy})	26.7%	45.9%
Fabric Price (p^f_{sbjy})	16.5%	59.5%
Marginal Cost (mc_{sbjy})	20.2%	44.8%
Markup Factor (μ_{sbjy})	17.9%	34.8%
Observations	8,109	

The table decomposes the variability in different seller-buyer-product-year outcomes. The underlying regression for the whole panel is $y_{sbjy} = \delta_b + \delta_{sjy} + \epsilon_{sbjy}$, where the δ terms collect buyer and seller-product-year fixed effects. The outcomes are: log quantities (q_{sbjy}), weighted averages of the price of garment (p_{sbjy}), price of fabric (p^f_{sbjy}), marginal costs (mc_{sbjy}) and markup factors (μ_{sbjy}). All outcomes are in logs and are aggregated from order level metrics as weighted averages using the size of the corresponding order. The cells in the panel show the percentage of the variance in each outcome that is explained by buyer (Column (1)), seller-product-year (Column (2)) effects: $Cov(\delta_x, y_{sbjy})/Var(y_{sbjy})$, for $x \in \{b, sjy\}$.

Table 5: Marginal Costs and Markups Paid by Largest Buyers

	Market Share %	Sellers per Year Average	Relational Ranking	Weighted Averages Across Orders Markup Factor	Weighted Averages Across Orders Marginal Cost
Top 25 Buyers	(1)	(2)	(3)	(4)	(5)
H&M Hennes And Mauritz	5.22	60.38	5	1.50	11.86
Wal Mart Stores	4.90	58.50	54	1.34	8.69
Vf Corporation	4.19	25.25	14	1.57	8.22
The Gap Inc	3.49	26.88	1	1.69	9.82
C & A Buying	3.21	43.00	41	1.70	9.99
PVH Corporation	3.18	39.75	32	1.41	10.30
K Mart Corporation	3.10	61.38	59	1.42	8.96
Levi Strauss & Co	2.26	7.50	3	1.47	9.27
J.C. Penney	2.00	25.75	44	1.29	10.73
Primark	1.46	22.88	29	1.47	8.27
Kik Textilen	1.31	51.38	421	1.45	7.65
Kohls Department Stores Inc	1.26	16.38	45	1.61	9.85
Asda	1.22	19.25	15	1.73	7.33
Marks & Spencer	1.18	10.00	35	1.46	10.11
Carrefour	1.15	26.63	33	1.63	7.97
Tesco	1.13	23.13	21	1.76	7.75
Tema Magazacilik	0.91	42.63	142	1.69	10.74
G. Gueldenpfennig Gmbh	0.87	31.88	373	1.28	8.94
Target Stores	0.87	19.13	52	1.48	9.27
Public Clothing Company Inc	0.85	24.88	355	1.38	8.32
Auchan S.A.	0.72	30.13	117	1.57	8.65
Charles Vogele	0.71	17.25	97	1.55	10.21
The Children's Place	0.69	11.00	13	1.57	10.10
IFG Corporation	0.64	14.00	215	1.55	5.87
Ospig Gmbh	0.63	5.50	34	1.11	12.03
Top 100 (Market Share = 66%)					
Mean	0.66	16.00		1.51	9.18
Median	0.29	13.00		1.49	9.22
St. Deviation	0.99	14.24		0.20	1.90
Coeff. Variation	1.50	0.89		0.14	0.21
All Buyers (N = 1,576)					
Mean	0.06	4.40		1.43	10.05
Median	<0.01	3.00		1.38	9.59
St. Deviation	0.30	5.84		0.36	3.53
Coeff. Variation	5.06	1.32		0.25	0.35

The top panel of the table reports market shares, number of sellers, ranking according to the buyer's relational characteristic and weighted average markup factors and marginal costs for the largest 25 buyers. Weights for averaging markups and marginal costs are given by the volume of the corresponding export orders, and aggregations are performed across all products and years the buyer is trading. The markup factor (Column (4)) is defined as Price / Marginal Cost and, as such, it is measured in units of marginal costs. The marginal cost (Column (5)) is expressed in dollars per kilogram of garment. Buyers are ranked using their market share according to traded volumes in the relevant product groups - trousers and shirts - throughout the sample. This information is collected in Column (1). Column (3) indicates the ranking of the buyer in the distribution of the relational characteristic. Column (2) averages (over years) the total number of sellers supplying the buyer in those product groups. For comparison, the bottom panels of the table reports summary statistics of the corresponding variables across the top 100 buyers and across all buyers.

Table 6: Prices, Costs and Markups with Relational Buyers

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	p_{sbjy}	mc_{sbjy}	μ_{sbjy}	p_{sbjy}	mc_{sbjy}	μ_{sbjy}	p_{sbjy}	mc_{sbjy}	μ_{sbjy}
$Relational_b^{0/1}$	0.043***	0.012	0.035***	0.065***	0.039***	0.027**	0.054***	0.026*	0.028**
	(0.012)	(0.015)	(0.012)	(0.012)	(0.015)	(0.012)	(0.010)	(0.013)	(0.012)
q_{sbjy}				-0.050***	-0.063***	0.019***	-0.040***	-0.051***	0.017***
				(0.003)	(0.004)	(0.003)	(0.003)	(0.004)	(0.003)
p^f_{sbjy}							0.538***	0.636***	-0.082***
							(0.014)	(0.019)	(0.017)
FEs	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d
R^2	0.56	0.51	0.40	0.58	0.54	0.40	0.69	0.63	0.40
Obs.	8,169	8,169	8,169	8,169	8,169	8,169	8,169	8,169	8,169

Standard errors in parentheses, clustered at the buyer-seller level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcomes correspond to the weighted average price, marginal cost and markup factor in the seller-buyer-product-year combination, p_{sbjy}, mc_{sbjy} and μ_{sbjy} respectively. The regressor of interest,$Relational_b^{0/1}$, is a dummy taking value one if the buyer belongs to the top 10 percent of the distribution of the relational characteristic, constructed as described in the main text. All specifications include seller-product-year and destination fixed effects, as well as a control for the overall buyer size, q_{b-s}, constructed summing over all the imports of the buyer in the main woven categories, excluding imports from the incumbent seller. Further controls are added sequentially, accounting for the volume of trade and the average price of fabric in the seller-buyer-product-year combination (both in logs), q_{sbjy} and p^f_{sbjy} respectively.

Table 7: Prices, Costs and Markups with Relational Buyers, Accounting for Quality

	Seasons			Fabric Varieties			Complexity		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	p_{sbjy}	mc_{sbjy}	μ_{sbjy}	p_{sbjy}	mc_{sbjy}	μ_{sbjy}	p_{sbjy}	mc_{sbjy}	μ_{sbjy}
$Relational_b^{0/1}$	0.055***	0.025*	0.029**	0.055***	0.027*	0.028**	0.048***	0.020	0.028**
	(0.010)	(0.013)	(0.012)	(0.010)	(0.014)	(0.012)	(0.010)	(0.013)	(0.012)
FEs	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d
R^2	0.70	0.64	0.43	0.73	0.67	0.47	0.70	0.63	0.40
Obs.	8,121	8,121	8,121	7,771	7,771	7,771	8,169	8,169	8,169

Standard errors in parentheses, clustered at the buyer-seller level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcomes correspond to the weighted average price, marginal cost and markup factor in the seller-buyer-product-year combination, p_{sbjy}, mc_{sbjy} and μ_{sbjy} respectively. The regressor of interest,$Relational_b^{0/1}$, is a dummy taking value one if the buyer belongs to the top 10 percent of the distribution of the relational characteristic, constructed as described in the main text. All specifications include seller-product-year and destination fixed and controls accounting for the volume of trade and the average price of fabric in the seller-buyer-product-year combination (both in logs), q_{sbjy} and p^f_{sbjy} respectively. A control for the size of the buyer is included (q_{b-s}) in all columns. Columns (1)-(3), in addition, include product-*season*-year fixed effects. In these regressions, we allow for four seasons in a year, but results are similar when we allow for six seasons in a year. Columns (4)-(6) include a large set of dummies collecting the variety of fabric most used in the orders traded in the buyer-seller-product-year combination. A variety of fabric is defined as the combination of a six-digit HS type of fabric and the country of origin of the fabric. Columns (7)-(9) condition on (log) number of types of fabric and origins of fabric, as a proxy for the complexity of the garment. For each buyer-seller-product-year combination, these complexity metrics are constructed as weighted averages over all orders.

Table 8: Prices, Costs and Markups with Relational Buyers

	(1)	(2)	(3)	(4)	(5)
	p_{sby}	mc_{sby}	μ_{sby}	q_{sby}	π_{sby}
$Relational_b^{0/1}$	0.065***	-0.001	0.063***	0.469***	0.752***
	(0.010)	(0.013)	(0.012)	(0.072)	(0.086)
FEs	sy,d	sy,d	sy,d	sy,d	sy,d
All Controls	Yes	Yes	Yes	Yes	Yes
R^2	0.66	0.58	0.33	0.54	0.53
Obs.	9,210	9,210	9,210	9,210	8,056

Standard errors in parentheses, clustered at the buyer-seller level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcomes correspond to the weighted average price, marginal cost and markup factor in the seller-buyer-year combination, p_{sby}, mc_{sby} and μ_{sby} respectively for the first three columns and, for the last two, the traded volumes, q_{sby} and profits, π_{sbt} constructed as the product between markup values and volumes. All outcomes are in logs. The regressor of interest, $Relational_b^{0/1}$, is a dummy taking value one if the buyer belongs to the top 10 percent of the distribution of the relational characteristic, constructed as described in the main text. All specifications condition on seller-year and destination fixed effects. In addition, the include controls for the price of fabric (p^f_{sby}) and the size of the buyer (q_{b-s}). Columns (1)-(3) also control for traded volumes, q_{sby}.

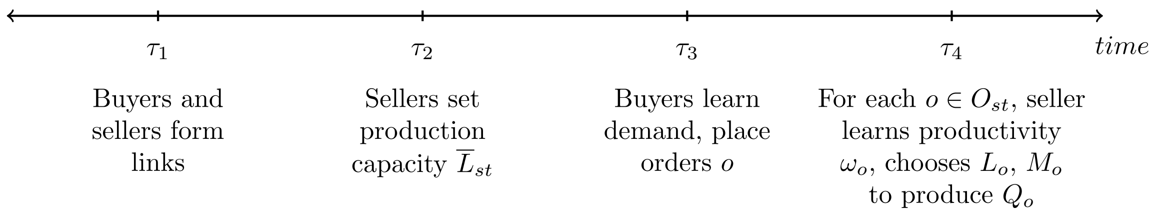

Figure 1: Timing of Events

The Figure illustrates the timing of events in the sourcing and production problem as described in the main text. The markers $\tau_1 \dots \tau_4$ correspond to different theoretical instances *within* a temporal unit, for example, a year. $\overline{L}_{st}$ corresponds to the amount of labor hired by seller s seller for year t; o indexes orders and O_{st} is the set of orders received by s for production in t; L_o, M_o and ω_o indicate the labor, materials and productivity operated in the production of o; finally, Q_o stands for the size of the order.

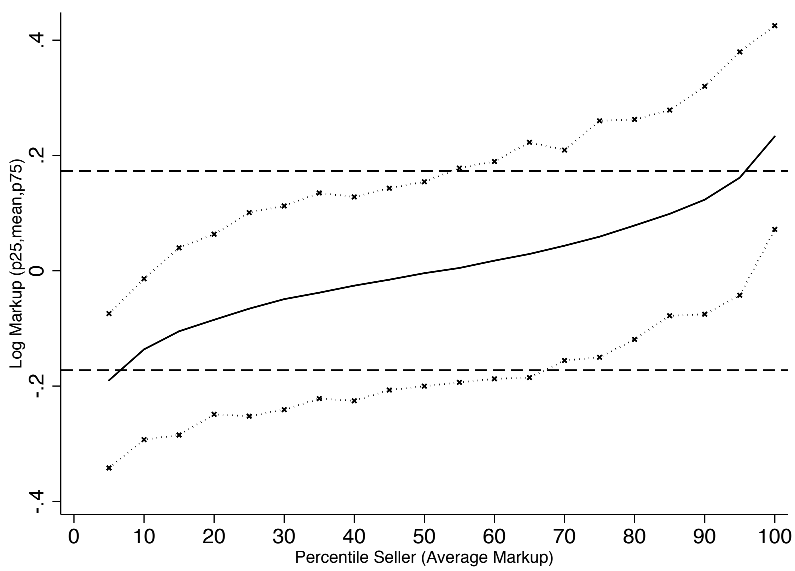

Figure 2: Dispersion in Markups Within Seller

We aggregate order-level log markup factors for each seller-buyer-product-year combination. We residualize these against product-year fixed effects. For each seller, we construct the simple average, 25^{th} and 75^{th} percentile residual markup across those residuals (discarding any seller with less than 10 data points). The horizontal axis arranges sellers ascendingly in percentiles according to their average markup. The solid line connects the average markup in bins of 20 sellers. The dotted lines represent the 25^{th} and 75^{th} percentiles. The dashed horizontal lines correspond to the average interquartile range across sellers, centered around the average residual markup of the median seller.

A Further Details on the Sample and Methodology

A.1 Sample Construction

In the empirical analysis, we focus on a subset of woven garment export products, comprising 17 six-digit HS codes, grouped more broadly into *trousers* and *shirts*. These account for approximately 86% of the exported volume in woven garments in our data. We discard export orders whose quality of underlying data is low, preventing a clean import-export matching exercise.[39] To mitigate sparseness in the data for our analysis, we only consider the top-500 exporters (out of almost 1,500), who jointly account for 87% of the exported volumes in the subsample. Our final sample of 22,445 export orders accounts for 37% of the Bangladesh's exports in the relevant product categories.

The order-level regressions that make use of the instrument for the size of the export order constructed as explained in the main text, necessitate a slightly more restrictive sample. In particular, the instrumentation strategy requires that the exporter is trading in the same quarter with *other* buyers, who in turn trade with *other* sellers. Together with use of fixed effects as granular as seller-product-year (where product is an six-digit HS code), this restriction renders a sample of 486 sellers with 16,561 export orders.[40]

Table A1 compares key shipment, buyer, seller, and relationship characteristics between the original sample and the two sub-samples described above.

Table A1: Sample Comparisons

	Panel A: Average Shipment Characteristics				
Shipments:		Count	Price (USD/kg)	Size (tonnes)	
Under UD System		598,423	17.00	2.70	
Outside UD System		5,267	15.35	1.72	
	Panel B: Firm and Relationship Characteristics				
Orders:	Buyer Vol. (tonnes)	N_b^s	Seller Vol. (tonnes)	N_s^b	Rel. Vol. (tonnes)
Used in Analysis	223.45	14.14	493.69	21.10	74.43
Used in Estimation	368.80	21.44	500.40	21.47	93.07

The top panel compares shipments from orders in the UD system and shipments outside the UD system for buyers and sellers active and relevant products in the sub-sample used in the empirical analysis. A test of equal means finds that both average price and shipment size are not significantly different across samples. The bottom panel compares buyer, seller, and relationship characteristics for the two sub-samples used in the paper. Volume measures the average yearly traded volume.

[39] Specifically, we work with export shipments that are channeled via the UD system, therefore ignoring isolated or stand-alone export transactions. Within this sub-sample, we exclude orders characterized as outliers (lower than 3% and larger than 97%) in the distribution of relevant observables: the buy-to-ship weight ratio, the output price, the input price, the cost share of fabric with respect to the order revenue. These conditions are satisfied for almost half of the volume exported in the relevant product categories.

[40] The instrument construction does not drop any individual seller, but discards some orders of these sellers, such that there is not enough variation within narrow clusters.

A.2 Methodology

Input versus Waste Ratio							
Factory No	Input Quantity (KG) [A_1]	Inspection Loss (KG) Point 1	Cutting Loss (KG) Point 2	Sewing Loss (KG) Point 3	Finishing Loss (KG) Point 4	Total Waste (KG) [A_2]	% of Waste (A_2 / A_1) x 100
1	700	35	50	20	10	115	16.25%
2	750	30	40	25	15	110	14.67%
3	780	40	50	15	10	125	16.03%
4	800	25	30	30	20	105	13.13%
5	820	20	45	30	15	110	13.42%
6	880	25	40	35	20	120	13.63%
7	910	50	70	30	25	175	19.24%
8	950	45	65	25	20	155	16.34%
9	990	25	35	35	15	110	11.12%
10	1000	50	50	30	10	140	14%
11	1100	25	40	25	5	95	8.64%
12	1900	100	100	50	40	290	15.27%
13	2000	80	60	30	50	120	6%
14	2300	110	100	50	20	280	12.18%
15	2500	25	20	10	5	60	2.4%
16	3000	20	40	30	10	100	3.34%
17	3200	60	35	20	20	135	4.26%
18	3600	50	30	10	15	105	2.9%
19	3900	90	35	30	20	175	4.49%
20	4000	80	30	25	25	160	4%
21	4100	40	25	50	20	135	3.30%
22	4250	35	30	30	10	105	2.48%
23	4400	55	25	50	5	135	3.06%
24	4700	70	30	30	5	135	2.89%
25	5000	65	25	50	10	150	3%
26	14000	50	120	20	45	235	1.68%
27	1100	25	15	25	10	75	6.8%
28	24200	220	200	50	40	470	2%
29	23100	140	180	45	30	385	1.6%
30	1600	10	10	25	5	50	3.1%
Total	**136930**	**1585**	**1325**	**930**	**540**	**4240**	

Table A2: Input versus Waste Ratio from Tanvir and Mahmood (2014)

This table shows data on fabric wastage from 30 Bangladeshi garment factories surveyed in Tanvir and Mahmood (2014).

Table A3: Input Choice and Input Prices

	(1) q_{sbjo}^{fabric}	(2) q_{sbjo}^{fabric}	(3) q_{sbjo}^{fabric}	(4) q_{sbjo}^{fabric}	(5) q_{sbjo}^{fabric}	(6) q_{sbjo}^{fabric}
$p_{m(o)}^{cotton}$	-0.021**		-0.075***	-0.027**		-0.079***
	(0.010)		(0.011)	(0.012)		(0.013)
$m(o) \geq MinWage$		0.086***	0.115***		0.083***	0.114***
		(0.010)	(0.011)		(0.012)	(0.013)
Size Order	Yes	Yes	Yes	Yes	Yes	Yes
Time Trend (Month)	Yes	Yes	Yes	Yes	Yes	Yes
FEs	sj	sj	sj	sj	sj	sj
R^2	0.95	0.95	0.95	0.94	0.94	0.94
Obs.	21,904	21,904	21,904	16,561	16,561	16,561

Standard errors in parentheses, clustered at the seller-product level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. All specifications have the log of the quantity of fabric used in the order, q_{sbjo}^{fabric}, as the outcome. All specifications include seller-product fixed effects, a control for the size of the order, in log kilos of garment, q_{sbjo} and a monthly linear time trend. $p_{m(o)}^{cotton}$ is the log of the international price of cotton in the first month of the order. $m(o) \geq MinWage$ is a dummy that takes value one if the order started after the implementation of the minimum wage increase in November 2010. The analogous exercise (not reported here) using the wage inflation update in November 2006 shows the same pattern, but with an effect on the outcome smaller in magnitude, consistent with the size of the wage increase. Columns (1) to (3) use the full sample of orders we use in the analysis. Columns (4) to (6) restricts the sample to those orders used for the estimation of the input-output elasticities.

B Firm Level Patterns

B.1 Sellers

To compare our findings with those of the literature, we aggregate order-level outcomes at the seller-product-year level.[41] We find that at this level of aggregation, the average markup factor is 1.44 and the median is 1.37, with approximately 8% of the observations having markup factors below one. Appendix Figure B1 displays the residual dispersion in markups and marginal costs across sellers. Consistent with the results of Atkin et al. (2015), the figure shows that markup values are significantly more dispersed than marginal costs. The corresponding coefficients of variation are 0.82 for markups and 0.36 for marginal costs.

Appendix Figure B2 plots markups and marginal costs for each seller-product-year combination against the quantity exported by the triplet (where all variables are logged and demeaned with product-year fixed effects). The figure shows that exported quantity is negatively correlated with marginal costs and positively correlated with markups. Hence, in line with the findings of De Loecker et al. (2016) for India and Atkin et al. (2015) for the soccer ball sector in Sialkot, Pakistan, we find that larger firms enjoy lower marginal costs and higher markups than smaller firms.

Finally, in Appendix Table B1, we investigate the core product hypothesis discussed in the literature on multi-product firms (see, e.g., Mayer et al., 2014). This hypothesis posits that firms will specialize in their core products—those for which their marginal costs are lowest and their markups are highest—and will lose profitability as they move away from these products. Table B1 explores this hypothesis using two specifications. First, we show that conditional on seller and product-year fixed effects, the share of a product in a seller's exports in a given year correlates negatively with marginal costs (column (1)) and positively with markups (column (4)). Second, we show that these relationships also hold when conditioning on seller-year and product fixed effects (columns (2) and (5)). Moreover, a dummy that indicates the seller's main product—that with the largest total exports across all years—correlates negatively with marginal costs and positively with markups (columns (3) and (6)). In sum, consistent with the core product hypothesis and as also found in De Loecker et al. (2016), Table B1 shows that multi-product firms have relatively lower marginal costs and command relatively higher markups in the products in which they specialize.

[41] The aggregation is a weighted average across orders in the seller-product-year triplet, where an order's weight is determined by its size in kilos.

Table B1: Markups, Costs and the Product Share within Seller-Year

	(1)	(2)	(3)	(4)	(5)	(6)
	mc_{sjy}	mc_{sjy}	mc_{sjy}	μ_{sjy}	μ_{sjy}	μ_{sjy}
$Share^j_{sy}$	-0.144***	-0.138***		0.066***	0.065***	
	(0.016)	(0.017)		(0.013)	(0.013)	
$Core^j_s$			-0.026**			0.017**
			(0.011)			(0.009)
FEs	s,jy	sy,j	sy,j	s,jy	sy,j	sy,j
R^2	0.41	0.58	0.58	0.23	0.47	0.47
Obs.	5,362	5,362	5,362	5,362	5,362	5,362

Standard errors in parentheses, clustered at the seller-product level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table shows OLS regressions whose outcomes are, the log of the marginal cost for the seller-product-year combination (Columns (1) to (3)) and its log markup factor (Columns (4) to (6)). Both outcomes are constructed as weighted averages of the order-level markups and marginal costs. Columns (1) and (4) include seller and product-year fixed effects, while the rest of the table includes seller-year and product effects. The share of the product in the seller-year's trade, $Share^j_{sy}$ is constructed as the ratio between the volume traded in the product and the total exports of the seller in the year. It is therefore $\in (0, 1]$. $Core^j_s$ is a dummy variable that takes value one if product j is the one with the largest share in the seller's trade.

Table B2: Variance Decomposition of Seller Level Outcomes

	Seller-Product-Year Markups	
	Specification	
Outcome: μ_{sjy}	Within Seller	Within Seller-Product
Fixed Effects:	(1)	(2)
Seller (δ_s)	17.6%	-
Seller-Product (δ_{sj})	-	38.1%
Main Buyer ($\delta_{B(sjy)}$)	13.2%	11.0%
Product-Year (δ_{jy})	4.6%	2.8%
Observations	6,103	5,404

The table decomposes the log average markup of seller-product-year combinations, into variation explained by different fixed effects in two specifications. The specification in Column (1) is $\mu_{sjy} = \delta_s + \delta_{B(sjy)} + \delta_{jy} + \epsilon_{sjy}$, where δ_s is a fixed effect of the seller, $\delta_{B(sjy)}$ is a dummy collecting an intercept for the main buyer in the seller-product-year combination according to volumes, and δ_{jy} is a product-year fixed effect. The specification in Column (2) is $\mu_{sjy} = \delta_{sj} + \delta_{B(sjy)} + \delta_{jy} + \epsilon_{sjy}$, with δ_{sj} seller-product fixed effects. The cells in the panel correspond to the percentage of the variance in the outcome μ_{sjy} that each component explains: $Cov(\delta_x, \mu_{sjy})/Var(\mu_{sjy})$, for $x \in \{B(sjy), s, sj, jy\}$.

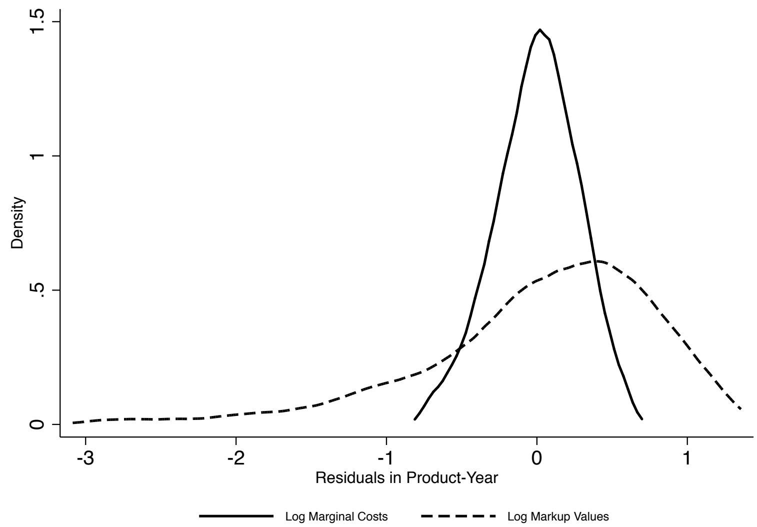

Figure B1: Residual Dispersion in Markups and Costs Across Sellers

We construct average markup values and average marginal costs for each seller-product-year combination. The aggregation is done weighting each order-level markup or marginal cost by the size of the order. We residualize the seller-product-year log markups and marginal costs against product-year means. We plot the density of these residuals. We consider all seller-product-year combinations for which markup values are non-negative and trim the top and bottom 1% of the densities.

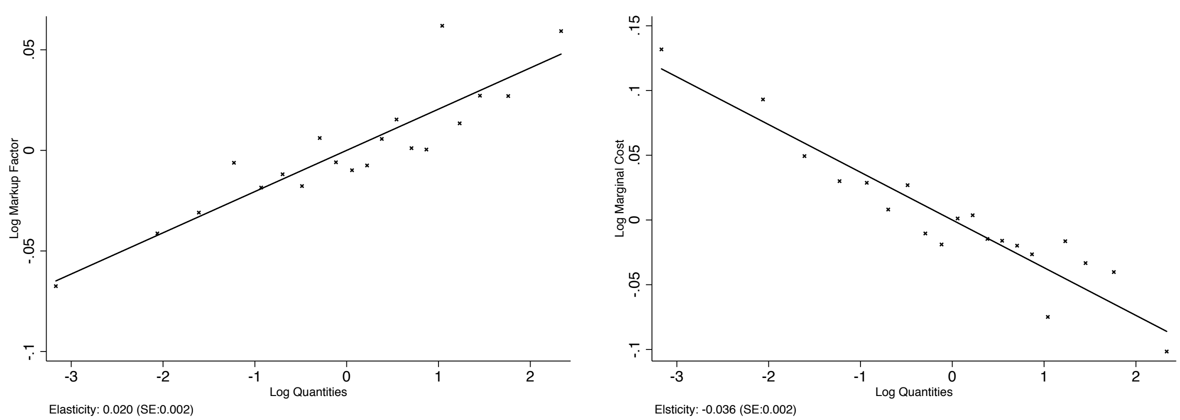

Figure B2: Markups, Marginal Costs and Quantities

The horizontal axis of each scatter plot measures the log quantity exported by seller-product-year combinations. The vertical axis corresponds to the log markup factor, on the left panel, and the average log marginal cost, on the right panel. Data along both sets of axis is arranged in equally sized bins, whose means are indicated with scatter-points. All variables are residualized against product-year fixed effects. The superimposed lines correspond to OLS regressions of the underlying data, comprising 6,496 seller-product-year combinations. Underneath each plot, the elasticity of each outcome -markups or marginal costs- with respect to traded quantities is reported.

B.2 Buyers

Table B3: Relational Buyers

Panel A: Destination Characteristics						
		(1) $Relational_b$	(2) $Relational_b$	(3) $Relational_b$	(4) $Relational_b$	
$Distance_d$		0.340*** (0.108)				
GDP_d			0.146*** (0.021)			
$Population_d$				0.134*** (0.023)		
$GDPpc_d$					0.329*** (0.072)	
R^2		0.01	0.05	0.03	0.03	
Obs.		1,045	1,045	1,045	1,045	
Panel B: Buyer Characteristics						
	(1) q_{by}	(2) $Med\ Share^s_{by}$	(3) $Max\ Share^s_{by}$	(4) $Count^o_{by}$	(5) $Count^{ship}_{by}$	(6) $Count^j_{by}$
$Relational_b$	1.301*** (0.038)	0.214*** (0.020)	0.122*** (0.009)	-0.143*** (0.017)	0.261*** (0.019)	-0.147*** (0.014)
R^2	0.36	0.55	0.41	0.69	0.87	0.51
Obs.	6,257	6,257	6,257	6,257	6,257	6,257
Panel C: Order Characteristics						
	(1) q_{sbjo}	(2) $\overline{q}^{ship}_{sbjo}$	(3) $Count^{ship}_{sbjo}$	(4) $Count^{ship}_{sbjo}$	(5) p^{fabric}_{sbjo}	
$Relational_b$	0.200*** (0.027)	-0.255*** (0.022)	0.455*** (0.023)	0.322*** (0.019)	0.009 (0.007)	
R^2	0.56	0.52	0.61	0.82	0.65	
Obs.	19,494	19,494	19,494	19,494	19,494	

Standard errors in parentheses, clustered at the buyer level in Panels B and C, heteroskedasticity-robust in Panel A. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. Panel A has the standardized metric on the relational characteristic of the buyer as the outcome and it runs in the cross-section of active buyers in 2009. All gravity variables are in logs and correspond to the distance from the buyer's county to Bangladesh ($Distance_d$), the GDP of the destination in the selected year (GDP_d), its population ($Population_d$) and GDP per capital (GDP_{pc}). Panel B regresses the standardized buyer-specific relational characteristic on the buyer's size of trade (q_{by}), the log share the median seller of the buyer has in the buyer's yearly trade ($Med\ Share^s_{by}$), the log share that the largest seller of the buyer has in the buyer's yearly trade ($Max\ Share^s_{by}$), the log number of orders the buyer has in the year ($Count^o_{by}$), the log number of shipments the buyer has in the year ($Count^{ship}_{by}$) and the log number of products the buyer purchases in the year ($Count^j_{by}$). All columns (1)-(6) include year fixed effects and columns (2)-(6) also control for the size of the buyer's trade, q_{by}. Panel C regresses the standardized relational characteristic on order-level outcomes: the log size of the export order (q_{sbjo}), the log average size of the shipments in the order ($\overline{q}^{ship}_{sbjo}$), the log number of shipments in the order ($Count^{ship}_{sbjo}$) and the log price of the fabric used for the order (p^{fabric}_{sbjo}). All specifications (1)-(5) include seller-product-year and destination fixed effects. They also control for the the size of the buyer's trade, q_{by}. Columns (4) and (5) further control for the size of the order (q_{sbjo}).

Table B4: Markups, Costs and the Product Share within Buyer-Year

	(1)	(2)	(3)	(4)	(5)	(6)
	mc_{bjy}	mc_{bjy}	mc_{bjy}	μ_{bjy}	μ_{bjy}	μ_{bjy}
$Share^j_{by}$	-0.228***	-0.239***		0.123***	0.125***	
	(0.022)	(0.024)		(0.018)	(0.019)	
$Core^j_b$			-0.060***			0.037***
			(0.013)			(0.011)
FEs	b,jy	by,j	by,j	b,jy	by,j	by,j
R^2	0.48	0.59	0.58	0.29	0.46	0.46
Obs.	4,979	4,979	4,979	4,979	4,979	4,979

Standard errors in parentheses, clustered at the buyer-product level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table shows OLS regressions whose outcomes are, the log of the marginal cost for the buyer-product-year combination (Columns (1) to (3)) and its log markup factor (Columns (4) to (6)). Both outcomes are constructed as weighted averages of the order-level markups and marginal costs. Columns (1) and (4) include buyer and product-year fixed effects, while the rest of the table includes buyer-year and product effects. The share of the product in the buyer-year's trade, $Share^j_{by}$ is constructed as the ratio between the volume traded in the product and the total imports in woven of the buyer in the year. It is therefore $\in (0, 1]$. $Core^j_b$ is a dummy variable that takes value one if product j is the one with the largest share in the buyer's (woven) trade.

Table B5: Prices, Costs and Markups with Large Buyers

	(1) p_{sbjy}	(2) mc_{sbjy}	(3) μ_{sbjy}	(4) p_{sbjy}	(5) mc_{sbjy}	(6) μ_{sbjy}	(7) p_{sbjy}	(8) mc_{sbjy}	(9) μ_{sbjy}	(10) p_{sbjy}	(11) mc_{sbjy}	(12) μ_{sbjy}
q_b	-0.004*	-0.006**	0.004**									
	(0.002)	(0.002)	(0.002)									
q_{b-s}				0.006***	0.009***	-0.003	0.005**	0.008***	-0.003			
				(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)			
q_{sb}				-0.023***	-0.035***	0.016***						
				(0.003)	(0.003)	(0.003)						
q_{sb-j}							0.000	0.000	0.000			
							(0.001)	(0.001)	(0.001)			
q_{sbj}							-0.032***	-0.040***	0.012***			
							(0.003)	(0.004)	(0.003)			
q_{by-s}										0.005***	0.003*	0.001
										(0.001)	(0.002)	(0.001)
q_{sby-j}										-0.000	0.000	-0.000
										(0.001)	(0.001)	(0.001)
q_{sbjy}										-0.048***	-0.060***	0.017***
										(0.003)	(0.004)	(0.003)
FEs	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d	sjy,d
R^2	0.56	0.51	0.40	0.57	0.52	0.40	0.57	0.53	0.40	0.58	0.54	0.40
Obs.	8,169	8,169	8,169	8,169	8,169	8,169	8,169	8,169	8,169	8,169	8,169	8,169

Standard errors in parentheses, clustered at the buyer-seller level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table shows OLS regressions whose outcomes are, the log of the price, marginal cost or markup factor for the seller-buyer-product-year combination. All outcomes are constructed as weighted averages of the order-level prices, markups and marginal costs. All specifications include seller-product-year and destination fixed effects. The regressors are defined as follows: q_b is the log volume imported by the buyer throughout the sample period (across all relevant products); q_{b-s} is the log volume imported by the buyer, excluding the volumes traded with the seller; q_{sb} is the log volume traded by the buyer-seller pair; q_{sb-j} is the log volume traded by the buyer-seller pair excluding the trade in the product; q_{sbj} is the log volume traded by the buyer-seller pair in the product category; q_{by-s} is the log volume imported by the buyer in the corresponding year, excluding the volume traded with the seller; q_{sby-j} is the log volume traded by the buyer-seller pair in the year, excluding the product; q_{sbjy} is to the log volume traded by the buyer-seller pair in the product category in the corresponding year. The variables defined over exclusions might be missing where the data is sparse - for example, a buyer-seller pair trading only one product would have q_{sb-j} missing. Such sparseness is not pervasive. We recode these cases as a zero and absorb the imputation in a dummy taking value one if such imputation was performed. The corresponding dummies are included in the specification whenever q_{b-s}, q_{by-s}, q_{sb-j} or q_{sby-j} feature as regressors.

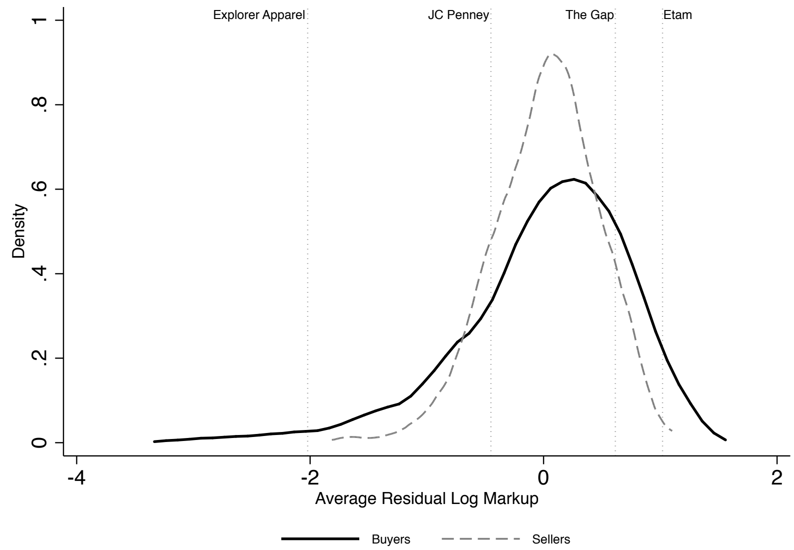

Figure B3: Dispersion in Markups Across Buyers and Sellers

We aggregate order level log markup factors for each buyer-product-year combination. We residualize these against product-year fixed effects. For each buyer, we construct the weighted average of these residual markups. The density over these buyer-level residual markups is represented by the solid line. The analogous density for sellers is overlaid using a dashed line. Selected buyers are labeled for reference.

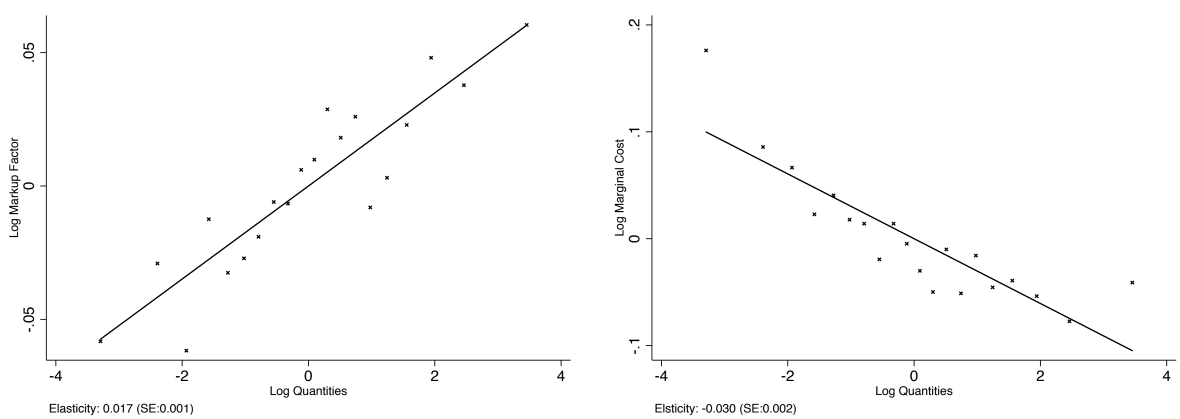

Figure B4: Buyers' Markups, Marginal Costs and Quantities

The figures correspond to data aggregated at the level of the buyer-product-year. The horizontal axis of each scatter plot measures the log quantity traded. The vertical axis corresponds to the log markup factor, on the left hand, and the log marginal cost, on the right hand. Data along both sets of axis is arranged in equally sized bins, whose means are indicated with scatter-points. Variables are residualized against product-year fixed effects. The superimposed lines correspond to OLS regressions of the underlying data, comprising 7,740 buyer-product-year triplets. Underneath each plot, the elasticity of each outcome -markups or marginal costs- with respect to traded quantities is reported.

C Robustness and Additional Results

C.1 Robustness

C.1.1 Seasonality and Specialization

In Section 4.3 we postponed the discussion of two important aspects of our analysis. In particular, we established our baseline result on relational sourcing in Table 6, leveraging the granularity of the panel to condition on fixed effects and relevant controls. However rich these specifications, the reader might not consider them exhaustive of factors, such as seasonality and specialization patterns, that could be systematically correlated with the sourcing strategy of the buyer. This is the set of concerns that we address here.

We augment the baseline specification in Table 6 at the seller-buyer-product-year level as follows

$$y_{sbjy} = \delta_{sjy} + \delta_d + \beta_1 x + \beta_2 Relational_b + \beta_3 q_{sbjy} + \beta_4 p^f_{sbjy} + \beta_5 q_{b-s} + \epsilon_{sbjy}$$

where we have included an additional control x, that varies with the specification. Results are presented in Appendix Table C1, where the left hand panel (Columns (1)-(3)) reports β_1 - the coefficient on the additional control - and the right panel (Columns (4)-(6)) reports β_2 - the coefficient on the relational regressor. The outcomes are $y = \{p, mc, \mu\}$ in columns (1)-(4), (2)-(5) and (3)-(6), respectively.

The top panel features six specifications, each adopting a different strategy to control for seasonality.[42] $\#Seasons_{sb}$ collects the number of seasons (1 to 4) the buyer and seller trade in, throughout the sample. $Share_{sb}^{\overline{Season}}$ and HHI_{sb}^{Season} collect the share of the largest season in the relationship and a Herfindahl index of concentration in seasons in the relationship. This index is normalized to range between zero and 1, where 1 is full concentration in one season. $Share_{sby}^{\overline{Season}}$ is the share of the largest season in the relationship's trade in the year. The last two specifications, 5 and 6, include dummies for the largest season the buyer and seller are trading in, in the corresponding year (or year-product). The baseline is the first season (low spring) so $\overline{Season}_{sby}$ and $\overline{Season}_{sbjy}$ are dummies for seasons 2, 3 or 4.

Looking at the left hand panel, relationships that spread out over multiple seasons and products enjoy higher markups. The more concentrated a buyer-seller pair are on trading in a particular season or a particular HS code the lower the prices and the lower the markups. As this always conditions on sjy effects, this is not down to the nature of the j per se. Instead, this is compatible with buyers having some suppliers with whom they trade regularly and 'supplementary' suppliers that they access in specific seasons to complement their mainstream supply. In this sense, for a given seller-product-year, those buyers that source over multiple seasons and products from the seller, pay higher markups, while those that are concentrated to specific seasons and products, pay lower prices and markups. Looking at the right hand panel, conditioning for this effect as collected by the various specification, the pattern we extracted on relational sourcing remains qualitatively unchanged.

The middle panel introduces a set of specifications that condition on relationship-specific product specialization in different ways. This is important, given the intuition we gave before

[42] In all cases, we divide the calendar year in four seasons. Results are very similar when we allow for six seasons.

for the 'core product' effects at the level of both the buyer and the seller. We want to make sure that the result we are putting forward on relational sourcing is not, at least entirely, driven by relational buyers sourcing specific products or bundles of products from their suppliers. We remind the reader at this point that the variation that is exploited in regressions like those presented in Table 6 operates across the different buyers a specific manufacturer selling a narrowly defined product in a calendar year trades with. This estimation strategy already accounts for simple(r) specialization mechanisms.

Specifications 1 to 5 introduce controls that account for scope or diversification effects. $\#Prod._{sb}$ ($\#Prod._{sby}$), collects the number of products the buyer and seller trade in, throughout the sample (in the year). $Share_{sb}^{\overline{Prod.}}$ and $HHI_{sb}^{Prod.}$ reflect the share of the largest product in the relationship and a (normalized) Herfindhal index of concentration in products in the relationship, respectively. $Share_{sby}^{\overline{Product}}$ is the share of the largest product in the relationship's trade of that year. Focusing on the left hand panel, across all specifications we observe that relationships that are more diversified across products tend to enjoy higher prices and markups, on average. Accounting for this scope effect, the relational result remains also unchanged, as shown in the right hand panel.

Specifications 6 and 7 reproduce the 'core product' analysis, but now at the level of the relationship. $Main_{sb}^{Prod.}$ takes value one if the product is the largest in the relationship, while $Share_{sby}^{Product}$ contains the share of the product in the relationship-year. Note that these specifications do not include seller-buyer fixed effects, so we cannot extract conclusions exploiting within relationship and across products comparisons. The key to these specifications is that the relational result presented in the main text does not seem to be driven by relational buyers sourcing only products that are core to their relationships. In the final two specifications of this block, we attempt to rule out the idea that the relational result is induced by relational buyers sourcing products that command higher markups due to differentiation. This needs to be a form of differentiation that exceeds specifics of the destination market (like in pricing-to-market hypothesis) and that is not fully accounted for by the price of fabric. This issue is discussed in depth in the main text. For now, we limit the discussion to showing that the inclusion of metrics collecting the complexity of the traded garment are positively related to marginal costs and prices and leave the relational result unchanged.[43]

The bottom panel of Table C1 presents the results of specifications where x contains information on the 'network' of relationships standing in the data: $\#Sellers_{bjy}$ count the number of different sellers the buyer is trading with in a year-product combination and $Share_{b}^{s}$ and $Share_{by}^{s}$ correspond to shares of the seller in the trade of the buyer and buyer-year combinations, computed using volumes. With these controls, we address concerns that might arise by considering that the higher markups that relational buyers offer respond to those buyers allocating most of their volume with one or few partners, who, in turn would enjoy a stronger bargaining position in the presence of relationship specific assets or switching costs. Columns (4) to (6) of the three specifications in this bottom panel suggest this mechanism is not a driver of the relational result.

[43]We construct the measures of complexity as follows. For each order in the relationship, we count the number of types (or origins) of fabric for producing this order. We weight each order by its size and generate a buyer-seller-product-year measure of the complexity of the garments traded, based on these average number of fabrics or origins.

Table C1: Relational Buyers, Seasonality, Specialization and Partners

			β_1			β_2	
		(1)	(2)	(3)	(4)	(5)	(6)
		p_{sbjt}	mc_{sbjy}	μ_{sbjy}	p_{sbjt}	mc_{sbjy}	μ_{sbjy}
	Seasonality:						
1	$\#Seasons_{sb}$	0.015***	0.002	0.015***	0.053***	0.025*	0.027**
		(0.004)	(0.005)	(0.004)	(0.010)	(0.013)	(0.012)
2	$Share_{sb}^{\overline{Season}}$	-0.074***	-0.028	-0.051**	0.052***	0.025*	0.027**
		(0.017)	(0.023)	(0.021)	(0.010)	(0.013)	(0.012)
3	HHI_{sb}^{Season}	-0.052***	-0.011	-0.045***	0.052***	0.025*	0.027**
		(0.012)	(0.016)	(0.014)	(0.010)	(0.013)	(0.012)
4	$Share_{sby}^{\overline{Season}}$	-0.105***	-0.033	-0.084***	0.052***	0.025*	0.027**
		(0.016)	(0.022)	(0.020)	(0.010)	(0.013)	(0.012)
5	$\overline{Season}_{sby}$=2	-0.003	-0.032***	0.029***	0.054***	0.027**	0.028**
		(0.008)	(0.011)	(0.011)	(0.010)	(0.013)	(0.012)
	$\overline{Season}_{sby}$=3	-0.029***	0.004	-0.028***			
		(0.008)	(0.011)	(0.010)			
	$\overline{Season}_{sby}$=4	0.014*	0.012	0.001			
		(0.008)	(0.012)	(0.011)			
6	$\overline{Season}_{sbjy}$=2	-0.006	-0.037***	0.031***	0.054***	0.026**	0.028**
		(0.008)	(0.012)	(0.011)	(0.010)	(0.013)	(0.012)
	$\overline{Season}_{sbjy}$=3	-0.030***	0.002	-0.029***			
		(0.008)	(0.011)	(0.010)			
	$\overline{Season}_{sbjy}$=4	0.013	0.018	-0.006			
		(0.008)	(0.012)	(0.010)			
	Specialization:						
1	$\#Prod._{sb}$	0.011***	0.003	0.008***	0.055***	0.026**	0.029**
		(0.002)	(0.003)	(0.002)	(0.010)	(0.013)	(0.012)
2	$Share_{sb}^{\overline{Prod.}}$	-0.073***	0.003	-0.078***	0.055***	0.026**	0.029**
		(0.019)	(0.025)	(0.023)	(0.010)	(0.013)	(0.012)
3	$HHI_{sb}^{Prod.}$	-0.031***	0.005	-0.036***	0.054***	0.025**	0.029**
		(0.010)	(0.013)	(0.012)	(0.010)	(0.013)	(0.012)
4	$\#Prod._{sby}$	0.020***	0.012**	0.009**	0.054***	0.026**	0.029**
		(0.004)	(0.005)	(0.005)	(0.010)	(0.013)	(0.012)
5	$Share_{sby}^{\overline{Prod.}}$	-0.071***	-0.025	-0.048**	0.054***	0.026**	0.029**
		(0.020)	(0.027)	(0.024)	(0.010)	(0.013)	(0.012)
6	$Share_{sby}^{Prod.}$	-0.022*	-0.004	-0.022	0.054***	0.026*	0.028**
		(0.012)	(0.016)	(0.015)	(0.010)	(0.013)	(0.012)
7	$Main_{sb}^{Prod.}$	0.000	0.035***	-0.036***	0.054***	0.026**	0.028**
		(0.008)	(0.011)	(0.010)	(0.010)	(0.013)	(0.012)
8	$Complex_{sbjy}^{fab.}$	0.059***	0.056***	0.002	0.048***	0.020	0.028**
		(0.009)	(0.012)	(0.011)	(0.010)	(0.013)	(0.012)
9	$Complex_{sbjy}^{orig.}$	0.059***	0.056***	0.002	0.048***	0.020	0.028**
		(0.009)	(0.012)	(0.011)	(0.010)	(0.013)	(0.012)
	Partners:						
1	$\#Sellers_{bjy}$	-0.001*	0.000	-0.001*	0.056***	0.025*	0.031***
		(0.001)	(0.001)	(0.001)	(0.010)	(0.013)	(0.012)
2	$Share_{b}^{s}$	0.003	-0.008	0.018	0.053***	0.026**	0.027***
		(0.026)	(0.035)	(0.031)	(0.010)	(0.013)	(0.012)
3	$Share_{by}^{s}$	-0.010	0.014	-0.022	0.054***	0.025*	0.030***
		(0.016)	(0.021)	(0.016)	(0.019)	(0.013)	(0.012)

Standard errors in parentheses, clustered at the buyer-seller level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table combines the results of multiple OLS regressions, all running on the same 8,169 observations.The outcomes correspond to the weighted average price, marginal cost and markup factor in the seller-buyer-product-year combination, p_{sbjy}, mc_{sbjy} and μ_{sbjy} respectively. The regressor of interest,$Relational_b^{0/1}$, is a dummy taking value one if the buyer belongs to the top 10 percent of the distribution of the relational characteristic, constructed as described in the main text. Its coefficient is labelled as β_2 and reported on the right hand side of the table. All specifications include seller-product-year and destination fixed effects, the volume of trade and the average price of fabric in the seller-buyer-product-year combination (both in logs), q_{sbjy} and p^f_{sbjy} and the size of the buyer, excluding the seller q_{b-s}. Each specification introduces a further control, one per row, whose coefficient is labelled as β_1 and reported on the left hand side of the table. Definitions for the <u>Seasonality</u> panel: In these regressions, we allow for four seasons in a year, but results are similar when we allow for six seasons in a year. $\#Seasons_{sb}$ collects the number of seasons (1 to 4) the buyer and seller trade in, throughout the sample. $Share_{sb}^{\overline{Season}}$ and HHI_{sb}^{Season} collect the share of the largest season in the relationship and a Herfindhal index of concentration in seasons in the relationship. This index is normalized to range between zero and 1, where 1 is full concentration in one season. $Share_{sby}^{\overline{Season}}$ is the share of the largest season in the relationship's trade of that year. The last two specifications, 5 and 6, include dummies for the largest season the buyer and seller are trading in, in the corresponding year (or year-product). The baseline is the January-March season so $\overline{Season}_{sby}$ and $\overline{Season}_{sbjy}$ are dummies for seasons 2, 3 or 4. Definitions for the <u>Specialization</u> panel: In these regressions, a product is a six-digit HS code. $\#Prod._{sb}$ ($\#Prod._{sby}$) collects the number of products the buyer and seller trade in, throughout the sample (year). $Share_{sb}^{\overline{Prod.}}$ and $HHI_{sb}^{Prod.}$ collect the share of the largest product in the relationship and a Herfindhal index of concentration in products in the relationship. This index is normalized to range between zero and 1, where 1 is full concentration in one product. $Share_{sby}^{\overline{Product}}$ is the share of the largest product in the relationship's trade of that year. $Main_{sb}^{Prod.}$ takes value one if the product is the largest in the relationship. The complexity measures are generated as follows. For each order in the relationship, we count the number of types (origins) of fabric for producing this order. We weight each order by its size and generate a buyer-seller-product-year measure of the complexity of the garments traded, based on these average number of fabrics or origins. Definitions for the <u>Partners</u> panel: $\#Sellers_{bjy}$ count the number of different sellers the buyer is trading with in a year-product combination. The shares of the seller in the buyer and buyer-year combinations are computed using volumes.

C.1.2 Alternative Measures of the Relational Characteristic

The results on relational sourcing in the main text made use of a synthetic metric collecting the (unobservable) buyer-level relational characteristic. We address the robustness of our statements to alternative ways of constructing this metric. We do so by retaining the specification used in Table 6, and subsequently changing the construction of the regressor of interest. The results of these exercises are collected in Appendix Table C2, whose rows correspond to those different constructions.

The relational characteristic as studied in Table 6 was constructed as a dummy variable taking value one if the negative of the weighted average (number of) sellers-to-shipments ratio for the buyer falls above the 90^{th} percentile of the distribution. In specifications 1 and 2 of C2, the cutoff for the relational dummy is set to be the 85^{th} percentile or the 95^{th} percentile. In specifications 3 and 4 we use a metric similar to the one used in Heise et al. (2017), both in a continuous and discrete version. We construct the ratio $\#Sellers_{bjy}/\#Shipments_{bjy}$, take logs and regress against jy fixed effects. We take these residuals, times -1 - for comparability with our metric - and aggregate at the buyer level as a weighted average using volumes. This gives the continuous metric. The dummy is generated using the 90^{th} percentile as cutoff. The final specification recomputes our baseline discrete metric (as in Table 6), but now excluding all transactions and trade partners in the relevant woven product categories, shirts and trousers. This includes all other garment products, in both woven apparel and knitwear.

The exercises in Appendix Table C2 show that the results presented in the main text are not driven by the arbitrary cutoff for the construction of the relational dummy, the inclusion of in-sample products or by favoring our metric over the one proposed in Heise et al. (2017).

Table C2: Robustness to Alternative Measures of Sourcing Strategy

			(1) p_{sbjy}	(2) mc_{sbjy}	(3) μ_{sbjy}
1	Baseline metric, cutoff 85^{th} pctile	$Relational_b^{0/1,85^{th}}$	0.042***	0.020*	0.020*
			(0.009)	(0.012)	(0.011)
2	Baseline metric, cutoff 95^{th} pctile	$Relational_b^{0/1,95^{th}}$	0.044***	0.019	0.027**
			(0.010)	(0.014)	(0.013)
3	Metric a la Heise, Pierce, Schaur Schott (2017), discrete	$Relational_b^{0/1,HPSS}$	0.039***	0.012	0.029**
			(0.009)	(0.013)	(0.011)
4	Metric a la Heise, Pierce, Schaur Schott (2017), continuous	$Relational_b^{Cont.,HPSS}$	0.037***	0.025***	0.013**
			(0.005)	(0.007)	(0.006)
5	Baseline metric, excluding relevant products, cutoff 90^{th} pctile	$Relational_b^{0/1,-J}$	0.058***	0.026**	0.029**
			(0.009)	(0.012)	(0.011)

Standard errors in parentheses, clustered at the buyer-seller level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table combines the results of multiple OLS regressions, all running on the same 8,169 observations.The outcomes correspond to the weighted average price, marginal cost and markup factor in the seller-buyer-product-year combination, p_{sbjy}, mc_{sbjy} and μ_{sbjy} respectively. All specifications condition on seller-product-year and destination fixed effects. They also include controls for traded volumes, the price of fabric and the size of the buyer (q_{sbjy}, p^f_{sbjy}, q_{b-s}). Specification 1 and 2 alter our baseline metric for the relational characteristic by setting the cutoff for the relational dummy ito be the 85^{th} percentile or the 95^{th} percentile, respectively. In rows 3 and 4 we construct the ratio $\#Sellers_{bjy}/\#Shipments_{bjy}$, take logs and regress against jy fixed effects. We take these residuals, times -1 and aggregate at the buyer level as a weighted average using volumes. On this continuous metric (specification 4), we take the 90^{th} percentile as cutoff to construct a dummy. These two correspond, as much as possible, to the constructions in Heise et al. (2017). In row 5, we recompute our baseline discrete metric (as in Table 6), but now exclude all transactions and trade partners in the relevant woven product categories, shirts and trousers. This includes all other garment products, in both woven apparel and knitwear. In this case, we drop 210 seller-buyer-product-year observations corresponding to buyers fully specialized in the products under study.

C.2 Additional Results

C.2.1 Relationship's Dynamics

$$y_{sby} = \delta_{sy} + \delta_{sb} + \beta_1 Age^{y}_{sb} + \beta_2 Relational_b \times Age^{y}_{sb} + \boldsymbol{\gamma}'\boldsymbol{X}_{sby} + \varepsilon_{sby}$$

The outcomes and controls mimic the specification presented in columns (7)-(9) of Table 6. The novelty in this exercise is that we condition on seller-buyer fixed effects and study the the evolution of the relationship over time, using the (log) age of the relationship in years.[44] Note that the relationship fixed effect absorbs the 'intercept' effect of the relational characteristic, as well as the common effect of the country of destination. We are interested in β_1, the average relationship between costs and markups with age, and β_2 capturing any additional shifter on the age correlation, corresponding to relational buyers.

Appendix Table C3 presents the results of this specification. Columns (1) to (3) show that, once we condition on relationship-specific and seller-year-specific means, prices, costs and markups do not seem to change significantly over the course of the relationship. Column (4), however, suggests that, conditional on survival, volumes tend to grow in all relationships, but in particular in those sustained with relational buyers. This is consistent with the promise of future trade acting as a mechanism to induce compliance of suppliers in relational arrangements.

[44]For the construction of the age variable, Age^{y}_{sb}, we take the first month we observe the buyer and seller trading (any product) in the data as the start of the relationship. For every buyer-seller-year combination we construct the age of the relationship as the log of the number of years (number of months / 12) elapsed from the start of the relationship to the start of the calendar year.

Table C3: Prices, Costs and Markups within Relationships

	(1)	(2)	(3)	(4)	(5)
	p_{sby}	mc_{sby}	μ_{sby}	q_{sby}	π_{sby}
Age_{sb}^{y}	-0.007	-0.007	0.002	0.062**	0.005
	(0.005)	(0.008)	(0.007)	(0.029)	(0.044)
$Relational_b^{0/1}$=1 × Age_{sb}^{y}	0.003	-0.004	0.005	0.071**	0.092**
	(0.004)	(0.008)	(0.007)	(0.031)	(0.045)
FEs	sb,sy	sb,sy	sb,sy	sb,sy	sb,sy
All Controls	Yes	Yes	Yes	Yes	Yes
R^2	0.90	0.83	0.72	0.87	0.84
Obs.	5,478	5,478	5,478	5,478	4,626

Standard errors in parentheses, clustered at the buyer-seller level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcomes correspond to the weighted average price, marginal cost and markup factor in the seller-buyer-year combination, p_{sby}, mc_{sby} and μ_{sby} respectively for the first three columns and, for the last two, the traded volumes, q_{sby} and profits, π_{sby} constructed as the product between markup values and volumes. All outcomes are in logs. The regressor $Relational_b^{0/1}$, is a dummy taking value one if the buyer belongs to the top 10 percent of the distribution of the relational characteristic, constructed as described in the main text. For the construction of the age variable, Age_{sb}^{y}, we take the first month we observe the buyer and seller trading (any product) in the data as the start of the relationship. For every buyer-seller-year combination we construct the age of the relationship as the log of the number of years (number of months / 12) elapsed from the start of the relationship to the start of the calendar year. All specifications condition on seller-year and buyer-seller fixed effects. In addition, the include controls for the price of fabric (p_{sby}^{f}), the size of the buyer (q_{b-s}). In addition, columns (1)-(3) also control for traded volumes, q_{sby}.

C.2.2 Relational Buyers and Firm's Profits

The overall effect that trading with relational buyers has on seller-level outcomes remains to be explored. We do so by offering a specification of the form:

$$y_{sy} = \delta_s + \delta_y + \beta_1 Share_{sy}^{Rel.} + \boldsymbol{\gamma}'\boldsymbol{X}_{sy} + \epsilon_{sy}.$$

Outcomes now are aggregated at the level of the seller-year combination and the fixed effects have been adjusted to control for all seller and year specific variation. The regressor of interest, $Share_{sy}^{Rel.}$, collects the share of the volume traded by the seller-year combination that corresponds to buyers classified as relational. With this, we study variation in prices, costs, markups and profits, for a seller, as the portfolio of buyers changes over time.

The results of this exercise are presented in Appendix Table C4. Columns (1) and (2) show that the higher the incidence of relational buyers in the seller's portfolio, the higher the prices with constant marginal costs. These specifications condition on the volume the seller exports, q_{sy}, and also on the average size across the buyers the seller is trading with. With these, those higher prices are not induced by the seller trading more, with larger buyers. Similarly, we condition of the average price of the fabric, p_{sy}^f, and as such, control for changes in the quality composition of the seller's trade.

The result of such variation in costs and prices reflects in (imprecisely estimated) higher markups when the seller trades more with relational buyers (Column (3)). This, together with volumes increasing in the incidence of relational sourcing (Column (4)), translates into higher profits (Column (5)). Taken altogether, the evidence suggests that when sellers trade more with relational buyers, they *perform better*.

Table C4: Sellers Trading with Relational Buyers

	(1) p_{sy}	(2) mc_{sy}	(3) μ_{sy}	(4) q_{sy}	(5) π_{sy}
$Share_{sy}^{Rel.}$	0.039**	0.010	0.023	0.257***	0.328**
	(0.017)	(0.023)	(0.022)	(0.095)	(0.131)
FEs	s,y	s,y	s,y	s,y	s,y
q_{sy}, p_{sy}^f	Yes	Yes	Yes	No	No
$Ave(Size_b)_{sy}$	Yes	Yes	Yes	Yes	Yes
R^2	0.77	0.64	0.37	0.61	0.55
Obs.	3,146	3,146	3,146	3,146	2,990

Standard errors in parentheses, clustered at the seller level. $*(p < 0.10)$, $**(p < 0.05)$, $***(p < 0.01)$. The outcomes are in logs in all cases, correspond to seller-year combinations and, in order, are: the marginal cost (mc_{sy}), the price of the garment (p_{sy}), the markup factor (μ_{sy}), the quantities traded (q_{sy}) and profits (π_{sy}). The regressor of interest, $Share_{sy}^{Rel.}$, is the share of the volume traded by the seller-year combination that corresponds to buyers classified as relational. All specifications include seller and year fixed effects and control for the average size of the buyer trading with the seller-year ($Ave(Size_b)_{sy}$). Columns (1) to (3) also control for the volume of trade, q_{sy} and all columns condition on the price of fabric (p_{sy}^f). The drop in the number of observations in the last column correspond to seller-year combinations with average markup factors below one.