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

Julia Cajal-Grossi† Rocco Macchiavello‡ Guillermo Noguera§

April 6, 2022

Abstract

We study how suppliers' markups vary across buyers adopting different sourcing strategies in the Bangladeshi garment sector. We distinguish between buyers with relational versus spot sourcing strategies. We show that a buyer's approach to sourcing is correlated across product-origin combinations and that buyer fixed effects explain most of the variation in sourcing strategies – suggesting that these depend at least in part on organizational capabilities. We build a sourcing model with imperfect contract enforcement and idiosyncratic shocks to suppliers. In equilibrium, ex-ante identical buyers adopt different sourcing strategies: relational buyers pay higher prices and markups than spot buyers and secure reliable deliveries while spot buyers occasionally suffer delivery failures. Consistent with the model's predictions, we find that Bangladeshi suppliers earn higher prices on export orders produced for relational buyers compared to those produced for spot buyers. Matching the inputs used to produce specific orders, we find no difference in the utilization or prices of fabric and labor between orders produced for relational and spot buyers. We derive the conditions under which the data allow us to recover within seller-product-time differences in markups across orders; we establish that relational buyers pay significantly higher markups relative to spot buyers for comparable orders. We discuss alternative mechanisms and policy implications.

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

---

*We are particularly indebted to Chris Woodruff with whom this project was started, and to Pol Antràs, Nathan Nunn and four anonymous referees for their insightful comments. We thank David Atkin, Johannes Boehm, Laura Boudreau, Davide Del Prete, Dave Donaldson, Bob Gibbons, Penny Goldberg, Amit Khandelwal, Kalina Manova, Nina Pavcnik, Tristan Reed, Tim Schimdt-Einselhor, John Sutton, Catherine Thomas, John VanReenen and Eric Verhoogen for very helpful discussions. This project benefited from excellent research assistance from Madhav Malhotra. We are grateful to participants at seminars, workshops and conferences at Aarhus, Bank of Italy, Barnard, Chicago, Columbia-NYU, Düsseldorf, FREIT-EITI, Harvard, Kellogg-Northwestern, London School of Economics, Mannheim, McGill, Munich, Oslo, Oxford, Sciences Po, Warwick, WIEN, World Bank and Yale. We gratefully acknowledge financial support from the Centre for Competitive Advantage in the Global Economy (CAGE) and ERC Funded Project *Sharing Gains* 818767.

†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

Firm-level decisions play a critical role in explaining aggregate productivity (Van Reenen, 2018; Goldberg et al., 2010), as well as the structure of trade flows (Antràs, 2016; Bernard et al., 2007; Gereffi et al., 2005). The diffusion of just-in-time inventory systems and outsourcing – including across borders – have turned firms' approaches to sourcing into a particularly important strategic decision (Dyer et al., 1998). Different ways of organizing sourcing must be coordinated with other operational processes (Cooper and Ellram, 1993), and require specific internal structures and suitable management practices (Milgrom and Roberts, 1990, 1995). Firms, even within narrowly defined industries, end up developing distinctive approaches to sourcing (Helper and Henderson, 2014). At one extreme – which we label *spot* sourcing – the buyer's purchases are spread among multiple suppliers to "improve the firm's bargaining power" (Porter, 1980, pp. 123). Buyers keep suppliers at arm's length, avoid any type of commitment, allocate short-term orders to the lowest bidders and bear the costs of suppliers' non-performance. At the other extreme – which we label *relational* sourcing – orders are allocated to few suppliers with whom the buyer develops long-term relationships to incentivize behavior that might otherwise be difficult to contract upon.

Buyers' sourcing can have far-reaching implications for suppliers. In particular, existing theories highlight the role of markups: under spot sourcing, suppliers' markups are squeezed by intense competition; under relational sourcing, buyers may pay higher markups to incentivize suppliers (see, e.g., Taylor and Wiggins, 1997). To our knowledge, this hypothesis has not been tested empirically. Do suppliers indeed earn higher markups from relational buyers? And, if so, are buyers' choices of sourcing socially efficient or is there scope for policy intervention? Answering these questions is as important – particularly so in developing countries, where buyers can act as potent vehicles for upgrading (World Bank, 2020) – as is challenging. The first challenge is that measuring the markups earned from different buyers requires knowledge of the prices obtained from – as well as the costs incurred from – supplying a specific buyer. While information on prices at the required level of detail is increasingly available, costs remain difficult to estimate since the amounts and prices of inputs used to produce for specific buyers are typically unobserved. The second challenge is that buyers' sourcing strategies are not directly observed either. As a result, one must construct proxies based on observable sourcing behavior which, however, may correlate with prices and markups through multiple channels.

This paper studies the prices and markups earned by Bangladeshi woven garment exporters supplying foreign buyers with different sourcing strategies. In addition to its intrinsic interest, unique features of the context allow us to make progress on the empirical front.[1] For woven garments we observe the type, prices and amounts of the main variable inputs (fabric and labor employed on sewing lines) used to produce specific export orders for different buyers, thereby overcoming the first challenge. In addition, many buyers source both woven garments and knitwear from Bangladesh. Due to differences in production processes, woven garments and knitwear are produced by different exporters. This allows us to use transactions in *knitwear* to characterize buyers' sourcing strategies and correlate those with prices and markups across *woven* garments export orders, thereby overcoming the second challenge.[2]

We find that Bangladeshi exporters earn higher prices and markups on otherwise identical export orders produced for relational buyers compared to spot buyers. Orders exported to relational buyers do not differ in the utilization and prices of fabric and sewing labor relative to those produced for spot buyers. We interpret these patterns – which are robust across a wide range of specifications – through the lens of a model in which suppliers are hit by idiosyncratic shocks and struggle to supply buyers reliably. Imperfect contract enforcement implies that spot contracts are effective in securing supply under 'business as usual' conditions, but fail to provide adequate incentives when suppliers are disrupted by shocks. This introduces a trade-off: relational buyers are able to secure reliable supplies, but pay higher prices; spot buyers pay lower prices but occasionally suffer delivery failures. Relative to the social optimum, the equilibrium features too few relational buyers, creating a rationale for policy intervention.

The paper proceeds in four sections. Section 2 characterizes international buyers' sourcing strategies in the garment sector. We introduce an intuitive proxy for buyers' sourcing strategies, building on the fact that relational buyers concentrate their sourcing amongst a small number of suppliers. Specifically, we compute the weighted average across product-year combinations of the number of suppliers a buyer sources from, normalized by scale. This yields a cross-sectional characterization of buyers' sourcing strategies that maps closely to qualitative accounts in the industry. In Bangladesh, buyers' sourcing strategies in knitwear and woven garments are strongly correlated. Computing the proxy for other sourcing origins reveals that a buyer's approach to sourcing is also correlated across the origins they source

[1] 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.

[2] In the paper we use *buyer* to refer to the firm that purchases ready-made garments from Bangladesh. When this creates no ambiguity, we sometimes refer to these firms as *retailers*. Similarly, we use *seller*, *supplier* or *exporter* to refer to the garment manufacturer that sells to international buyers. Finally, we use *upstream supplier* in passages where we refer to (foreign) firms that sell fabric to the Bangladeshi exporters.

from, which in turns means that buyer-level fixed effects explain a large share of the variation in sourcing strategies across products and origins. This observation, which is consistent with organizational capabilities underpinning buyers' global approaches to sourcing, justifies our buyer-level – as opposed to buyer-seller-level – characterization of sourcing strategies. Three further motivating facts shape our theoretical model: (i) there is substantial unexplained variation in buyers' choice of sourcing strategies; (ii) exporters tend to supply both relational and spot buyers; and (iii) unexpected disruptions to shipping – due for instance to hartals – are common in the industry, and exporters appear to prioritize orders for relational buyers during difficult times.

Building on these facts, Section 3 introduces a model in which buyers seek to secure reliable supplies from sellers that facing idiosyncratic shocks. Spot sourcing adequately ensures delivery under normal conditions but fails to incentivize sellers to undertake costly actions and avoid delivery failures when shocks occur. This mechanism creates a rationale for relational sourcing. In equilibrium, ex-ante identical buyers choose different sourcing strategies: some buyers invest in organizational capabilities and become relational – i.e., are able to make clear and credible promises of higher prices and markups in exchange of reliable deliveries from their suppliers; other buyers do not invest, source through spot contracts at low prices, but are unable to secure reliable supplies. The equilibrium of the model is consistent with the three motivating facts described above and with our empirical proxy for buyers' sourcing strategies.

The main prediction of the model – tested and confirmed in Section 4 – is that relational buyers pay higher prices and markups to suppliers relative to spot buyers sourcing the same product, from the same supplier, at the same time. Within seller-product-year combinations, the pattern that relational buyers indeed pay higher prices is indeed extremely robust. We take advantage of the unique features of our data to investigate differences in variable costs that suppliers incur when producing orders for relational buyers compared to spot buyers. In addition to standard information on the output side (quantity, prices and product type), we observe the amount, price and type of fabric used in the production of each export order. Conditional on seller-product-year fixed effects, the buyer's sourcing strategy does not correlate with the order-level buy-to-ship ratio (a measure of fabric efficiency) nor with the price of fabric. For a sample of factories, we also observe labor utilization and efficiency on the sewing lines – the most labor intensive step in garment production – and can exploit workers' surveys that contain information on labor characteristics and wages. These data confirm that orders produced for relational and spot buyers are sewed by workers of comparable skills, earning similar wages and working with similar efficiency.

The evidence on fabric and sewing labor is consistent with the model's prediction that the higher prices paid by relational buyers reflect a higher markup. However, there might be unobserved order-level variable costs that systematically vary between orders produced for relational and spot buyers. We develop an empirical framework, compatible with our theoretical model, that clarifies the conditions under which, given the available data, we can recover within seller-product-time *differences* in markups across orders. This enables us to conduct a precise test of the model's main prediction.[3] We show that exporters earn higher markups from orders produced for relational buyers relative to spot buyers.

Section 5 unpacks multiple facets of our analysis, discusses alternative mechanisms, provides a quantification of the value of supplying relational buyers, and sets forth policy implications of our study. First, we revisit our approach to characterizing the sourcing strategy at the buyer – as opposed to the buyer-seller – level. We then complement our across-buyers analysis with an event study around the shift from spot sourcing to relational sourcing rolled out in the global supply chain of VF Corporation – a large buyer of garments. Both exercises confirm that a buyer's relational approach to sourcing is associated with higher supplier's markups. For concreteness, the model in Section 3 focuses on reliability as the mechanism that gives rise to relational sourcing in our context. While we believe this particular mechanism to be important, we certainly do not contend that it is the only mechanism at play. We thus discuss the reliability mechanism in greater detail as well as alternative mechanisms (including differences in market power, search behavior, product quality and demand assurance) that could result in differences in markups earned from different buyers. The value of supplying relational buyers is substantial. A conservative estimate is that a shift in sourcing strategy from the average buyer in the sample to the relational sourcing adopted by The Gap (a shift of about one standard deviation in our empirical proxy) is associated with an 11% increase in the average markup value. A back-of-the-envelope calculation suggests that the net present value of supplying a relational buyer is equal to at least 30% of the yearly profits in the relationship.

Finally, we discuss policy implications. In our model, prices do not adjust in such a way that the buyers' choice of sourcing strategy reflects their marginal contribution to overall surplus – i.e., a relational buyer exerts a positive pecuniary externality on other market participants. Relative to the social optimum, too few buyers become relational. As a result, a planner may want to subsidize their entry.[4] In light of this, our conceptualization of the

[3] The main condition is a production function that features (log-)separability of fabric relative to other costs – an assumption justifiable in light of the two-step production process for garments. Other than that, the framework allows for an elasticity of output with respect to fabric that varies at the seller-product-time level and for an arbitrary number of other inputs that sellers might be able to chose freely (e.g., casual labor, bribes) or subject to capacity constraints (e.g., managerial labor and attention).

[4] In our model, limits to ex-ante transfers between buyers and sellers imply that suppliers earn rents. In such a case, even an export promotion agency that only cares about exporters profits might want to subsidize

sourcing strategy as a buyer-level attribute – as opposed to an exclusive emphasis on the relational nature at the buyer-seller pair – is of practical relevance. Even though organizational level capabilities underpin a buyer's ability to establish long-term relationships with suppliers, the relational contract with a particular supplier remains deeply rooted in both parties' specific circumstances (Gibbons and Henderson, 2012; Baker et al., 2002). It is thus unlikely that policy makers can improve specific relationships between exporters and buyers. If certain buyers possess organizational capabilities that make them valuable relational partners, and if such capabilities generate benefits for suppliers, an actionable margin for policy opens up. It might be possible to attract such buyers, e.g., by subsidizing visits to the country or targeting factors that favor their entry. In our model, the benefits come from the relational buyers' ability to overcome contracting problems that hinder reliable supply. Exploring other sources of benefits (e.g., a stable demand; upgrading) is a priority area for future research. Section 6 concludes.[5]

Related Literature. Differences in sourcing strategies within narrow industries appear in many contexts, including automotive (Helper and Sako, 1997; Nishiguchi, 1994; Richardson, 1993); electronics and machinery (De Toni and Nassimbeni, 2000); aerospace (Masten, 1984); and apparels (Gereffi, 1999). These differences echo those in the adoption of lean management practices (Bloom and Van Reenen, 2010; Bloom and Van Reenen, 2007). Unlike the evidence on management practices, accounts of sourcing strategies are mostly qualitative.[6] We contribute novel data, a quantitative characterization of sourcing strategies and how they correlate with suppliers' prices and markups as well as suggestive evidence of the importance of organizational capabilities in enabling relational contracting (Gibbons and Henderson, 2012). Taylor and Wiggins (1997) model spot and relational sourcing strategies.[7] We highlight reliability as a mechanism that induces ex-ante identical buyers to choose different sourcing strategies in equilibrium – thereby offering a rationale for within-industry variation in sourcing strategies.

Sociologists (see, e.g., Ponte et al., 2019 and Gereffi, 1999) and economists (see Antràs,

the entry of relational buyers.

[5]We collect supplementary material in several online appendices. Online Appendix A describes the data sources; Online Appendix B extends our econometric approach to the estimation of markups in levels; Online Appendix C discusses various robustness exercises pertaining to our main empirical findings; Online Appendix D collects additional tables and figures.

[6]Macchiavello and Morjaria (2021) show that the adoption of relational sourcing practices with supplying farmers is associated with better performance across mills in the Rwanda coffee sector.

[7]Heise et al. (2021) characterize sourcing systems using a measure inspired by Taylor and Wiggins (1997). Our measure is similar to theirs and also consistent with the model. Startz (2021) structurally estimates a model along the lines of Taylor and Wiggins (1997) using tailored surveys on the sourcing decisions of Nigerian importers.

2016, Antràs, 2020 and Macchiavello, 2022 for reviews) alike have emphasized the relational nature of global value chains. Macchiavello and Morjaria (2015) show that Kenyan rose exporters hit by an unanticipated shock prioritized buyers with whom they had valuable relationships. Exploiting unanticipated surges in international coffee prices, Blouin and Macchiavello (2019) show that suppliers' opportunism – as opposed to *force majeure* – causes many delivery failures. These papers provide a test of the reliability mechanism in our model and *infer* the value of relationships from observed responses to shocks. We borrow from these papers the idea that buyers' concerns over reliability are important drivers of relational contracting in our context as well. In contrast to these papers, we directly *measure* the higher markups earned when supplying relational buyers and we characterize the sourcing strategy at the buyer level, rather than focusing on buyer-seller relationships. Besides the novelty, this distinction has practical relevance for policy design. Recent contributions on buyers' role in global value chains include Amengual and Distelhorst (2019)'s study of the impact of a change in the global sourcing approach at The Gap Inc. on suppliers' compliance; Boudreau (2020)'s evaluation of a buyer-driven initiative aimed at enforcing worker-manager safety committees in Bangladeshi garment factories; and Macchiavello and Miquel-Florensa (2019)'s analysis of a buyer-driven quality upgrading program in the Colombian coffee chain.[8]

A vast body of work studies firms' upgrading from exporting and FDI in developing countries (see Verhoogen, 2021 for a review). For example, Atkin et al. (2017) show that randomly assigned export orders induced quality upgrading among Egyptian rug producers (see also Chor et al., 2021; Pavcnik, 2002). Alfaro Urena et al. (2021) find that Costa Rican suppliers increase sales, employment and productivity after starting to supply multinational corporations. We highlight buyers' sourcing strategies as an upgrading dimension. We do so by relaxing data constraints that have hindered progress in the estimation of markups in multi-product firms (Garcia-Marin and Voigtländer, 2019; De Loecker et al., 2016).[9]

[8]The literature has also studied vertical integration in global value chains (see e.g., Boehm and Sonntag, 2019; Antràs and Chor, 2013; Costinot et al., 2011; Alfaro and Charlton, 2009; and Antràs, 2003). Macchiavello and Miquel-Florensa (2018) compare integrated and relational sourcing in the coffee sector. We abstract from vertical integration as it is virtually nonexistent in our context.

[9]Recent contributions that use within firms data and are closely related include Brandt et al. (2020)'s analysis of vertical integration in the Chinese steel industry; Adhvaryu et al. (2020)'s study of workers allocation across production lines in a large Indian garment exporter; De Roux et al. (2020)'s analysis of the relationship between product quality and markups in a large Colombian coffee exporter; Atkin et al. (2015)'s survey based study of markups in the Sialkot soccer ball cluster in Pakistan.

2 Buyers' Sourcing Strategies

This section describes the sourcing strategies of international garment buyers. Section 2.1 builds upon the supply-chain literature in management as well as on case studies to motivate our approach that characterizes the sourcing strategy as a buyer-level attribute. We conceptualize sourcing as a bundle of complementary practices that are supported by organization-wide systems. This leads buyers to adopt consistent sourcing strategies across their supplier base. Section 2.2 introduces our measure of relational sourcing. Section 2.3 presents novel evidence on international sourcing of garments, provides a formal test of our approach, and describes further motivating evidence that guides the model we develop in Section 3.

2.1 Conceptualizing Sourcing

The introduction of lean management practices and just-in-time inventory systems has enhanced the importance of sourcing as a key strategic function (Dyer et al., 1998). Reliability of supply – the ability to guarantee supply in due time and form under most contingencies – has become a key part of firms' competitive advantage in many industries, including garments, as already discussed in Section 1. Reliability, however, is hard to contract upon – especially in developing countries and in the context of international sourcing. Buyers generally deal with this limitation by pursuing either one of two stylized approaches to sourcing: a *spot* sourcing strategy at one end or a *relational* sourcing strategy at the other. This conceptual distinction has its origin in the literature comparing the sourcing practices of American versus Japanese car manufacturers. The two respective models are sometimes referred to as 'adversarial' or 'American-style' sourcing in contrast to 'collaborative' or 'Japanese-style' sourcing (see, e.g., Helper and Sako, 1997; McMillan, 1990). Under *spot* sourcing, buyers purchase from multiple suppliers, with whom trade relationships tend to be short-lived, ending as a result of out-bids from cheaper suppliers. Procurement orders tend to be large and either one-off or sporadic. Under *relational* sourcing, buyers concentrate orders on a small number of suppliers with whom they develop relational contracts — defined as "*informal agreements sustained by the value of future relationships*" (Baker et al., 2002).

Adopting one or the other sourcing strategy requires an array of compatible structures and practices within the buyer's organization. The literature on supply chain management highlights how the main operational processes (source, make, deliver) must be coordinated across functions within the firm (Cooper and Ellram, 1993). This coordination hinges on complementarities between management practices and internal structures and processes (Milgrom and Roberts, 1990). The ability to implement relational sourcing thus depends on organization-wide capabilities (Helper and Henderson, 2014). For example, specific systems of (inward and outward) communication and knowledge diffusion are needed to foster "clarity" in the relational contracts with suppliers (Gibbons and Henderson, 2012). In the car industry, relational sourcing is characterized by a deep integration between the procurement, production and design functions. Such integration goes hand-in-hand with adequate human resources policies. Collaborative relationships between the sourcing department and other functions can be achieved by rotating personnel across functions and by avoiding excessively high powered incentives that create conflict between functions. Similarly, relationships with suppliers are fostered by avoiding frequent rotation of purchasing agents. By contrast, spot sourcing systems do not require integration between the sourcing and design functions. Purchasing agents are given high powered incentives and rotations are frequently made to avoid capture by suppliers. Accordingly, post-procurement functions, such as in-house quality control and product cycle integration, are critical.

These considerations are relevant in the context of the international sourcing of garments (see, e.g., Gereffi, 1999). Several case studies that document organization-wide restructuring of sourcing strategies have illustrated the point. For example, VF, a multi-brand U.S. apparel retailer, shifted its approach to sourcing *globally* from a spot-style of procurement in favor of a relational approach – the *Third Way* – in the mid-2000s (see Pisano and Adams, 2009). Quoting from the case study: "*Historically, apparel companies and apparel suppliers showed little loyalty to one another. Contracts were short-term (typically one season). In their aggressive pursuit of low costs, apparel companies drove hard bargains on pricing and freely shifted production from one supplier to another. There were no guarantees in either direction. Every year, suppliers had to bid to get new business from a company and never guaranteed production capacity beyond a very short time horizon [...] They also took on products from as many companies as possible (often competitors) to diversify their risks.*" Similarly, Nike also shifted towards a more relational approach to sourcing. This culminated in 2009 in a company-wide reorganization in which a new corporate division merged the Social Compliance Team into the Global Sourcing Department (see Nien-he et al., 2019). Again, quoting from the case study: "*Sourcing decisions are often decoupled from the enforcement of private regulation [...], resulting in a tension between the two functions*" and it is "*not uncommon to hear complaints from [Social Compliance] managers that their mission is not taken seriously by their colleagues in purchasing departments*". The merging of two previously distinct functions at the headquarters level is an organizational change that impacts sourcing across Nike's global supplier base.

The organizational complementarities that support one or the other sourcing strategy thus create economies of scale and scope in the formation of relationships: a buyer with an organizational structure conducive to relational sourcing tends to trade relationally across its entire supply chain. That is, relative to spot buyers, companies like VF and Nike commonly source relationally across products and origins. We therefore posit that the sourcing strategy is a buyer-level attribute, rather than a relationship-specific one. We contend that this approach not only accommodates insights from the aforementioned management literature, but it also has practical policy relevance. We now introduce our buyer level measure of sourcing strategy before testing this hypothesis in the data.

2.2 Measuring Buyers' Sourcing Strategies

Relational contracts – informal arrangements sustained by the value of future interactions – are not directly observable in the data (see Macchiavello, 2022 for a review). Much of the existing empirical work, thus, resorts to relationship's age – which is instead observable – as a proxy for relational trade. There are, however, a number of drawbacks to this approach. First, repeated trade does not imply relational trade which, instead, critically hinges on the promise of future rents to induce parties to resist temptations to engage in opportunist behavior. Second, measures based on relationship's age require the researcher to deal with censoring (which is typically severe and not exogenous) and make assumptions about the demand structure across buyers. For example, using calendar time to measure relationship's age ignores that the frequency of interactions may entail different implicit commitments; on the other hand, using transaction counts risks confounding relational contracting with, e.g., buyer-specific seasonal patterns.

A distinctive feature of relational buyers is the concentration of sourcing in a small number of suppliers. We thus pursue an alternative approach and measure sourcing according to how concentrated a buyer's sourcing is on a small number of suppliers. To lend intuition, Table 2 examines the 25 largest buyers of woven garments in Bangladesh. Column (1) ranks buyers according to their market shares in the country. H&M, Walmart, and the multi-brand apparel company VF Corporation lead the board with market shares of 5.22%, 5% and 4.14% respectively, more than 500 times larger than the median buyer in the sample. Even among these large buyers, there are large differences in their approach to sourcing. For example, Levi Strauss & Co. has a reputation for developing long-term collaborative relationships with suppliers; J.C. Penney has traditionally adopted 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). Column (2) in Table 2 shows that Levi Strauss & Co. and J.C. Penney have similar market shares (2.21% and 1.96% respectively) but differ in the number of suppliers they source from: in a typical year the former only sources from 7.4 suppliers while the latter does so from 25.8 suppliers.

Guided by the observations above, we construct the buyer-level measure of relational sourcing as the weighted average of the negative of the number of sellers divided by the number of shipments at the product-year level. Specifically, we define

$$Relational_b = \sum_{jt \in \mathcal{J}_b} \left[\frac{Q_{bjt}}{Q_b} \times Relational_{bjt} \right] \quad \text{and} \quad Relational_{bjt} = -\frac{\#Sellers_{bjt}}{\#Shipments_{bjt}}, \tag{1}$$

where Q_{bjt} is the overall volume of garment sourced by buyer b in product j in year t and $\mathcal{J}_b$ is the set of all product-year combinations jt sourced by buyer b. We normalize the number of sellers by the number of shipments so that, other things equal, a buyer with a higher number of shipments per seller, or equivalently, with fewer sellers per shipment is assigned a higher value in the relational metric. Relative to a normalization based on volumes or values, the number of shipments is observed with less error and is better aligned to our model in Section 3, as well as to the model in Taylor and Wiggins (1997), which in turn is used by Heise et al. (2021) to build a similar measure. Online Appendix C.2 shows that our results are robust to alternative definitions of a buyer's sourcing strategy, including those that normalize the number of partners using traded volumes or values, as well as measures based on the average duration of a buyer's relationships.

The $Relational_b$ metric produces a sensible ordering of buyers. Column (3) in Table 2 ranks the largest 25 buyers according to their sourcing strategies (the first one being the most relational buyer). The ranking maps closely to qualitative accounts in industry publications. For example, Levi Strauss & Co. ranks second, close to other large buyers known for their relational approach to sourcing, such as The Gap and H&M, ranked first and third, respectively. Large European discount retailers (e.g., Kik Textilen and JCK), known for a spot sourcing strategy, appear lower in the ranking. Zara's owner Inditex is ranked a bit lower: during the sample period, Zara sourced relationally from suppliers located near its headquarters in North-Western Spain and sourced from Bangladeshi suppliers through the traditional spot approach (see Ghemawat and Nueno Iniesta, 2006).

In the empirical analysis in Section 4 we correlate buyers' sourcing strategies with order level outcomes, such as prices and markups. A potential source of concern is that our proxy for sourcing strategy relies on features of the buyers' transactions that might be correlated with these outcomes of interest other than through the relational or spot nature of sourcing. To assuage such concerns, the empirical analysis takes advantage of the fact that garment exports in Bangladesh are concentrated in two distinct sets of products – woven garments and knitwear (see Online Appendix A.1). The production process of the two types of garments is radically different and the sets of exporters in the two sub-sectors are largely disjoint. Our analysis focuses on woven garments (for which we can match inputs and output at the order level). We thus construct our metric separately on two sets of products: products *included* in our analysis, $\mathcal{J}^+$, and products *excluded* from our analysis, $\mathcal{J}^-$, with $\mathcal{J} = \mathcal{J}^+ \cup \mathcal{J}^-$ and $\mathcal{J}^+ \cap \mathcal{J}^- = \emptyset$. The metric of relational sourcing that we take to the data is then,

$$Relational_b = \sum_{jt \in \mathcal{J}^-} \left[\frac{Q_{bjt}}{Q_{b\mathcal{J}^-}} \times Relational_{bjt} \right]. \tag{2}$$

Figure C1 in the Online Appendix shows that the sourcing of the buyer is strongly correlated between included and excluded products. This is not only reassuring for our approach, but it also responds to a broader pattern we document in the next subsection.

Appendix Figure D1 presents the distribution of the proxy for sourcing strategies across woven garment buyers in Bangladesh, computing the metric defined in equation (2) on excluded products. By construction, the measure ranges the interval $[-1, 0)$, with -1 the most spot sourcing and $\rightarrow 0$ the most relational. The median of this distribution is -0.344, meaning that the median buyer has just over 34 suppliers for every 100 transactions. The most relational buyers, conversely, have one supplier for every 100 transactions on average. To fix ideas, we consider two large international buyers presented in Table 2. Across the main woven product categories, H&M trades with 157 different sellers throughout our sample period, allocating an average of 847 shipments to each. Instead, KiK trades with 206 sellers, allocating an average of 26 shipments to each of them. As a result, H&M is located at the top end of the distribution of our sourcing metric ($Relational_{H\&M} = -0.021$), while KiK is almost one (0.83) standard deviation below ($Relational_{Kik} = -0.241$). We will return to this comparison, as an illustrative example in our quantification in Section 5.4.

Our empirical analysis studies how export-order level outcomes vary with the sourcing strategy of the buyer. By construction, the proxy for buyers' sourcing strategies includes several potential forms of measurement error that may lead to attenuation bias – i.e., making differences across buyers harder to detect in the data. First, buyers may tailor their sourcing practices to specific suppliers. In Section 5.1 we introduce bilateral proxies for relationalness and confirm that their inclusion strengthens our baseline results.[10] Second, buyers might change their sourcing strategy over time. If that is the case, our reliance on a time-invariant measure makes identification elusive. Indeed, in Section 5.2 we present an event study leveraging a large buyer's change in sourcing strategy and find significantly larger estimates

[10] Concerns that proxies for relational sourcing might correlate with unobserved features of the transaction that also correlate with outcomes of interest – such as prices and markups – arise, a *fortiori*, when proxying relational trade at the buyer-seller level. This provides a further rationale for our choice to focus on buyer-level proxies for relational sourcing.

than in our baseline. Third, in some cases, buyers might source through specialized sourcing intermediaries (such as Li & Fung). From collaborations with some large retailers, we know that sometimes buyers use *both* direct sourcing and intermediaries. These intermediaries are paid a commission and very rarely take ownership of the good – particularly so in Bangladesh due to the widespread use of back-to-back letters of credit – and thus might not appear in the customs data. This introduces measurement error: a spot buyer that uses an intermediary that has relationships with the exporters will look similar to a relational buyer according to our measure – making it harder to detect any difference. Finally, our strategy to focus on excluded products alleviates endogeneity concerns, but also introduces measurement error.

2.3 Motivating Evidence on Sourcing Patterns

Having introduced an empirical measure for buyers' sourcing strategies, this subsection first provides a formal test of the idea that buyer-level capabilities are important drivers of sourcing strategies. It then presents three motivating facts that guide our model in Section 3.

2.3.1 Sourcing Strategies as a Buyer-Specific Characteristic

We begin by describing novel evidence on buyers' sourcing strategies using transaction-level customs records of garment exports (defined at the HS6 level – like in the rest of the paper) from Bangladesh, Ethiopia, India, Indonesia, Pakistan and Vietnam. Taken together, these countries account for 36% of garments exports from developing countries into the U.S. and the EU.[11] Our working sample contains approximately 16.5 million transactions, across the six countries, corresponding to almost 10 thousand buyers and 29 thousand sellers. For each buyer-product-country, we construct our proxy $Relational_{bjc}$ analogously to the definition in equation (1) above, where $c \in \{BD, ET, IN, ID, PK, VN\}$.

Appendix Figure D2 shows (stylized) scatter plots of $Relational_{bjc}$ and $Relational_{bjc'}$ for all pair-wise combinations c and c'. A positive slope indicates that a buyer b that sources product j relationally in country c tends to do so in country c' as well. This is the slope that we find in all but one pairs of countries – the sole exception being the pair Vietnam (the most advanced garment producer in our sample) and Ethiopia (which only recently began exporting large volumes of garments). For example, in HS 610442, H&M is classified at the 99^{th} percentile of the relational metric in both Bangladesh and Pakistan, and at the 96^{th} percentile in Indonesia. In HS 620459, H&M is above the 95^{th} percentile in Bangladesh,

[11] The data comes from ongoing work in Cajal-Grossi et al. (2022) and it is described in detail in Online Appendix A.3. We are grateful to Davide Del Prete for letting us use the data in this paper. We focus on the years 2018 and 2019 to avoid overlap with the onset and early development of the Covid-19 pandemic.

Indonesia, Pakistan and Vietnam. In contrast, J.C. Penney, is below the 25^{th} percentile in HS 610520 in Vietnam and Ethiopia and in HS 620429 in Bangladesh and Vietnam.

We now formally test the hypothesis that buyers' sourcing strategies are largely driven by buyer-level capabilities, thereby justifying our approach to model them at the buyer-level. Our test is inspired by, and generalizes, Monteverde and Teece (1982)'s classic study of vertical integration for 133 components used by Ford and GM. The authors test, and find empirical support for, the transaction costs economics theory of vertical integration by showing that car assemblers integrate components whose production processes generate quasi-rents in the form of specialized, non-patentable, know-how. A perhaps less appreciated finding in this classic study, however, is that the buyer's dummy accounts for a substantial share of the observed variation in vertical integration across components. This suggests that – holding a component's technical specification constant – Ford and GM differ in their overall approach to sourcing.[12]

Returning to our context, transaction cost economics (Williamson, 1971, 1975, 1985) predicts that the choice of governance form – in our case, the choice between spot vs. relational sourcing – is driven by characteristics of the product and the market in which it is being sourced. For example, products that are more differentiated (Rauch, 1999), have different fashion cycles (Woodruff, 2002), or are sourced from countries in which contracts are harder to enforce (Antràs and Foley, 2015), are more likely to be sourced relationally. Similarly, conditions in the downstream market might also influence the choice of sourcing strategy. This logic implies that origin-product fixed effects and destination-product fixed effects should account for most of the observed variation in sourcing strategy. If, instead, organizational capabilities – as opposed to transaction costs – are a key driver of sourcing strategy choices, a buyer's sourcing strategy should be correlated across the different products and origin countries the buyer sources from (as seen in Appendix Figure D2) and, quantitatively, a buyers' identity should explain a significant proportion of the variation in how sourcing is organized.

We implement a loss-of-fit exercise to quantify the relative importance of buyer fixed effects versus other factors in driving variation in sourcing strategies $Relational_{bjc}$. Appendix Table D1 reports the results. We regress $Relational_{bjc}$ on a set of fixed effects $\{\delta_i\}_{i \in I}$ and obtain the loss in model fit from removing each component from model I. Denote by b, j, c and d the buyer, product, country of origin and destination respectively. Starting from the most saturated specification with $I = \{b, jc, jd\}$, we find that buyer fixed effects account for

[12]More recently, Helper and Munasib (2021) use U.S. customs data on the imports of car parts and find that – controlling for detailed product fixed effects – Japanese owned importers source parts more relationally than American and European companies.

over 40% of the explained variation in sourcing strategies, vis-à-vis 16% and 14% explained by product-country (the origin of the garment) and product-destination (the country of the buyer), respectively. Organizational capabilities at the buyer level appear to play a key role in driving a buyers' approaches to sourcing in the industry.

2.3.2 There is Unexplained Variation in Sourcing Strategies

A question beckoned by these findings is the extent to which buyers' characteristics correlate with sourcing strategies. Due to data limitations, this paper can only provide a partial answer to this question. Appendix Table D3 shows that some of the observed dispersion in sourcing strategies can indeed be explained by buyers' characteristics. In addition to characteristics observed in the customs records (such as the buyer's main product and main destination), for up to 34% of buyers representing 53% of woven volume in our data, we can study characteristics from the Bureau van Dijk's ORBIS database, including firm size category, main activity, main domestic country, year of incorporation, assets, sales, and number of employees (as of 2010). We find suggestive evidence that relational buyers are larger, tend to operate mainly in retail and wholesale rather than in manufacturing, but are neither older or younger than spot buyers. While these patterns are informative, the richer specifications that control for main destination, main product or activity, size, and cohort year fixed effects only explain between 25 and 36% of the observed variation in sourcing strategies using buyer characteristics from our data and from data from ORBIS. In other words, there appears to be substantial unexplained variation in sourcing strategies across firms within narrowly defined industries. In light of this, and the fact that information is not available for many of the buyers in our sample, the model in Section 3 considers ex-ante identical buyers that endogenously sort into ex-post different sourcing strategies. The key insights of the model extend to a setting with ex-ante heterogeneous buyers. Exploring buyers' choices of sourcing strategy is an important avenue for future research, beyond the scope of this paper.

2.3.3 Sellers Supply Both Relational and Non-Relational Buyers

We now describe the composition of exporters' portfolios in terms of their buyers' sourcing strategies. For the sake of simplicity, we label buyers in the top decile of the distribution of the sourcing strategy proxy in equation (2) as relational. Figure 2 plots the distribution of the share of exports delivered to relational buyers across sellers. For purposes of consistency, we focus on the sample of sellers used in the empirical analysis in Section 4. Two facts emerge. First, 10% of exporters supply exclusively spot buyers, and 25% of exporters trade no more than 10% of their volumes with relational buyers. Second, among the exporters that do supply relational buyers, very few do so *exclusively.* Thus, most suppliers serve a mix of relational and spot buyers. Appendix Table D2 shows that this pattern is not driven by the specific cutoff used to define relational buyers: partitions in the top 5%, 25% and 50% yield a picture similar to the 10% cutoff. Importantly, the table shows that exporters trade with an average of 21 buyers throughout the sample period. There is substantial variation in buyers' sourcing strategies within exporters. The within-seller range in buyers' sourcing strategies – the absolute value of the difference in the sourcing metric between the seller's most and least relational buyers – is 0.72, out of a theoretical maximum range strictly less than one.

This mix of buyers is sellers' portfolios is central to our empirical analysis, which focuses on within-seller variation across buyers with different sourcing strategies. Given our focus on within-seller variation, the model in Section 3 features *ex-ante* identical sellers. In equilibrium, and consistent with Figure 2, some sellers supply only spot buyers; while others supply both spot and relational buyers. The main insights of the model extend to a scenario with *ex-ante* heterogeneous sellers. Across sellers, Appendix Table D5 in Online Appendix D shows that exporters that sell to relational buyers are larger. Conditional on size, however, these sellers *do not* export more products or to more destination countries. Exploring drivers of selection into the supply chain of buyers with different sourcing strategies remains beyond the scope of this paper.

2.3.4 Supply Shocks and Reliability towards Relational Buyers

Relational sourcing is generally used to incentivize suppliers to undertake costly actions that are hard to contract upon. In our context, several, not mutually exclusive, mechanisms are potentially at play: relational sourcing could emerge as a solution to the problem of quality control; it could improve coordination and ensure swift responses to changes in demand; or it could incentivize suppliers to deliver in a timely and reliable fashion. It is not the primary objective of this paper to contend that a particular mechanism is the *only* one at play. Numerous conversations with factories and buyers, however, suggest that shocks to suppliers' ability to deliver orders on time appear to be common in the industry. For concreteness, therefore, the model in Section 3 focuses on the reliability mechanism.

Hartals – (often violent) national political protests that are common in Bangladesh – are a frequently cited source of supply disruptions. Hartals may force exporters to reschedule shipments and/or scale back production due to the blockage of roads and ports, the closure of bureaucratic offices and (sometimes) the coercive shut down of factories. Using Bangladeshi data, Ahsan and Iqbal (2015) show that when hartals take place, garment exports drop by approximately 3%. Combining their time series on hartals with our customs records, we find similar effects in our sample. Furthermore, Appendix Table D6 shows that orders that are produced during hartals take longer to be completed – an indication that exporters may have to delay shipments to cope with the hartals' disruptive effects on production and shipping. Interestingly, however, exporters' behavior during hartals differs depending on the buyer for which the export order is produced. Delays in export shipments are significantly mitigated when the order is produced for a relational buyer. This evidence – which holds across orders produced within seller-product-time combinations as per our baseline specification in Section 4 – suggests that exporters may be able to undertake certain actions to prioritize orders for relational buyers at difficult times.

While useful for illustrative purposes, it is worth noting that hartals do not provide an ideal 'natural experiment' to test for reliability as a driver of relational sourcing. First, hartals are frequent – almost half of the export orders in our sample are affected – but very imprecisely measured. Hartals' impact on exporters is highly heterogeneous and depends on how far in advance they are announced, as well as on their intensity and location – dimensions which are not observed to us. Second, while relational buyers may be prioritized during hartals; it could also be the case that relational buyers give slack to some of their suppliers depending on circumstances that are unobservable to us. This ambiguity – noted in other contexts (Macchiavello and Morjaria, 2015) – underscores the suggestive nature of the evidence presented here. Finally, it is unlikely that hartals are the primary source of supply disruption in our context. Localized episodes of labor unrest at specific factories, for example, are also common and have stronger impacts on workers' absenteeism and productivity than hartals (Ashraf et al., 2015). Given this last observation, and the likely heterogenous impact of hartals on exporters, Section 3 models how a desire to ensure deliveries at times of idiosyncratic supply shocks provides a rationale for relational sourcing. We discuss this and other mechanisms in Section 5.3.

3 Model

We present a simple model in which buyers source from sellers in either a relational or a spot fashion. The main features of the model are as follows. First, sellers are hit by idiosyncratic shocks that affect their capacity: in some states of the world, a seller's capacity is scarce and not all buyers can be prioritized. Second, while buyers and sellers can transact at market prices under normal conditions, it is not possible to formally enforce contracts that prevent delivery failure when shocks occur. However, buyers can invest in organizational capabilities to be able to make clear and credible promises to suppliers. In equilibrium, ex-ante identical buyers choose different sourcing strategies: some do not invest in organizational capabilities, pay low prices, and suffer delivery failures; others invest, promise higher prices and markups to suppliers and are reliably supplied.

The model is consistent with our empirical proxy for relational sourcing (defined in Section 2.2) and with the sourcing patterns of Section 2.3. The main prediction – which we test in Section 4 – is that relational buyers pay higher prices and markups to suppliers relative to spot buyers sourcing the same product, from the same supplier, at the same time. Section 5 discusses alternative mechanisms and policy implications from this framework.

3.1 Setup

Players. Consider a setting with a measure $B = 1$ of ex-ante identical buyers indexed by b and a measure $S < 1$ of sellers indexed by s. Time is an infinite sequence of periods $t = 0, 1, ...$ All parties have the same discount factor $\delta \in (0, 1)$. In period $t = 0$ buyers choose whether to invest in organizational capabilities, at cost F. The investment allows buyers to make credible promises to sellers, as described below. We say that a buyer is *relational* if they make the investment and is *spot* otherwise. We denote with ρ the (endogenous) share of buyers that make the investment to become relational in period $t = 0$.

Demand. From $t = 1$ onward, sellers produce garments and trade with buyers. In each period, buyers need to source a quantity $q = 1$ of identical garments. We refer to this as an order. All orders are identical, indivisible, and, when fulfilled, yield a gross payoff v to the buyer. A non-fulfilled order yields zero to the buyer. Sellers can sell completed orders to an external market at price $\underline{v} < v$.

Production. Idiosyncratic shocks affect sellers' production capacity over time. Shocks are i.i.d. over time and across sellers. Each seller is hit by a shock with probability $\alpha \in (0, 1)$ in each period. We use the indicator $\alpha_t^s \in \{0, 1\}$ to denote whether seller s is hit ($\alpha_t^s = 1$) or not hit ($\alpha_t^s = 0$) by a shock in period t. If $\alpha_t^s = 1$, seller s can produce only one order at t, at cost c_1. If instead $\alpha_t^s = 0$, seller s can produce two orders at t, at cost $c_0 < c_1$ each.

Assumption 1. $c_0 < \underline{v} < c_1 < v$.

This assumption implies that the efficient trade is for a seller to serve one buyer when hit by a shock, and to supply two buyers in the absence of a shock. The seller however would not produce for the external market when hit by a shock.[13]

[13]These assumptions can be microfounded in a way consistent with the empirical production model presented in Section 4.3 and estimated in Online Appendix B.2 to recover order-level markups. Specifically,

Timing and Matching. Following investments at $t = 0$, the sequence of events is as follows. At the beginning of period $t = 1$, pairs of buyers and sellers are formed and negotiate contracts. As explained below, some pairs start long-term relationships and continue trading in future periods $t = 2, 3, ...$ Other pairs only trade through spot contracts. There is random matching for everyone in $t = 1$, and then random matching at the beginning of every subsequent period between buyers and sellers that are not in a relationship. Since $S < B = 1$, a seller not in a relationship finds a buyer with probability one in all periods. A buyer not in a relationship finds a seller able to deliver through a spot contract with a probability μ that is determined in equilibrium.

Once pairs are formed and contracts negotiated, idiosyncratic production shocks are realized; sellers decide which orders to produce among the set of sourcing contracts agreed to with buyers; and, finally, order delivery and payments take place.

Contracts. A sourcing contract is an exchange of promises: the buyer promises to pay a certain price upon delivery; the seller promises to deliver at the agreed price. We rule out ex-ante transfers between parties. This assumption is made for both simplicity and realism. Spencer (2005) discusses why export contracts typically do not include lump-sum payments. Coupled with the contracting problems described below, the assumption implies that relational buyers cannot extract the entire surplus that their investment generates from the sellers they are matched with. We return to the role of this assumption for policy implications in Section 5.5.

There are two types of contracts: *spot* and *relational*. A spot contract is simply a price p_t^S to be paid to the seller upon order delivery in the current period. Buyers and sellers with a spot contract do not expect to continue trading in the future, so these contracts are offered at the beginning of each period. Relational contracts, instead, are plans that specify what the buyer and the seller are expected to do in each period of the relationship. Relational contracts are thus negotiated at the beginning of period $t = 1$ only.

Sourcing contracts are not perfectly enforced by courts: sellers cannot be penalized for

assume that orders are produced according to a production function $q = \omega F^{\beta} L^{(1-\beta)}$ in which F is materials (fabric), L is labor and ω is productivity. Denote with $c^n_{\alpha^s_t} = c(P^f, W; \omega, \alpha^s_t, n)$ the cost of producing one order when n other orders are produced. The solution to the firm's cost minimization yields a cost that depends on input prices, P^f and W, productivity ω, as well as on the shock realization α^s_t and, potentially, on the number of orders produced. In the absence of a shock ($\alpha^s_t = 0$), sellers optimally set F and L to produce each order. Since the two orders are identical and separately produced, the solution yields an order cost denoted $c_0 = c(P^f, W; \omega, 0, n)$ that is independent of n. If the seller is hit by a shock ($\alpha^s_t = 1$), instead, the firm faces a capacity constraint and must produce using at most $\overline{L}$ units of labor. A binding capacity constraint implies that the cost of producing each order is higher and depends on the number of orders produced. We assume that $\overline{L}$ is sufficiently low that $c^2_1 > v$ (i.e., it is not efficient to produce more than one order when hit by a shock). Furthermore, if the capacity constraint $\overline{L}$ also binds when only one order is produced, then $c^1_1 > c_0$, as in the text.

failing to deliver and buyers cannot be fully penalized for withholding payment. We assume that if a buyer does not pay the promised price after delivery, a court is unable to adjudicate whether the order was appropriately delivered or not. The seller is, however, able to prevent the buyer from withholding payment completely, e.g., because of a letter of credit (see Antràs and Foley, 2015 on payment terms in international trade). For simplicity, we assume that the court enforces a payment that corresponds to the market value of the order, $\underline{v}$.

3.2 Spot and Relational Trade

Spot Trade. Given our assumptions, a buyer would always renege on a spot contract that promises a price $p_t^S > \underline{v}$. Thus, omitting time subscripts, all spot contracts specify $p^S = \underline{v}$. This implies that sellers do not produce orders to fulfill spot contracts when their cost is $c_1 > \underline{v}$, and therefore buyers with spot contracts can only source from sellers that have not been hit by a shock. This captures the idea that spot contracts are useful under business as usual (no shock) but are ineffective to induce sellers to undertake costly actions during unusual, difficult to contract upon, circumstances. In equilibrium, buyers that source through spot contracts may thus suffer delivery failures.

Relational Trade. A relational contract between buyer b and seller s is a plan that specifies $\{\mathrm{I}_t(\alpha_t^s), p_t^R(\alpha_t^s)\}$ for each period $t = 1, 2, ...$ as a function of the past history of play. The function $\mathrm{I}_t(\alpha_t^s) \in \{0, 1\}$ indicates whether the seller is supposed to deliver in state α_t^s, while $p_t^R(\alpha_t^s)$ is the price the buyer promises to pay in state α_t^s. We focus on stationary relational contracts, in which $\mathrm{I}_t(\alpha_t^s)$ and $p_t^R(\alpha_t^s)$ are time-independent. We thus drop the time subscript in what follows.

Gibbons and Henderson (2012) point out that *credibility* and *clarity* are two necessary features of successfully managed relationships. Credibility refers to self-enforcement: a relational contract is self-enforcing if it constitutes a subgame-perfect equilibrium of the repeated game between the buyer and the seller. In equilibrium, it must be that parties do not want to deviate from the agreed plan. Credibility is thus captured by dynamic incentives constraints for the buyer and the seller as described below. Clarity is about understanding (and selecting) what equilibrium is played. Even when dynamic incentive constraints can be satisfied, different equilibria, including inefficient ones, can emerge. As noted in Section 2, a large management literature has argued that certain capabilities are necessary to establish successful relational contracts (see, e.g., Helper and Henderson, 2014). These investments are needed to create a shared understanding about the nature of the proposed relational contract and to persuade suppliers to trust the buyer and coordinate on the preferred equilibrium. We thus assume:

Assumption 2. *Only relational buyers can offer relational contracts.*

We construct an equilibrium in which sellers always deliver to relational buyers regardless of the shock (i.e., $\mathrm{I}(1) = \mathrm{I}(0) = 1$) and relational buyers pay the promised price $p^R(\alpha_t^s)$. We assume that if a seller fails to deliver, no buyer will source relationally from her in the future. Similarly, we assume that if a relational buyer reneges on a promised payment, no seller will ever believe their promises in the future. Thus, upon reneging, the buyer would only be able to source through spot contracts. These rather drastic assumptions capture the idea that there must be some reputational loss from reneging on a relational contract. No relational equilibrium could be sustained if, following a deviation, buyers and sellers could immediately re-match with a partner offering an identical relational contract. Less drastic assumptions do not alter the key insights but come at the cost of a more cumbersome setup.

Along the equilibrium path of a self-enforcing relational contract, the seller produces for the relational buyer regardless of shocks, and also produces for a spot buyer when not hit by a shock. The seller's expected per-period payoff along the equilibrium path is given by $\pi_s^R = \alpha p^R(1) + (1-\alpha)(p^R(0) + \underline{v}) - [\alpha c_1 + 2(1-\alpha)c_0]$, whereas her expected per-period payoff following a deviation is $\underline{\pi}_s^R = 2(1-\alpha)(\underline{v} - c_0)$. Let $\Delta\pi_s^R \equiv \pi_s^R - \underline{\pi}_s^R$. The seller's incentive compatibility constraint in state $\alpha_t^s \in \{0,1\}$ is

$$\frac{\delta}{1-\delta}\Delta\pi_s^R \geqslant \alpha_t^s\left(c_1 - p^R(1)\right) + (1-\alpha_t^s)\left(\underline{v} - p^R(0)\right). \qquad (\mathrm{DICS}_{\alpha_t^s})$$

The left-hand side is the net present value of the relationship for the seller, namely the discounted difference between her expected payoff along the equilibrium path and off the equilibrium path. Since shocks are i.i.d., this value does not depend on α_t^s. The right-hand side is the temptation to deviate. When hit by a shock ($\alpha_t^s = 1$), the seller may be tempted to not produce the order and save the cost c_1. Her temptation to deviate is thus $(c_1 - p^R(1))$. When not hit by a shock ($\alpha_t^s = 0$), the seller may be tempted to sell the order on the spot market at price $\underline{v}$, which is more profitable than not producing it. In this case her temptation to deviate is $(\underline{v} - p^R(0))$.

Consider next the incentive compatibility constraint for the buyer. The buyer's expected per-period payoff along the equilibrium path of the relational contract is given by $\pi_b^R = v - \alpha p^R(1) - (1-\alpha)p^R(0)$, whereas his expected per-period payoff following a deviation is $\underline{\pi}_b^R = \mu(v - \underline{v})$ (recall that μ denotes the probability with which a buyer finds a seller not hit by a shock in the spot market). Let $\Delta\pi_b^R \equiv \pi_b^R - \underline{\pi}_b^R$. The buyer's incentive compatibility constraint in state $\alpha_t^s \in \{0,1\}$ is

$$\frac{\delta}{1-\delta}\Delta\pi_b^R \geqslant p^R(\alpha_t^s) - \underline{v}. \qquad (\mathrm{DICB}_{\alpha_t^s})$$

As in the case of spot contracts, we maintain the assumption that a defaulting buyer is forced by the court to pay a price $\underline{v}$. A court however would not enforce any additional promised payment $(p^R(\alpha_t^s) - \underline{v})$.

In every period, the relationship creates an expected joint surplus given by $\Delta\pi_s^R + \Delta\pi_b^R$. The relationship can be sustained if the net present value of the expected surplus is sufficient to resist the sum of the seller's and buyer's temptations to deviate in each state. Note that the relational price $p^R(\alpha_t^s)$ drops out from the sum of the temptations to deviate as it is simply a transfer between parties. When the seller is not hit by a shock, the sum of these temptations is then zero. Intuitively, in the absence of a shock, the incentive constraint cannot be binding if the relationship generates value. In the presence of a shock, instead, the relationship's incentive constraint is given by

$$\frac{\delta}{1-\delta}(\Delta\pi_s^R + \Delta\pi_b^R) \geqslant c_1 - \underline{v}. \tag{DICR$_1$}$$

This is a necessary condition for the relational contract to be sustained. Note that this condition does not pin down prices $p^R(\alpha_t^s)$ nor how the surplus generated by the relationship is split between the buyer and the seller. For example, the relational contract could specify a fixed contractible price $\underline{v}$, and an additional non-contractible bonus to be paid to the seller when there is a shock. Alternatively, the relational contract could specify a price p^R that does not depend on the shock. A large literature has built upon the idea that prices cannot easily be adjusted in response to production conditions (see, e.g., Dana (1998) and Carlton (1978) for theoretical models and Carlton (1986) for a discussion of the empirical relevance of such assumptions in a variety of industries). In this spirit, we assume:

Assumption 3. *The relational price is shock-independent:* $p^R(1) = p^R(0) = p^R$.

Note that Assumption 3 does not alter the conditions under which (DICR$_1$) can be satisfied. Moreover, conditional on (DICR$_1$) being satisfied, there exists a continuum of shock-independent prices p^R that satisfy the incentive compatibility constraints. Let $\overline{p}^R \equiv \delta[v - \mu(v - \underline{v})] + (1-\delta)\underline{v}$ be the highest such price, i.e., the price that makes the buyer's incentive constraint (DICB$_1$) hold with equality. Analogously, let $\underline{p}^R \equiv (1-\delta)c_1 + \delta[\alpha c_1 + (1-\alpha)\underline{v}]$ be the lowest such price, i.e., the price that makes the seller's incentive constraint (DICS$_1$) hold with equality.

Observation 1. *A seller that (contempouraneously) supplies both a relational buyer and a spot buyer earns higher prices (and markups) on orders produced for the relational buyer:* $\underline{p}^R \in (\underline{v}, c_1)$. *Moreover, a seller earns higher profits from a relational buyer than from spot* *buyers:*

$$\begin{aligned}
\underline{p}^R - \alpha c_1 - (1-\alpha)c_0 &= c_1(1-\delta) + \delta[\alpha c_1 + (1-\alpha)\underline{v}] - \alpha c_1 - (1-\alpha)c_0 \\
&= (1-\alpha)[(1-\delta)c_1 + \delta \underline{v} - c_0] \\
&> (1-\alpha)(\underline{v} - c_0),
\end{aligned}$$

where the inequality follows from Assumption 1. The higher markups from the relational buyer under no shock more than compensate for the cost of delivering under the shock.

We thus find that with shock-independent relational prices, even the lowest such price $\underline{p}^R$ leaves rents to the supplier. In this sense, Assumption 3 is not inconsequential. If prices could be made contingent on the shock, the buyer would be able to extract all rents from the supplier (e.g., by paying a price that just compensates the seller for its costs). This, however, is a consequence of the perfect monitoring in our model and would not be the case if suppliers' deviations were observed only imperfectly or if the shock was the supplier's private information. In such cases, even if prices could condition on the shock, the supplier would earn rents, as in efficiency wage models (MacLeod and Malcomson, 1998). Assumption 3 thus captures in a parsimonious way the common insight that rents must be paid to induce suppliers to undertake costly, non-contractible, actions. The assumption that parties cannot make ex-ante transfers at the contracting stage implies that suppliers keep those rents in equilibrium.

Buyers' Investments in Capabilities. We complete our description of the equilibrium by studying the buyers' decisions to become relational or remain spot in period $t = 0$. Note that for an equilibrium with relational trade to exist, it must be that spot buyers fail to secure supply with positive probability; that is, we must have $\mu < 1$. If $\mu = 1$, the incentive constraint for a relational buyer requires $p^R = \underline{v}$, which violates the seller's incentive constraint.

In equilibrium, a mass ρ of buyers invest to become relational, and the remaining mass $(1-\rho)$ are spot buyers. If $\rho > S$, then a mass $\rho - S$ of relational buyers are not able to find a supplier to start a relationship. These buyers only source through spot contracts from sellers that are not hit by shocks. There is a mass $(1-\alpha)S$ of such sellers. A mass $(1-\alpha)\max\{S-\rho, 0\}$ can supply two units, while a mass $(1-\alpha)\min\{\rho, S\}$ can only supply one unit. In equilibrium, the probability of being able to source with a spot contract is thus given by

$$\mu(\rho) = \min\left\{\frac{(1-\alpha)(2\max\{S-\rho,0\} + 1\min\{S,\rho\})}{1-\rho+\max\{\rho-S,0\}}, 1\right\}.$$

Assumption 4. $\dfrac{1}{2-\alpha} < S < \dfrac{1}{2(1-\alpha)}$.

The assumption $S < \frac{1}{2(1-\alpha)}$ states that the spot market is not sufficient to achieve efficiency. If $S \geq \frac{1}{2(1-\alpha)}$, there would always be sufficient capacity from sellers that were not hit by shocks to supply all buyers. The assumption $S > \frac{1}{2-\alpha}$ allows us to simplify the equilibrium analysis and avoid a taxonomy of cases that yields limited additional insight. In particular, this assumption yields that $\mu(\rho) = 1$ for any $\rho \geq S$. Given our discussion above, it thus follows that there cannot be equilibria with $\rho \geq S$.

We thus consider equilibria in which $\rho < S$. In such equilibria there are more sellers than relational buyers. It is therefore natural to assume that the relational price is the lowest possible, namely $p^R = \underline{p}^R$. To pin down the equilibrium value of ρ, note that for $\rho \in (0, S)$, the probability that a spot buyer finds a supplier is equal to $\mu = (1-\alpha)(2S-\rho)/(1-\rho)$, which is increasing in ρ. An equilibrium with $\rho \in (0, S)$ is thus defined by $p^{R^*} = \underline{p}^R$ and ρ^* such that the following buyer entry condition holds:

$$\delta\left[v - p^{R^*} - \mu(\rho^*)(v - \underline{v})\right] = (1-\delta)F. \tag{BE}$$

Condition (BE) states that in an equilibrium with relational buyers ($\rho > 0$), ex-ante identical buyers must be indifferent between entering as relational or spot. The equilibrium would feature no relational buyers ($\rho = 0$) if $\delta\left[v - p^{R^*} - \mu(0)(v - \underline{v})\right] < (1-\delta)F$. As we are interested in equilibria with relational trade, we assume that F is sufficiently small that the equality can hold.

Figure 1 illustrates the equilibrium. The figure reports on the horizontal axis the share of relational buyers ρ, and on the vertical axis the relational price p^R. The buyer entry (BE) condition traces a curve that slopes down. A higher ρ increases μ and thus the payoff of being spot; to restore the equilibrium, this must be compensated by a decline in p^R. Note that given an equilibrium price p^{R^*}, a reduction in the cost F of becoming relational would imply a shift of the buyer-entry curve and an increase in ρ^*. We discuss policies that reduce F in Section 5.5.

3.3 Summary of Takeaways

Relational buyers trade with fewer sellers than spot buyers. The model is consistent with the empirical proxy for buyers' sourcing strategy introduced in Section 2.2. Relational buyers always buy from the same seller (only one seller given our simplifying assumptions on unit demands), while spot buyers end up being matched with different sellers

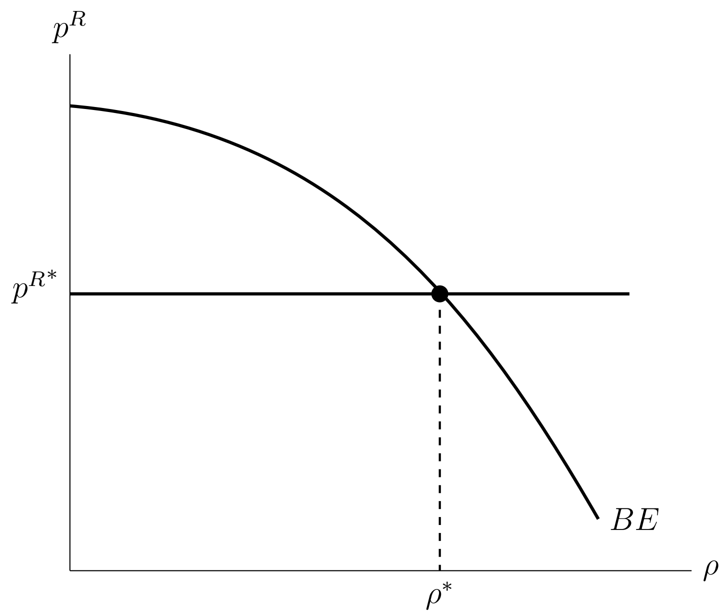

Figure 1: Equilibrium share of relational buyers

along the equilibrium path. The reason is that spot buyers' orders are fulfilled in the spot market only by sellers who are not affected by a shock. As such, spot buyers switch between partners over time. In addition to sourcing from fewer sellers, relational buyers also trade more frequently than spot buyers: they trade in all periods, whereas spot buyers trade only if they find an available supplier that is not hit by a shock.

Within a narrow industry, buyers differ in their sourcing strategies and sellers do not specialize. The model rationalizes the evidence in Sections 2.3.2 and 2.3.3. While buyers are ex-ante identical, they endogenously make different decisions with regards to becoming relational or spot. Scarce capacity and contracting frictions imply that sourcing strategies are *strategic substitutes*. While we do control for buyer characteristics in the empirical analysis that follows, the model provides a rationale as to why a significant share of the observed variation in buyers' sourcing strategies cannot be explained by observable characteristics.[14] This matches the evidence in Section 2.3.2. Furthermore, in equilibrium, a share ρ/S of sellers supplies a mix of relational and spot buyers while the remaining sellers only supply spot buyers. No seller only supplies relational buyers as it is not possible to

[14]Consider an extension in which buyers demand two units per period instead of only one. In this case, relational buyers will form relationships with two sellers, sourcing one unit from each of them in each period. This follows from the fact that the cost of becoming relational is a fixed cost, and in equilibrium relational trading must give buyers a higher ex-post payoff than spot trading. Hence, the extension would imply that relational buyers source relationally from all their suppliers.

promise reliability to all buyers. The model thus also rationalizes the evidence in Section 2.3.3.

Relational buyers pay higher markups. The main prediction of the model – in 1 – is that relational buyers pay a higher price, and thus a higher markup, relative to spot buyers. The prediction that relational sourcing uses higher markups to incentivize suppliers is *not* unique to our model. Testing the prediction, however, is difficult: the main challenge is that one needs to compare the costs that a given supplier incurs when producing the same product for buyers adopting different sourcing strategies. These costs are difficult to estimate in standard datasets since the amounts and prices of the variable inputs used to produce for a particular buyer are typically not observed. The next section leverages unique data that allow us to overcome this challenge and directly test the prediction. We leave a discussion of alternative mechanisms and policy implications of this finding to Section 5.

4 Evidence

This section tests the main prediction of the model: relational buyers pay higher markups. Before presenting the main results, we describe garments' production process and our data. The customs data reveal that, within seller-product-year combinations, orders produced for relational buyers earn higher prices. The higher prices may reflect higher markups (as predicted by the model) or higher costs of producing for relational buyers. Disentangling the two is difficult as the allocation of inputs to output is not typically observed. Our customs data and the internal records from factories, however, allow us to link inputs to specific orders and reveal that within seller-product-year combinations, orders produced for relational buyers *do not* differ in the type, price and efficiency of fabric and sewing labor. We derive conditions under which the data recover differences in markups across orders produced for different buyers and confirm the model's main prediction.

4.1 Buyer-Specific Inputs and Outputs

4.1.1 Garment Production

Ready-made garment manufacturers in Bangladesh, who are entirely export-oriented, make production decisions based on the orders they receive from international buyers. Buyers provide suppliers with a design and a set of technical specifications on the items to be produced. Unlike cut-make-trim systems in which buyers provide fabric and other material inputs to the manufacturer (e.g., China, Mexico and Myanmar), Bangladeshi exporters source fabric and inputs, then proceed to cut, sew and package the garments according to their buyers' specifications.

Fabric and labor employed on sewing lines are the two main variable inputs utilized in the production of a garment export order and jointly account for 85-90% of the variable costs of producing a typical garment. Fabric utilization choices are made order by order. Once the fabric is available at the manufacturing plant, two sequential production stages take place: (i) inspection and cutting, and (ii) sewing and finishing (see Online Appendix B.1 for details). Fabric efficiency is tracked by two performance indicators. The *buy-to-cut ratio* – the ratio of purchased fabric to cut fabric that is fed to the sewing lines – measures performance at the inspection and cutting stage. The *cut-to-ship ratio* – the ratio of cut fabric to shipped garments – measures performance at the sewing and finishing stage. The product of these two metrics, the *buy-to-ship ratio*, is a commonly used performance indicator. Lower values represent lower levels of waste and, thus, higher efficiency over the two stages of production.

Labor employed in the sewing section of the factory is the other main variable input in the production of garments. Like the buy-to-ship ratio, labor efficiency is a standard performance indicator in the industry. It is measured as the ratio between the minutes-equivalent output of the production line and the minutes of labor input. On a given day, the input minutes on a line are given by the number of sewing operators multiplied by the line's runtime. The output minutes are calculated as the product between the garment's Standard Minute Values (SMVs) and the number of pieces produced by the line. The SMV is a measure computed by the factory's industrial engineers – often based on international libraries of SMVs of elemental sewing processes – and captures the amount of time required to sew a particular garment.

4.1.2 Data and Sample

Our main source of data consists of transaction-level export and import customs records from Bangladesh over the period 2005-2012. We complement these data with internal production records and workers surveys from a sample of factories. These additional data were collected as part of a series of RCTs (see Macchiavello et al., 2020; Ashraf et al., 2015; and Macchiavello and Woodruff, 2014). The main novelty of the data is that they allow us to explore differences in the price and efficiency of the two main variable inputs – fabric (in the customs data) and labor (in the production line data) – across export orders produced for different buyers. We offer a brief description here and refer to Online Appendices A.1 and A.2 for details.

Customs Records. We focus on woven garments. Two features of the Bangladeshi woven garment sector enable us to link the use of material inputs to output at the export order level. First, unlike other major garment exporters including 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., woven cotton fabric) exclusively, as there are no suitable domestic substitutes. Second, to participate in a customs bonded warehouse regime that allows duty free import of material inputs, exporters must indicate the export order for which the imported fabric will be used. Specifically, after receiving an order from an international buyer, the manufacturer submits a utilization declaration (UD) to the Bangladesh Garment Manufacturers and Exporters Association. A unique UD identifier is assigned to all export and import transactions belonging to that export order. These two features enable us to identify the material inputs that correspond to an export order. We aggregate transaction-level records at the order (i.e., UD) level, producing a single entry for each order that denotes the following information: 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. To illustrate, a hypothetical observation in our dataset looks 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 to fulfill an order subsequently exported to *Walmart Inc.* of 450 kg of men's or boys' woven cotton shirts (HS620520) on 03/01/2018 at \$10 per kg.

We focus on woven garment orders channeled through the UD system in the 17 six-digit HS codes in the two largest woven apparels: shirts and trousers. Across orders variation within export-product-time combinations is needed to test the model's prediction. We thus restrict our analysis to the 500 largest exporters, accounting for 78% of the relevant sample. Appendix Table A1 compares the analysis sample with the broader population.

Table 1 provides descriptive statistics. Panel A reveals that the average order has a buy-to-ship ratio of 0.87 – similar to $400/450 \approx 0.89$ for our hypothetical order exported by *Nice Apparel Co. Ltd.* to *Walmart Inc.* The buy-to-ship ratio is often less than one because it is computed using the net export volumes (kilos) that, on top of fabric, also include accessories and packaging (garments are folded in plastic envelopes and then stored in carton boxes). Our results are robust to controlling for accessories and packaging characteristics. Buy-to-ship ratios at the order level are also quite dispersed, with a coefficient of variation of 0.33. This dispersion is consistent with differences in efficiency – at the inspection and cutting and/or at the sewing and finishing stages of production – and in the substitution between fabric and other inputs (see Online Appendix B.1 for a discussion and evidence). Panels B, C and D provide descriptive statistics at the exporter, buyer and buyer-seller pair level. Our baseline specification explores differences across orders within seller-product-year – denoted sjt – combinations. There are 6,872 seller-product-year sjt combinations. Across these, the median (mean) number of buyers in the triplet is 2 (2.91). There are 4 (6) buyers at the 75th (90th) percentile.

Internal Plant Production Records. We complement the customs records with daily production data on approximately 1,300 sewing lines from 51 garment factories. Sewing lines are observed for approximately 340 days. The data record the utilization, composition and efficiency of labor, including the Standard Minute Values (SMVs) defined above. Record keeping varies across plants and also within plants over time. The buyer for whom the line is producing on a specific day is observed for 46% of the observations. Appendix Table A3 shows that there are no significant differences between observations with and without information on the buyer. We observe the buyer whose order is being produced, for almost 200 thousand production line-day combinations (see Panel A of Appendix Table A2). This allows us to compare labor usage for buyers with different sourcing strategies. Appendix Table A5 reports summary statistics in the labor data, and shows that there is significant variation (coefficient of variation of 0.5) on the sourcing characteristic of the buyers lines are producing for. The production records do not contain information on the skills and wages of workers on the lines. We thus complement the data with surveys of over one thousand workers employed at these plants (Panel C of Appendix Table A2), as well as internal HR records for over 35 thousand workers in eleven factories (Panel B of Appendix Table A2).

4.2 Relational Buyers and Export Prices

The model predicts that relational buyers pay higher prices than spot buyers for otherwise identical orders produced by a given supplier under identical (shock) conditions. To test this prediction, we estimate

$$p_{sbjo} = \delta_{sjt} + \beta Relational_b + \varepsilon_{sbjo}, \tag{3}$$

where p_{sbjo} is the log unit price of garment order o, of product j (six digits HS code), manufactured by seller s for buyer b and δ_{sjt} is a fixed effect that absorbs seller-product-year variation. These fixed effects allow us to compare differences across orders produced for different buyers as in the model. The regressor of interest, $Relational_b$, is our baseline metric of buyers' sourcing defined in Section 2. Throughout the analysis we use the metric in *excluded* products to assuage endogeneity concerns and avoid mechanical correlations with order-level outcomes.

Table 3 reports the results. Column (1) shows that a standard deviation increase in the sourcing metric (i.e. the more relational the buyer is) is associated with 2% higher prices. Columns (2) to (4) sequentially add controls that are buyer-, relationship- and order-specific. Across all specifications, the estimated coefficient remains quantitatively and qualitatively unchanged, ranging from 1.9% to 2.3%.

Relational sourcing is unconditionally correlated with the buyer's size (see Panel B of Appendix Table D4). Column (2) controls for these buyer-level characteristics as well as destination fixed effects, δ_d, to absorb differences explained by characteristics common to all buyers in a given country. Column (3) adds buyer-seller controls: the age and cohort of the relationship, its size, the share of the seller in the buyer's trade and share of the buyer in the seller's trade. Finally, relational buyers place more frequent, smaller orders (Panel C of Appendix Table D4) and might demand garments of different quality. Column (4) controls for the size of the order and the price of the fabric used in its production.

The pattern is extremely robust. We explore robustness of the results along two dimensions: one that relaxes the controls, and one that uses alternative time horizons. We consider all combinations that (i) let the set of covariates to feature none, some or all sets of controls (i.e., buyer-, relationship- and/or order-level controls); (ii) include one, two and three way combinations of fixed effects (s for seller, j for product, d for destination and t for period); and (iii) define the period t at either the month m, quarter q or year y. Figure 3 reports estimates from the 522 resulting specifications. All point estimates fall in the interval $[0.005, 0.046]$, with our baseline specification (corresponding to column (4) of Table 3) below the midpoint.[15]

The baseline specifications likely *under*estimates differences in prices paid by relational and spot buyers. Leaving aside the evidence in Figure 3, $Relational_b$ only exploits across buyers variation. To the extent that buyers tailor sourcing behavior to suppliers' circumstances (e.g., some relational buyers might source spot from some suppliers), or buyers change sourcing strategies over time; our approach induces attenuation bias. We return to both issues in the next section. Furthermore, Online Appendix C explores additional departures from our baseline specifications, including the use of 15 alternative operational definitions of relational sourcing (Appendix Table C3), as well as different estimation samples (Appendix Tables C4 and C5). Results are robust and often larger in magnitude than in our baseline specification.

[15] The 36 specifications with coefficients not significantly different from zero (albeit positive) correspond specifications that include either seller-month (or seller-product-month) fixed effects or destination-seller fixed effects alongside product-month fixed effects, and no order-level controls. These fixed effects leave insufficient variation either because not enough exporters ship multiple orders of the same product within a month or have multiple buyers with different sourcing strategy within the same destination market.

4.3 Relational Buyers and Variable Inputs

Having established that relational buyers pay higher prices, we now turn to the two main variable inputs – fabric and labor on the sewing lines. The main takeaway is that, conditional on exporter-product-time fixed effects, we *do not* detect any difference in the type, efficiency, price or utilization of the two variable inputs across orders produced for relational and non-relational buyers.

Input Usage: Fabric. We use the specification in equation (3) and consider three outcomes: the price of the fabric used in the production of the order; the order level buy-to-ship ratio; and a proxy for product complexity given by the number of different types of fabric used to produce the order. Table 4 reports the results. Odd columns estimate the specification in column (1) of Table 3; even columns include buyer-, relationship- and order-level controls as in column (4) of the same table. Columns (1) and (2) show that the price of the fabric does not correlate with the sourcing strategy adopted by the buyer for whom the order is being produced. Columns (3) and (4) show that there is also no correlation between fabric efficiency – as measured by the order's buy-to-ship ratio – and whether the order is produced for a relational buyer. Finally, Column (5) shows a small positive correlation between a proxy for product complexity (the number of different fabric types used in the order) and the buyer's sourcing strategy, but the correlation vanishes once controls (in particular, the size of the order) are included in column (6).

Taken together, these results suggest that the higher export prices paid by relational buyers are unlikely to reflect differences in the type of fabric used, or differences in the efficiency with which suppliers sew fabric into garments, when producing for buyer with different sourcing strategies. Before turning our attention to labor on the sewing lines, we note that the extent to which labor and fabric can be substituted also does not differ across buyers adopting different sourcing strategies. Exploiting time variation in cotton prices (the main input to produce fabric) and a large increase in the minimum wage, Appendix Table B2 shows that when the price of fabric (labor) increases, exporters use less (more) fabric to produce orders of a given size. These substitution patterns, however, do not differ across orders produced for buyers with different sourcing strategies. See Online Appendix B.1 for details.

Input Usage: Labor. We now turn to labor employed on the sewing lines. To the extent possible, we would like to study labor usage using the same specification in equation (3). Differences in the nature of the customs and production data, however, impose certain adjustments. To fix notation, let τ, s, l and b denote a calendar day, seller, production line and buyer respectively. We estimate specifications

$$y_{slb\tau} = \delta_{sm(\tau)} + \delta_\tau + \delta_{sl} + \beta Relational_b + \varepsilon_{slb\tau}, \tag{4}$$

where $y \in \{\#Workers, Share\ Helpers, SMVs, Efficiency\}$ is the outcome and the main regressor of interest is $Relational_b$ – the sourcing strategy of the buyer for which the line is producing on a given day. Denoting with $m(\tau)$ the calendar month of date τ, fixed effects $\delta_{sm(\tau)}$ absorb factory-month specific variation common across all production lines and buyers. δ_τ is a day fixed effect collecting common shocks that could affect production in all plants, such as hartals or festive days, and δ_{sl} are production line fixed effects.

There are four main differences in the structure of the production line data relative to the customs records, and these induce small discrepancies between equations (3) and (4). First, the production line data do not include information on the product. It is thus not possible to directly include the product j dimension and replicate the seller-product-time fixed effects in equation (3). The fixed effects $\delta_{sm(\tau)}$ in equation (4) are thus akin to st fixed effects in the customs data. In practice, this might not be a significant departure because (i) most factories specialize in a narrow range of products with exporters typically using multiple factories to offer a broader range of products to their buyers; (ii) Figure 3 above and Online Appendix C show that the customs data results are robust to specifications with st – as opposed to sjt – fixed effects. A second potential departure pertains to the inclusion (or not) of production line fixed effects. Since the allocation of orders to lines is part of the cost minimization problem solved by the sellers and production lines are not observed in the customs records, consistency would suggest to exclude line fixed effects in the production data. At the same time, it is interesting to explore whether orders from relational buyers are systematically allocated to more/less efficient lines. Furthermore, factories that produce multiple products tend to assign dedicated lines to specific product types. Including production line fixed effects therefore helps dealing with the unobserved j dimension in the production data. For these reasons, we report results both with and without line fixed effects. A third potential source of discrepancy is that the higher frequency of the production data – a day, as opposed to a more sporadic export order – suggests a narrower definition of time period relative to the customs data (a month $m(\tau)$ instead of a quarter, or a year). Nevertheless, Online Appendix C shows that the results are robust to different definitions of time period in either of the two datasets. Finally, the specification includes the buyer-level controls added in column (2) of Table 3. As we cannot precisely match the factories in the production data with the customs records, we cannot include relationship- and order-level controls.

Table 5 reports the results for the four outcomes of interest excluding (odd columns) and including (even columns) the line fixed effects. Columns (1)-(2) and (3)-(4) show that orders produced for relational buyers are not characterized by higher SMVs (a commonly used measure of the garment's complexity) or by a lower efficiency on the sewing lines.[16] Columns (5) and (6) show that orders produced for relational buyers do not have a significantly different number of operators working on the sewing line. Finally, columns (7) and (8) show that factories use similar shares of helpers (relative to more skilled sewing operators) when producing orders for relational buyers. In sum, the table reveals no significant differences in the efficiency or type of labor when producing orders for buyers with different sourcing strategies.

The production data do not contain information on the skills or pay of workers on the production lines. If orders produced for relational buyers employ more skilled workers that earn higher wages, physical efficiency will fail to detect differences in labor costs. To investigate this possibility, we leverage workers surveys and administrative HR data from (some of) the factories. A sample of 704 workers was asked whether they have rotated lines on a temporary or permanent basis (in the entirety of their job history at the factory). Considering permanent rotations, 93.6% of the workers respond never having rotated and 98% report at most one rotation in their entire work history at the plant. Considering temporary rotations, 73.5% of the workers answer to never have rotated and 95% of the workers had at most 3 rotations in their entire history. In other words, workers do not rotate frequently across lines. These patterns are confirmed by HR records on production line assignment for almost 20,000 operators in a more limited set of factories. While it is possible that these records underreport the movement of workers across lines, 79% of the workers are always on the same line; 99.5% are assigned to at most two lines. In light of this, line fixed effects in Table 5 effectively control for differences in workers' composition across orders.

Furthermore, workers surveys provide information on demographics for 1,500 line operators, line supervisors and line chiefs. A subset of the workers (approximately 700) were also asked about wages and pay. Knowledge of the line on which the worker was working at the time of the survey allows us to construct a variable, $Relational_{sl}$, that measures the share of days during which that worker was likely producing for a relational buyer. We are able to construct this variable for approximately 1,000 workers in the overall sample and for 560 workers in the sample for which we have information on compensation. Although the construction of $Relational_{sl}$ inevitably entails measurement error, Appendix Table D7 shows that, conditional on factory and workers' position fixed effects, $Relational_{sl}$ does not correlate with the wage, whether the worker is paid piece rates, quality bonuses, or other types

[16] Column (3) estimates a negative coefficient that is statistically different from zero at conventional levels. The estimated coefficient is, however, economically small and corresponds to an increase in variable costs of approximately 0.25% – nearly a tenth of the (rather conservative) estimate for prices.

of bonuses (Panel A). Furthermore, $Relational_{sl}$ does not correlate with the gender, education level, experience and a measure of cognitive skills (Raven Test) – this latter test being completed by line chiefs only (Panel B). This evidence assuages concerns that differences in skills could be driving differences in prices.

A final concern is that, if relational buyers expect timely shipments, production workers may need to work overtime shifts to ensure that production is ready for shipment. To determine whether this is the case, we use the production data to construct the share of the time (days) that sewing lines in the factory are producing for relational buyers in a certain month, $Relational_{sm}$. We correlate this variable with overtime and absences reported in the HR data. Appendix Table D8 finds no correlation. This is also the case when we restrict our attention to non-line workers ('upstream' workers such as cutters and fabric markers, and 'downstream' workers, such as packers) and managers separately. We are nevertheless aware that overtime is a very sensitive issue on which factories might misreport and thus interpret this evidence with caution. As an additional piece of evidence, we check whether the daily runtime of the production line correlates with the type of buyer the line produces for and, again, find no correlation (see Appendix Table D9). To the extent that our data allow, we find no differences across relational and non-relational buyers in terms of labor costs.

4.4 Relational Buyers and Markups

The evidence is thus consistent with the model's prediction that the higher prices paid by relational buyers reflect a higher markup. However, there might be costs unobserved to us that are variable at the order level and systematically vary across orders produced for relational and non-relational buyers. We develop an empirical framework that clarifies under which conditions the available data recover within seller-product-time *differences* in markups across orders thereby allowing for a precise test of the model's main prediction.

It turns out that these conditions are quite mild: they boil down to a production function that features (log-)separability of fabric use relative to other costs. This condition appear justifiable in light of the two-step production process for garments described in Section 2 and further studied in Online Appendix B.1. Other than that, the framework allows for an elasticity of output with respect to fabric that varies at the seller-product-time level and for an arbitrary number of other inputs that sellers may choose freely (e.g., casual labor) or subject to capacity constraints (e.g., managerial labor and attention). Under more stringent functional form assumptions the framework also recovers the *level* of markups at the order level. When we do so, we estimate an output elasticity with respect to fabric and overall returns to scale that are in line with industry reports. We sketch here the main elements of the framework. Online Appendix B.2 provides the details; Online Appendix B.3 estimates markup levels.

4.4.1 Estimating Differences in Order-Level Markups

Following the model in Section 3 (see in particular footnote 13), let the timing of events within a period t be as follows. First, buyers b and sellers s form links (and sellers choose production capacity). Second, each buyer's demand is realized, buyers place orders and shocks to sellers' capacity are realized. 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 denote the set of orders placed to seller s in period t (by all buyers and in all products) by O_{st}. Order o is seller-buyer-product-time specific (i.e., $sbjt$ specific). Each order specifies a volume Q_o and a unit output price P_o.

To produce an order o, a seller combines labor L_o^z (of potentially different types $z \in \{1, 2..., Z\}$) with fabric F_o. We allow orders to vary in the way they combine the different types of labor and have idiosyncratic productivity ω_o. We denote θ_o the output elasticity with respect to fabric and $\boldsymbol{L}_o = \{L_o^1, L_o^2, ..., L_o^Z\}$ the available capacity in labor of type z.

Seller s in period t chooses $\{\boldsymbol{L}_o, F_o\}_{o \in O_{st}}$ to minimize costs, subject to technology and capacity constraints, taking order characteristics and prices as given. Denote the wages for labor of type z and the price of fabric with W_o^z and P_o^f respectively. The order-specific Lagrange multiplier λ_o represents the increase in cost associated with producing one additional unit of output in order o, that is, the short-run marginal cost for order o. Denoting the order-level markup factor M_o, the (order-specific) first order condition w.r.t. fabric F_o yields

$$\lambda_o = \frac{P_o^f F_o}{\theta_o Q_o} \text{ and } M_o \equiv \frac{P_o}{\lambda_o} = \theta_o \frac{P_o Q_o}{P_o^f F_o}. \tag{5}$$

Equation (5) implies that the order-level markup M_o depends on the buy-to-ship ratio F_o/Q_o, the unit price of garment P_o and fabric P_o^f and the output fabric elasticity θ_o. The unique feature of our data is that F_o/Q_o, P_o and P_o^f are directly observed. The output fabric elasticity θ_o, however, is not. Denote $\psi_o = \frac{P_o Q_o}{P_o^f F_o}$ the term that is directly observed in the data. We can write the *difference* in (log) markups factors between two orders o and o' as:

$$\Delta_{oo'} \equiv \ln(M_o) - \ln(M_{o'}) = \underbrace{(\ln(\psi_o) - \ln(\psi_{o'}))}_{\text{Directly Observed in the Data}} + \underbrace{(\ln(\theta_o) - \ln(\theta_{o'}))}_{\text{Not Observed in the Data}}. \tag{6}$$

The data allow us to directly observe *differences* in markups across orders that share the same fabric elasticity. We assume that the output-to-fabric elasticity vary at the seller-product-time, i.e., $\theta_o = \theta_{sjt}$. Under this assumption, we can directly explore differences in (log) markups factors across buyers within a seller-product-time combination using ψ_o as the dependent variable in our baseline regression in equation (3). Denote with μ and mc the log markup factor and marginal cost, respectively, and note that $\mu_{sbjo} \equiv p_{sbjo} - mc_{sbjo}$. Seller-product-time fixed effects, δ_{sjt}, flexibly control for differences in the (log of the) output-to-fabric elasticity $\ln(\theta_{sjt})$ – the only unobserved component of markups. A potential concern is that the fabric elasticity might vary across orders produced for buyers that use different sourcing practices. Online Appendix B.1 and B.3 present evidence that assuages such concerns.

4.4.2 Relational Buyers Pay Higher Markups

Figure 4 replicates Figures 3 and reports results from over 500 regressions using μ as dependent variable. All point estimates in the markups regressions are bounded in $[0.009, 0.048]$, with our baseline specification – explored in Table 6 – estimating a coefficient of 0.026. For ease of comparison, column (1) replicates the result on prices reported in column (4) of Table 3. Columns (2) and (3) decompose the difference in prices into marginal costs and markup factors. Orders produced for relational buyers do not have higher marginal costs and, therefore, the price difference follows from higher markups. Only 11 of the 522 estimates reported in Figure 4 are not significantly different from zero at 10% significance.[17] Online Appendix C confirms the robustness of these results to different estimation samples and proxies for relational sourcing.

In sum, consistent with the main prediction of the model, we find robust evidence that relational buyers pay higher prices and markups.

5 Discussion and Further Evidence

This section revisits our approach that characterizes the sourcing strategy at the buyer – as opposed to the buyer-seller – level; complements the across buyers comparison with an event study exploiting VF Corporation's shift in its *global* approach to sourcing from spot to relational; discusses the reliability mechanism, alongside alternative channels; explores the quantitative implications of our estimates and, finally, returns to the model to discuss policy implications.

[17] Like with prices, these 11 specifications generally include either seller-product-month or destination-seller-product fixed effects that absorb most of the across orders variation in the data.

5.1 Relational Buyers vs. Relational Relationships

A distinctive feature of our analysis is that we consider the sourcing strategy as a buyer-level characteristic, as opposed to a relationship feature. We subject this approach to further empirical scrutiny. We start by quantifying the role of buyer- vis-à-vis relationship-specific effects in explaining variation in prices and markups and then revisit our baseline specification introducing a measure of "relationalness" at the buyer-seller level.

Section 2.3 argued that buyer-level factors (capabilities) are key drivers of observed sourcing patterns. This does not preclude a buyer's sourcing behavior to vary across different suppliers. For example, a relational buyer might source relationally from a core set of suppliers but also source spot from a fringe of suppliers that are used at times of especially high demand. As already noted, this would imply that our approach *under*estimates the influence of relational sourcing on prices and markups.

We borrow from the employer-employee literature (see Card et al., 2012) to assess the relative importance of buyer and buyer-seller effects in explaining order level prices and markups. We compare the explanatory power of a model that includes buyer fixed effects with one that includes buyer-seller fixed effects, conditional on the baseline fixed effects δ_{sjt}. Appendix Table D10 shows that between 95.5% (87.8%) and 98.5% (90.1%) of the fit in a model of prices (markups) with bilateral buyer-seller fixed effects, can be explained by a model with buyer effects only. In both cases, we cannot reject the null hypothesis that the model with buyer effects explains as much variation as the one with relationship effects.

We also study buyer-seller specific sourcing explicitly. To do so in a manner that is consistent with our measure of relational sourcing at the buyer level, we construct an analogous metric of relationalness at the buyer-seller level. Taylor and Wiggins (1997)'s model implies that, holding constant sourced volumes, relational trade is associated with more frequent shipments. This suggests to proxy the relational nature of buyer-seller pairs using traded volumes (in kilos) per shipment. Formally,

$$Relational_{sb} = -1 \times \sum_{jt \in sb} \left[\frac{q_{sbjt}}{q_{sb}} \times \frac{1}{\#Shipments_{sbjt}} \right]. \tag{7}$$

This measure has three appealing features. First, it is consistent with our approach to measuring sourcing practices in the rest of the paper. Second, it is empirically correlated with measures of relational sourcing at the buyer-seller level (such as relationship duration, see Appendix Table D11). Finally, it can be aggregated over relationships to give a buyer-level measure of sourcing comparable to our baseline, as

$$\widetilde{Relational}_b = \sum_{s \in b} \left[\frac{q_{sb}}{q_b} \times Relational_{sb} \right]. \tag{8}$$

This feature allows us to write the relationship-level metric as $\widetilde{Relational}_{sb} = -1 \times (|Relational_{sb}| - |\widetilde{Relational}_b|)$ and decompose the sourcing strategy into a relationship-specific component and a buyer-level one.

Appendix Table D12 reports results from our baseline specification that include seller-product-year (δ_{sjt}) fixed effects and buyer and order-level controls. Note that we omit controls at the buyer-seller pair level – bilateral volumes, shares in partner's trade and relationship age – that we have included in the baseline specification since they also proxy for relationalness. For comparability with the rest of the table, column (1) reports the baseline relational sourcing metric, $Relational_b$, in levels (i.e., unlike the rest of the paper, not standardized). Column (2) reports the buyer's relational metric, $\widetilde{Relational}_b$ in excluded products. This metric benchmarks the estimates once we introduce buyer-seller specific metrics. Results confirm both positive correlations. Column (3) uses the relationship-specific measure, $Relational_{sb}$. There is a positive correlation between the bilateral metric and both prices and markups. Of course endogeneity concerns prevent us from interpreting these correlations: the bilateral sourcing metric could be correlated with unobserved aspects that are also correlated with our outcomes of interests – a further reason why we favor our buyer-level approach.

Column (4) provides our main test. It includes *both* the buyer-level metric and the bilateral metric, centered around the buyer's mean ($\widetilde{Relational}_b$). The buyer level metric has a positive coefficient in both the prices and markups regressions that is, if anything, larger relative to column (2). This is expected: as noted above, omitting the relationship level metrics introduces measurement error and biases our estimates towards zero.

Note also that the bilateral proxy for relationalness is positively correlated with markups but not with prices. This suggests that sellers incur lower marginal costs to produce for buyers with whom they trade relationally. This result, however, is difficult to interpret. On the one hand, suppliers might learn and become more efficient – either through learning by doing, or from transfers of capabilities from the buyer – when producing orders for their more stable partners. On the other hand, it might be those orders for which the supplier has lower costs that are more likely to result in more stable matches with buyers. The difficulty in interpreting results with proxies for relationalness at the buyer-seller level provides further justification for our approach that considers the buyer-level proxies for relational sourcing computed in excluded products. Removing these concerns facilitates interpretation.

5.2 An Event Study

Our evidence is identified out of cross-sectional variation in sourcing strategies across buyers. We leverage an event to probe whether the patterns are robust to within-buyer changes in sourcing strategies over time. We zoom in on the Bangladeshi supplier base of VF, the large apparel buyer mentioned in Section 2. In 2004, VF begun a shift in its global sourcing from a spot strategy to a relational approach (see Pisano and Adams, 2009 for details). The approach was called *Third Way* "*because it represented an alternative to both in-house manufacturing and traditional sourcing*". VF used to source internally from its own plants and externally from suppliers through short-term contracts. The transition was slow: within a global supply network of over 1,000 suppliers, by 2009 there were only five *Third Way* suppliers (none among VF's Bangladeshi suppliers in our product categories). The new approach ramped up globally in 2010.

The data confirm the profound supply-chain restructuring brought about by the transition to the *Third Way*. The transition induced a significant degree of churning – with the termination of many suppliers and fewer new ones being added after the transition. At the same time, VF expanded the volumes sourced from suppliers: the number of orders per supplier increased from around 160 in 2005 to close to 400 in 2012. Before the transition, VF accounted for 27-28% of the volumes exported by its suppliers in 2007 and 2008 on average. That share jumped to 44-47% in 2010 and 2011. This is consistent with VF's (continuing) suppliers, dropping buyers they had previously supplied. On average, VF's continuing suppliers dropped 2.6 more buyers and begun supplying 2.1 fewer new buyers after VF's transition relative to the pre-period In line with our metric for relational sourcing, VF consolidated its supplier base in Bangladesh. The transition, however, also implied a re-organization of suppliers' own set of buyers.

We compare the evolution of the markups earned in orders sold to VF relative to other buyers within a difference-in-differences framework,

$$\mu_{sbjo} = \delta_{sjt} + \delta_b + \sum_{r=2005}^{2012} \beta_r VF_b \times I_{t(o)=r} + \gamma Z_{sbjo} + \varepsilon_{sbjo}, \tag{9}$$

where VF_b is an indicator that takes value one if the buyer in the order is VF and zero otherwise, while $I_{t(o)=r}$ is a dummy for year r ($r = 2009$ is the excluded year). We include seller-product-year fixed effects, δ_{sjt}, as in our baseline specification and thus compare *changes* in differences in order-level markups between VF and other buyers. The inclusion of buyer fixed effects, δ_b, accounts for unobservable, time-invariant buyer characteristics.

The churning of trade partners resulting from the transition implies that a simple before-and-after comparison of orders is marred by selection effects. The most likely form of selection is that in the post period continuing suppliers likely dismiss (and avoid forming new relationships) with buyers that are less profitable, i.e., those on which they earn lower markups. This implies that the DID coefficient estimated on the entire sample would be biased *downward*. We restrict the sample to only include the main buyers of the supplier – defined as those accounting for at least 20% of the sellers' non-VF exports in each year and control for relationship cohort fixed effects as these buyers change over time.

The difference-in-differences analysis confirms the cross-sectional evidence in Section 4. Figure 5 reveals no differential trend in markups in orders sold to VF relative to other buyers before VF's transition. After the transition, orders produced for VF start earning significantly higher markups relative to comparable orders produced for other buyers. The pattern persists until the end of our sample period. Appendix Table D15 shows that the pattern in Figure 5 is driven by an increase in prices following VF's change in its approach to sourcing, rather than changes in marginal costs. Furthermore, alternative samples that deal with selection restricting attention to continuing buyers and to surviving suppliers estimate similar results. Including all buyers, however, estimates a lower and non statistically different from zero effect that is consistent with the selection effect discussed above.

5.3 Mechanisms

Relational buyers pay higher markups. Building on motivating evidence in Section 2.3.4, our model rationalizes this fact through a particular mechanism: the buyer's need to ensure (non-contractible) reliable deliveries. While we believe this mechanism to be relevant in our context, we do not contend that it is the only mechanism that might be at play. We discuss in greater detail the reliability mechanism before turning to other potential alternative explanations to our main finding.

Reliability The model conceptualizes reliability as a costly action that is difficult to contract upon – akin to a pure moral hazard model. This is motivated by suggestive evidence in Section 2.3.4, indicating that the disruptive effects of hartals – as evidenced by delayed orders conditional on size – are mitigated in the case of orders for relational buyers. Indeed, as shown in Appendix Table D13, a shorter order throughput time is associated with both higher markups and relational buyers, suggesting that on-time deliveries are an important aspect of relational sourcing (see also Taylor and Wiggins, 1997 endogenizing the frequency of shipments).[18]

[18]In the presence of demand shocks, *flexibility* – intended as the supplier's ability to accelerate production or allocate additional production capacity at short notice – can also be important. In such cases, we also expect the order lead-time (the time elapsed between the incoming shipment of fabric and the outgoing

An interesting question is whether reliability could instead be considered as a (possibly hidden) type, whereby only some suppliers are able to be reliable. Our empirical analysis includes seller-(product-time) fixed effects. These fixed effects thus control for the type of the seller – whether observed or not. A model in which reliability is purely a type is thus difficult to reconcile with the different prices and markups charged by the same seller to different buyers. An alternative formulation in which reliability is the result of both hidden types and actions is also possible. In such a reputation model, buyers' beliefs matter for how the seller responds to shocks. For example, Macchiavello and Morjaria (2015) develop and test such a model. In their model, uncertainty over the seller's type (whether she is reliable or not) and the seller's actions (whether she exerts effort or not to prioritize the buyer) influence buyers' beliefs about the value of future interactions with the seller. A common feature of such models is that uncertainty over types is needed to preserve reputational incentives. They find an inverted-U pattern in sellers' responses to an unanticipated supply shock: sellers prioritize relationships that are neither too young, nor too old. Young relationships are not valuable enough; in old relationships, there is nothing left to prove. The findings in Macchiavello and Morjaria (2015) suggest that, during hartals, relational buyers may be prioritized, but may also give slack to the exporter depending on circumstances that are unobservable to us. Furthermore, unlike the shock in Macchiavello and Morjaria (2015) – which is large, unanticipated and observable – the hartals in Section 2.3.4 are relatively small, frequent and measured with significant error. These considerations limit our ability to replicate their analysis to untangle a pure moral hazard model from a model with both moral hazard and hidden types.[19]

To close, we do not contend that reliability is the *only* mechanism that generates a rationale for relational contracting in this industry – nor that reliability should be conceptualized within a pure moral hazard framework with no hidden types. Reliability appears to be an empirically plausible mechanism that, once formalized, is consistent with the motivating facts we put forward, as well as the main empirical result in the paper. We now turn to discussing alternative mechanisms that are consistent with some – but not all – the facts in the paper.

shipment of garments) (i) to correlate (negatively) with relational sourcing, (ii) to correlate (positively) with higher markups, and (iii) sourcing to still display a positive correlation with markups, conditional on lead-times. Appendix Table D13 finds support for these three patterns and thus supports flexibility as an additional mechanisms.

[19] Macchiavello and Morjaria (2015) point out that a further implication of a model with hidden types is that the relationship dynamics are not stationary – a prediction tested looking at the age profile of contractual outcomes over the course of a relationship. In our context, unreported results reveal that, conditional on relationship fixed effects, prices increase over time in the first few years of the relationships – but this effect is not different between relational and non-relational buyers. Later in the relationship, a (weak) differential dynamic effect on markups between relational buyers and non-relational buyers appears.

Demand Assurance Not all models with relational contracts imply that relational buyers pay higher prices than spot buyers. With demand uncertainty (see, e.g., Carlton, 1978 and Dana, 1998), for instance, suppliers face uncertain capacity utilization. Relational buyers may promise reliable capacity utilization and offer relational rents to sellers in the form of lower costs. Indeed, in industries in which demand uncertainty is important, prices tend to be *lower* in long-term relationships (see, e.g., Macchiavello and Miquel-Florensa, 2018 and Pirrong, 1993). Furthermore, in such industries, buyers might adopt a dual sourcing strategy in which they keep a few "reliable" suppliers to serve the stable part of demand and then use a fringe of spot suppliers to cover unforeseen spikes. Based on our understanding of the sector as well as interviews in the field, demand assurance is likely also an important aspect of relationships in this context. Yet, our results suggest that this is quantitatively outweighed by alternative mechanisms – such as the reliability mechanism in our model – that imply higher prices. Our estimates thus understate the value of relational buyers to exporters if demand assurance is at play.[20]

Costly, seller-specific, capabilities. Sellers might need to undertake, and be compensated for, specific investments to supply relational buyers. To the extent that our data allows it, we do not find evidence for such differences. However we cannot rule out other unobservable costs, fixed from the perspective of an export order, that are necessary to build capabilities to supply relational buyers (see, e.g., footnote 22 below for a quantification of such costs, unlikely to outweigh the higher markups paid by relational buyers). If the ability to supply relational buyers was a seller capability, however, the higher markups paid by relational buyers would induce sellers that have acquired such capabilities to specialize in supplying relational buyers. In contrast, the evidence in Section 2.3.3 reveals that sellers supply a mix of relational and spot buyers – a pattern that is naturally explained by the mechanism in our model.

Product Quality. Differences in the physical quality of products are unlikely to account for the observed differences in markups across buyers. The –rather limited– existing empirical evidence suggests that higher quality products are associated with higher markups (see, e.g., De Roux et al., 2020 and Atkin et al., 2015). However, buyers with different sourcing strategies do not appear to differ in the quality of the garments they source. Table 5 found no

[20]Our model could be extended to consider flexibility as a response to demand assurance incentives. At full capacity, flexibility requires an exporter to divert resources from other orders: flexibility towards a buyer compromises reliability towards another one. An exporter that only supplies relational buyers thus cannot simultaneously guarantee flexibility and reliability to *all* of them. This logic is consistent with our motivating fact in Section 2.3.3.

differences in SMVs – a direct measure of a garment's technical complexity – between orders produced for relational and spot buyers. While this is suggestive, other dimensions of quality are unobserved. Two pieces of evidence assuage such concerns. First, higher quality garments are produced using higher quality inputs (see, e.g., Kugler and Verhoogen, 2012) – they are made with better fabric and sewed by more skilled operators. We found no differences in the price and type of both fabric and labor across orders produced for buyers with different sourcing strategies. Furthermore, Table 6 reports results from additional specifications that further control for proxies for physical quality, including specialization, seasonality, proxies for product complexity, and dummies for the type and origin of the fabric most used in the order. Results are robust to the inclusion of these different proxies for product quality.

Bargaining Power. Differences in bargaining power are unlikely to explain our results. Before we discuss the evidence for this assertion, we introduce an important distinction between *ex-ante* and *ex-post* bargaining power. The former refers to the relative strength of parties as they negotiate their initial agreement. Note that in our model, relational buyers do have *ex-ante* bargaining power and indeed negotiate a relational price that is the most favorable price that still satisfies the supplier's incentive compatibility constraints. Despite this, they still pay higher prices and markups relative to spot buyers. To rationalize the evidence without introducing incentive compatibility constraints, an alternative model would thus need to assume that relational buyers have *lower* ex-ante bargaining power. When we control for common proxies for bargaining power, however, we find that our results remain robust. All specifications in Table 6 control for the buyer's size in the market, the age of the buyer-seller relationship and traded volumes between parties. In addition, they control for the share of the buyer (seller) in the seller's (buyer's) trade. Furthermore, column (7) (respectively (8)) discards orders sold to (bought from) the main buyer (supplier) and finds that the result remain largely unchanged in the sample of orders in secondary relationships. These patterns suggest that the higher markups paid by relational buyers are unlikely to solely reflect a weaker ex-ante bargaining position of these buyers *vis-à-vis* their suppliers.

Ex-post bargaining power refers instead to parties' relative negotiating position once the relationship has been formed. By design, relational buyers *chose* to have lower *ex-post* bargaining power *vis-à-vis* their suppliers: quoting from Helper and Sako (1998) "*a deliberate strategy of locking oneself into a relationship, thus raising switching costs, may facilitate the creation and maintenance of trust*". In other words, lower *ex-post* bargaining power should not be considered an alternative explanation to be ruled out – it is a quintessential feature of relational sourcing systems.

Search and Switching Costs. Analogously to the previous discussion, it is useful to distinguish between search and switching costs. Search costs are relevant at the *ex-ante* stage, i.e., when a buyer is searching for, and negotiating with, adequate suppliers. Differences in search costs are unlikely to explain our evidence. Buyers may differ in their costs of searching for a supplier. Certain buyers may be more patient (or have lower search costs) and thus be "pickier": they search for longer to find a suitable supplier and, when they find one, they establish long lasting relationships, thus mimicking relational behavior. In standard models, however, more patient buyers have stronger bargaining power and negotiate a *lower* price – we find instead that relational buyers with more stable relationships pay *higher* prices. The prediction, however, could be reversed if "picky" buyers attain higher value matches. These buyers would form lasting relationships that generate more surplus – potentially shared with the supplier in the form of higher prices. The evidence in Section 5.1, however, suggests that match-specific components are unlikely to quantitatively account for our patterns and that controlling for proxies at the buyer-seller pair strengthens our results.[21]

Switching costs, instead, refer to the cost of finding alternative suppliers *ex-post*, i.e., once the relationship has been established. As with bargaining power, it is a deliberate strategy to introduce higher switching costs to support the relationship. Switching costs are thus not an unobserved confounder to be ruled out, but rather an attribute of relational sourcing systems.

Pricing to Market and Rent Sharing. In our context, higher markups could also stem from sellers' discriminating pricing across markets. By controlling for destination fixed effects, Table 6 accounts for average differences across destinations. Appendix Table D14 further explores related confounders. For ease of comparison, column (1) reproduces column (3) of Table 6. Column (2) includes destination-product-year fixed effects while column (3) controls for seller-destination fixed effects. Column (4) includes country-product-year fixed effects, where country corresponds to where the order is shipped to (which could differ from the main destination of the buyer). These fixed effects control for differences in markups following sellers' pricing-to-market behavior and from heterogeneous consumers' tastes across countries, products and time. These mechanisms do not explain the markup differentials across buyers, which remain robust throughout the exercise.

Relational buyers might also have higher market power downstream and pass-through some of their profits to upstream suppliers. In this case, the higher markups paid by relational

[21] In the same context of this paper, Cajal-Grossi (2021) develops, and tests, a model in which buyers that are sensitive to the risk of reputational losses due to suppliers' misconduct may have both higher search costs *and* values from being matched with suppliers of a higher type. She finds that these buyers experiment with different suppliers, but do so to a lesser extent when reputational risks are high.

buyers might reflect profit sharing. To explore this possibility, we match the buyers in our sample with data from Euromonitor, which capture the sales of the buyer in the destination market (Euromonitor International, 2015). We find 53 buyers for which the downstream market share is observed for every year in our sample. Columns (5) and (6) of Appendix Table D14 shows that our results are robust to controlling for the buyer's sales in the downstream market within this restricted sample.

5.4 How Valuable are Relational Buyers?

We now explore the quantitative implications of our estimates – as well as their limitations – through a back of the envelope calculation. The estimated correlation between markups and relational sourcing is quantitatively sizable. Our baseline specification reveals that a one standard deviation increase in the buyer's measure of relationalness is associated with a 0.026 increase in the (log) markup factor. To interpret this magnitude, consider the average markup factor (1.44) and marginal cost ($10.35) estimated in Online Appendix B.5. The estimated coefficient implies that a shift in sourcing strategy from a spot approach like Kik's to relational sourcing like H&M's is associated with an additional $0.32 per kilogram of garments, equivalent to a 9.8% increase over the average markup value ($3.32). Put otherwise, a change in sourcing strategy from the average buyer to The Gap (a shift of about one standard deviation) yields an increase in markup of approximately 11%. Comparing the 25^{th} (10^{th}) to 75^{th} (90^{th}) percentiles in the distribution of buyers' relational metric gives a 15.3% (30.6%) increase over the average markup value.[22]

To our knowledge, there are no other estimates of markups earned from specific buyers in the literature – it is thus difficult to benchmark our results. Macchiavello and Morjaria (2015) and Blouin and Macchiavello (2019) estimate the net present value of a relationship through a revealed preference approach. From a seller's point of view, the relationship with a buyer is worth at least as much as the seller's "temptations to deviate" – which is directly observed in those papers. Both studies find that the relationships with buyers are highly valuable. For example, Macchiavello and Morjaria (2015) find that to the typical Kenyan rose dealer, the average relationship is worth 161% of their weekly turnover . Assuming a profit margin of 10%, this translates into a net present value of $161\%/(10\% \times 52) \approx 30\%$

[22] Relational contracts are plausibly more time-consuming for contract managers/procurement managers, as they are expected to offer a more personalized service to relational buyers. While these administrative costs are not variable at the order level, they could potentially erode the extra profits that suppliers earn from relational buyers. The average annual gross salary of *Managing Directors and Chief Executives* in manufacturing in Bangladesh's Labor Force Survey of 2017 is 4,755 USD (34,133 BDT per month, for 12 months). The average size of orders from relational buyers in our analysis sample is over 64.5 thousand kilograms. The extra 0.32 USD markup, thus, amounts to almost 20 thousand dollars in the average order, almost four times the annual gross pay of a managing director in manufacturing.

of the yearly profits from that relationship. To benchmark these estimates to ours, we need to discount the estimated markup increase. Conditional on the buyer and the seller trading at least one year, the average duration of relationships with relational buyers is $D = 3.71$ years. Assuming an annual interest rate of 15% yields an effective discount factor $\delta = 1/(1 + 0.15) \times (1 - 1/D) \approx 0.635$. This yields a net present value in the range $9.8\%/(1 - 0.635) \approx 26.8\%$ to $11\%/(1 - 0.635) = 30.15\%$.

The reduced form results in our paper, however, likely underestimate the value of relational buyers. First, our proxy for the buyer's relational strategy is conservative and suffers from attenuation bias. Using alternative definitions (Appendix Table C3) often yields higher estimates; controlling for a bilateral proxy for relationalness increases our buyer level estimate by a tenth (Appendix Table D12). Second, we are not taking into account two potentially important sources of value from supplying relational buyers: higher volumes and demand assurance. Relational buyers source larger volumes than spot buyers from their suppliers – a typical supplier thus earns significantly higher variable profits when supplying relational buyers relative to spot buyers. Relational buyers are also likely to provide a more stable demand, thereby allowing for better capacity planning and utilization - as well as lower costs. On the other hand, in our model suppliers incur a loss when delivering to relational buyers while hit by the shock. The within-seller-time comparison of our baseline, thus, might overestimate the difference in average markups between relational and spot buyers. Given this limitations, we see the development of structural models to estimate the value of relationships as an important avenue for future research.

5.5 Policy Implications

We now return to our model to discuss policy implications. Due to contracting problems, the spot market is not efficient: when suppliers are hit by shocks, they sell to relational buyers but not to spot buyers, despite the fact that their cost is lower than buyers' valuation ($c_1 < v$). Some capacity thus remains inefficiently unutilized, and overall market efficiency is increasing in the share of relational buyers ρ. In deciding whether to become relational, however, a buyer only takes into account his private returns but not the rents that his investment generates for other market participants. As a result, there is insufficient entry of relational buyers in equilibrium. In such circumstances, a planner may want to intervene and subsidize the entry cost of relational buyers.

Formally, consider a subsidy τ to be paid to relational buyers. The planner chooses τ in order to maximize welfare $W(\tau)$ subject to the buyer entry condition (BE), where

$$
\begin{aligned}
W(\tau) =& \rho(\tau) \underbrace{\delta \left[p^R - \alpha c_1 - (1-\alpha)c_0 + (1-\alpha)(\underline{v} - c_0)\right]}_{\text{Profits of sellers in relationships}} + (S - \rho(\tau)) \underbrace{\delta 2(1-\alpha)(\underline{v} - c_0)}_{\text{Profits of sellers not in relationships}} \\
&+ \lambda \underbrace{\left\{\rho(\tau) \left[\delta(v - p^R) - (1-\delta)(F-\tau)\right] + (1-\rho(\tau))\delta\mu(\rho(\tau))(v - \underline{v})\right\}}_{\text{Buyers' profits}} \\
&- \rho(\tau)(1-\delta)\Psi\tau.
\end{aligned}
$$

Here $\lambda \in [0,1]$ is the welfare weight assigned to buyers, and $\Psi \geq 1$ is the marginal cost of public funds (Stiglitz and Dasgupta, 1971) for the planner. The case $\lambda = 0$ captures the preferences of an export promotion agency that is exclusively concerned with exporters' profits and the cost of public funds. The optimal subsidy is increasing in λ. To show that a planner may want to subsidize entry of relational buyers, it is thus sufficient to focus on the case $\lambda = 0$ to obtain a lower bound to the optimal subsidy. The (BE) condition defines a relationship between the subsidy τ and the share of relational buyers ρ: its implicit differentiation yields $\frac{d\rho}{d\tau} = \frac{1-\delta}{\delta\mu'(\rho)(v-\underline{v})} > 0$. Taking the derivative of $W(\tau)$ with respect to τ, a positive subsidy is optimal when $\lambda = 0$ if

$$
\frac{dW(\tau)}{d\tau}\bigg|_{\tau=0} = \frac{d\rho}{d\tau}\delta\left[p^R - \alpha c_1 - (1-\alpha)\underline{v}\right] - \rho(0)(1-\delta)\Psi > 0.
$$

Substituting with $p^R = \underline{p}^R$, one can verify that $\left[p^R - \alpha c_1 - (1-\alpha)\underline{v}\right] > 0$. Therefore, provided that the equilibrium share of relational buyers under no subsidy, $\rho(0)$, is sufficiently small, a subsidy is justified. This is the case for a planner that only cares about sellers' profits and the cost of public funds and – a fortiori – for a planner that also values buyers' profits.

We have assumed away ex-ante lump-sum transfers between sellers and buyers. This implies that sellers earn rents in equilibrium. While the assumption provides a rationale for policy intervention when $\lambda = 0$, the assumption is not needed to rationalize a subsidy to the entry of relational buyers. It can be shown that, even if buyers could capture all the rents from relational trade by charging suppliers an ex-ante lump fee, a planner with $\lambda = 1$ would subsidize entry of relational buyers if the equilibrium share of relational buyers and cost of public funds are sufficiently low. The reason is that spot buyers are better off when there are more relational buyers in the market ($\mu'(\rho) > 0$), and thus a buyer investing in relational capabilities exerts a positive externality on spot buyers.

These observations provide a practical justification for our approach to consider the sourcing strategy as a buyer-level – as opposed to a buyer-seller pair level – attribute. Even though organizational level capabilities are important to build relational arrangements with suppliers, a particular relational contract between a buyer and one of their suppliers will still be rooted in a mutual understanding of the specific circumstances of that particular pair (Gibbons and Henderson, 2012; Baker et al., 2002). It is thus difficult for policy makers – e.g., export promotion agencies in developing countries – to directly improve specific relationships between exporters and buyers. On the other hand, if certain buyers possess organizational capabilities that make them valuable relational partners, an actionable margin for policy opens up. It might be possible to attract such buyers, e.g., by subsidizing visits to the country or understanding the specific factors that favor their entry.

6 Conclusions

This paper studied how order-level prices, variable costs and suppliers' markups vary with the sourcing strategies of international buyers in the Bangladeshi garment sector. We contributed novel evidence that sourcing strategies are largely driven by buyer-level capabilities, leading us to propose a model in which ex-ante identical buyers endogenously chose different sourcing strategies in equilibrium. The main prediction of the model is that, to induce suppliers' reliable deliveries under bad contingencies, relational buyers pay higher markups relative to spot buyers for otherwise identical orders from the same supplier. We tested and found empirical support for this prediction by leveraging original data that allow for the direct measurement of utilization and prices of the main variable inputs (fabric and labor) used for producing orders for different buyers.

Interpreted through the lens of the model, the empirical results have policy implications for export promotion agencies, particularly in developing countries. The results provide quantitative support to the view that international buyers' sourcing strategies are a potentially important dimension of upgrading for exporting firms in developing countries (see, e.g, Gereffi, 1999 and Egan and Mody, 1992). Similarly to models that distinguish between "good jobs" – in which workers earn rents – versus "bad jobs" (see, e.g., Acemoglu, 2001), the laissez-faire equilibrium generates too few relational buyers relative to the social optimum. This gives rise to the possibility that export promotion agencies might want to target programs to assist exporters in establishing relationships with relational buyers.

This paper provides a first step towards a more systematic understanding of the implications of sourcing practices for economic performance and international trade. Much research remains to be done, however, and we hope that our results will spur further work on this important topic. Two areas appear to be particularly pressing. First, while we have focused on suppliers' markups, buyers' sourcing strategies likely impact other important aspects of supply chains' performance, e.g., their resilience to, and transmission of, shocks; the transfers of capabilities to suppliers – especially in developing countries. Second, we have documented, rationalized, and then taken as given, substantial unexplained variation in sourcing strategies across firms within a narrowly defined sector. The discussion of the policy implications of our results, however, suggests that exploring drivers of buyers' choices of sourcing strategy should be a priority in future research.

References

Acemoglu, Daron, "Good Jobs versus Bad Jobs," *Journal of Labor Economics*, 2001, *19* (1), 1–21.

Adhvaryu, Achyuta, Vittorio Bassi, Anant Nyshadham, and Jorge A Tamayo, "No line left behind: Assortative matching inside the firm," Technical Report, National Bureau of Economic Research 2020.

Ahsan, Reshad N and Kazi Iqbal, "Political strikes and its impact on trade," *IGC Working Paper*, 2015.

Alfaro, Laura and Andrew Charlton, "Intra-industry Foreign Direct Investment," *American Economic Review*, December 2009, *99* (5), 2096–2119.

Amengual, Matthew and Greg Distelhorst, "Can Sourcing Help Enforce Global Labor Standards? Evidence from the Gap Inc Supply Chain," *mimeo*, 2019.

Antràs, Pol, "Firms, Contracts, and Trade Structure," *The Quarterly Journal of Economics*, 2003, *118* (4), 1375–1418.

_ , *Global Production: Firms, Contracts, and Trade Structure*, Princeton University Press, 2016.

_ , "Conceptual aspects of global value chains," *The World Bank Economic Review*, 2020, *34* (3), 551–574.

_ **and Davin Chor**, "Organizing the Global Value Chain," *Econometrica*, 2013, *81* (6), 2127–2204.

_ **and Fritz Foley**, "Poultry in Motion: A Study of International Trade Finance Practices," *Journal of Political Economy*, 2015, *123* (4), 853–901.

Ashraf, Anik, Rocco Macchiavello, Atonu Rabbani, and Christopher Woodruff, "The effect of political and labour unrest on productivity," Technical Report, International Growth Center Working Paper 2015.

Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit Khandelwal, and Eric Verhoogen, "Markup and Cost Dispersion across Firms: Direct Evidence from Producer Surveys in Pakistan," *American Economic Review*, May 2015, *105* (5), 537–544.

Baker, George, Robert Gibbons, and Kevin J. Murphy, "Relational Contracts and the Theory of the Firm," *The Quarterly Journal of Economics*, 02 2002, *117* (1), 39–84.

Bernard, Andrew B, J Bradford Jensen, Stephen J Redding, and Peter K Schott, "Firms in International Trade," *Journal of Economic Perspectives*, 2007, *21* (3), 105–130.

Bloom, Nicholas and John Van Reenen, "Measuring and explaining management practices across firms and countries," *Quarterly Journal of Economics*, 2007, *122* (4), 1341 to 1408.

_ and _ , "Why Do Management Practices Differ across Firms and Countries?," *Journal of Economic Perspectives*, March 2010, *24* (1), 203–24.

Blouin, Arthur and Rocco Macchiavello, "Strategic Default in the International Coffee Market," *The Quarterly Journal of Economics*, 01 2019, *134* (2), 895–951.

Boehm, Johannes and Jan Sonntag, "Vertical Integration and Foreclosure: Evidence from Production Network Data," CEP Discussion Papers dp1641, Centre for Economic Performance, LSE August 2019.

Boudreau, Laura, "Multinational enforcement of labor law: Experimental evidence from Bangladesh's apparel sector," *mimeo*, 2020.

Brandt, Loren, Gueorgui Kambourov, and Kjetil Storesletten, "Barriers to entry and regional economic growth in China," 2020.

Cajal-Grossi, Julia, "Searching for Trade Partners in Developing Countries," *mimeo*, August 2021.

_ , Davide Del Prete, and Rocco Macchiavello, "Sourcing Strategies and Resiliance of GVCs in Garments," 2022.

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez, "Inequality at work: The effect of peer salaries on job satisfaction," *American Economic Review*, 2012, *102* (6), 2981–3003.

Carlton, Dennis W, "Market behavior with demand uncertainty and price inflexibility," *The American Economic Review*, 1978, *68* (4), 571–587.

_ , "The rigidity of prices," 1986.

Chor, Davin, Kalina Manova, and Zhihong Yu, "Growing like China: Firm performance and global production line position," *Journal of International Economics*, 2021, *130*, 103445.

Cooper, Martha C and Lisa M Ellram, "Characteristics of supply chain management and the implications for purchasing and logistics strategy," *The international journal of logistics management*, 1993.

Costinot, Arnaud, Lindsay Oldenski, and J. Rauch, "Adaptation and the Boundary of Multinational Firms," *Review of Economics and Statistics*, 2011, *93* (1), 298–308.

De Loecker, Jan, Pinelopi K Goldberg, Amit Khandelwal, and Nina Pavcnik, "Prices, Markups, and Trade Reform," *Econometrica*, 2016, *84* (2), 445–510.

De Roux, Nicolas, Rocco Macchiavello, Josepa Miquel-Florensa, and Eric Verhoogen, "Markups and Quality: Evidence from Colombian Coffee Exports," *mimeo* December 2020.

De Toni, Alberto and Guido Nassimbeni, "Just-in-Time Purchasing: an Empirical Study of Operational Practices, Supplier Development and Performance," *Omega-International Journal of Management Science*, December 2000, *28* (6), 631–651.

Dickerson, Kitty G, *Textiles and Apparel in the Global Economy*, Prentice Hall, 1999.

Dyer, Jeffrey H, Dong Sung Cho, and Wujin Cgu, "Strategic supplier segmentation: The next "best practice" in supply chain management," *California management review*, 1998, *40* (2), 57–77.

Egan, Mary Lou and Ashoka Mody, "Buyer-Seller Links in Export Development," *World Development*, March 1992, *20* (3), 321–334.

Euromonitor International, "GBO and GBN, sales, Apparel," *Passport Database, Accessed 2015: http://www.euromonitor.com/research-methodology-market-analysis* 2015.

Garcia-Marin, Alvaro and Nico Voigtländer, "Exporting and Plant-Level Efficiency Gains: It's in the Measure," *Journal of Political Economy*, 2019, *127* (4), 1777–1825.

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

_ **, John Humphrey, and Timothy Sturgeon**, "The governance of global value chains," *Review of international political economy*, 2005, *12* (1), 78–104.

Ghemawat, Pankal and Jose Luis Nueno Iniesta, "Zara: Fast Fashion," *Harvard Business School Case 9-703-407*, 2006.

Gibbons, Robert and Rebecca Henderson, "Relational Contracts and Organizational Capabilities," *Organization Science*, 2012, *23* (5), 1350–1364.

Goldberg, Pinelopi Koujianou, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova, "Imported intermediate inputs and domestic product growth: Evidence from India," *The Quarterly journal of economics*, 2010, *125* (4), 1727–1767.

Heise, Sebastian, "Firm-to-Firm Relationships and Price Rigidity: Theory and Evidence," *mimeo*, August 2021.

— , **Justin R Pierce, Georg Schaur, and Peter K Schott**, "Tariff rate uncertainty and the structure of supply chains," Technical Report, Working Paper, Federal Reserve Bank of New York 2021.

Helper, Susan and Abdul Munasib, "Economies of scope and relational contracts: Exploring global value chains in the automotive industry mimeo," *mimeo*, 2021.

— **and Mari Sako**, "Supplier Relations in Japan and the United States: Are They Converging?," *MIT-Japan Program*, 1997.

— **and** — , "Determinants of trust in supplier relations: Evidence from the automotive industry in Japan and the United States," *Journal of Economic Behavior & Organization*, 1998, *34* (3), 387–417.

— **and Rebecca Henderson**, "Management Practices, Relational Contracts, and the Decline of General Motors," *Journal of Economic Perspectives*, February 2014, *28* (1), 49–72.

Jr, James D Dana, "Advance-purchase discounts and price discrimination in competitive markets," *Journal of Political Economy*, 1998, *106* (2), 395–422.

Kugler, Maurice and Eric Verhoogen, "Prices, Plant Size, and Product Quality," *Review of Economic Studies*, February 2012, *79* (1), 307–339.

Macchiavello, Rocco, "Relational Contracts and Development," *Annual Review of Economics*, 2022, *14*.

— **and Ameet Morjaria**, "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports," *American Economic Review*, 2015, *105* (9), 2911–45.

— **and** — , "Competition and relational contracts in the Rwanda coffee chain," *The Quarterly Journal of Economics*, 2021, *136* (2), 1089–1143.

— **and Christopher Woodruff**, "Supervisory Management and Productivity Dispersion in the Bangladeshi Garment Sector," Technical Report 971, Society for Economic Dynamics 2014.

— **and Josepa Miquel-Florensa**, "Vertical Integration and Inter-Firm Relationships in the Costa Rica Coffee Chain," *mimeo*, 2018.

— **and** — , "Improving Export Quality: The case of the Sustainable Quality Program in Colombian Coffee," *mimeo*, 2019.

— , **Andreas Menzel, Atonu Rabbani, and Christopher Woodruff**, "Challenges of change: An experiment promoting women to managerial roles in the bangladeshi garment sector," Technical Report, National Bureau of Economic Research 2020.

MacLeod, W Bentley and James M Malcomson, "Motivation and markets," *American Economic Review*, 1998, pp. 388–411.

Martin, Julien, Isabelle Mejean, and Mathieu Parenti, "Relationship stickiness, international trade, and economic uncertainty," *CEPR Discussion Paper No. DP15609*, 2020.

Masten, Scott, "The Organization of Production: Evidence from the Aerospace Industry," *Journal of Law and Economics*, 1984, *27* (2), 403–17.

Mayer, Thierry, Marc J. Melitz, and Gianmarco I. P. Ottaviano, "Market Size, Competition, and the Product Mix of Exporters," *American Economic Review*, February 2014, *104* (2), 495–536.

McMillan, John, "Managing Suppliers: Incentive Systems in Japanese and U.S. Industry," *California Management Review*, July 1990, *32* (4), 38–55.

Milgrom, Paul and John Roberts, "The economics of modern manufacturing: Technology, strategy, and organization," *The American Economic Review*, 1990, pp. 511–528.

_ **and** _ , "Complementarities and fit strategy, structure, and organizational change in manufacturing," *Journal of accounting and economics*, 1995, *19* (2-3), 179–208.

Monteverde, K. and D. Teece, "Suppplier switching costs and vertical integration in the automobile industry," *The Bell Journal of Economics*, 1982, *13* (1), 206 to 213.

Morlacco, Monica, "Market Power in Input Markets: Theory and Evidence from French Manufacturing," *mimeo*, 2019.

Nien-he, Hsieh, Michael Toffel, and Olivia Hull, "Global Sourcing at Nike," *Harvard Business School Case 619-008*, June 2019.

Nishiguchi, Toshihiro, *Strategic Industrial Sourcing: The Japanese Advantage*, Oxford University Press, 1994.

Pavcnik, Nina, "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants," *Review of Economic Studies*, 2002, *69*, 245–276.

Pirrong, Stephen Craig, "Contracting practices in bulk shipping markets: a transactions cost explanation," *The Journal of Law and Economics*, 1993, *36* (2), 937–976.

Pisano, G. and P. Adams, "VF Brands: Global Supply Chain Strategy," *Harvard Business School Case 610-022*, November 2009.

Ponte, Stefano, Gary Gereffi, and Gale Raj-Reichert, *Handbook on global value chains*, Edward Elgar Publishing, 2019.

Porter, Michael E, *Competitive strategy*, New York Free Press, 1980.

Rauch, James E, "Networks versus markets in international trade," *Journal of international Economics*, 1999, *48* (1), 7–35.

Reenen, John Van, "Increasing differences between firms: Market power and the macro-economy," 2018.

Richardson, James, "Parallel Sourcing and Supplier Performance in the Japanese Automobile Industry," *Strategic Management Journal*, 1993, *14* (5), 339–350.

Spencer, Barbara J, "International outsourcing and incomplete contracts," *Canadian Journal of Economics/Revue canadienne d'économique*, 2005, *38* (4), 1107–1135.

Startz, Meredith, "The Value of Face-to-Face: Search and Contracting Problems in Nigerian Trade," *mimeo*, 2021.

Stiglitz, Joseph E and Partha Dasgupta, "Differential taxation, public goods, and economic efficiency," *The Review of Economic Studies*, 1971, *38* (2), 151–174.

Tanvir, Saiful Islam and Tarique Mahmood, "Solid Waste for Knit Fabric: Quantification and Ratio Analysis," *Journal of Environmental and Earth Science*, 2014, *4* (12), 68–80.

Taylor, Curtis R and Steven N Wiggins, "Competition or compensation: Supplier incentives under the American and Japanese subcontracting systems," *The American Economic Review*, 1997, pp. 598–618.

Urena, Alonso Alfaro, Isabela Manelici, and Jose Vasquez, "The effects of multinationals on workers: evidence from Costa Rican microdata," Technical Report, Princeton University, Department of Economics, Center for Economic Policy ... 2021.

Verhoogen, Eric, "Firm-Level Upgrading in Developing Countries," Working Paper 29461, National Bureau of Economic Research November 2021.

Williamson, Oliver E, "The vertical integration of production: market failure considerations," *The American Economic Review*, 1971, *61* (2), 112–123.

_ , "Markets and hierarchies: analysis and antitrust implications: a study in the economics of internal organization," *University of Illinois at Urbana-Champaign's Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship*, 1975.

_ , *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*, New York: Free Press, 1985.

Woodruff, Christopher, "Non-contractible investments and vertical integration in the Mexican footwear industry," *International Journal of Industrial Organization*, 2002, *20* (8), 1197–1224.

World Bank, *World Development Report 2020: Trading for Development in the Age of Global Value Chains*, 0 International Bank for Reconstruction and Development, The World Bank, 2020.

Tables and Figures

Table 1: Summary Statistics

	Obs.	Mean	Std. Dev.	P10	P25	P50	P75	P90
Panel A: Orders								
$Buy-to-Ship_o$	22,741	0.87	0.29	0.51	0.67	0.86	1.04	1.22
$Length_o$ $(months)$	22,741	4.24	3.25	1.47	2.17	3.3	5.23	8.03
Panel B: Sellers								
$Count^o_{st}$	3,165	14.60	13.07	3	6	11	19	29
$Count^o_{sjt}$	6,872	6.03	7.53	1	2	3	7	14
$Count^j_{st}$	3,165	3.27	1.88	1	2	3	4	6
$Share^j_{st}$	3,165	57.76	34.92	6.40	24.29	62.30	92.67	100
$Count^b_s$	500	20.97	17.10	4	8.5	17	28	42.5
$Count^b_{st}$	3,165	5.88	4.86	1	2	5	8	12
$Count^b_{sjt}$	6,872	2.91	2.91	1	1	2	4	6
$Share^b_{st}$	3,165	43.98	36.91	1.71	8.50	34.67	82.18	100
$Length_s$ $(years)$	500	6.65	1.54	4.08	5.75	7.63	7.75	7.75
Panel C: Buyers								
$Count^o_{bt}$	4,478	13.37	29.75	1	2	5	13	27
$Count^o_{bjt}$	8,070	5.75	11.54	1	1	2	5	12
$Count^j_{bt}$	4,478	4.24	3.83	1	2	3	5	9
$Share^j_{bt}$	4,478	59.47	35.71	5.96	26.41	63.58	100	100
$Count^s_b$	2,529	54.40	50.06	9	18	37	72	137
$Count^s_{bt}$	7,569	22.05	20.52	4	7	14	30	58
$Count^s_{bjt}$	11,942	8.80	9.07	1	3	5	12	21
$Share^s_{bt}$	4,478	48.62	37.97	0	11.72	42.08	92.28	100
$Length_b$ $(years)$	1,578	5.48	2.42	1.58	3.58	6.42	7.67	7.75
Panel D: Relationships								
$Count^o_{sbt}$	10,448	3.38	4.58	1	1	2	4	7
$Count^o_{sbjt}$	12,858	2.52	3.14	1	1	1	3	5
$Count^j_{sbt}$	10,448	1.46	0.85	1	1	1	2	2
$Length_{sb}$ $(years)$	5,658	1.87	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, t to years. $Count^x_y$ is the number of x per y. For example, $Count^o_{sjt}$ 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 are observed trading in the dataset, respectively. A value of 7.75 in these variables implies censoring, given the time span of our dataset. 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 $Share^s_{bt}$, the average seller's share in buyer's trade in a year is 48.62%. The column under the heading 'Obs.' reports the count of *cells* relevant to the level of aggregation of the variable in the row. For example, the first row of Panel C, corresponding to $Count^o_{bt}$ shows that there are 4,478 buyer-year combinations in the data; across these, the average number of orders is 13.37.

Table 2: Buyers' Concentration and Sourcing

	Market Share %	Sellers per Year Average	Relational Ranking	Price (Residuals) Ranking
Top 25 Buyers	(1)	(2)	(3)	(4)
H&M Hennes And Mauritz	5.22	55.25	3	2
Wal Mart Stores	5.00	57.50	17	16
VF Corporation	4.14	23.75	5	17
The Gap Inc	3.44	26.13	1	1
C & A Buying	3.17	41.00	8	9
K Mart Corporation	3.08	59.25	16	14
PVH Corporation	3.11	39.00	7	15
Levi Strauss & Co	2.21	7.38	2	7
J.C. Penney	1.96	25.75	11	10
Primark	1.42	22.75	10	24
Kik Textilen	1.32	49.88	25	22
Tesco	1.25	23.00	12	19
Kohls Department Stores Inc	1.25	16.13	13	5
Asda	1.21	19.50	6	8
Marks& Spencer	1.15	9.88	4	11
Carrefour	1.13	26.38	14	18
G. Gueldenpfennig Gmbh	0.87	30.88	24	20
Tema Magazacilik	0.91	41.63	21	4
Public Clothing Company Inc	0.84	24.75	23	23
Target Stores	0.85	19.38	15	12
Inditex (Zara)	0.81	32.25	20	3
Auchan S.A.	0.71	29.00	19	21
Charles Vogele	0.69	17.25	18	13
The Children's Place	0.68	11.13	9	6
IFG Corporation	0.65	14.13	22	25
Top 100 (Market Share = 66%)				
Mean	0.66	17		
Median	0.29	12		
St. Deviation	0.99	13.74		
Coeff. Variation	1.49	0.81		
All Buyers (N = 1,578)				
Mean	0.06	4.55		
Median	0.01	3		
St. Deviation	0.30	5.95		
Coeff. Variation	5.04	1.31		

The top panel lists the largest 25 buyers in descending order based on their imports of woven garments (trousers and shirts). For each of them, it reports the buyer's market share (column (1)), the number of sellers the buyer trades with on average every year (column (2)), the ranking according to the buyer's relational characteristic in woven products (column (3)) and the ranking of the buyer according to the average price it pays for its orders, residualized against the size of the order and seller-product-year fixed effects (column (4)). The bottom panels of the table report summary statistics of the corresponding variables in columns (1) and (2) across the top 100 buyers and across all buyers.

Table 3: Buyers' Sourcing and Prices

	(1)	(2)	(3)	(4)
		p_{sbjo}		
$Relational_b$	0.020***	0.023***	0.019**	0.021***
	(0.008)	(0.008)	(0.009)	(0.008)
FEs	sjt	sjt,d	sjt,d	sjt,d
Controls	.	B	B,R	B,R,O
R^2	0.57	0.59	0.63	0.73
Obs.	18,664	18,513	15,647	15,647

Standard errors in parentheses, clustered at the buyer level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcome in all regressions is the log price of an order between a seller and a buyer in a given product category, p_{sbjo}. The main regressor in all cases is the baseline, buyer-specific metric of relational sourcing and it is standardized. Column (1) includes our baseline fixed effect, defined at the level of the seller-product-year triplet. Columns (2) to (4) sequentially add buyer-, relationship- and order-level covariates, as follows. Buyer controls (B): fixed effect for the main destination of the buyer, cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship controls (R): Cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order controls (O): size of order (log volume), log price of fabric of the order.

Table 4: Buyers' Sourcing and Input Usage

	(1)	(2)	(3)	(4)	(5)	(6)
	p^f_{sbjo}		$(F/Q)_{sbjo}$		$Complex_{sbjo}$	
$Relational_b$	0.008	0.003	-0.004	-0.007	0.021*	-0.003
	(0.006)	(0.007)	(0.006)	(0.007)	(0.012)	(0.011)
FEs	sjt	sjt,d	sjt	sjt,d	sjt	sjt,d
Controls	.	B,R,O	.	B,R,O	.	B,R,O
R^2	0.64	0.69	0.36	0.44	0.46	0.58
Obs.	18,664	15,647	18,664	15,647	18,664	15,647

Standard errors in parentheses, clustered at the buyer level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The main regressor in all cases is the baseline, buyer-specific metric of relational sourcing and it is standardized. Outcomes are: the log weighted average price of fabric in the order, p^f_{sbjo} (columns (1) and (2)), the buy-to-ship ratio of the order, $(F/Q)_{sbjo}$ (columns (3) and (4)) and a measure of complexity of the garment order (the log of the number of fabric types used for producing the order), $Complex_{sbjo}$ (columns (5) and (6)). All columns feature seller-product-year fixed effects. In addition, even numbered columns also include buyer-, relationship- and order-level controls, s, as follows. Buyer controls (B): fixed effect for the main destination of the buyer, cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship controls (R): Cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order controls (O): size of order (log volume), log price of fabric of the order.

Table 5: Buyers' Sourcing and Labor Usage

	$SMV_{slb\tau}$		$Efficiency_{slb\tau}$		$\#Workers_{slb\tau}$		$Share\ Helpers_{slb\tau}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Relational_b$	-0.084	-0.024	-0.010*	-0.005	0.428	0.396	0.001	-0.002
	(0.352)	(0.317)	(0.006)	(0.005)	(0.534)	(0.335)	(0.002)	(0.001)
FEs	sm(τ),τ	sm(τ),sl,τ	sm(τ),τ	sm(τ),sl,τ	sm(τ),τ	sm(τ),sl,τ	sm(τ),τ	sm(τ),sl,τ
R^2	0.77	0.85	0.20	0.23	0.80	0.92	0.92	0.94
Obs.	155,723	155,713	116,905	116,896	125,940	125,932	125,940	125,932

Standard errors in parentheses, clustered at the level of the buyer and production line. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. Across all specifications, the regressor of interest is the metric on relational sourcing, standardized and increasing in the relational characteristic of the buyer. The outcome in columns (1) and (2) is the Standard Minutes Value (SMV), defined as the amount of time a particular garment is supposed to take to be sewed together computed by the factory's industrial engineers (often based on international libraries of SMVs of elemental sewing processes). Columns (3) and (4) study labor efficiency of a particular line in a plant, producing for a buyer on a given day, $Efficiency_{slb\tau}$. Labor efficiency is constructed as the ratio between the minutes-equivalent of the output and the minutes of labor input. In turn, the output is calculated as Standard Minute Values times the number of pieces and the input is calculated using the number of workers times the runtime. See main text for a comprehensive description. The outcome in columns (5) and (6) is the number of workers active on the line, $\#Workers_{slb\tau}$, and in columns (7) and (8) it is the share of such workers that are line helpers, $Share\ Helpers_{slb\tau}$. The discrepancies in sample size across columns are due to the fact that not all plants keep administrative records of all labor usage metrics studied here. All specifications include as controls for relevant buyer characteristics, its size as a garment importer in Bangladesh, whether the buyer is a signatory of the compliance Accord as of 2019 and the cohort of the buyer. Odd numbered columns condition on fixed effects corresponding to the seller-month ($sm(\tau)$) and the day (τ). Even numbered columns, in addition, include a fixed effect for the production line of the seller (sl).

Table 6: Buyers' Sourcing, Markups and Costs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	p_{sbjo}	mc_{sbjo}	μ_{sbjo}	μ_{sbjo}	μ_{sbjo}	μ_{sbjo}	μ_{sbjo}	μ_{sbjo}
$Relational_b$	0.021***	-0.005	0.026***	0.026***	0.026***	0.032***	0.031***	0.027**
	(0.008)	(0.007)	(0.007)	(0.007)	(0.007)	(0.010)	(0.007)	(0.012)
FEs	sjt,d	sjt,d	sjt,d	sjt,d	sjt,d	sjfot,dq	sjt,d	sjt,d
Controls	B,R,O	B,R,O	B,R,O	B,R,O	B,R,O	B,R,O	B,R,O	B,R,O
Robustness	.	.	.	Season	Product	Quality	Small b	Small s
R^2	0.73	0.62	0.41	0.41	0.41	0.59	0.41	0.44
Obs.	15,647	15,647	15,647	15,647	15,647	10,103	15,144	7,479

Standard errors in parentheses, clustered at the buyer level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcome in column (1) is the log price of an order between a seller and a buyer in a given product category, p_{sbjo}. The outcome in column (2) is the estimated log marginal cost of the order, mc_{sbjo}. In all other columns, the outcome is the log markup factor, μ_{sbjo}. The main regressor in all cases is the baseline, buyer-specific metric of relational sourcing and it is standardized. All columns but (6) include seller-product-year and destination fixed effects. As such, column (1) simply reproduces the results of column (4) in Table 3 and columns (2) and (3) use the same specification to study marginal costs and markups. All remaining columns report the results of different robustness exercises, for brevity, shown only on μ_{sbjo}. All columns in the table include buyer-, relationship- and order-level covariates, as follows. Buyer controls (B): fixed effect for the main destination of the buyer, cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship controls (R): Cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order controls (O): size of order (log volume), log price of fabric of the order. Columns (4) to (6) include rich sets of controls to condition on seasonality patterns, product specialization, and the physical quality of the order. These additional controls are as follows. *Season*: Herfindhal index describing how concentrated the trade in a relationship is in one season, the share of the largest season in the seller-buyer-year combination and an indicator that picks up orders in such season. *Product*: defined analogously to controls described for seasonality. *Quality*: measure of complexity of the garment order (the log of the number of fabric types used for producing the order, elsewhere labeled as $Complex_{sbjo}$) and a fixed effect for the seller-product-year-fabric-type-origin ($sjfot$), exploiting the type and origin of fabric to define the variety of the order; a category is as specific as *Nice Ltd.'s men's shirts made of wov. fab. containing > 85% cotton, printed, plain weave, weighing more than 100g/m2 but not more than 200g/m2 sourced from India in 2010*. Column (7) trims the sample to drop all the orders of the largest buyer of the seller-year-product. Column (8), analogously, drops all orders of the largest seller of the buyer in the product-year combination.

Figure 2: Share of Volumes Sold to Relational Buyers

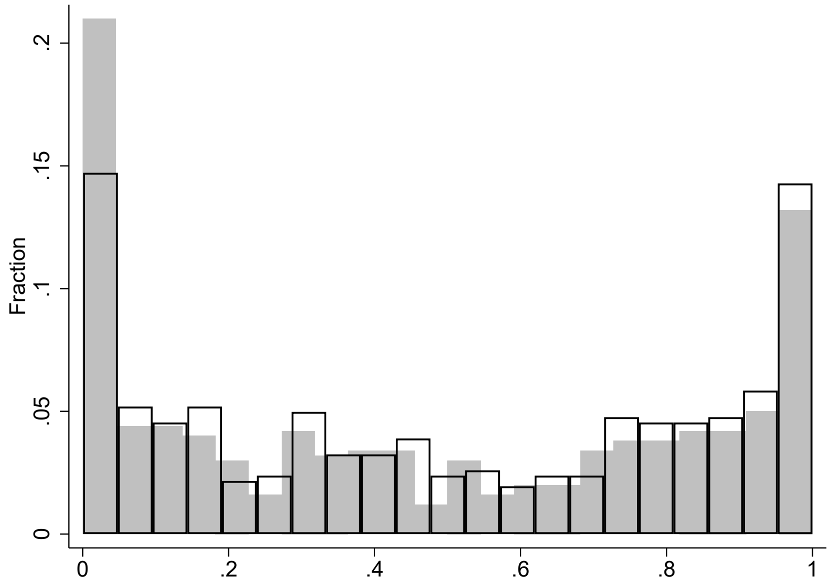

The histogram shows the fraction of the volume of woven garments that each seller sells to relational buyers. Relational buyers are defined as those located in the top 10% of the distribution of the relational sourcing metric. The grey bars represent data for the 500 sellers in our sample. Across them, the average (median) share of volumes sold to relational buyers is 0.47 (0.43). The bars with black contours represent data for the subset of 462 that sell at least some volume to relational buyers. On this subsample, the average (median) share is 0.50 (0.49). Further trimming to this subsample, restricts the histogram to the 442 sellers that trade both with relational and non-relational buyers. In that subsample, the average (median) is 0.48 (0.44). The histogram on this subset of observations is not reported in the graph for visual clarity.

Figure 3: Robustness of Price Result to Alternative Fixed Effects and Controls

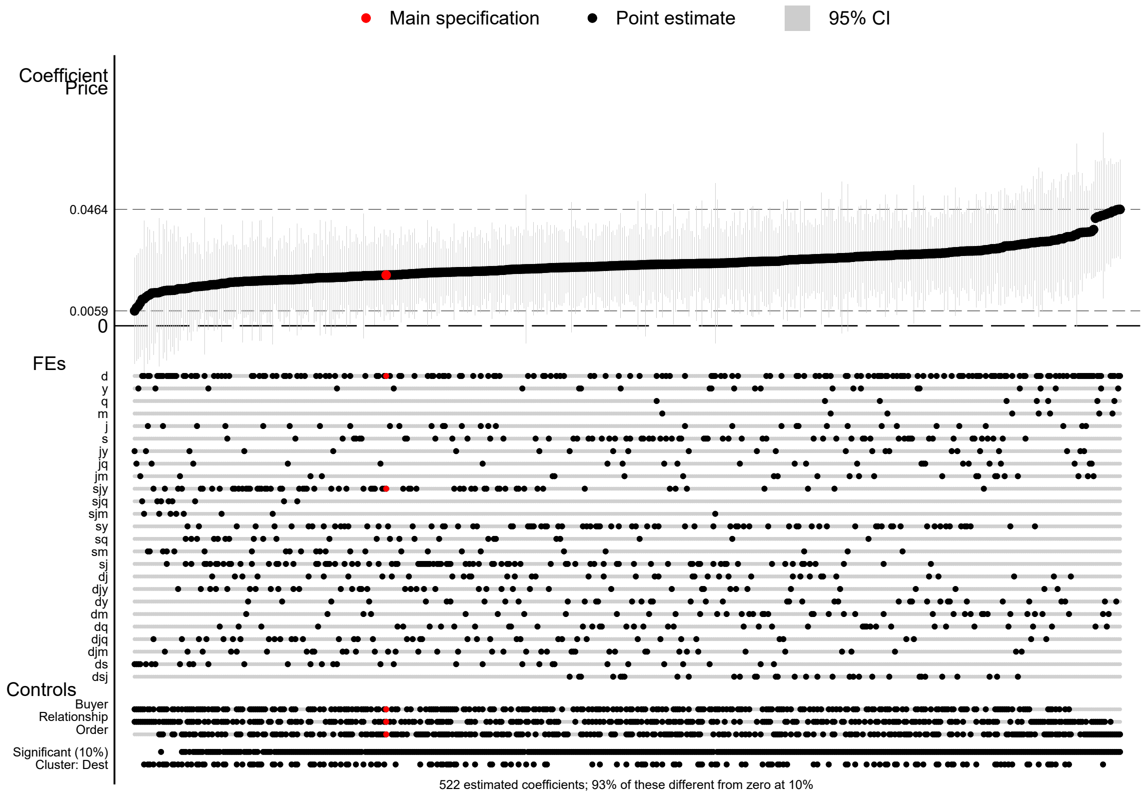

The graph presents 522 estimates of the coefficient on the buyer-specific relational metric in the regression of order prices following specifications with alternative controls and fixed effects. Our baseline, highlighted in red in the graph, includes seller-product-year fixed effects, destination fixed effects, and buyer-, relationship- and order-level controls. These controls are as follows. Buyer: cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship: cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order: size of order (log volume), log price of fabric of the order. The fixed effects are labeled following the notation of the paper: s for seller, j for product, y for year, d for destination, m for month, q for quarter. The scatter marks in black present the point estimates and the bars in grey show 95% confidence intervals. The bottom panel reflects the set of fixed effects and controls used for the corresponding estimation. For example, a point estimate that has a black marking in dy, sjq and $Buyer$ corresponds to a price regression on the relational metric, with destination-year and seller-product-quarter fixed effects, as well as buyer-level controls. All possible combinations of fixed effects and controls give an intractably large set of estimates to report. The specifications presented here exclude: (i) redundant combinations (for example, seller-product-quarter and year in the same specification), (ii) combinations with more than two sets of additive fixed effects and three multiplicative effects. We note that the number of observations may vary across specifications, as the change in the fixed structure gives rise to different singleton nests. The average specification runs on 17,453 orders. The largest sample runs on 21,577 orders and the smallest sample includes only 6,483 orders. This sample reduction is associated to the use of seller-product-month fixed effects (alongside different forms of destination fixed effects and controls). Standard errors are clustered at the level of the buyer in most specifications. In 148 cases, very granular fixed effects (namely, destination-month, destination-product-month, destination-quarter, destination-quarter-month) and buyer-level clustering give variance matrices that are close to singular. In these cases, the standard errors are clustered by destination (which is a conservative solution in general). This is indicated at the bottom of the figure. There are 36 out of the 522 point estimates (6.8% of all estimations) that are not significantly different from zero. Of these, two thirds correspond to computations of the standard errors clustering at the level of the destination. The remaining 14 (out of the 36) non-significant point estimates correspond to buyer-level clustering of the standard errors. The 36 specifications with coefficients not significantly different from zero (albeit positive), in general correspond to two types of specifications: (i) seller-product-month, seller-month or seller-product-quarter fixed effects, alongside destination fixed effects and no order-level controls; (ii) destination-seller fixed effects, alongside product-month, product-quarter or product-year fixed effects and no order-level controls.

Figure 4: Robustness of Markups Result to Alternative Fixed Effects and Controls

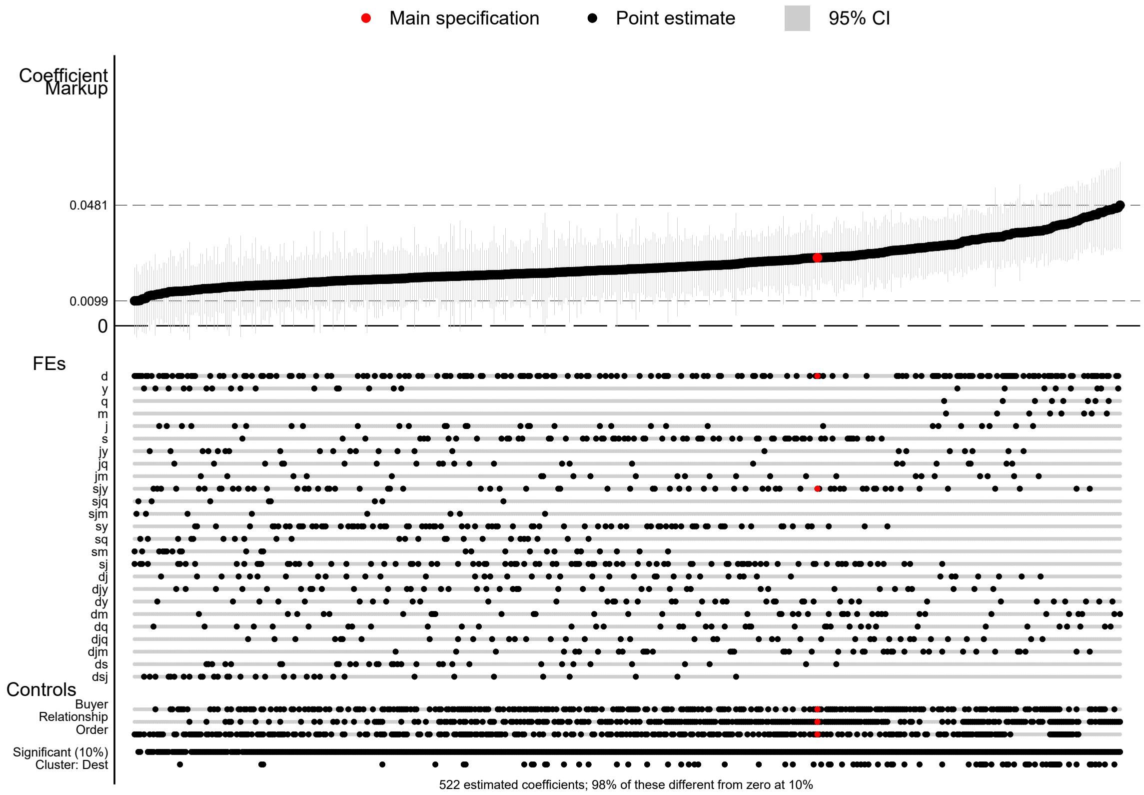

The graph presents 522 estimates of the coefficient on the buyer-specific relational metric in the regression of order-level markup factors following specifications with alternative controls and fixed effects. Our baseline, highlighted in red in the graph, includes seller-product-year fixed effects, destination fixed effects, and buyer-, relationship- and order-level controls. These controls are as follows. Buyer: cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship: cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order: size of order (log volume), log price of fabric of the order. The fixed effects are labeled following the notation of the paper: s for seller, j for product, y for year, d for destination, m for month, q for quarter. The scatter marks in black present the point estimates and the bars in grey show 95% confidence intervals. The bottom panel reflects the set of fixed effects and controls used for the corresponding estimation. For example, a point estimate that has a black marking in dy, sjq and $Buyer$ corresponds to a price regression on the relational metric, with destination-year and seller-product-quarter fixed effects, as well as buyer-level controls. All possible combinations of fixed effects and controls give an intractably large set of estimates to report. The specifications presented here exclude: (i) redundant combinations (for example, seller-product-quarter and year in the same specification), (ii) combinations with more than two sets of additive fixed effects and three multiplicative effects. We note that the number of observations may vary across specifications, as the change in the fixed structure gives rise to different singleton nests. The average specification runs on 17,453 orders. The largest sample runs on 21,577 orders and the smallest sample includes only 6,483 orders. This sample reduction is associated to the use of seller-product-month fixed effects (alongside different forms of destination fixed effects and controls). Standard errors are clustered at the level of the buyer in most specifications. In 75 cases, very granular fixed effects (namely, destination-month, destination-product-month, destination-quarter, destination-quarter-month) and buyer-level clustering give variance matrices that are close to singular. In these cases, the standard errors are clustered by destination (which is a conservative solution in general). This is indicated at the bottom of the figure. There are 11 out of the 522 point estimates (2.1% of all estimations) that are not significantly different from zero. The 11 specifications with coefficients not significantly different from zero (albeit positive), in general correspond to two types of specifications: (i) seller-product-month or seller-month fixed effects, alongside destination fixed effects; (ii) destination-seller-product fixed effects, alongside product-year or year fixed effects.

Figure 5: A Change in Sourcing Strategy

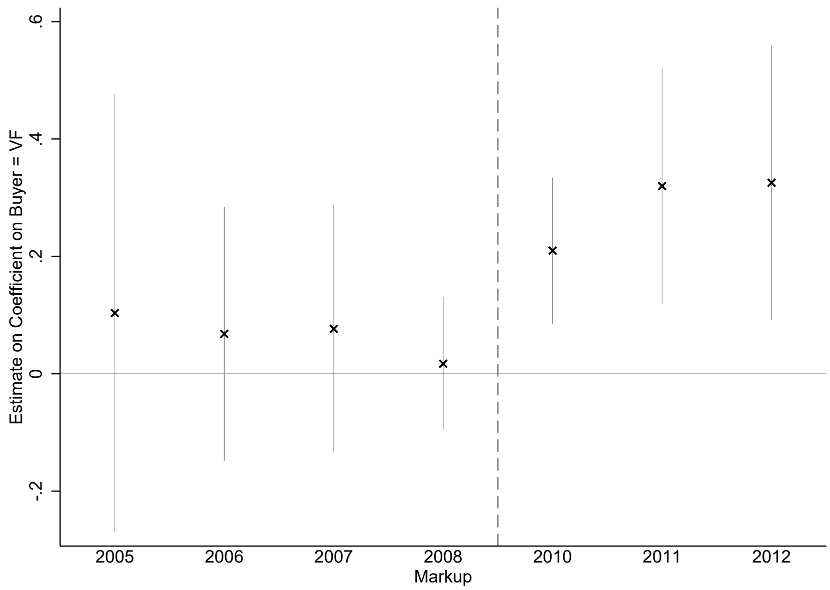

The figure plots estimated year-specific coefficients, β_r, on a dummy that takes value one when the buyer is VF, following specification (9). The excluded category corresponds to $VF \times I_{r=2009}$. We focus on export orders manufactured by sellers that traded at some point with VF. Among those, we consider the orders placed by VF or by another main buyer of the seller. A main buyer is either the largest buyer (in volumes) of the supplier over the entirety of the sample period, before 2010 or after 2010. The regression includes seller-product-year fixed effects. This controls already on the first difference (time) in order level markups. The specification also includes buyer fixed effects, which absorb all buyer level controls included elsewhere (see, for example, Table 3) and the first difference in markups, comparing buyers with VF. Finally, we include relationship- and order-level covariates, defined as follows. Relationship: cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order: size of order (log volume), log price of fabric of the order. The vertical bars correspond to 95% confidence intervals, when standard errors are clustered at the buyer-year level.

A Data Sources and Working Sample

A.1 Customs Records: Analysis Sample

Ready-made garment products can be classified into *woven* garments and *knitted* garments, in accordance with the interlacing of their fibers. Over the period 2005-2012, 45.8% of the garment exports from Bangladesh correspond to woven products, the focus of our study. We restrict attention to this subset of products to leverage the fact that there is virtually no domestic production of woven fabrics in Bangladesh, making this subsector completely dependent on imported fabric.[23] Most of the woven garment exports are concentrated in a few product codes. We do not consider small woven product categories, such as the ones corresponding to *parts of garments or clothing* (all codes in 6217), *brassieres, corsets, suspenders, etc.* (all codes in 6212), *track suits, ski suits and swimwear, etc.* (all codes in 6211), among other technical or specialized garments. The remaining product codes, all in four four-digit HS codes, account for 92% of all exported volumes in woven garments. In turn, 90% of the volumes in these main woven products correspond to trousers or shirts of some description. We focus our analysis on these two products types, which we label as *included* products, to differentiate them from *excluded* products, which are all other product categories, in either knit or woven. These cover 17 different HS codes at six digits of disaggregation.[24]

While some outcomes (such as prices and volumes) can be studied transaction by transaction, the analysis of fabric usage and the recovery of markups require that transactions are grouped into their correspondent *export orders*. This is possible only when trade happens through a recorded *Utilization Declaration* (UD) procedure. 80% of the volume exported in the products that we study are associated to a UD. Of this data (which corresponds to 56 thousand orders, approximately), we remove those orders whose quality of underlying data is low (missing values or extreme outliers), preventing a clean import-export matching exercise.[25]

To mitigate sparseness in the data for our analysis, our baseline sample only considers orders of the top 500 exporters (out of approximately 1,500), who jointly account for 78% of the exported volumes in the subsample. Our final sample of 22 thousand export orders accounts for approximately 45% of the Bangladesh's exports in the UD system, in the relevant product categories.

Table A1 compares the data across the different sample trimmings described here. For this comparison, we consider: (1) all shipments in woven exports throughout the entire sample period, (2) all shipments in the selected products categories, (3) all shipments in the selected product categories, that also have a UD associated to them, (4) all shipments in the

[23] As described in the main text, we use the products excluded from our analysis for the construction of measures of relational sourcing.

[24] The included HS codes are: 620341, 620342, 620343, 620349, 620461, 620462, 620463, 620469, 620510, 620520, 620530, 620590, 620610, 620620, 620630, 620640, 620690. Altogether, they account for 83% of all woven exports.

[25] Specifically, we exclude orders that either have missing values or outliers (lower than 3% and larger than 97%) in 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 59% of the exported woven volumes in the UD system, corresponding to 34 thousand export orders.

relevant product categories, with UDs with high-quality data, and (5) the analysis sample. For each buyer and each seller in the data, we compute the total traded volume (in kilos) and the (weighted) average price across all of their transactions. We report the distribution of volumes and prices for buyer and sellers across the five samples in the comparison exercise.

The comparison across samples reveals very intuitive patterns, given the nature of the trimmings described above. First, restricting attention to included products (comparing samples (1) and (2)) removes small product categories. In this case, the size distribution of buyers and sellers shifts upwards and the price distributions compress (among the excluded product, one finds cheap items like handkerchiefs, as well as expensive items per unit of weight such as corsets). Second, moving from all shipments in included products to shipments with a UD (comparing samples (2) and (3)), consequently removes all buyers and sellers with isolated shipments. This discards small firms on either side of the market, shifting the distribution of size upwards and further compressing the distribution of prices, particularly on the bottom tail. Reassuringly, studying the sample with UDs (column (3)), alongside the sample that discards orders with outliers or missing values in any variable that we will use for analysis (column (4)) does not significantly shift the distribution of prices for buyers or sellers. The size distribution of sellers remains similar, while that of buyers shifts upwards. Finally, our analysis sample, by definition, retains the largest sellers and, consequently selects on large buyers: the median buyer (10^{th} percentile seller) in the analysis sample is larger than a buyer (seller) in the 75^{th} percentile of the buyers' (sellers') size distribution in all woven. Buyers' prices do not appear significantly different, when comparing the analysis sample with the sample with UDs in included products (comparing samples (3) and (4)), while sellers' prices are on average one dollar higher.

As we show in Table C4, the trimming criteria we follow is extremely conservative and our main results on both prices and markups are larger in magnitude and significant when we reproduce the estimations in the less restrictive samples discussed above. This is the case when both transaction- and order-level outcomes are studied. We favor our conservative sample in the body of the paper, to ensure that the sample on all outcomes is the same and that our results stem from adequate variation in the data, rather than unbalancedness in the hierarchical panel, or the presence of outliers.

Table A1: Comparisons of Volumes and Prices Across Samples

Sample	(1) All woven	(2) Included	(3) With UDs	(4) Suitable UDs	(5) Analysis Sample
Share of volume in all woven	100%	78%	63%	40%	29%
Share of volume in Included	.	100%	80%	48%	37%
Share of volume in Incl. w. UDs	.	.	100%	59%	45%
Share of volume in Incl. w. suitable UDs	.	.	.	100%	78%
Count of shipments	1,369,609	1,085,535	815,926	418,998	353,580
Count of orders	.	.	76,366	34,307	22,741
Count of buyers	5,436	4,748	3,114	2,106	1,578
Count of sellers	5,898	4,705	2,340	1,528	500
Buyers' trade volumes (kg)					
Mean	1,403,435	1,586,402	2,319,611	3,054,738	3,881,589
p10	3,703	6,403	18,052	30,541	37,763
p25	24,426	34,785	72,026	108,447	148,728
p50	121,933	152,094	274,749	445,647	593,775
p75	508,114	604,100	1,123,805	1,672,577	2,364,826
p90	1,982,118	2,332,416	3,744,288	5,750,063	7,287,804
Buyers' average price (USD/kg)					
Mean	11.99	11.66	11.86	11.94	12.09
p10	4.22	5.19	7.04	7.87	8.13
p25	7.93	8.21	9.06	9.48	9.67
p50	10.80	10.87	11.32	11.50	11.65
p75	14.15	13.95	13.98	13.90	14.00
p90	18.53	17.88	17.10	16.49	16.51
Sellers' trade volumes (kg)					
Mean	1,252,599	1,499,374	2,352,346	2,176,370	4,565,052
p10	2,283	6,885	102,013	156,499	1,506,784
p25	19,500	51,388	372,517	418,006	2,013,343
p50	237,776	377,894	1,134,213	1,088,230	3,271,812
p75	1,177,703	1,525,823	2,754,710	2,552,031	5,451,742
p90	3,315,640	3,867,229	5,700,036	5,327,198	8,698,080
Sellers' average price (USD/kg)					
Mean	10.37	10.20	11.56	11.77	12.54
p10	2.78	3.60	7.75	8.13	9.42
p25	6.21	7.10	9.40	9.64	10.54
p50	9.74	10.01	11.16	11.26	12.08
p75	12.52	12.57	13.45	13.58	14.11
p90	15.51	15.13	15.64	15.89	16.56

The table compares five samples in the customs data over the period 2005-2012: (1) all shipments in woven exports throughout the entire sample period, (2) all shipments in the selected products categories, (3) all shipments in the selected product categories, that also have a UD associated to them, (4) all shipments in the relevant product categories, with UDs with high-quality data, and (5) the analysis sample. The top panel shows counts of shipments, orders, buyers and sellers in each sample, as well as the share over the exported volume over different denominators. The rest of the table shows the distribution of volumes and prices, for buyers and sellers in the different samples. The volumes are constructed by aggregating all volumes (in kilos) traded by the buyer or seller throughout the sample period. The prices (in USD per kilo) are computed as weighted averages across all the shipments of the buyer or the seller.

A.2 Production Lines and Labor Data

For the analysis of labor usage, this paper combines buyer-level data from the customs records (described in Online Appendix A.1) with production line records from factories in Bangladesh, as well as worker-level surveys. These data were collected as part of a series of RCTs led by a team of researchers who kindly lent their data to this study. The manuscripts of reference are Macchiavello and Woodruff (2014), Ashraf et al. (2015) and Macchiavello et al. (2020).

Table A2 lists the different datasets available to us employed in our analysis. The information presented here is organized in *Phase 1* and *Phase 2* which correspond to different RCTs performed by the research team. For the purpose of this paper, this distinction is not relevant, except when some data is only available for units in one or the other phase of the study. This is explicitly reported when necessary. We leverage three different types of data, which are described in turn.

Panel A of Table A2 reports the structure of the production line data. A unit of observation in this dataset is a factory-line-day triplet. There are almost 460 thousand such triplets in the data, distributed across 51 plants and 1,344 lines observed for an average of 341 days. The variables in these data contain information pertaining to the operators working on the line on a given day, as well as the efficiency of production.[26] For 46% of the entries, we also observe the buyer for whom the line is producing on that specific day. For the purposes of this paper, only the line-day combinations for which the buyer is observed are suitable for analysis. We are not aware of any systematic aspects of the data collection protocol driving the availability of this information. To our knowledge, the structure of production records varied widely across plants and the time of the month during which the data were requested. In Table A3 we show that there are no significant differences in most of our key outcomes, when we compare observations that have an identified buyer with those that do not. The factory-line-day triplets that record the buyer can be matched with the buyer identities in the customs records.[27] Through this matching, we can study the composition of labor and the efficiency in labor usage, vis-a-vis buyer-level characteristics, including how relational the buyer is, the volume they import, etc., as collected from the customs records.

Panel B of Table A2 describes a dataset compiled from the Human Resources records of eleven factories. The data include the pay (including overtime) of all workers employed at the plants on a monthly basis. There are a total of almost 37 thousand workers in the data, who were observed over multiple months (for an average of 10 months per factory). The records include both production line operators, as well as workers in non-line occupations (including managers, workers in cutting stations, workers on finishing, tagging, boxing, needle replacement, spot washers, quality control, ironing and folding etc.). Of the over 250 thousand worker-month combinations in the data, 140 thousand correspond to workers assigned to production lines, and the rest are workers upstream or downstream to sewing lines. For this reason, we are able to characterize the pay and overtime of workers when the plant is serving relational buyers. Specifically, we use the production data to construct the

[26] Not all factories and lines record all variables, so the size of the sample varies slightly across different outcomes.

[27] We note that we cannot compellingly match the plants in the RCTs with customs records, which identify sellers using tax identification codes. We discuss this issue in detail in Online Appendix C.

share of the time (days) that sewing lines in the factory are producing for relational buyers in a given month, $Relational_{sm}$.

Finally, Panel C of Table A2 reports the structure of data assembled from worker-level surveys. In total, we have access to 1,538 surveys. Of these, 1,035 (67%) report the production line they are assigned to. Table A4 shows that the workers in our analysis (those matched with production lines and, hence information on the buyers) are no different than those excluded, when compared on all the outcomes we analyze: demographics (gender, experience, education and ability) and pay (wage, bonuses and piece-rate pay). Since the production line on which the worker was assigned at the time of the survey is known to us, we can match the worker to the mix of relational and spot buyers for which the production line was producing, whenever this information is available. We thus create a variable, $Relational_{sl}$, measuring the share of days during which the sewing line of the worker produced for a relational buyer.

We present descriptive statistics on all the variables we analyze (both for the main text and appendices), in our sample, in Table A5.

Table A2: Data on Production Lines and Worker Surveys

	Phase 1	Phase 2	Total
Panel A: Production Data			
Study Period	Jan 2012 - Jul 2014	Jan 2013 - Mar 2015	-
Observations (factory-line-day)	270,725	188,180	458,905
Factories	34	17	51
Lines	811	533	1,344
Days (total)	938	820	1,758
Days per line (average)	334	353	341
Observations reporting buyer	137,321 (51%)	74,761 (40%)	212,082 (46%)
Panel B: Human Resources Data			
Observations (worker-month)	-	250,510	250,510
Observations assigned to lines	-	142,977	142,977
Observations matched to buyer	-	90,791	90,791
Observations matched to buyer-month	-	53,457	53,457
Factories	-	11	11
Workers	-	36,997	36,997
Lines	-	260	260
Lines matched to production	-	255	255
Months (total)	-	15	15
Months per factory (average)	-	10	10
Panel C: Worker Surveys			
Observations (worker)	708	830	1,538
Factories	26	24	50
Lines	129	223	352
Observations with buyer characteristics	566 (79%)	469 (56%)	1035 (67%)

The table describes the structure of three types of data describing labor usage in garments plants. The underlying data comes from a set of RCTs performed by Macchiavello and Woodruff (2014), Ashraf et al. (2015) and Macchiavello et al. (2020). Full documentation of the experiments and data collection efforts is provided in those papers.

Table A3: Observations with and without buyer characteristics - Production Line Data

	(1) $SMV_{slb\tau}$	(2) $Efficiency_{slb\tau}$	(3) $\#Workers_{slb\tau}$	(4) $Share\ Helpers_{slb\tau}$
In Sample	-0.601	0.048**	0.572	-0.011
	(0.695)	(0.018)	(1.669)	(0.009)
FEs	sm(τ),τ	sm(τ),τ	sm(τ),τ	sm(τ),τ
R^2	0.74	0.23	0.76	0.87
Obs.	236,618	156,290	171,327	171,327

Standard errors in parentheses, clustered at the level of the seller. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. Across all columns, the key regressor is an indicator that takes value one if the observation is in our sample of analysis (i.e. if it is matched to a buyer). The outcomes correspond to those analyzed in the paper. $SMV_{slb\tau}$ correspond to the Standard Minute Values in the line on day τ; $Efficiency_{slb\tau}$ the ratio between the minutes-equivalent of the output of the line and the minutes of labor input; the outcome in column (3) is the number of workers active on the line, $\#Workers_{slb\tau}$; and $Share\ Helpers_{slb\tau}$ is the share of such workers that are line helpers. The discrepancies in sample size across columns can be attributed to the the fact that not all plants keep administrative records of all labor usage metrics studied here. The regressions include fixed effects corresponding to the seller-month ($sm(\tau)$) and the day (τ), in line with the least restrictive specification in our analysis.

Table A4: Observations with and without buyer characteristics - Survey Data

Panel A: Outcomes in top panel of Table D7				
	(1) $Wage_{isl}$	(2) $Piece\ Rate_{isl}$	(3) $Quality_{isl}$	(4) $Other_{isl}$
In Sample	-0.038	0.024	-0.016	0.047
	(0.064)	(0.050)	(0.020)	(0.033)
R^2	0.00	0.00	0.00	0.01
Obs.	696	705	706	708

Panel B: Outcomes in bottom panel of Table D7				
	(1) $Female_{isl}$	(2) $Experience_{isl}$	(3) $Educated_{isl}$	(4) $Ability_{isl}$
In Sample	-0.027	-2.182	-0.019	-0.110
	(0.035)	(4.770)	(0.027)	(0.246)
R^2	0.00	0.00	0.00	0.00
Obs.	1,538	1,535	1,538	428

Standard errors in parentheses, clustered at the level of the seller. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table compares outcomes of regressions using survey data, for workers in the sample and those excluded from the regressions. As shown in Panel C of Table A2 a minority share of the observations (workers) in the survey data cannot be used for the purpose of the regressions presented in Table D7. The excluded workers correspond to cases in which a worker has not been assigned to a production line or when the production line does not have enough information about the buyers it is producing for. We only consider production lines for which there are at least 30 days of active production for which a buyer is reported. The data used for the regressions in the top panel of Table D7 is based on surveys of Phase 1 and excludes 20% of the observations on these grounds. The comparison of *In Sample* relative to excluded observations is presented in Panel A here. The data used for the regressions in the bottom panel of Table D7 is based on a combination of surveys from Phase 1 and Phase 2 and, in this case, 33% were discarded. The comparison of *In Sample* against excluded observations in these regressions is presented in Panel B here. In all cases, the outcomes defined at the level of worker i assigned to line l of seller (factory) s In the top panel the outcomes correspond to the log basic salary (*Wage*), and indicators for whether the worker reports being paid piece rate (*Piece Rate*), quality bonuses (*Quality*), or other bonuses (*Other*). In the bottom panel, the outcomes are the gender of the worker (*Gender*), whether they have completed secondary education (*Educated*), the months of experience in the garment industry (*Experience*), and the overall score of the worker's Raven Test (*Ability*). The Raven Test was only completed by supervisors and chiefs in Phase 1 of the study.

Table A5: Summary Statistics of Production, Workers and HR variables (In Sample)

	Obs.	Mean	Std. Dev.	P10	P25	P50	P75	P90
Panel A: Variables in analysis with production lines								
$SMV_{slb\tau}$	155,727	13.83	10.19	4.95	6.16	9.16	20.31	27.35
$Efficency_{slb\tau}$	116,905	0.54	0.34	0.28	0.40	0.52	0.66	0.82
$\#Workers_{slb\tau}$	125,940	47.24	23.62	23	28	38	67	83
$Share\ Helpers_{slb\tau}$	125,940	0.29	0.15	0	0.2	0.31	0.40	0.48
$Relational_b$	188,411	1.15	0.57	0.64	1.10	1.32	1.47	1.47
Panel B: Variables in analysis with worker surveys								
$Wage_i$	556	8,584	4,383	4,100	4684	8,217	11,500	14,540
$Piece\ Rate_i$='No'	563	0.97						
$Quality_i$='No'	564	0.96						
$Other_i$='No'	566	0.95						
$Female_i$='Yes'	1,035	0.33						
$Educated_i$='Yes'	1,035	0.47						
Age_i	1,032	27.59	5.63	21	24	27	30	35
$Experience_i$	1,033	93.27	52.87	36	57	84	120	168
$Ability_i$	345	2.42	2.15	0	1	2	4	5
$Relational_{sl}$	1,035	0.74	0.31	0.14	0.58	0.86	1	1
Panel C: Variables in analysis with HR data								
$Wage_{ism}$	195,672	6,531	2,752	5,300	5,510	6,420	6,670	6,949
$Overtime_{ism}$	195,699	31.34	21.73	0	8	38	48	52
$Absenteeism_{ism}$	195,699	0.76	2.6	0	0	0	0	2
$Relational_{sm}$	195,699	0.56	0.300	0.15	0.28	0.50	0.86	1

The table shows summary statistics for all the variables used in the analysis of data from production lines (Panel A), worker surveys (Panel B) and Human Resources records (Panel C). The summary statistics are computed over units of observation in the analysis sample (see Tables A3 and A4 for a comparison with excluded observations. The production line data (Panel A) is disaggregated at the level of the production line and day, thus observations are indexed by s (for seller), l (for line), b (for buyer) and τ (day). The variables are defined as follows. $SMV_{slb\tau}$ corresponds to the Standard Minute Values in the line on day τ; $Efficiency_{slb\tau}$ the ratio between the minutes-equivalent of the output of the line and the minutes of labor input; $\#Workers_{slb\tau}$ is the number of workers active on the line; $Share\ Helpers_{slb\tau}$ is the share of such workers that are line helpers; $Relational_b$ is a buyer level characteristic and it is available for 188 thousand observations. It is computed in the customs records following the baseline sourcing metric (ratio of sellers to shipments) and standardized. The survey data (Panel B) is disaggregated at the level of the individual i. By definition, an individual is assigned to a line l in seller s, so these indices are redundant. The variables of interest are the basic salary ($Wage$), and indicators for whether the worker reports being paid piece rate (*Piece Rate*), quality bonuses (*Quality*), or other bonuses (*Other*) and demographics including the gender of the worker (*Female*), whether they have completed secondary education (*Educated*), the months of experience in the garment industry (*Experience*), their age (*Age*), and the overall score of the worker's Raven Test (*Ability*), which is only available for the workers in one of the phases. $Relational_{sl}$ corresponds to the share of production (days) in which the line operates for buyers classified as relational (a buyer in the top 10% percent of the distribution of the sourcing variable). The observations in the HR data (Panel C) are specific to an individual i (working for seller s) in month m. The variables that we study are the wage that worker i is paid on month m by their employer, seller s, as reported in the HR records ($Wage_{ism}$), $Overtime_{ism}$ which is the hours of overtime recorded for the worker and $Absenteeism_{ism}$ which is the number of days the worker is absent during the month. For each factory-month combination, we compute the share of line-day pairs that are producing for a relational buyer. This is labeled as $Relational_{sm}$.

A.3 Global Sourcing Data

We combine customs records from six garment-exporting, developing countries: Bangladesh, Indonesia, India, Vietnam, Pakistan and Ethiopia. These countries account for 36% (58%) of all exports of garments –including (excluding) China– to the U.S. and Europe.[28] Each transaction has an identifier for the exporter (seller), and a name and address of the importer (buyer). We exploit string-matching routines, transaction-level imports data from the U.S. and manual procedures to homogenize buyers' denominations. In all of the countries we restrict our working sample to the years 2018 and 2019, to avoid overlap with the onset and early development of the Covid-19 pandemic. We drop any transactions in which a buyer name is available, but our diagnostics on the quality of the cleaning routine are pessimistic. This trimming drops 0.55% of the exported values across the data. While we continue to improve on the cleaning routines, we are confident this trimming is not inducing significant distortions in the data. We also drop buyers that have less than 100 shipments in the global data. The discarded observations account for less than 6% of the value exported in the data. Our working sample contains approximately 16.5 million transactions, across the six countries, corresponding to almost 10 thousand buyers and 29 thousand sellers. In Table A6 we present summary statistics for the global data.

Table A6: Summary Statistics in Global Data

	Obs.	Mean	Std. Dev.	P10	P25	P50	P75	P90
Panel A: All buyers								
$Count_b^c$	9,852	1.51	0.95	1	1	1	2	3
$Count_b^j$	9,852	27.0	23.5	6	12	21	34	54
$Count_b^{jc}$	9,852	30.0	34.3	6	12	21	36	59
$Count_b^s$	9,852	16.7	38.6	2	3	7	16	35
$Count_b^o$	9,852	1658.8	21209.2	118	157	284	697	1894
Panel B: Buyers active in multiple countries (68% of exported volumes)								
$Count_b^c$	2,915	2.73	0.98	2	2	2	3	4
$Count_b^j$	2,915	38.4	30.8	11	18	30	49	78
$Count_b^{jc}$	2,915	48.8	52.3	11	19	33	57	102
$Count_b^s$	2,915	33.8	63.4	5	8	16	34	71
$Count_b^o$	2,915	3860.2	38456.7	134	211	476	1444	4631
$Count_{bc}^j$	7,965	17.8	21.5	1	4	11	23	44
$Count_{bc}^s$	7,965	12.4	30.8	1	2	5	12	27
$Count_{bc}^o$	7,965	1412.7	16824.0	4	21	117	451	1660

Super- and sub-scripts are as follows: o corresponds to *shipments*, b to buyers, s to sellers, j to HS6 product categories, c to countries, with $c \in \{Bangladesh, India, Indonesia, Vietnam, Pakistan, Ethiopia\}$. $Count_y^x$ is the number of x per y. For example, $Count_b^{jc}$ is the number of product-country combinations within a buyer. The column under the heading 'Obs.' reports the count of *cells* relevant to the level of aggregation of the variable in the row. For example, the first row of Panel B, corresponding to $Count_b^c$ shows that there are 2,915 buyers that are active in multiple countries; across these, the average number of countries is 2.73.

[28] Calculated with data on export value flows in 2018-2019 from UN Comtrade, and supplemented with aggregated Customs Records where required.

References

Ashraf, Anik, Rocco Macchiavello, Atonu Rabbani and Christopher Woodruff, "The effect of political and labour unrest on productivity," Technical Report, *International Growth Center* Working Paper, 2015.

Macchiavello, Rocco and Christopher Woodruff, "Supervisory Management and Productivity Dispersion in the Bangladeshi Garment Sector," Technical Report 971, *Society for Economic Dynamics*, 2014.

Macchiavello, Rocco, Andreas Menzel, Atonu Rabbani, and Christopher Woodruff, "Challenges of Change: An Experiment Training Women to Manage in the Bangladeshi Garment Sector," *National Bureau of Economic Research*, 2020.

B Discussion on the Estimation of Markups and Costs

B.1 Modeling Garment Production

The production of woven garments takes place in two sequential stages: (i) inspection and cutting, and (ii) sewing and finishing. In the first stage, manufacturers inspect the fabric, plan fabric utilization, and then proceed to mark the fabric according to patterns, cutting, ticketing and bundling. In the second stage, cut fabric is sent to the sewing department, where the pieces are sewn on 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.

Many decisions in these two stages have a direct impact on fabric efficiency and relative labor, capital and fabric usage. For example, manufacturers may use fabric inspection machines to check for fabric and print defects and shading. The markers for cutting can be completed either by hand or through the use of 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 misalignment in the spread. Similarly, losses of fabric at the sewing stage can occur due to quality defects such as stains or incorrectly sewn garments. These losses can be reduced by introducing 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 at the finishing section when the garment is pressed or ironed, finished, and packed.

To guide our estimation framework, we further elaborate upon important characteristics of the garment production process and discuss how to model them. We first present evidence on existing dispersion in buy-to-ship ratios, and then show that labor and fabric are substitutes in production, at least to some extent.

Dispersion in buy-to-ship ratios. We describe further evidence from within-firm studies documenting the sources of variation in buy-to-cut and cut-to-ship ratios, which jointly determine buy-to-ship ratios.

The engineering study of Tanvir and Mahmood (2014) examines 30 Bangladeshi factories producing single jersey standard shirts. This study finds that fabric waste 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 Table B1).

Using data on reject rates and other defects from Macchiavello et al. (2020), we find that there is also variation in the sewing and finishing stage, namely in the cut-to-ship ratios. We examine a subsample of 6,000 line-day combinations in the daily production records of the 51 factories described in Online Appendix A.2. For these subsamples, we have information on rejection rates. The rejection rates, which occur at the final inspection point on the sewing line vary from 0% to 5% across these observations. This figure, however, is only a lower bound for the actual dispersion in cut-to-ship ratios. This can be explained by the fact that while rejections lead to a complete waste of the garment's fabric, there is an additional waste of fabric created by 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 some (but usually not all) of the sewing lines. Other factories, conversely, 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 the flexibility to substitute, to a certain extent, between fabric and other inputs. This can also contribute to the variation in buy-to-ship ratios 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. To assess whether this verifies empirically, we relate the amount of fabric imported at the order level, q^f_{sbjo}, to two exogenous sources of variation in input prices. First, we study the effects of changes in the international price of cotton, the most commonly used material in garment production in Bangladesh. Second, we consider the effects of a significant increase in the minimum wage in Bangladesh in November 2010. The specification is given by:

$$q^f_{sbjo} = \delta_{sj} + m(o) + \beta_1 Shock_{m(o)} + \beta_2 Relational^D_b + \beta_3 Shock_{m(o)} \times Relational^D_b + \beta_4 q_{sbjo} + \varepsilon_{sbjo}, \tag{B1}$$

which includes seller-product fixed effects, δ_{sj}, and a time trend (linear, quadratic and cubic) corresponding to the (calendar) month the order was made, $m(o)$. All specifications control for the size of the order, q_{sbjo}. The two exogenous input price shifters are captured by the term $Shock_{m(o)} \in \{p^{cotton}_{m(o)}; m(o) \geq Nov2010\}$. With this structure, we assess whether the use of fabric responds to exogenous shifts in input prices, according to expected substitution patterns. A second goal of our analysis if to explore whether these substitution patterns differ when the seller serves a relational buyer. To this end, we interact the shock variables with a dummy indicating whether the buyer is in the top decile of the distribution of the sourcing metric, $Relational^D_b$.[29] As an alternative to studying the volume of fabric conditional on the size of the order, we analyze the buy-to-ship ratio as the outcome of our regressions.

Column (1) in Table B2 shows that increases in cotton prices translate into lower import volumes of fabric to produce orders of a given size. Conversely, column (2) shows that a significant increase in the minimum wage (which resulted into significant increases in the wages of sewing operators) results in higher volumes of fabric being used to produce orders of a given size. Columns (3) and (4) show that the two patterns hold within the same specification and with linear, as well as higher order time trends. This evidence lends support to the hypothesis that fabric and labor can be substituted, to some extent, in response to changes in the price of inputs.

The degree to which fabric and labor can be substituted for one another does not appear

[29] Table **??** shows robustness of the interaction results to the use of different cutoffs of the relational dummy.

to differ across buyers who adopt different sourcing practices. This is illustrated in the results in columns (5), (6) and (7). Column (5) interacts the two price shocks in column (4) with the dummy for relational buyers. Both interaction terms render coefficients that are small and non-significant. The exercise is repeated in columns (5) and (6) on the sample of orders, adding the controls of our baseline specifications, to the effect of showing that results remain unchanged.[30]

In sum, we find that a model of garment production should accommodate three important characteristics: (1) a production process operating at the order level (see description of the UD system in Section 4.1.2), (2) variations in fabric efficiency across orders, and (3) a technology that allows for substitution across inputs, likely not different across buyers. The framework we propose in Section B.2 incorporates these characteristics into a technology that transforms material fabric inputs and labor into garments. We address (1) by specifying this production function at the order level. We address (2) by allowing for a productivity shock at the order level. Finally, we address (3) with a flexible specification in which the fabric enters production in a log additive separable manner, can be substituted for labor of different types, subject to capacity constraints accounting for the possibility of fixed or quasi-fixed factors of production.

These restrictions lead to a parsimonious cost minimization problem that allows us to recover *relative differences* in marginal costs and markups, across buyers. These differences are the objects of interest in the body of the paper and we devote Section B.2 to the presentation of our approach. In Section B.3 we derive conditions under which we can recover the *levels* of marginal costs and markups. To that end, we restrict the production function to be Cobb-Douglas for the purpose of estimating the elasticity of output to fabric, which is needed for the recovery of markups in levels. While there are different formulations that could be used to fit characteristics (1)-(3), the Cobb-Douglas appears appropriate for the purposes of the empirical context, as well as convenient for our estimation strategy. In a nutshell, and as we explain in detail in the sections to follow, to obtain the levels of markups and marginal costs, we first need to estimate output elasticities. The Cobb-Douglas production functional form assumes constant output elasticities for a given disaggregation level of the production function parameters. This allows us to complete our estimation even though we observe the use of fabric and not the usage of labor or capital, which would be necessary if allowing for a more flexible production function like the translog. It must be stressed that we require the elasticity of output to fabric only to compute the levels of markups. Our main results, which focus on exploring difference in markups across buyers within seller-product-time combinations, *do not* rely on the 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, for any production function.

[30]Note that in column (7) the coefficients on input prices are statistically insignificant. This is because the inclusion of seller-product-year effects absorbs nearly all the variation in the input price shocks. The purpose of the specification in column (7) is to study the interaction terms.

B.2 Framework: Markup Differences

We model trade between buyers indexed by b and sellers indexed by s. Sourcing and production over any period t are modeled as follows. First, buyers b and sellers s form links and sellers choose their production capacity. Second, each buyer's demand is realized and buyers place product orders. We impose no restrictions on the mechanism through which orders are allocated to sellers. Finally, each seller s produces the orders they 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 and a unit output price P_o.

Set Up. The production of woven garments is organized at the level of the order and comprises two sequential stages: (i) inspection and cutting, and (ii) sewing and finishing. In the first stage, the fabric is cut into pieces in preparation for the second labor intensive stage, in which the garments are sewn and finished. The inspection and cutting stage generates most of the variation in fabric waste. Conditional on the fabric fed onto the sewing lines, however, labor can also be used to reduce defects and fabric waste at the sewing and finishing stage.

To capture the two stages of production we assume an order-level production function that features log additive separability in the two main inputs, labor and fabric. Specifically, to produce an order o, a seller combines labor L_o^z of different types $z \in \{1, 2..., Z\}$ with fabric F_o. The different types of labor z capture the fact that orders are produced using workers of different skills, such as helpers, operators, supervisors and managers. We allow orders to vary in the way they combine the different types of labor and have idiosyncratic productivity ω_o. The production function can thus be written as:

$$Q_o = F_o^{\theta_o} H_o(\boldsymbol{L}_o, \omega_o) \tag{B2}$$

where θ_o is the output elasticity with respect to fabric and $\boldsymbol{L}_o = \{L_o^1, L_o^2, ..., L_o^Z\}$.

The seller may face capacity constraints in labor type z. Specifically, seller s chooses how much labor of type z to use in each order $o \in O_{st}$ subject to the capacity constraint:

$$\overline{L}_{st}^z = \sum_{o \in O_{st}} L_o^z. \tag{B3}$$

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

Seller s in period t chooses $\{\boldsymbol{L}_o, F_o\}_{o \in O_{st}}$ to minimize costs, subject to the technology constraint in (B2) and capacity constraint (B3), and taking order characteristics and prices as given. Denote the wages for labor of type z and the price of fabric with W_o^z and P_o^f respectively. We assume that fabric prices P_o^f do not depend on the size of the order.[31]

Cost Minimization. The Lagrangian for the seller's problem is

[31] We discuss the empirical validity of this assumption below.

$$\mathcal{L}_{st} = \sum_o \Big(\sum_z (W_o^z L_o^z) + P_o^f F_o\Big) + \sum_o \lambda_o\Big(Q_o - F_o^{\theta_o} H_o(\boldsymbol{L}_o, \omega_o)\Big) + \sum_z \lambda_{st}^z\Big(\overline{L}_{st}^z - \sum_o L_o^z\Big).$$

The Lagrange multipliers λ_{st}^z reflect the value of relaxing the capacity constraint for labor of type z. Having an extra unit of labor of type z to be allocated across orders allow the seller to reduce fabric input use and thus costs. Note that orders $o \in O_{st}$ are interrelated only via the capacity constraints, as captured by the Lagrange multipliers λ_{st}^z. Naturally, the analysis also applies if labor of type z can be adjusted freely (in which case the multiplier λ_{st}^z is equal to zero).

The order-specific first order condition with respect fabric F_o yields

$$F_o = \theta_o \frac{Q_o}{P_o^f}\lambda_o, \tag{B4}$$

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 short-run marginal cost for order o.

Knowledge of the marginal cost allows us to compute order-level markup factor M_o as the ratio between the order price P_o and the marginal cost λ_o:

$$M_o \equiv \frac{P_o}{\lambda_o} = \theta_o \frac{P_o Q_o}{P_o^f F_o}. \tag{B5}$$

Equation (B5) implies that the order-level markup M_o depends on the buy-to-ship ratio F_o/Q_o, the unit price of garment P_o and fabric P_o^f and the output fabric elasticity θ_o. The unique feature of our data is that F_o/Q_o, P_o and P_o^f are directly observed. The output fabric elasticity θ_o, however, is not. Denote $\alpha_o^{-1} = \frac{P_o Q_o}{P_o^f F_o}$ the term that is directly observed in the data. We can write the *difference* in (log) markups factors between two orders o and o' as:

$$\Delta_{oo'} \equiv \ln(M_o) - \ln(M_{o'}) = \underbrace{(\ln(\alpha_o^{-1}) - \ln(\alpha_{o'}^{-1}))}_{\text{Directly Observed in the Data}} - \underbrace{(\ln(\theta_o) - \ln(\theta_{o'}))}_{\text{Not Observed in the Data}}. \tag{B6}$$

The data thus allow us to directly observe differences in markups across orders that share the same fabric elasticity.

Taking the Model to the Data. Section 4 investigates differences in (log) markups factors across buyers with different sourcing strategies. The baseline specification allows for the output-to-fabric elasticity to vary at the seller-product-time level through the inclusion of the corresponding fixed effects. That is, in the empirical analysis, we assume $\theta_o = \theta_{sjt}$.

A fabric elasticity that varies at the seller-product-time level is more flexible than typically allowed for in the literature. A potential concern, however, is that within seller-product-time combinations, the fabric elasticity might also vary across orders produced for buyers adopting different sourcing practices. Two considerations assuage this concern. First, in Section B.1, we presented reduced form evidence that, conditional on seller-product-time fixed effects, substitution patterns between fabric and labor *do not* vary across buyers adopting different sourcing practices. Second, in Online Appendix B.3 we develop a framework to estimate the fabric elasticity. We find no evidence of it differing across buyers employing different sourcing strategies.

As is standard in the literature, our framework requires that fabric is flexibly chosen at the order level, taking its price as given. The assumption is consistent with the context of our analysis. As noted in Section 4.1.2, through the UD system, Bangladeshi garment exporters import, on an FOB basis, fabric to produce a specific order. This ensures that fabric is sourced flexibly for each order. Still, two potential concerns arise. First, the price of fabric could depend on the amount of fabric purchased - e.g., if the seller has market power over upstream suppliers of fabric or is able to negotiate discounts. Second, the price of fabric may be dependent upon the buyer for whom the order is produced. Even though buyers do not provide material inputs to their suppliers, they could influence its price - e.g., through relational sourcing with upstream suppliers of fabric.

Table B3 mitigates these concerns. First, the table shows that, conditional on seller-product-year fixed effects, the amount of fabric purchased in the order does not affect the fabric unit price. As expected, there is a negative correlation (all else equal, a garment manufacturer purchases more fabric when it is cheaper). Instrumenting for the amount of fabric purchased, however, we find no statistically significant relationship between the size of the fabric order and its price (see Online Appendix B.3 for details of the IV strategy). Second, if (relational) buyers negotiated price discounts with fabric manufacturers, we would expect that the price of fabric correlates with the buyer's sourcing practices and potentially with volumes traded between the buyer and the garment manufacturer in the past. We find no evidence for either of these hypotheses.

B.3 Econometric Approach: Markup Levels

In Section B.2 we develop a parsimonious model of garment production that allows us to recover deviations in markups, across orders, within seller-product-time combinations. We can directly map the components of these markup deviations to readily available information in our data. In this section we extend this framework with the purpose of recovering the level of markups and marginal costs in each order from our data.

Naturally, the estimation of orders' markups and marginal costs in levels requires an estimate of θ, the elasticity of output to fabric. This section derives a structural input demand equation that identifies the fabric elasticity. We assume a Cobb-Douglas production function with fabric (labor) output elasticity θ (β).[32]

We introduce two additional 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 $W_o = W_{sjt}$).[33] This significantly relaxes assumptions commonly made in the literature. Second, we require that the first order condition for the

[32] For the purpose of clarity, the rest of this section presents derivations for the case $Z = 1$, i.e. one type of labor z only. The extension to any number of Cobb-Douglas inputs is immediate.

[33] In the presence of this assumption, it continues to be possible to extend the model to multiple production factors, flexible or subject to capacity constraints, without altering the structural equation derived below. 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.

labor input, L_o, also holds exactly. This takes the form

$$L_o = \frac{\beta}{\widetilde{W}_{sjt}} Q_o \lambda_o, \tag{B7}$$

with $\widetilde{W}_{sjt} \equiv W_{sjt} + \lambda^L_{st}$.

We combine the Cobb-Douglas structure in production, the first order condition for fabric (see Section B.2) and (B7), and solve for the observable buy-to-ship ratio, F_o/Q_o. Taking logs, we obtain a structural equation that relates an order's buy-to-ship ratio to the order's size, the price of fabric used for producing the order, and two additional terms:

$$\ln \frac{F_o}{Q_o} = \frac{1-\beta-\theta}{\beta+\theta} \ln Q_o - \frac{\beta}{\beta+\theta} \ln P^f_o + \frac{\beta}{\beta+\theta} \ln\left(\frac{\theta \widetilde{W}_{sjt}}{\beta}\right) - \frac{1}{\beta+\theta}\omega_o. \tag{B8}$$

In principle, the framework allows for flexible production function parameters θ_o and β_o. In practice, in estimating (B8) we are constrained by the amount of variation in the data and we obtain more precise estimates when we restrict the elasticity to be common across all orders.[34] Exercises in which we allow for further disaggregation reveal nearly identical estimates.

The following relabeling is convenient: $\gamma_1 \equiv \frac{1-\beta-\theta}{\beta+\theta}$, $\gamma_2 \equiv -\frac{\beta}{\beta+\theta}$. The third term in (B8) reflects a seller-product-time-specific shifter, whenever the production function elasticities vary at most at that level of disaggregation, i.e. $\theta_o = \theta_{sjt}$ and $\beta_o = \beta_{sjt}, \forall o \in O_{sjt}$. Let this shifter be denoted $\delta_{sjt} \equiv -\gamma_2 \ln(\theta \widetilde{W}_{sjt}/\beta)$, 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 (B8) by means of the proposed notation yield the estimating equation:

$$\ln \frac{F_o}{Q_o} = \gamma_1 \; q_o + \gamma_2 \; p^f_o + \delta_{sjt} + \varepsilon_o, \tag{B9}$$

where lowercase letters denote logged variables.

The dependent variable in (B9) 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 can also be observed in our data; these are the order size q_o and the price of fabric p^f_o. Instead, the third explanatory term, δ_{sjt}, is not observable in our data. It 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^L_{st}, which varies at the seller-time level. We remind the reader that this multiplier is a sufficient statistic capturing the interdependence in input choices across orders arising from prices and capacity constraints. 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 control not only for the interdependence across orders but also for unobservable factors and productivity shocks that affect buy-to-

[34] Note that this assumption, i.e. $\theta_o = \theta \; \forall o$, is significantly stronger than the assumption $\theta_o = \theta_{sjt}$ $\forall o \in O_{sjt}$ used to recover differences in markups in the body of the paper, as explained in Online Appendix B.2.

ship ratios and are common across orders at the sjt level.[35] Finally, the fourth explanatory variable on the right-hand side of (B9) includes an order-specific productivity deviation, which is not observable.

Estimating equation (B9) allows us to construct our variables of interest in levels: specifically, from 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, $\widehat{\lambda}_o = P_o^f F_o/(\widehat{\theta} Q_o)$ and $\widehat{M}_o = P_o/\widehat{\lambda}_o$. We next discuss the approach that we use for estimating equation (B9).

B.4 Estimation of Elasticities

The recovery of elasticity θ by means of estimation of equation (B9) poses a number of challenges. Our baseline approach is an OLS estimation of a unique θ across all orders. We explore alternatives to this approach, to accommodate variations to our modeling assumptions that would lead to specification problems in equation (B9). We discuss each of these concerns before presenting the alternative estimation approaches.

First, since quantities q_o are obtained from customs records, measurement error is likely present in our data. In its classical form, measurement error would bias our estimate of $\gamma_1 \equiv \frac{1-\beta-\theta}{\beta+\theta}$ towards zero, thus yielding $\beta + \theta = 1$ even when the production technology does not exhibit constant returns to scale. Second, a similar concern applies to measurement error in the price of fabric p_o^f after which, other things equal, θ would be biased upwards towards $(1 + \gamma_1)^{-1}$. Third, we derived equation (B9) 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. In particular, misspecified productivity can overstate the scale coefficient and bias our estimates of θ upwards. Similarly, a fourth and related concern arises when the price of fabric is correlated with the error term. Such a scenario appears relevant in the presence of omitted inputs whose prices vary from order to order concomitantly with the cost of fabric or if bargaining power upstream is not fully captured by seller-product-time effects (e.g., if fabric prices are negotiated by the buyer).

We address these issues by performing a range of different estimation exercises. Across various specifications described momentarily, the estimate of the output-to-fabric elasticity θ is always around 0.6. All specifications also yield nearly constant returns to scale at the order level. The estimate of θ is thus remarkably consistent with industry reports and costing sheets, which show that fabric represents roughly two thirds of variable unit costs in garment production. We also find that the availability of detailed information on the heterogeneous input prices and varying allocation of fabric across orders is crucial for the recovery of θ: estimating equation (B9) ignoring these features of our data yields implausible large output fabric elasticities.

An IV strategy mitigates issues arising from measurement error and/or endogeneity of the order size with respect to unobservables governing the buy-to-ship ratio. The center of Panel

[35] Note also that the inclusion of sjt fixed effects in the estimating equation allows us to recover the relevant elasticity even in the presence of exporters' market power upstream as described in Morlacco (2019).

A in Table B4 presents the estimate of θ after instrumenting the size of the order in (B9) with volumes traded by third parties connected through the network of buyers and sellers. A full description of the construction of this instrument is included below. Diagnostics indicating a strong first stage are presented in column (1) of Table B5. The estimated elasticity is 0.615, very close to the point estimate under OLS, 0.623 (leftmost column in Panel A of Table B4). The similarity between the OLS and IV estimates suggests that productivity shocks that correlate with the buy-to-ship ratio and with the size of the order are well captured by the seller-product-time fixed effects. Additional order-specific productivity shocks (e.g., worker absenteeism due to hartals, power cuts, etc.) are plausibly *ex-post*, this is, revealed after the size of the order has been determined.

As described above, another set of concerns arises from the fact that there might be (unobservable) factors that correlate with both the price of fabric and with the buy-to-ship ratio, even conditional on seller-product-year fixed effects. Of particular interest is the possibility that garment buyers are able to exercise market power upstream. The evidence presented in Table B3 suggests that this mechanism is not supported by our data. Here we propose two approaches to address this issue directly. First, we instrument for fabric prices in the structural equation (B9). We pursue this strategy exploiting data on international prices of cotton in the countries from which the fabric is sourced. The exclusion restriction requires that the unobservables in equation (B9) are uncorrelated with shifts in the international price of cotton and with exchange rates (details on the construction of the instrument are included in the text below). The rightmost columns of Panel A in Table B4 present the estimate of θ after instrumenting for both the size of the order (as described above) and the price of fabric. The elasticity is slightly lower (0.544) than the one obtained via the OLS approach, but is accompanied by much higher standard errors, inherited from a borderline first stage (see column (3) of Table B5).

The inclusion of the sjt fixed effects, however, yields a weak first stage. In light of this, we also address the potential endogeneity of fabric prices to unobserved buyer-specific characteristics via a second approach. We augment equation (B9) to include buyer-specific fixed effects (see equation (B10) below). Panel B of Table B4 presents the elasticities obtained from estimating this augmented equation for the buy-to-ship ratio, by OLS (left side) and IV on quantities (right side). The elasticities obtained under these are 0.591 and 0.583 respectively (first stage diagnostics for the IV included in column (2) of Table B5). These estimates are very close to those obtained in estimations without buyer fixed effects. This assuages the concern that buyers of garments possess market power two tiers upstream or that choices of fabric at the order-level (e.g., with respect to fabric type) are influenced by the buyer in ways that correlate with order-level efficiency.

Our main interest is the study of dispersion in marginal costs and markups across orders sold to different buyers. The estimation approaches considered in this Appendix constrain the elasticity of output to materials to take a unique value across all orders, products, buyers and sellers. If the true parameter was not constant along all those dimensions, our estimation would understate the amount of dispersion in the level of marginal costs and markups by underestimating heterogeneity in technology. Given our focus, a relevant concern is that the elasticities vary with the sourcing strategy of the buyer. Panel C of Table B4 presents estimates of elasticities that are specific to whether the buyer is relational or not (i.e. is, spot). These follow the specification in equation (B11) described below, which is estimated by OLS and IV. The inclusion of fixed effects at the $sjt\bar{b}$, with $\bar{b} \in \{Relational, Spot\}$ depending on the sourcing strategy of the buyer, yields a weak first stage: including seller-product-time-sourcing fixed effects absorb most of the relevant identifying variation in the instrument.[36] The elasticities that we obtain are 0.592 for relational buyers and 0.638 for spot buyers in the OLS and, respectively, 0.590 and 0.618 in the IV. Very close to each other, we cannot reject the null hypothesis that the elasticity of output to fabric is the same across the two types of buyers. This result reinforces the earlier evidence suggesting that suppliers do not employ significantly different technologies when producing for buyers of different sourcing characteristics.

The following paragraphs develop in detail the alternative estimation approaches described here. All relevant results are presented in Tables B4 and B5.

Instrumenting for quantities. To assuage measurement error and endogeneity concerns in the regressor capturing quantities, we instrument for the size of the order, q_o. Our IV strategy leverages the observed network of trade partnerships. The key identifying assumption is that buyers cannot adjust their orders in response to shocks that are realized after orders have been allocated with sellers. Put differently, buyers take into account any information they have on the demand and seller-product-year characteristics when placing their orders, but they cannot respond to ex-post production shocks (for example, unexpected disruptions 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.

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', which we can label $q_{s'b'}$, as an instrument for q_{sb}. The argument for relevance 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_o is the

[36] As explained in this Appendix, the mechanics of the construction of the instrument are such that, while the potentially endogenous regressor, q_{sbjo} varies with each order, the instrument is only seller-buyer-time specific. To the extent that a seller might be trading with one buyer of each type $\bar{b}$ at a given time, the fixed effect absorbs the instrument.

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. Note that while the instrumented regressor q_o is an order-level variable, the instrument is constructed 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.

We note that the construction of this instrument necessitates a slightly more restrictive sample, relative to our analysis sample. 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 16,500 export orders.[37] Table B6 compares key shipment, buyer, seller, and relationship characteristics between the original sample and the two sub-samples described above.

Instrumenting for the price of fabric. To address measurement error and potential endogeneity of input prices in equation (B9), we instrument for the price of the fabric, p_o^f. To this end, we leverage the rich information we have on the dates of import shipments relevant to the order and the origins of each fabric shipment. Specifically, we use the international price of cotton in the month of the order, converted from dollars to the relevant currency using the exchange rate between the main country of origin of the fabric and the US dollar, in the corresponding month.[38]

Buyer Fixed effects. Equation (B9) includes a seller-product-time-specific term, capturing the market and shadow prices of inputs other than fabric. The baseline estimation absorbs this term, which is unobservable to us, in a fixed effect that removes variability in the error term at that level of aggregation. This device mitigates several concerns with the specification at hand. In particular, it allows for the unobservable productivity to be specific to a seller-product-time combination. This nests the standard assumptions the literature puts in place when estimating production functions in manufacturing, usually at higher levels of aggregation. The fixed effect also captures rich bargaining protocols upstream. For example, it allows for garment manufacturers negotiating prices with an upstream supplier of textiles for all the orders to be produced in, say, a product-year combination. A remaining concern specific to our context, however, is that the international buyer negotiates the fabric price directly with the foreign upstream supplier of fabric. To overcome this concern we estimate a version of (B9), including buyer fixed effects:

[37] 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.

[38] In practice, we also include an interaction between the price of cotton in local currency and an indicator that takes value one if the order uses fabric from a single origin. This allows for the slope of the international price of cotton in the first stage regression to differ for orders in which the main origin is the only relevant one, relative to orders sourcing from multiple origins and for which the currency conversion might be noisier.

$$\ln \frac{F_o}{Q_o} = \gamma_1\ q_o + \gamma_2\ p_o^f + \delta_{sjt} + \delta_b + \varepsilon_o, \tag{B10}$$

Buyer's Sourcing Strategy and Elasticities. In principle, our framework presented allows for flexible production function parameters θ and β. In practice, when estimating elasticities in (B9) we are constrained by the amount of variation in the data. In the baseline specification, we fix these elasticities to be common across all orders in the data. We introduce one relevant extension to this specification, following

$$\ln \frac{F_o}{Q_o} = \gamma_{1\bar{b}}\ q_o + \gamma_{2\bar{b}}\ p_o^f + \delta_{sjt\bar{b}} + \varepsilon_o, \tag{B11}$$

where the two coefficients of interest, γ_1 and γ_2, are allowed to vary at a disaggregation level of $\bar{b}$. Given the structural components collected in δ, allowing the elasticities to vary at level $\bar{b}$ requires that we introduce richer fixed effects of the form $\delta_{sjt\bar{b}}$. Specifically, $\bar{b}$ reflects the relational characteristic of the buyer of the order, i.e. $\bar{b} \in \{Relational_b, Spot_b\}$. As is the case in the rest of the paper, we define a buyer to be *Relational* if it falls in the top 10^{th} percentile of the distribution of the relational characteristic. All other buyers are defined as *Spot*, for the purpose of this discrete classification. The extension here, allowing for sourcing-specific elasticities is particularly relevant, given the focus of the analysis in the main text.

Naïve estimation with insufficient data. To illustrate the importance of the data on input utilization at the order-level, we conclude by estimating equation (B9) by OLS ignoring the available information. First, we ignore the information on order-specific input prices. In this exercise, we assume that all sellers pay the average price of fabric (across all orders and sellers) when producing a product (HS6 code) in a given month. This severely underestimates the responses of buy-to-ship ratios to changes in the price of fabric and, and thus, overstates the elasticity of output to materials, which is now estimated to be 0.99. In the second exercise, we ignore the information on the allocation of inputs to outputs. We assume that we observe the total volume of fabric purchased by the seller-buyer combination in a year, F_{sbt}, but not the amount of fabric assigned to each order, F_o. We split the volume in F_{sbt} across orders proportionally to the share of the order in the $-sbt$ combination, i.e. Q_o/Q_{sbt}. As expected, the loss of informative identifying variation produces coefficients on the order size and the price of fabric that approach zero; θ is biased towards one and estimated to be 0.86. When combined with cost shares, these elasticities would result in significantly larger estimates of markups.

B.5 Levels of Markups and Marginal Costs

Online Appendix B.3 estimates the *level* of markups for each export order. While the analysis in Section 4 *does not* require this estimation, the exercise allows us to benchmark our environment against other papers in the literature.

Across various specifications, the estimate of the output-to-fabric elasticity θ (our key outcome of interest) falls in the range $0.55 - 0.62$. All specifications also yield nearly constant returns to scale at the order level. These estimates are consistent with industry reports and costing sheets that show that fabric represents roughly two thirds of variable unit costs in garment production. Furthermore, in Online Appendix B.3 we find no statistically significant differences in the fabric elasticity in orders produced for relational buyers relative to spot buyers.

Table B7 presents our estimates of the order-level marginal costs and markups, $\widehat{\lambda}_o$ and $\widehat{M_o}$. The table shows that, on average, the price per kilo of garment paid by buyers is \$13.65. This average price is composed by \$3.30 of markup and \$10.35 of marginal cost, where the latter is in turn composed by \$7.57 of fabric and a reminder of labor and other costs. The implied average markup factor is 1.44. 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.[39] Table B7 shows that both markups and marginal costs are highly dispersed. We find that order-level markup values are more dispersed than order-level marginal costs: the interquartile ratio is 6.29 for markups and 1.80 for marginal costs.

The *within*-seller dispersion in markups (across buyers) is similar in magnitude to the dispersion *across* sellers. Appendix Figure B1 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 in ascending order in percentiles according to their average markup. Across the full range of sellers, the within-seller interquartile range is everywhere wide. Moreover, the average within-seller interquartile range in markups is of comparable magnitude to the interquartile range observed across sellers.

After taking product and time variation into account, the buyer rather than the destination appears to account for the sizable within-seller dispersion in markups. In an unreported exercise, we decompose seller-buyer-product-year markups into a seller-product-year component and either a buyer or a destination component. Buyer effects account for about 30% of the total variation in markups explained by the decomposition. The alternative specification, replacing the buyer fixed effects with country fixed effects, shows that destinations account for less than 5% of the total explained variation in markups.[40]

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 in the broader literature. To this end, we aggregate order-level outcomes at the seller-product-year level and find that at this level, (i) markups are more dispersed than marginal costs, as in Atkin et al. (2015); (ii) exported quantities are negatively correlated with marginal costs and positively correlated with markups,

[39]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, reported 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 B7.

[40]For concreteness, we estimate specifications of the form $\mu_{sbjy} = \delta_\iota + \delta_{sjy} + \varepsilon_{sbjy}$, where δ_ι with $\iota \in \{b, d\}$ are fixed effects for the buyer or the destination. The decomposition on buyers gives a share over total explained variation of 0.296, computed as 16.80%/(16.80% + 39.89%), where 16.80% corresponds to the variation explained by the buyer fixed effect and 39.89% that accounted for by seller-product-year effects. The decomposition on destinations gives 0.047, as a result of 1.93%/(1.93% + 41.83%).

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 (iii) 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).[41]

We close this discussion with a note on the importance of observing, rather than extrapolating, input usage and heterogeneous input prices, even when the relevant elasticities can be obtained without error. We conduct two exercises that explore the implications of forms of mismeasurement that are common in standard datasets. We consider the mismeasurement in the allocation of fabric across orders (i.e., pretend that such an allocation is not observed) and order-specific input prices (i.e., pretend that price of fabric is not observed at the order level). We find that each of these two types of measurement error would lead to significant over or under-estimations of order-level markups. Standard proportionality imputations of input allocations induce mismeasurements of at least 19% in half of the orders in our data and of 50% in 10% of the orders. These translate to under or over-estimation of order level markups of 20% and 53%, respectively. Similar errors follow the setting of common input prices across orders.

We first study the incidence of mismeasurement as a result of allocating inputs proportionally to outputs, based on the share of outputs on total sales or production. We do so by constructing a 'naive' weight ratio of fabric to garments for each order, $\widetilde{WR}_o \equiv \frac{\widetilde{F_o}}{Q_o}$. We leave the denominator as observed in the data (i.e. the true size of the order). We assume that we can observe the total amount of fabric (in kilos) bought by the seller in a given year (i.e. we observe F_{sy}). We impute this volume proportionally across garment orders, based on the share of each order in the year's output (Q_o/Q_{sy}). This gives the imputed fabric $\widetilde{F_o}$ that is used in the numerator of the naive weight ratio $\widetilde{WR}_o$. We label WR_o the true weight ratio, and study the mismeasurement induced by the imputations described above, using the absolute value of the difference between the naive and true weight ratios, as a share the true weight ratio – $|(\widetilde{WR}_o - WR_o)/WR_o|$. For the 22 thousand orders in our analysis sample, Appendix Figure B2 shows that half of the orders under or overstate the weight ratios by 19% or more, with the top 10% of the distribution inidicating mismeasurement of more than 50%.

Next, we study mismeasurement that can result from imputing common input prices across orders. To do so, we assume that the researcher observes the (weighted) average price of fabric paid in the industry when producing product j in year y. That is, for all orders of garment product j produced by any seller in year y we define a 'naive' input price $\widetilde{Pf}_o \equiv P^f_{jy}$. We study the mismeasurement induced by the imputation of common input prices, using the absolute value of the difference between the naive and true prices, as a share of the true price – $|(\widetilde{Pf}_o - P^f_o)/P^f_o|$. For the 22 thousand orders in our analysis sample, Appendix Figure B3 shows that half of the orders under or overstate the input prices by 17% or more, with the top 10% of the distribution featuring mismeasurement of more than 45%.

Both the weight ratio and the input price enter multiplicatively in the markup factor expression, so any mismeasurement in these transfers proportionally to markup factors, with a proportionality rate equal to the elasticity of fabric to output. For concreteness, with an

[41] These additional results are available upon request.

elasticity of 0.6 (as obtained in earlier subsections) and an output-to-input price ratio of 1.79 (the median in the data) the 19% mismeasurement in weight ratios translates to markups mismeasured by approximately 20%.

References

Atkin, David, Amit Khandelwal and Adam Osman, "Exporting and Firm Performance: Evidence from a Randomized Experiment," *The Quarterly Journal of Economics*, 2017, 132 (2), 551-615.

De Loecker, Jan, Pinelopi Goldberg, Amit Khandelwal, and Nina Pavcnik, "Prices, Markups, and Trade Reform," *Econometrica*, 2016, 84 (2), 445-510.

Macchiavello, Rocco, Andreas Menzel, Atonu Rabbani, and Christopher Woodruff, "Challenges of Change: An Experiment Training Women to Manage in the Bangladeshi Garment Sector," *National Bureau of Economic Research*, 2020.

Mayer, Thierry, Marc J Malitz and Gianmarco Ottaviano, "Market Size, Competition, and the Product Mix of Exporters," *American Economic Review*, 2014, 104 (2), 495-536.

Morlacco, Monica, "Market Power in Input Markets: Theory and Evidence from French Manufacturing," *mimeo*, 2019.

Tanvir, Saiful Islam and Tarique Mahmood, "Solid Waste for Knit Fabric: Quantification and Ratio Analysis," *Journal of Environmental and Earth Science*, 2014, *4* (12), 68–80.

Table B1: Fabric Input and Waste over Production Stages

Factory Number	Input Quantity (KG)	Inspection Loss (KG)	Cutting Loss (KG)	Sewing Loss (KG)	Finishing Loss (KG)	Total Waste (KG)	% of waste
	$A1$	(1)	(2)	(3)	(4)	$A2$	$(A1/A2) \times 100$
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	1,000	50	50	30	10	140	14
11	1,100	25	40	25	5	95	8.64
12	1,900	100	100	50	40	290	15.27
13	2,000	80	60	30	50	120	6
14	2,300	110	100	50	20	280	12.18
15	2,500	25	20	10	5	60	2.4
16	3,000	20	40	30	10	100	3.34
17	3,200	60	35	20	20	135	4.26
18	3,600	50	30	10	15	105	2.9
19	3,900	90	35	30	20	175	4.49
20	4,000	80	30	25	25	160	4
21	4,100	40	25	50	20	135	3.3
22	4,250	35	30	30	10	105	2.48
23	4,400	55	25	50	5	135	3.06
24	4,700	70	30	30	5	135	2.89
25	5,000	65	25	50	10	150	3
26	14,000	50	120	20	45	235	1.68
27	1,100	25	15	25	10	75	6.8
28	24,200	220	200	50	40	470	2
29	23,100	140	180	45	30	385	1.6
30	1,600	10	10	25	5	50	3.1
Total	136,930	1,585	1,325	930	540	4,240	

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

Table B2: Buyers' Sourcing and Input Substitution

Panel A: Fabric volumes as outcome

	(1)	(2)	(3)	(4) q^f_{sbjo}	(5)	(6)	(7)
$p^{cotton}_{m(o)}$	-0.020**		-0.074***	-0.060***	-0.059***	-0.072***	0.014
	(0.010)		(0.011)	(0.013)	(0.017)	(0.021)	(0.025)
$m(o) > Nov2010$		0.088***	0.117***	0.115***	0.127***	0.135***	-0.006
		(0.010)	(0.011)	(0.017)	(0.020)	(0.024)	(0.035)
$Relational^D_b\!=\!1 \times p^{cotton}_{m(o)}$					0.000	0.013	0.004
					(0.022)	(0.026)	(0.031)
$Relational^D_b\!=\!1 \times m(o) > Nov2010$					-0.020	-0.028	-0.029
					(0.020)	(0.023)	(0.031)
q_{sbjo}	0.946***	0.946***	0.946***	0.946***	0.944***	0.942***	0.941***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
FEs	sj	sj	sj	sj	sj	sj	sjt,d
Time Trend	Linear	Linear	Linear	Cubic	Cubic	Cubic	Cubic
R^2	0.95	0.95	0.95	0.95	0.95	0.95	0.96
Obs.	21,986	21,986	21,986	21,986	20,841	15,647	15,595

Panel B: Weight ratio as outcome

	(1)	(2)	(3)	(4) $(F/Q)_b$	(5)	(6)	(7)
$p^{cotton}_{m(o)}$	-0.028***		-0.070***	-0.056***	-0.053***	-0.058***	0.026
	(0.009)		(0.010)	(0.011)	(0.014)	(0.017)	(0.017)
$m(o) > Nov2010$		0.064***	0.091***	0.081***	0.084***	0.083***	-0.017
		(0.008)	(0.009)	(0.013)	(0.016)	(0.019)	(0.018)
$Relational^D_b\!=\!1 \times p^{cotton}_{m(o)}$					-0.002	0.006	0.001
					(0.019)	(0.023)	(0.023)
$Relational^D_b\!=\!1 \times m(o) > Nov2010$					-0.005	-0.007	-0.008
					(0.017)	(0.020)	(0.020)
FEs	sj	sj	sj	sj	sj	sj	sjt,d
Time Trend	Linear	Linear	Linear	Cubic	Cubic	Cubic	Cubic
R^2	0.20	0.21	0.21	0.21	0.21	0.24	0.39
Obs.	21,986	21,986	21,986	21,986	20,841	15,647	15,595

Standard errors in parentheses, clustered at the seller-product level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. All specifications in Panel A have the log of the quantity of fabric used in the order, q^f_{sbjo}, as the outcome. In Panel B, the outcome is the buy-to-ship ration F/Q of the order. Specifications in the top and bottom panel are identical, except for the control for the size of the order, in log kilos of garment, q_{sbjo}, which is only included in Panel A (in levels, this is the denominator of the outcome variable in Panel B). Columns (1) to (6) include seller-product fixed effects (sj) and column (7) uses the baseline seller-product-year and destination effects (sjt, d). In addition, columns (1) to (3) include a linear time (month) trend, and columns (4) to (7) have linear, quadratic and cubic time trends. $p^{cotton}_{m(o)}$ is the log of the international price of cotton in the first month of the order, $m(o)$. $m(o) > Nov2010$ 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. The richer fixed effects in column (7) restrict the size of the sample. Column (6) reproduces the specification of column (5), in the restricted sample of column (7). $Relational^D_b$ is a dummy that takes value one if the buyer is in the top 10% of the distribution of the relational souring metric.

Table B3: Price of Fabric, Relationship Dynamics and Fabric Quantities

Panel A: Price and Quantity of Fabric

	(1) p^f_{sbjo}	(2) q^f_{sbjo}	(3) p^f_{sbjo}	(4) p^f_{sbjo}	(5) q^f_{sbjo}	(6) p^f_{sbjo}
q^f_{sbjo}	-0.050***		0.020	-0.039***		0.049
	(0.004)		(0.030)	(0.002)		(0.031)
$z_{sb\tau}$		0.109***			0.092***	
		(0.016)			(0.012)	
FEs	sjt	sjt	sjt	sjt,fo	sjt,fo	sjt,fo
Fabric	Single	Single	Single	All	All	All
Specification	OLS	First Stage	IV	OLS	First Stage	IV
KP F-Stat	.	44.188	.	.	57.880	.
R^2	0.66	0.53		0.71	0.55	
Obs.	6,754	6,754	6,754	16,209	16,209	16,209

Panel B: Price of Fabric and Relationship Dynamics

	(1)	(2)	(3)	(4)
		p^f_{sbjo}		
$Relational_b$	0.010			
	(0.007)			
$Past\ Trade_{sbo}$	0.003	0.002	0.006	0.005
	(0.003)	(0.003)	(0.005)	(0.004)
$Relational_b \times Past\ Trade_{sbo}$	-0.001	0.002	0.000	
	(0.002)	(0.003)	(0.004)	
$Relational^D_b{=}1 \times Past\ Trade_{sbo}$				-0.002
				(0.004)
FEs	sjt,d	sjt,sb	sjt,sb	sjt,sb
Controls	B,R,O	B,R,O	B,R,O	B,R,O
Relationships	All	All	Main	All
R^2	0.69	0.78	0.79	0.78
Obs.	18,261	16,002	7,373	16,002

Standard errors in parentheses, clustered at the buyer level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. Panel A presents OLS and IV estimations of regressions of the price of fabric, p^f_{sbjo}, on the quantity of fabric used in the order, q^f_{sbjo}, both in logs. Columns (1), (2) and (3) use a sample of orders that use a unique fabric type (a unique HS code), as the specification in this trimmed sample more closely resembles an inverse-demand relationship. The specifications in these first three columns include seller-product-year fixed effects. Columns (4), (5) and (6) use all orders (including multi-fabric orders) and augment the specification to include fabric-origin fixed effects. The specifications correspond to the OLS (columns (1) and (4)), the first stage regression of q^f_{sbjo} on the excluded instrument $z_{sb\tau}$ (columns (2) and (5)) and the second stage in the 2SLS (columns (3) and (6)). Please, see Section B.3 for details on the construction of the instrument. For columns (2) and (5) the Kleibergen-Paap F statistic is reported. In Panel B, the outcome of all specifications is the log price of the fabric used in the order, p^f_{sbjo}. Columns (1), (2) and (4) correspond to regressions in the entire sample, while column (3) performs robustness of column (2) to restricting the sample to orders the buyer places with the seller with the highest share of its imports in the product-year combination, i.e. its *main* partner. The key regressors in all specifications are: the baseline, buyer-specific metric of relational sourcing and it is standardized, $Relational_b$; the experience in the relationship measured as the log cumulative traded volumes until the date of the order, $Past\ Trade_{sbo}$; the interaction between the two. All columns include buyer-, relationship- and order-level controls, as described in the notes of Table 3. Column (1) includes the baseline seller-product-year and destination fixed effects, therefore exploiting variation across buyers. Columns (2), (3) and (4) include seller-product-year and seller-buyer effects (so the coefficient on the buyer-level variable $Relational_b$ is not identified). Column (4) simply reproduces the exercise in columns (2), replacing the continuous relational metric with a dummy variable indicating the buyers in the top 10% of the relational characteristic, $Relational^D_b$. These specifications using the dummy are included for ease of interpretation and to match our practice in other tables that feature interactions in this paper.

Table B4: Elasticities and Returns to Scale

Panel A: OLS and IV, sjt fixed effects

	OLS		IV: Quantities		IV: Quantities, Fab. Price	
	Coeff	SE	Coeff	SE	Coeff	SE
Materials: θ	0.623	0.016	0.615	0.016	0.544	0.28
Labor: β	0.445	0.016	0.343	0.026	0.453	0.273
RTS: $\theta + \beta$	1.068	0.003	0.958	0.025	0.998	0.013

Panel B: OLS and IV, sjt, b fixed effects

	OLS		IV: Quantities	
	Coeff	SE	Coeff	SE
Materials: θ	0.591	0.017	0.583	0.017
Labor: β	0.48	0.017	0.398	0.073
RTS: $\theta + \beta$	1.071	0.004	0.981	0.076

Panel C: OLS and IV, $sjt\bar{b}$ fixed effects, sourcing-specific elasticities ($\bar{b} \in \{R, S\}$)

	OLS		IV: Quantities	
	Coeff	SE	Coeff	SE
Materials: θ^R	0.592	0.028	0.59	0.028
Labor: β^R	0.471	0.029	0.301	0.029
RTS: $\theta^R + \beta^R$	1.064	0.005	0.891	0.005
Materials: θ^S	0.638	0.023	0.618	0.023
Labor: β^S	0.438	0.022	0.341	0.022
RTS: $\theta^S + \beta^S$	1.076	0.005	0.961	0.005
Test $\theta^R = \theta^S$ (χ^2)	1.70; pval: 0.192		0.73; pval: 0.392	

Panel D: Naïve Estimations, OLS, , sjt fixed effects

	OLS: Naïve Allocations		OLS: Naïve Prices	
	Coeff	SE	Coeff	SE
Materials: θ	0.865	0.011	0.999	0.024
Labor: β	0.164	0.011	0.058	0.023
RTS: $\theta + \beta$	1.029	0.002	1.057	0.003

The table reports detailed results of the main estimation strategies used for computing the elasticities of output to materials and labor, θ and β, respectively, on the sample of 16,500 garment orders. Panel A shows the elasticities resulting from the estimation of equation (B9) using our data. The underlying specification includes seller-product-time fixed effects. The leftmost panel performs the estimation using OLS. The central block reports the results of the IV strategy when only the size of the order, q_{sbjo}, is instrumented for. The rightmost block presents results from the IV strategy when both quantities and the price of fabric, p^f_{sbjo} are instrumented for. The first stages of all IV (2SLS) procedures, in all panels of this table, are reported in Table B5. Panel B shows the elasticities using the augmented specification in equation (B10), which to (B9) adds buyer-specific fixed effects. The estimation is again performed via OLS (left) and IV instrumenting the size of the order (right). Panel C presents elasticities that are specific to the sourcing strategy of the buyer, obtained via the OLS and IV estimation of equation (B11). At the bottom of this panel we include the test statistic for the null hypothesis that the elasticity of output to fabric is no different across buyers with the different sourcing strategies. Panel D presents the elasticities obtained via the OLS estimation of equation (B9), with data that we artificially restrict to mimic limitations present in commonly available datasets: the unobservability of the allocation of inputs to output and the setting of input prices to be common to all orders and manufacturers. We present the results on this under the headings of "OLS: Naïve Allocations" and "OLS: Naïve Prices". Please refer to the text in Online Appendix B.3 for further details. The standard errors in all panels are bootstrapped drawing, with replacement, the entire vector of export orders for each seller (in all products and time periods).

Table B5: First Stage Regressions and Diagnostics

Panel A: Unique Elasticity

	(1)		(2)		(3)			
Instrumentation:	Quantities Only		Quantities Only		Quantities and Fabric Price			
Equation:	q_{sbjo}		q_{sbjo}		q_{sbjo}		p^f_{sbjo}	
	Coeff	SE	Coeff	SE	Coeff	SE	Coeff	SE
$z_{sb\tau}$	0.095	0.010	0.055	0.012	0.085	0.010	0.004	0.002
p^f_{sbjo}	-0.568	0.051	-0.509	0.049				
p^c_{sbjo}					0.125	0.014	0.012	0.002
$p^c_{sbjo} \times \mathbf{1}\{\#fabric = 1\}$					-0.223	0.013	-0.001	0.002
Fixed effects	sjt		sjt,b		sjt			
First Stage K-P (F weak)	113.55		23.16		9.04			
First Stage K-P (LM underid)	103.77		24.06		26.87			

Panel B: Sourcing-specific Elasticities

	(1)			
Instrumentation:	Quantities Only			
Equation:	q_{sbjo}		$q_{sbjo} \times Relational^D_b$	
	Coeff	SE	Coeff	SE
$z_{sb\tau}$	0.091	0.014	0.000	0.000
$z_{sb\tau} \times Relational^D_b$	-0.002	0.030	0.089	0.027
p^f_{sbjo}	-0.472	0.073	0.000	0.000
$p^f_{sbjo} \times Relational^D_b$	-0.108	0.107	-0.581	0.080
Fixed effects	sjt$\bar{b}$			
First Stage K-P (F weak)	6.27			
First Stage K-P (LM underid)	13.00			

The table reports results of the first stage estimations corresponding to the IV strategies used for recovering elasticities, as reported in Panels A, B and C of Table B4. In all cases, the first stage equations include seller-product-year fixed effects. The specifications whose header read 'Quantities only' treat the price of fabric, p^f_{sbjo}, as exogenous and the size of the order, q_{sbjo}, as endogenous; those reading 'Quantities and Price of Fabric' instrument both variables. The instruments are the competitors' trade variable constructed using the trade network instrument described in the text of Online Appendix B.3, $z_{sb\tau}$ and the international price of cotton in the month of the order, p^c_{sbjo}, as the exogenous shifter (see text for further details). The coefficient on this instrument is allowed to vary when the order uses only one type of fabric. Panel A reports the first stages corresponding to the exercises that recover a unique elasticity for all orders. In this panel, column (1) is the IV estimation instrumenting for quantities only, column (2) augments the specification to include buyer-specific effects and column (3) corresponds to the IV of both quantities and fabric prices. Panel B reports the first stages corresponding to the exercises that recover elasticities specific to the souring strategy of the buyer. $Relational^D_b$ corresponds to a dummy taking value one if the buyer is in the top 10^{th} percentile of the distribution of the sourcing characteristic. In the specification of the fixed effects, $\bar{b} \in \{Relational^D_b = 1, Relational^D_b = 0\}$ such that $-sjt\bar{b}$ correspond to seller-product-time-sourcing fixed effects. For each estimation in this table we report test statistics for underidentification (LM) and weak instruments (F), allowing for clustering of the standard errors. The LM test corresponds to the Kleibergen-Paap rank test and in all cases all exogenous regressors, including the seller-product-year (and, when suitable, buyer) fixed effects) are partialled out (χ^2-distributed). The standard errors reported in the table are bootstrapped drawing, with replacement, the entire vector of export orders for each seller.

Table B6: Sample Comparisons

Panel A: Average Shipment Characteristics			
Shipments:	Count	Price (USD/kg)	Size (tonnes)
Under UD System	613,826	16.99	2.70
Outside UD System	5,181	15.35	1.76

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	228.96	13.71	504.83	20.09	75.40
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 in relevant products of 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. Volumes are constructed as averages of yearly traded volume.

Table B7: Order-Specific Markups and Marginal Costs

	But-to-Ship Ratio (Kg/Kg)	Price Garment (USD/Kg)	Price Fabric (USD/Kg)	Marginal Cost (USD/Kg)	Markup Factor (Units of Mc)	Markup Value (USD/Kg)
	(1)	(2)	(3)	(4)	(5)	(6)
Mean	0.87	13.65	7.57	10.35	1.44	3.30
Median	0.86	13.06	7.25	9.52	1.31	2.94
10^{th} Percentile	0.51	8.62	4.64	5.55	0.95	-0.64
25^{th} Percentile	0.67	10.43	5.64	7.13	1.08	0.86
75^{th} Percentile	1.04	16.32	9.15	12.83	1.67	5.38
90^{th} Percentile	1.22	19.77	11.03	16.36	2.14	7.80
St. Deviation	0.29	4.21	2.41	4.30	0.47	3.32
Coeff. Variation	0.33	0.31	0.32	0.42	0.33	1.01
$90^{th}/10^{th}$ Ratio	2.39	2.29	2.38	2.95	2.25	-12.24
$75^{th}/25^{th}$ Ratio	1.56	1.57	1.62	1.80	1.55	6.29
Number of orders			22,741			

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 body of the text and presented in Table B4, Panel A, unique θ estimated by OLS. The markup factor is defined as Price/Marginal Cost while the markup value is (Markup Factor - 1) $\times$ Marginal Cost.

Figure B1: Dispersion in Markups across Buyers

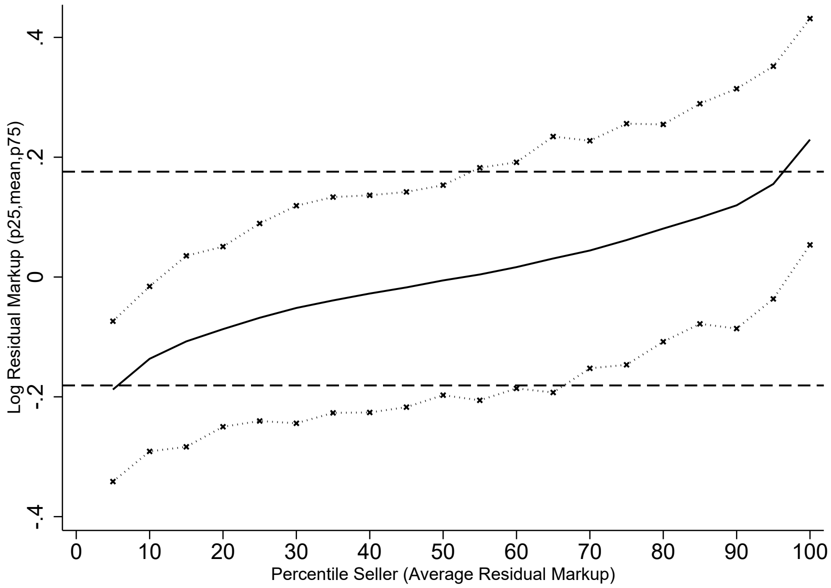

We aggregate order-level log markup factors for each seller-buyer-product-year combination, as weighted averages, where the weights are given by order volumes. We residualize these against product-year fixed effects. For each seller, we construct the simple average, 25^{th} and 75^{th} percentile markup 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 residualized 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 of residualized markups across sellers, centered around the average residualized markup of the median seller.

Figure B2: Comparison of Weight Ratios with Imputed Input Allocations

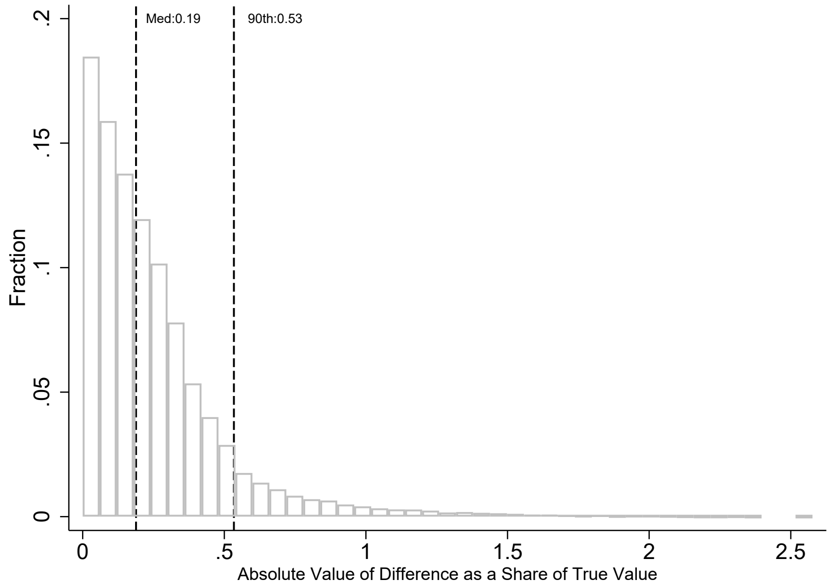

The figure plots the incidence of mismeasurement as a result of allocating inputs proportionally to outputs, based on the share of outputs on total sales or production. A'naive' weight ratio of fabric to garments for each order is constructed, $\widetilde{WR}_o \equiv \frac{\widetilde{F}_o}{Q_o}$. We leave the denominator as observed in the data (i.e. the true size of the order). We assume that we observe the total amount of fabric (in kilos) bought by the seller in a given year (i.e. we observe F_{st}). We impute this volume proportionally across garment orders, based on the share of each order in the year's output (Q_o/Q_{st}). This gives the imputed fabric $\widetilde{F_o}$ that is used in the numerator of the naive weight ratio $\widetilde{WR}_o$. We label WR_o the true weight ratio, and study the mis-measurement induced by the imputations described above, using the absolute value of the difference between the naive and true weight ratios, as a share the true weight ratio – $|(\widetilde{WR}_o - WR_o)/WR_o|$. For the 22 thousand orders in our analysis sample, the histogram shows that half of the orders under or overstate the weight ratios by 19% or more, with the top 10% of the distribution featuring mis-measurement of more than 50%.

Figure B3: Comparison of Fabric Price with Imputed Common Price

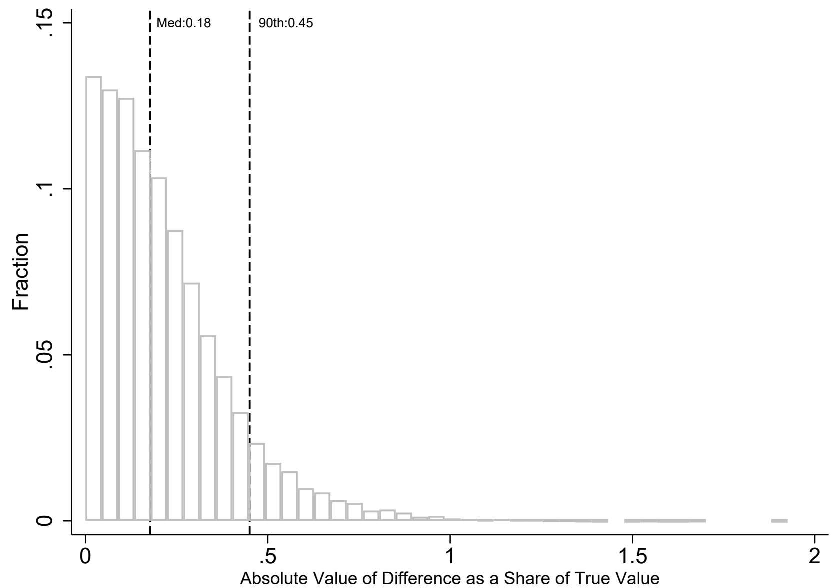

The figure plots the incidence of mismeasurement as a result of imputing common input prices across orders. We assume that the researcher observes the (weighted) average price of fabric paid in the industry when producing product j in year t. This is, for all orders of garment product j produced by any seller in year t we define a 'naive' input price $\widetilde{P^f}_o \equiv P^f_{jt}$. We study the mis-measurement induced by the imputation of common input prices, using the absolute value of the difference between the naive and true prices, as a share of the true price – $|(\widetilde{P^f}_o - P^f_o)/P^f_o|$. For the 22 thousand orders in our analysis sample, the histogram shows that half of the orders under or overstate the weight ratios by 17% or more, with the top 10% of the distribution featuring mis-measurement of more than 45%.

C Robustness of the Main Results

In this Online Appendix we conduct a systematic exploration of the robustness of our findings. We focus on issues of *measurement* and study the robustness of our main price and markup results to changes in specification, in the operational definition of relational sourcing and in the sample. We leave the discussion of alternative *mechanisms* to Section 5.3.

C.1 Robustness to Specifications

Our analysis focuses on differences in prices and markups across orders, accounted for by buyers' adoption of different sourcing strategies. Accordingly, the specifications we study in Sections 4.2 and 4.4 use variability *within* seller-product-year combinations, and condition on buyer-, relationship-, and order-level controls (including destination fixed effects). It is instructive to explore two types of departures from this structure: one that relaxes the controls, thereby allowing for the coefficient of interest to collect different selection and confounding forces, and one that uses alternative time horizons. We perform this exercise in a systematic manner, by considering all combinations that (i) let the set of covariates to feature none, some or all sets of controls in the paper (i.e., buyer-, relationship- and/or order-level controls); (ii) include one, two and three way combinations of fixed effects (s for seller, j for product, d for destination and t for period); and (iii) define the period t at either the month m, quarter q or year y. This approach leads to over 500 different specifications for each of the two outcomes.

The results of these estimations are presented in Figure 3, which studies order prices as outcomes, and Figure 4, which studies markups. We highlight three takeaways from these figures.

First, the over 500 specifications that we study produce results of (positive) sign and magnitude consistent with our baseline findings. All point estimates in the price regressions fall in the interval $[0.005, 0.046]$, with our preferred specification (corresponding to column (4) of Table 3) falling below the midpoint. All point estimates in the markups regressions are bounded in $[0.009, 0.048]$, while our baseline specification estimates a coefficient of 0.026 (column (3) of Table 6). Only 2.1% (6.8%) of the estimations are not significantly different from zero at 10% when the outcome is markups (prices).

Second, the specifications that render coefficients non-statistically different from zero correspond to those in which the set of fixed effects leave limited *within* variation for estimating the coefficient of interest. In most cases, these are specifications in which the time dimension t is set at the level of the month (or quarter) and there are not enough orders for seller-time or seller-product-time combinations – this is consistent with the fact that, at sufficiently disaggregated time, trade is lumpy.[42] At the other end of the spectrum, we obtain point

[42] In particular, the 36 specifications with coefficients not significantly different from zero (albeit positive) in the price regressions, in general correspond to two types of specifications: (i) seller-product-month, seller-month or seller-product-quarter fixed effects, alongside destination fixed effects and no order-level controls; (ii) destination-seller fixed effects, alongside product-month, product-quarter or product-year fixed effects and no order-level controls. The 11 specifications with coefficients not significantly different from zero (albeit positive) in the markups regressions, in general correspond to: (i) seller-product-month or seller-month fixed effects, alongside destination fixed effects; (ii) destination-seller-product fixed effects, alongside product-year

estimates that are twice as large as our baseline in both the price and markups regressions. Most of these specifications correspond to structures that pool together *within* and *across* sellers' variation. Thus, the sorting of heterogeneous sellers to buyers with different sourcing strategies pushes coefficients upwards through selection.

Third, we show that the choice of time horizon does not significantly change our main findings concerning input usage. We remind the reader that the data on material inputs are obtained from customs records. The data on labor usage, instead, come from internal production line records from factories. The unit of observation is an export order being shipped in the customs records, while it is a production line-day combination in the records from the factories. Garment plants produce every day, so we are able to exploit variation in labor usage within line and across days. Plants accumulate the output produced over several days (often months) and ship sporadically to fulfill orders. This leads to baseline specifications that account differently for time in the regressions that use customs records, and those that use production line records.

In Table 4 of Section 4.3 we show that the price and efficiency in the use of fabrics does not differ across buyers adopting different sourcing strategies, under our baseline specification. This structure implicitly suggests that the relevant time horizon for decisions pertaining to fabric usage is the year. In Table C1 we reproduce our results on fabric usage, changing the baseline time horizon (a year) for quarters (Panel A) and months (Panel B). We present the results with and without buyer, relationship and order-level controls. By and large, all patterns in our baseline specification remain unchanged. As in the price and markup regressions with very narrow time definitions, we note that variation across orders within a seller-product-month combination is more limited. In a similar vein, in Table 5 of Section 4.3 we document no differences in the use of labor across different types of buyers, within seller-month combinations. Table C2 reproduces the results on labor usage, changing the baseline time horizon (a month) to a quarter. Note that the average production line in the data is observed for around 350 days. Aggregating the time period at the year level thus does not control for anything more than the line fixed effects included in the specifications. Results appear to be robust.

C.2 Robustness to Definitions of Relational Sourcing

A key measurement challenge (see Macchiavello, 2022 for a review and a discussion) is that relational contracts are not typically observed in the data. Much of the empirical literature uses metrics of relationship age to proxy for relational trade. The advantage of this type of measure of relational trade is that it is observable in the data; the disadvantage is that repeated trade does not imply relational trade (which instead relies on future rents to provide incentives to parties). Specifications in Section 4 reveal that our results are robust to the inclusion of relationships' age as a control. As discussed in Section 2, we measure relational sourcing using the ratio between the number of suppliers and the number of shipments of a buyer. We compute the sourcing measure in *excluded* product categories. This allows for a more cautious interpretation of our results, while maintaining explanatory power (see Appendix Figure C1). The measure of relational sourcing that we introduce in the paper

or year fixed effects.

captures the idea that relational buyers concentrate most of their trade on a small group of suppliers, as formalized in our model of sourcing of Section 3. Naturally, this idea can be captured by other measures.

Table C3 reports the results of re-estimating our baseline regressions of price and markups, using 15 alternative definitions of the buyer's sourcing. We highlight here the most relevant findings. First, our results are robust to sourcing measures that use alternative normalizations, other than the number of shipments. In particular, we study constructions using values and volumes to find reassuring results. Relative to a normalization based on the number of shipments, normalizations based on volumes or values are more sensitive to outliers. This is because those variables are measured with some error, while the number of shipments is always correct by the structure of the Asycuda customs system. On both the price and markup estimations, when we use these alternative normalizations, the results are positive, significant and larger than our baseline. Second, other proxies for relational sourcing produce similar patterns. In particular, we study measures based on relationship duration [43] to find comparable results. Finally, we exploit information on the share of sellers' sales and the number of sellers to construct four metrics that capture concentration in a normalized measure along the lines of our baseline, and three direct proxies for concentration. Specifically, definitions 9 and 10 look at the ratio of sellers to shipments in excluded products, including only the largest sellers accounting for 50% or 70% of the buyer's shipments. Definitions 11 and 12 construct analogous metrics, but using volumes rather than shipments to measure concentration. Finally, definitions 13, 14 and 15 directly use the share of shipments in excluded products accounted for by the largest one, three or five sellers. Across all of these, our results on prices and markups remain very robust. In particular, when we restrict our baseline metric to cover only the largest sellers (the sellers that account of the majority of the buyer's trade), our point estimates are significantly larger than our baseline. This is intuitive, as focusing on the few suppliers that explain most of buyers' trade removes the fringe of small partners that make spot and relational buyers look alike.

C.3 Robustness to Samples

Broader Samples. In Online Appendix A.1 we explain the construction of our working sample in the customs records. This follows from a number of sequential trimmings over the universe of garment exports in Bangladesh. Relevant broader samples are described in Table A1, which considers five samples: (1) all shipments in woven exports throughout the entire sample period, (2) all shipments in the selected products categories, (3) all shipments in the selected product categories, that also have a UD associated to them, (4) all shipments in the relevant product categories, with UDs with high-quality data, and (5) the analysis sample. Here we show that our main results on both prices and markups can be recovered robustly in these broader samples.

Panel A of Table C4 studies the relationship between prices and relational sourcing in all knitted and woven garments, and in all woven garments (not only the 'included' products). In these samples, we cannot group transactions at the order-level, so this presents results for

[43] Measures of relationship duration as proxies for stickiness or relationalness have been used in Heise (2019) and Martin et al. (2020).

transaction-level prices. Columns (1) and (2) correspond to the samples of all garments and all woven garments, respectively, and have a coefficient of interest larger than the one we find in the transactions in included products (column (3)). The subsequent columns show increasingly trimmed samples, to reach our analysis sample in column (6). For the sample definitions of columns (4) to (6), the analysis can be performed both at the level of the transaction and the order in the case of prices (Panels A and B), and at the level of orders in the case of markups (Panel C). Across all of these samples our main results remain robust. As expected, the sample trimming that we need in order to (i) observe inputs used on export orders, and (ii) remove endogeneity concerns from using included products, attenuate the size of the coefficients we estimate.

Narrower Sample. As discussed in the main text, our results on material inputs, prices and markups are obtained using customs records. Instead, the analysis on labor usage is performed over a sample of factories whose production line records were collected in a series of RCTs. These data are described in detail in Online Appendix A.2. A relevant question is whether the results on material inputs, prices and markups also hold for the subset of factories in the production data.

The identifiers for exporters in the customs records are based on the Business Identification Number (BIN) – a tax identification code assigned by the National Board of Revenue in Bangladesh. BINs thus identify firms. Production records were directly collected, as part of field projects, from *plants* (or factories), and are identified by a name (and sometimes address) and a project ID. With this, there is no straightforward matching between plants in the production dataset and exporters in the customs records. For the purpose of examining the validity of our results in the overlapping sample, we exploit information on the *buyers* in the two data sources.

In both the customs records and the production line records, buyers are identified by their name, typically taken from a proforma invoice. This proves a consistent identifier across both data sources. A unit of observation in the production dataset is a factory-line-day triplet. There are almost 460 thousand such triplets in the data, distributed across 51 plants and 1,344 lines observed for an average of 341 days (see Table A2). For the purpose of this paper, only the line-day combinations for which the buyer is observed are suitable for analysis. Record keeping varies across plants and also within plants over time. The buyer for whom the line is producing on a specific day is observed for 46% of the observations. We are not aware of any systematic aspects of the data collection that drives the availability of this information. Table A3 shows that there are no significant differences between observations with and without information on the buyer.

There are a total of 164 buyers observed in the factory-line-day triplets and present in our customs records sample. We assess whether the patterns in the broader sample hold for the restricted sample of sellers that trade with these buyers. As these buyers are large brands, it turns out that almost all sellers in our analysis sample trade with these buyers at some point: of the 15,647 orders that feature in most of our regressions, 15,374 are exported by sellers trading at least once with these buyers. Naturally, the results in this alternative sample are almost identical to those in our baseline sample.

To impose a more meaningful restriction, we keep sellers trading at least 50% of their volumes with buyers in the production data. We first study orders in our analysis sample (characterized in column (5) of Table A1 in the data Appendix) that belong to these sellers. As this mechanically restricts attention to the top 500 sellers, we also explore a broader sample with sellers of any size (characterized in column (4) of Table A1). We reproduce our regressions on input usage, prices and markups for the two samples, respectively in Panel A and Panel B of Table C5. Across both exercises, we find results strongly consistent with those in the main analysis of the paper.

References

Heise, Sebastian, " Firm-to-Firm Relationships and Price Rigidity: Theory and Evidence," *mimeo*, August 2019.

Macchiavello, Rocco, "Relational Contracts and Development," *Annual Review of Economics*, 2022, 14.

Martin, Julien, Isabelle Mejean and Mathieu Parenti, "Relationship stickiness, international trade, and economic uncertainty," *CEPR Discussion Paper No. DP15609*, 2020.

Table C1: Buyers' Sourcing and Input Usage: Alternative Time Horizons

Panel A: Quarters

	(1)	(2)	(3)	(4)	(5)	(6)
	p^f_{sbjo}		$(F/Q)_{sbjo}$		$Complex_{sbjo}$	
$Relational_b$	0.008	0.002	-0.008	-0.006	0.010	-0.002
	(0.006)	(0.007)	(0.006)	(0.008)	(0.012)	(0.012)
FEs	sjq,d	sjq,d	sjq,d	sjq,d	sjq,d	sjq,d
Controls	.	B,R,O	.	B,R,O	.	B,R,O
R^2	0.72	0.75	0.49	0.55	0.58	0.67
Obs.	13,847	11,403	13,847	11,403	13,847	11,403

Panel B: Months

	(1)	(2)	(3)	(4)	(5)	(6)
	p^f_{sbjo}		$(F/Q)_{sbjo}$		$Complex_{sbjo}$	
$Relational_b$	0.010	0.002	-0.008	-0.005	-0.006	-0.024*
	(0.007)	(0.009)	(0.008)	(0.009)	(0.013)	(0.013)
FEs	sjm,d	sjm,d	sjm,d	sjm,d	sjm,d	sjm,d
Controls	.	B,R,O	.	B,R,O	.	B,R,O
R^2	0.75	0.78	0.55	0.60	0.63	0.71
Obs.	7,969	6,481	7,969	6,481	7,969	6,481

Standard errors in parentheses, clustered at the level of the buyer. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table reproduces the exercises on the usage of fabric, varying the time horizon on the fixed effects. The baseline specification (not in this table) uses seller-product-year fixed effects (sjt), alongside destination fixed effects. Panel A of this table replaces the sjt fixed effect for seller-product-quarter fixed effects (sjq) and Panel B uses seller-product-month fixed effects (sjm). All other aspects of the specifications here follow the table in the main text. In particular, odd columns do not include any other controls (aside from the fixed effects) and even columns control for buyer-, relationship- and order-level controls. These are as follows. Buyer: cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship: cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order: size of order (log volume), log price of fabric of the order (omitted when this is the outcome).

Table C2: Buyers' Sourcing and Labor Usage: Alternative Time Horizons

	(1) $SMV_{slb\tau}$	(2) $Efficency_{slb\tau}$	(3) $\#Workers_{slb\tau}$	(4) $Share\ Helpers_{slb\tau}$
$Relational_b$	0.017	-0.005	0.238	-0.002**
	(0.328)	(0.006)	(0.338)	(0.001)
FEs	sq(τ),sl,τ	sq(τ),sl,τ	sq(τ),sl,τ	sq(τ),sl,τ
R^2	0.84	0.22	0.92	0.92
Obs.	155,714	116,896	125,932	125,932

This table reproduces the results of even columns of Table 5, with the difference that instead of having $sm(\tau)$ (seller-month) fixed effects, we have $sq(\tau)$ (quarter) fixed effects. Standard errors in parentheses, clustered at the level of the buyer and production line. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. Across all specifications, the regressor of interest is the metric on relational sourcing, standardized and increasing in the relational characteristic of the buyer. The outcome in column (1) is the Standard Minutes Value (SMV), defined as the amount of time a particular garment is supposed to take to be sewed together computed by the factory's industrial engineers (often based on international libraries of SMVs of elemental sewing processes). Column (2) study labor efficiency of a particular line in a plant, producing for a buyer on a given day, $Efficiency_{slb\tau}$. Labor efficiency is constructed as the ratio between the minutes-equivalent of the output and the minutes of labor input. In turn, the output is calculated as Standard Minute Values times the number of pieces. The input is calculated using the number of workers times the runtime. See main text for a comprehensive description. The outcome in column (3) is the number of workers active on the line, $\#Workers_{slb\tau}$, and in column (4) it is the share of such workers that are line helpers, $Share\ Helpers_{slb\tau}$. The discrepancies in sample size across columns are due to the fact that not all plants keep administrative records of all labor usage metrics studied here. All specifications include as controls for relevant buyer characteristics, its size as a garment importer in Bangladesh, whether the buyer is a signatory of the compliance Accord as of 2019 and the cohort of the buyer.

Table C3: Alternative Definitions of Buyers' Relational Sourcing

Panel A: Variations of the Baseline Metric		p_{sbjo}	μ_{sbjo}
1	Weighted average across excluded products and all years of (minus) sellers-to-shipments ratio	0.038*** (0.007)	0.031*** (0.007)
2	Ratio of (minus) sellers to volumes (1,000 kg.), across excluded products and all years	0.170** (0.070)	0.214*** (0.068)
3	Ratio of (minus) sellers to values (1,000 USD), across excluded products and all years	0.116*** (0.033)	0.115*** (0.033)
4	Weighted average across excluded products and all years of (minus) sellers-to-shipments ratio, excluding the largest seller in each product-year combination of the buyer	0.032*** (0.006)	0.030*** (0.006)
5	Weighted average across excluded products and all years of (minus) sellers-to-shipments ratio, excluding the first six months the buyer is observed in the product category	0.040*** (0.007)	0.026*** (0.006)
6	Weighted average across excluded products and all years of (minus) sellers-to-shipments ratio, excluding the first year the buyer is observed in the product category	0.041*** (0.007)	0.027*** (0.006)
Panel B: Metrics based on Duration		p_{sbjo}	μ_{sbjo}
7	Weighted average across all excluded sellers in excluded products of the duration of the relationship, defined as the count of months of trade between the buyer and the seller	0.005 (0.005)	0.014*** (0.004)
8	Buyer-specific fixed effect after residualizing relationship duration (as defined in 7) on seller fixed effects, buyer's cohort fixed effects and buyer size	0.013** (0.005)	0.016*** (0.004)
Panel C: Metrics based on Concentration		p_{sbjo}	μ_{sbjo}
9	Ratio of (minus) sellers to shipments in excluded products including only the largest sellers accounting for 50% of the buyer's shipments	0.037*** (0.009)	0.030*** (0.007)
10	Ratio of (minus) sellers to shipments in excluded products including only the largest sellers accounting for 70% of the buyer's shipments	0.038*** (0.008)	0.031*** (0.007)
11	Ratio of (minus) sellers to shipments in excluded products including only the largest sellers accounting for 50% of the buyer's volumes	0.034*** (0.008)	0.032*** (0.006)
12	Ratio of (minus) sellers to shipments in excluded products including only the largest sellers accounting for 70% of the buyer's volumes	0.034*** (0.008)	0.032*** (0.006)
13	Share of shipments in excluded products accounted for by the largest one seller	0.020*** (0.005)	0.015*** (0.003)
14	Share of shipments in excluded products accounted for by the largest three sellers	0.017*** (0.005)	0.015*** (0.004)
15	Share of shipments in excluded products accounted for by the largest five sellers	0.016*** (0.005)	0.015*** (0.004)

The table reports the results of re-estimating our baseline regressions of price, p_{sbjo}, and markups, μ_{sbjo}, with alternative definitions of the buyer's relational metric. In all cases, the regressions include seller-product-year fixed effects, destination fixed effects, and buyer-, relationship- and order-level controls. These controls are as follows. Buyer: cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship: cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order: size of order (log volume), log price of fabric of the order. Standard errors in parentheses, clustered at the level of the buyer. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The construction of all alternative measures is described in the second column of the table. All measures are standardized, for comparability. Panel A reports results using six alternative constructions for the relational metric using seller-to-shipment ratios. Panel B focuses on metrics based on relationship duration. Panel C presents measures that exploit information on the share of sellers' sales and the number of sellers. The majority of the 32 regressions reported in the table run on the same number of observations. There are a few exceptions, due to specifics of the construction of alternative metrics. The smallest sample consists of 13,067 orders, out of the 15,647 used elsewhere in the paper. The discrepancy in sample size arises in three instances: (i) in definition 8, where the residualization of relationship duration does not always allow for retrieving a buyer fixed effect; (ii) in definition 6, where excluding transactions that happen within the first year of entering a product category removes buyer-product combinations that are active for less than a year; (iii) in definitions 2 and 3, where outliers (top and bottom 2%) are removed after normalizing with values or volumes.

Table C4: Prices in Transactions and Orders Across Samples

Panel A: Transaction level prices

	(1)	(2)	(3)	(4)	(5)	(6)
			p_{sbji}			
$Relational_b$	0.092***	0.094***	0.081***	0.079***	0.051***	0.042***
	(0.018)	(0.021)	(0.014)	(0.014)	(0.013)	(0.014)
FEs	sjt,d	sjt,d	sjt,d	sjt,d	sjt,d	sjt,d
Controls	B,R,O	B,R,O	B,R,O	B,R,O	B,R,O	B,R,O
Sample	All Garments	All Woven	Included	With UDs	Clean UDs	Analysis
R^2	0.55	0.58	0.58	0.58	0.49	0.49
Obs.	1,939,554	881,404	594,371	558,793	363,536	314,853

Panel B: Order level prices

	(4)	(5)	(6)
		p_{sbjo}	
$Relational_b$	0.044***	0.028***	0.027***
	(0.012)	(0.010)	(0.010)
FEs	sjt,d	sjt,d	sjt,d
Controls	B,R,O	B,R,O	B,R,O
Sample	With UDs	Clean UDs	Analysis
R^2	0.72	0.72	0.72
Obs.	25,996	17,353	15,607

Panel C: Order level markups

	(4)	(5)	(6)
		μ_{sbjo}	
$Relational_b$	0.019	0.024**	0.028***
	(0.023)	(0.011)	(0.008)
FEs	sjt,d	sjt,d	sjt,d
Controls	B,R,O	B,R,O	B,R,O
Sample	With UDs	Clean UDs	Analysis
R^2	0.43	0.42	0.40
Obs.	25,996	17,353	15,607

The table reports the main results on markups and prices, at the level of the transaction or the order, across different samples. In all cases, the standard errors are reported in parentheses, and clustered at the level of the buyer. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. Six different samples are considered. The sample in column (1) corresponds to all garment exports; column (2) corresponds to all woven exports; column (3) consists of all exports in the included products categories (in short, woven shirts and trousers - see Online Appendix A.1); column (4) restricts attentions to exports in included produced, so long as they are channeled through the UD system; column (5) restricts the sample of column (4) to remove observations in orders that have at least one missing or outlier value in any of the relevant variables for our analysis (essentially, prices and volumes of both inputs and outputs); column (6) further restricts attention to the largest 500 sellers of the sample. Online Appendix A.1 describes these samples in detail. Panel A focuses on transaction (or shipment) level prices, p_{sbji}, as an outcome, which can be observed in all samples. Panels B and C focus on prices and markups, P_{sbjo} and μ_{sbjo} respectively, aggregated at the level of the order, only possible where a UD is present. With this, there results are only available for the samples of columns (4), (5) and (6). In all cases, we include the fixed effects and controls used elsewhere in the paper. The regressor of interest is the measure of relational sourcing, $Relational_b$. This is constructed (as in the baseline in the paper) in excluded product categories in the case of columns (3) to (6). For columns (1) and (2) all product categories are used for the construction of the relational metric, because these samples contain what we define as *excluded* products. All regressions include seller-product-year fixed effects, destination fixed effects, and buyer-, relationship- and order-level controls. These controls are as follows. Buyer: cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship: cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order: size of order (log volume), log price of fabric of the order.

Table C5: Input Usage, Prices and Markups: Robustness in the Production Subsample

Panel A: Matching by Buyer - Among top 500 sellers					
	(1)	(2)	(3)	(4)	(5)
	p^f_{sbjo}	$(F/Q)_b$	$Complex_{sbjo}$	p_{sbjo}	μ_{sbjo}
$Relational_b$	0.007	-0.014	0.015	0.026**	0.036***
	(0.009)	(0.010)	(0.016)	(0.013)	(0.011)
FEs	sjt,d	sjt,d	sjt,d	sjt,d	sjt,d
Controls	B,R,O	B,R,O	B,R,O	B,R,O	B,R,O
R^2	0.71	0.44	0.51	0.73	0.40
Obs.	9,136	9,136	9,136	9,136	9,136

Panel B: Matching by Buyer - Sellers of any size					
	(1)	(2)	(3)	(4)	(5)
	p^f_{sbjo}	$(F/Q)_b$	$Complex_{sbjo}$	p_{sbjo}	μ_{sbjo}
$Relational_b$	0.009	-0.004	0.025	0.031**	0.030**
	(0.009)	(0.009)	(0.015)	(0.013)	(0.014)
FEs	sjt,d	sjt,d	sjt,d	sjt,d	sjt,d
Controls	B,R,O	B,R,O	B,R,O	B,R,O	B,R,O
R^2	0.70	0.39	0.49	0.72	0.43
Obs.	11,200	11,200	11,200	11,200	11,200

Standard errors in parentheses, clustered at the level of the buyer. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table reproduces the baseline regressions of Table 4 on material inputs usage and Table 6 on prices and markups in three alternative samples. Panel A considers the orders of all sellers in our analysis sample, so long as the seller trades at least 50% of their volumes with buyers present in the production data (see Appendix A.2 for details). Panel B considers the orders of all sellers for which there are clean UDs, so long as the seller trades at least 50% of their volumes with buyers present in the production data. The analysis sample is presented in detail in Table A1 and corresponds to column (5) and the orders included in Panel B here corresponds to those in column (4) in Table A1. In both cases, we intersect these samples with the condition trades at least 50% of their volumes with buyers present in the production data. Removing this trimming leaves the analysis sample largely unchanged. Please, refer to Appendix A for a discussion. The main regressor in all cases is the baseline, buyer-specific metric of relational sourcing and it is standardized. Outcomes are: the log weighted average price of fabric in the order, p^f_{sbjo}, the buy-to-ship ratio of the order, $(F/Q)_{sbjo}$, a measure of complexity of the garment order (the log of the number of fabric types used for producing the order), $Complex_{sbjo}$, the price of the order, p_{sbjo} and the order's markup, μ_{sbjo}. All columns include seller-product-year fixed effects and buyer-, relationship- and order-level controls, as follows. Buyer controls (B): fixed effect for the main destination of the buyer, cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship controls (R): Cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order controls (O): size of order (log volume), log price of fabric of the order (except when this is the outcome).

Figure C1: Sourcing in Excluded and Included Products

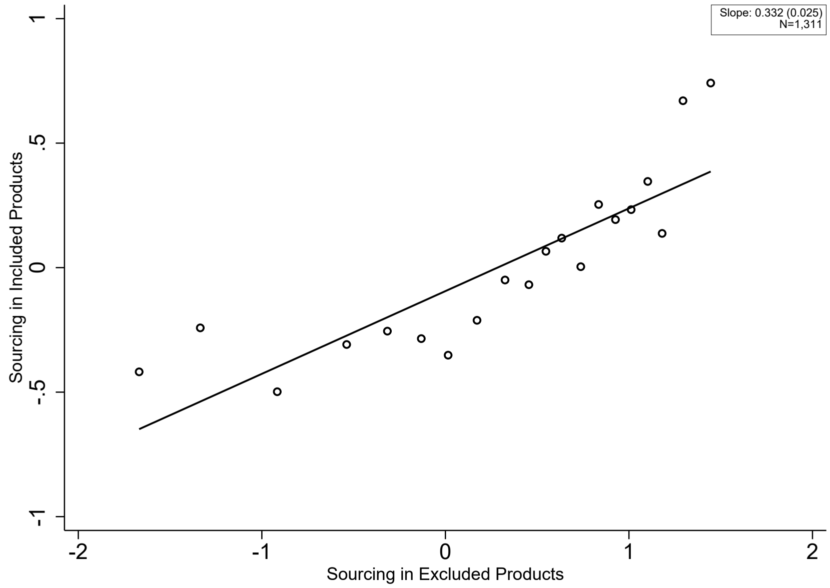

The figure presents the correlation between the sourcing metric computed using excluded products (horizontal axis) and using only included products (vertical axis) for 1,311 buyers in our analysis sample. The metrics are standardized in sample and the scatter markers denote averages across all buyers within each of 20 equally-sized bins. The overlayed line corresponds to a linear fit of the underlying data and has an estimated slope coefficient of 0.33.

D Appendix Tables and Figures

Table D1: Sources of Variability in Sourcing Strategies

Decomposition based on loss of fit (% of R^2)					
Fixed effects set:	I^1	I^2	I^3	I^3	I^4
Destination	16.90				
Buyer		68.33	61.10	57.80	41.64
Product	64.95	26.20			
Country	11.50	3.05			
Product-country			36.83	45.28	16.01
Product-destination					13.61
Sample		All		Multi-country	
Observations	300,660	300,660	300,531	141,135	137,239

Each entry reflects the loss of fit resulting from removing the fixed effects in the rows, from a linear projection of $Relational_{bjc}$ on the set of fixed effects in each I specification (columns). The specifications are as follows: $I^1 = \{destination, product, country\}$, $I^2 = \{buyer, product, country\}$ and $I^3 = \{buyer, product - country\}$ and $I^4 = \{buyer, product - country, product - destination\}$. The loss of fit is computed as a share over the fit in the full model: $(R_I^2 - R_{I-i}^2)/R_I^2$. The first three columns of the table use all buyer-product-country triplets available in the global data. The last two columns restrict attention only to buyers that are present in two countries or more.

Table D2: Within Seller Variation in Buyers' Sourcing Strategies

	Mean	Std. Dev.	P10	P25	P50	P75	P90
Count of buyers	21.0	17.1	4	8.50	17	28	42.5
Range of sourcing	0.72	0.28	0.27	0.56	0.85	0.94	0.97
Share of trade with relational (top 5%)	0.28	0.33	0	0	0.12	0.52	0.86
Share of trade with relational (top 10%)	0.47	0.37	0.007	0.08	0.43	0.84	0.98
Share of trade with relational (top 25%)	0.64	0.33	0.10	0.37	0.72	0.95	1.00
Share of trade with relational (top 50%)	0.82	0.23	0.48	0.73	0.92	0.99	1.00

The table shows summary statistics of different seller-level variables, computed for each of the 500 sellers in our analysis sample. The count of buyers is the number of different buyers the seller trades with, throughout the period of study. The range of sourcing is the difference (in absolute values) between the the relational sourcing metric of the most relational buyer and the least relational buyer of the seller. Note that this range can take values from zero to strictly less than one. The rest of the table shows statistics for the share of the seller's trade that is shipped to the most relational buyers in the sample. The different rows adopt different cutoffs for the binary definition of relational buyers. These are defined to be those whose metric of relational sourcing is in the top 5%, 10%, 25% and 50% of the distribution of relational sourcing among the buyers in the sample. The baseline discrete definition of relational sourcing used in this paper (elsewhere labeled as $Relational_b^D$) uses the 10% cutoff.

Table D3: Sourcing Strategies and Destination and Buyer Characteristics

Panel A: Destination Characteristics (Cross section for 2010)

	(1)	(2)	(3)	(4)	(5)	(6)
			$Relational_b$			
q_b	0.049***					0.049***
	(0.004)					(0.005)
GDP_d		0.002				
		(0.007)				
$Distance_d$			-0.074***			
			(0.027)			
$Population_d$				-0.002		
				(0.007)		
$GDPPC_d$					0.025	
					(0.018)	
Fixed Effects	y	y	y	y	y	y
R^2	0.14	0.03	0.04	0.03	0.04	0.25
Obs.	913	913	913	913	913	913

Panel B: Buyer Characteristics (Cross section for 2010)

	(1)	(2)	(3)	(4)
		$Relational_b$		
q_b	0.035***	0.034***	0.034***	0.030**
	(0.008)	(0.009)	(0.010)	(0.012)
$Assets_b$		0.011		
		(0.007)		
$Sales_b$			0.003	
			(0.009)	
$Employees_b$				0.011
				(0.009)
Fixed Effects	y	y	y	y
R^2	0.36	0.49	0.61	0.50
Obs.	353	244	177	160

Robust standard errors in parentheses. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. Panel A regresses the baseline buyer sourcing metric of relational sourcing on buyer's size of trade (q_b) in logs and gravity variables for the cross-section of active buyers in 2010. All gravity variables are in logs and correspond to the distance from the buyer's country to Bangladesh ($Distance_d$), the GDP of the destination country in the selected year (GDP_d), its population ($Population_d$) and GDP per capita ($GDPPC_d$). Specifications (1)-(5) include cohort year fixed effects (defined using the first year with observed trade in customs data), not shown. Specification (6) further includes main destination and main product fixed effects, not shown. Panel B regresses the baseline buyer sourcing metric of relational sourcing on buyer's size of trade (q_b) in logs and buyer characteristics for the selected year from Bureau van Dijk's ORBIS database. These include total assets ($Assets_b$), operating revenue ($Sales_b$), and number of employees ($Employees_b$) and are measured in logs. All specifications include fixed effects for main domestic country of the buyer, main activity (manufacturing, retail, wholesale, and services), size category (small, medium, large, and very large), and cohort year (winsorized at years 2004 and 2012 for dates of incorporation outside our timeframe), all calculated from Bureau van Dijk's ORBIS data, not shown.

Table D4: Sourcing Strategies, Time-Varying and Order Characteristics

Panel A: Buyer Characteristics

	(1) q_{bt}	(2) $Med\ Share^s_{bt}$	(3) $Max\ Share^s_{bt}$	(4) $Count^o_{bt}$	(5) $Count^{ship}_{bt}$	(6) $Count^j_{bt}$
$Relational_b$	0.639***	-0.123***	-0.041***	0.011	0.148***	0.082***
	(0.054)	(0.020)	(0.009)	(0.016)	(0.016)	(0.013)
R^2	0.11	0.56	0.40	0.70	0.87	0.52
Obs.	5,569	5,569	5,569	5,569	5,569	5,569

Panel B: Order Characteristics

	(1) q_{sbjo}	(2) $\overline{q}^{ship}_{sbjo}$	(3) N^{ship}_{sbjo}	(4) N^{ship}_{sbjo}
$Relational_b$	-0.042**	-0.164***	0.122***	0.150***
	(0.018)	(0.013)	(0.017)	(0.012)
R^2	0.56	0.52	0.61	0.83
Obs.	18,399	18,399	18,399	18,399

Standard errors in parentheses, clustered at the buyer level. $*(p < 0.10), **(p < 0.05), * * *(p < 0.01)$. Panel A regresses on the standardized buyer-specific sourcing characteristic, a number of outcomes: the buyer's size of trade (q_{bt}), the log share the median seller of the buyer has in the buyer's yearly trade ($Med\ Share^s_{bt}$), the log share that the largest seller of the buyer has in the buyer's yearly trade ($Max\ Share^s_{bt}$), the log number of orders the buyer has in the year ($Count^o_{bt}$), the log number of shipments the buyer has in the year ($Count^{ship}_{bt}$) and the log number of products (HS6 codes) the buyer purchases in the year ($Count^j_{bt}$). All columns (1)-(6) include year fixed effects and columns (2)-(6) also control for the size of the buyer's trade, q_{bt}. Panel B's main regressor of interest is again the standardized sourcing characteristic of the buyer on order-level and the outcomes are: the log size of the export order (q_{sbjo}), the log average size of the shipments in the order ($\overline{q}^{ship}_{sbjo}$) and the log number of shipments in the order (N^{ship}_{sbjo}). All specifications (1)-(4) include seller-product-year and destination fixed effects. They also control for the size of the buyer's trade, q_{bt}. Column (4) further controls for the size of the order (q_{sbjo}).

Table D5: Sourcing Strategies and Sellers' Characteristics

Panel A: Seller Characteristics

	(1) q_{st}	(2) $Count^j_{st}$	(3) $Count^d_{st}$	(4) $Count^b_{st}$	(5) $Count^o_{st}$
$Trades\ w/Relational_s$	0.371***	0.090	0.226***	0.271**	0.316***
	(0.125)	(0.058)	(0.086)	(0.115)	(0.103)
R^2	0.03	0.01	0.02	0.01	0.02
Obs.	3,248	3,248	3,241	3,248	3,248

Panel B: Seller Characteristics (conditional on size)

	(1) $Count^j_{st}$	(2) $Count^d_{st}$	(3) $Count^b_{st}$	(4) $Count^o_{st}$	(5) $Med\ Share^b_{st}$	(6) $Max\ Share^b_{st}$
$Trades\ w/Relational_s$	0.023	0.154	0.174	0.117	-0.202**	-0.109*
	(0.064)	(0.094)	(0.126)	(0.105)	(0.101)	(0.060)
R^2	0.14	0.12	0.14	0.44	0.03	0.03
Obs.	3,248	3,241	3,248	3,248	3,248	3,248

Panel C: Order Characteristics

	(1) q_{sbjo}	(2) $\overline{q}^{ship}_{sbjo}$	(3) N^{ship}_{sbjo}	(4) N^{ship}_{sbjo}
$Trades\ w/Relational_s$	-0.229***	0.009	-0.239***	-0.073**
	(0.064)	(0.040)	(0.059)	(0.036)
R^2	0.49	0.57	0.56	0.84
Obs.	18,030	18,030	18,030	18,030

Standard errors in parentheses, clustered at the seller level in Panels A and B and seller-product-year level in Panel C. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. Across all panels, the regressor of interest is a dummy that takes equal one if the seller trades at least once, in any product, at any point in time, with a buyer classified as relational (i.e. in the top 10 percentile of the distribution of the continuous sourcing characteristic). Panel A studies a number of outcomes at the level of the seller-year: the seller's log export volume in the products of interest (q_{st}), the log count of products, destinations, buyers and orders of the seller-year combination (respectively $Count^j_{bt}$, $Count^d_{bt}$, $Count^b_{bt}$, $Count^o_{bt}$). In all cases, Panel A conditions on year fixed effects. In columns (1) to (4), Panel B repeats the exercises on the counts of products, destinations, buyers and orders, but in addition to the year fixed effects it also controls for the seller-year's size (q_{st}). Columns (5) and (6) retains the same controls, and studies as outcomes the log share the median buyer of the seller in its yearly trade ($Med\ Share^b_{st}$) and the log share that the largest buyer of the seller has ($Max\ Share^b_{st}$). Panel C studies order-level outcomes: the log size of the export order (q_{sbjo}), the log average size of the shipments in the order ($\overline{q}^{ship}_{sbjo}$) and the log number of shipments in the order (N^{ship}_{sbjo}). All specifications (1)-(4) include buyer-product-year fixed effects and control for the size of the seller's trade, q_{st}. Column (4) further controls for the size of the order (q_{sbjo}).

Table D6: Orders Affected by Hartals

	(1)	(2)	(3)
		$Duration_{sbjo}$	
$Hartal_{sbjo}$=1	0.812***	0.761***	0.735***
	(0.034)	(0.029)	(0.039)
q_{sbjo}	0.144***	0.247***	0.231***
	(0.021)	(0.017)	(0.017)
$Count^{ship}_{sbjo}$	0.043***	0.033***	0.036***
	(0.005)	(0.004)	(0.005)
$Hartal_{sbjo}$=1 × $Relational^D_b$=1		-0.131***	-0.164***
		(0.049)	(0.059)
FEs	j,q	j,q,b	sjt,q,b
R^2	0.38	0.49	0.65
Obs.	18,010	16,824	13,964

Standard errors in parentheses, clustered at the buyer level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. This table studies the correlation between order duration, $Duration_{sbjo}$ and the occurrence of hartals. $Duration_{sbjo}$ is measured as the log count of days elapsed between the first and last shipment observed in an order, and we drop the top and bottom 5% tails in the duration distribution (orders with one shipment only and orders of more than 260 days). $Hartal_{sbjo}$ is a dummy that takes value one if at least one hartal occurs between the first and last shipment in the order (+/- seven days). The dates of hartals are obtained from Ahsan and Iqbal (2015). Across all specifications, we control for product fixed effects (j), as well as the calendar quarter in which the order started (q). In addition, we control for the size of the order (q_{sbjo}) and the number of shipments in the order ($Count^{ship}_{sbjo}$). Columns (2) and (3) also add buyer fixed effects (b) and in column (3) we introduce seller-product time (sjt) fixed effects, to mimic the baseline specifications we will introduce in Section 4. In columns (2) and (3) the coefficient of interest is the interaction between hartal occurrences and the relational characteristic of the buyer. For ease of interpretation of the interaction term, we use the discrete measure of relational sourcing ($Relational^D_b$), which takes value one if the buyer is in the top 10% of the distribution of the relational sourcing metric.

Table D7: Workers' Wages, Bonuses and Demographics

	(1) $Wage_{isl}$	(2) $Piece\ Rate_{isl}$	(3) $Quality_{isl}$	(4) $Other_{isl}$
$Relational_{sl}$	-0.052	-0.045	0.032	0.029
	(0.134)	(0.037)	(0.069)	(0.071)
FEs	s,r(i)	s,r(i)	s,r(i)	s,r(i)
R^2	0.54	0.24	0.14	0.12
Obs.	556	563	564	566
	(1) $Female_{isl}$	(2) $Educated_{isl}$	(3) $Ability_{isl}$	(4) $Experience_{isl}$
$Relational_{sl}$	-0.096	0.024	-0.434	-12.808
	(0.108)	(0.092)	(1.972)	(7.676)
FEs	s,r(i)	s,r(i)	s,r(i)	s,r(i)
R^2	0.55	0.47	0.06	0.47
Obs.	1,035	1,035	345	1,030

Standard errors in parentheses, clustered at the level of the seller. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. For the top panel, the data on workers' pay is only available for surveys conducted in Phase 1. With the exception of column (3), the bottom panel uses surveys of both Phase 1 and Phase 2. In all cases, the outcomes are defined at the level of worker i assigned to line l of seller (factory) s. The regressions include fixed effects at the level of the factory and role: $r(i)$ corresponds to whether the worker is a line operator, a supervisor or a chief. The regressor of interest $Relational_{sl}$ is the share of days that the line l is observed producing for relational buyers, over all days the line is active in the production data. The outcomes in the top panel correspond to the log basic salary ($Wage$), and indicators for whether the worker reports being paid piece rate (*Piece Rate*), quality bonuses (*Quality*), or other bonuses (*Other*). The bottom panel reports outcomes on demographics: the gender of the worker ($Female$), whether they have completed secondary education (*Educated*), the overall score of the worker's Raven Test (*Ability*) and the months of experience in the garment industry (*Experience*) - the latter, conditional on time invariant demographics. The Raven Test was only performed by supervisors and chiefs in Phase 1 of the study.

Table D8: Overtime and Pay when Producing for Relational Buyers

	$Wage_{ism}$			$Overtime_{ism}$			$Absenteeism_{ism}$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$Relational_{sm}$	0.001	-0.012	-0.006	-0.297	-0.763	-0.704	0.001	-0.073	-0.099
	(0.011)	(0.012)	(0.020)	(0.269)	(0.543)	(0.678)	(0.067)	(0.071)	(0.126)
Type Worker	All	Non-line	Managers	All	Non-line	Managers	All	Non-line	Managers
FEs	i	i	i	i	i	i	i	i	i
R^2	0.98	1.00	1.00	0.89	0.88	0.83	0.33	0.31	0.36
Obs.	191,935	16,170	3,074	191,941	16,170	3,074	191,941	16,170	3,074
Workers	25,313	2,654	512	25,315	2,654	512	25,315	2,654	512

Standard errors in parentheses, clustered at the level of the seller(factory)-month. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. For each factory-month combination, we compute the share of line-day pairs that are producing for a relational buyer (a buyer in the top 10% percent of the distribution of the sourcing variable). We study three outcomes. First, $Wage_{ism}$ is the log wage that worker i is paid on month m by its employer, seller s, as reported in the HR records. $Overtime_{ism}$ is the log hours of overtime recorded for the worker. $Absenteeism_{ism}$ is the log number of days the worker is absent in the month. All specifications include worker fixed effects. We study three samples. In columns (1), (4) and (7) we pool all workers in the HR records. In columns (2), (5) and (8), we focus on the subsample of workers that have designations that are not as line operators. This includes workers in cutting stations, workers on finishing, tagging, boxing, needle replacement, spot washers, quality control, ironing and folding. In columns (3), (6) and (9) we study workers with managerial designations. This includes production managers, assistant production managers, managing directors, general managers, HR managers, IE managers, in-house or external trainers, clerks (various types, including inventory).

Table D9: Buyers' Sourcing and Runtime

	(1) $Runtime_{slb\tau}$	(2) $Runtime_{slb\tau}$
$Relational_b$	0.022	0.029
	(0.035)	(0.030)
FEs	sm(τ),τ	sm(τ),sl,τ
R^2	0.66	0.67
Obs.	121,195	121,195

Standard errors in parentheses, clustered at the level of the buyer and production line. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. Across both specifications, the regressor of interest is the metric on relational sourcing, standardized and increasing in the relational characteristic of the buyer. The outcome in both column is the runtime of the line, i.e. the number of hours that the line was active (running) on a given day. All specifications include as controls for relevant buyer characteristics, its size as a garment importer in Bangladesh, whether the buyer is a signatory of the compliance Accord agreement as of 2019 and the cohort of the buyer. Odd numbered columns condition on fixed effects corresponding to the seller-month ($sm(\tau)$) and the day (τ). Even numbered columns, in addition, include a fixed effect for the production line of the seller (sl).

Table D10: Explanatory Power of Relationship Match Values

	p_{sbjo}		μ_{sbjo}	
	(1)	(2)	(3)	(4)
Explained by restricted model (buyer effects)	98.52%	95.55%	87.83%	90.17%
F statistic (saturated = restricted)	0.680	0.599	0.468	0.415
Controls	B,O	B,R,O	B,O	B,R,O
R^2 in saturated model	0.81	0.81	0.51	0.52
Obs.	16,979	15,183	16,979	15,183
Buyer effects (restricted)	792	679	792	679
Buyer-seller effects (saturated)	2,232	1,913	2,232	1,913

The table shows the explanatory power of a model with additive buyer and seller effects, relative to a specification that accounts for relationship-specific shifters in price and markups regressions. Columns (1) and (2) correspond to regressions of log order-level prices, p_{sbjo}, while columns (3) and (4) study order-level markups, μ_{sbjo}. Saturated specifications include seller-product-year (sjt) and seller-buyer fixed effects (sb). Restricted specifications include seller-product-year (sjt) and buyer fixed effects (b). The explained variation is computed as $R^2_{restricted}/R^2_{saturated}$ and expressed in percentages. For example, the entry in column (1) reads as follows: 98.52% of the variability in prices explained in a model with relationship-specific effects, can be accounted for by a model with additive buyer and seller effects. The second row presents an F-statistic for the null hypothesis for joint restrictions on the relationship fixed effects (i.e., $sb = b$). It follows a standard construction: $F \equiv [(R^2_{saturated} - R^2_{restricted})/J]/[(1 - R^2_{saturated})/(N - K)]$, with J the number of buyer-seller fixed effects constrained to a buyer effect. Odd columns include buyer and order controls, while even columns, in addition, control for relationship-specific characteristics. Time-invariant controls defined at the level of the buyer drop in all specifications and time-invariant relationship controls drop in saturated specifications. The full list of controls is as follows. Buyer: cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship: cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order: size of order (log volume), log price of fabric of the order.

Table D11: Correlation of Sourcing Measures and the Buyer and Relationship Levels

Panel A: Correlations across constructions of the buyer-level relational metric

		(1)	(2)	(3)	(4)	(5)	(6)
$Relational_b$ (baseline)	(1)	1.0000					
All Products	(2)	0.6104	1.0000				
Average over relationships, all products	(3)	0.6197	0.8128	1.0000			
Weighted avg. over relationships, all products	(4)	0.6102	0.9959	0.8070	1.0000		
Average over relationships, excl. products	(5)	0.7018	0.6986	0.8693	0.6950	1.0000	
Weighted avg over relationships, excl. products	(6)	0.7512	0.7866	0.7160	0.7880	0.8002	1.0000

Panel B: Correlations across constructions of the relationship-level relational metric

		(1)	(2)	(3)	Intra-cluster (buyer) corr.
Kilos per shipment, all products	(1)	1.0000			0.473
Kilos per shipment, excluded products	(2)	0.6774	1.0000		0.481
Duration, all products	(3)	0.4887	0.4442	1.0000	0.313

This table shows correlations across various sourcing metrics, defined at the level of the buyer (Panel A) and at the level of the relationship (Panel B). Panel A presents correlations over 1,311 buyers and the measures are as follows: (1), $Relational_b$, corresponds to the baseline relational metric used throughout the paper, constructed as (minus) the ratio of sellers to shipments in excluded products; (2) uses the same definition as the baseline metric (see equation (1), but exploits all products (including those in the analysis); (3) is the measure of sourcing as a simple average across relationship-specific shipment concentration, using all products; (4) is an alternative to (3) using volumes to construct a weighted average (see equation (8)); (5) and (6) are analogous to (3) and (4), respectively, but are constructed over excluded products only. Panel B studies correlations between relationship-specific measures of relational sourcing. In particular, (1) corresponds to the kilos per shipment in the relationship, across all products (see equation (7)); (2) follows the same construction, but excludes the products in the analysis; (3) is the duration of the relationship, in terms of months of effective trade interaction. The rightmost column of Panel B shows the intra-cluster (within buyer) correlation in the relationship specific measures.

Table D12: Relational Buyers and Relationships

Panel A: Prices

	(1)	(2)	(3)	(4)	(5)
			p_{sbjo}		
$Relational_b$	0.078***				
	(0.029)				
$\widetilde{Relational}_b$		0.141***		0.141***	
		(0.039)		(0.041)	
$Relational_{sb}$			0.069**		
			(0.034)		
$\widetilde{Relational}_{sb}$				-0.001	0.062
				(0.036)	(0.039)
FEs	sjt,d	sjt,d	sjt,d	sjt,d	sjt,b
Controls	B,O	B,O	B,O	B,O	O
R^2	0.73	0.73	0.73	0.73	0.77
Obs.	15,647	15,647	15,647	15,647	15,381

Panel B: Markups

	(1)	(2)	(3)	(4)	(5)
			μ_{sbjo}		
$Relational_b$	0.093***				
	(0.024)				
$\widetilde{Relational}_b$		0.132***		0.145***	
		(0.038)		(0.040)	
$Relational_{sb}$			0.064*		
			(0.035)		
$\widetilde{Relational}_{sb}$				0.063*	0.118**
				(0.037)	(0.047)
FEs	sjt,d	sjt,d	sjt,d	sjt,d	sjt,b
Controls	B,O	B,O	B,O	B,O	O
R^2	0.41	0.41	0.41	0.41	0.46
Obs.	15,647	15,647	15,647	15,647	15,381

Standard errors in parentheses, clustered at the level of the buyer. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. Panel A shows regressions on order-level prices (p_{sbjo}), while Panel B shows markups (μ_{sbjo}). All other aspects of the specifications in the top and bottom panels are the same. Columns (1) to (4) include seller-product-year and destination fixed effects, as well as buyer and order-level controls. Column (5), instead, includes seller-product-year and buyer fixed effects, as well as order-level controls. These controls are as follows. Buyer: cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Order: size of order (log volume), log price of fabric of the order. Standard errors in parentheses, clustered at the level of the buyer. The main regressors of interest vary across columns. Column (1) studies the baseline relational sourcing metric, $Relational_b$, constructed in excluded products. For comparability with the rest of the table, and unlike the rest of the paper, this measure is presented in levels (not standardized). Column (2) studies the buyer's relational metric, constructed as a weighted average across the sourcing in all relationships of the buyer in excluded products (see the definition of $\widetilde{Relational}_b$ in equation (8). Column (3) uses the relationship-specific measure of sourcing, $Relational_{sb}$ as defined in equation (7). Column (4) combined the buyer-level metric (as in column (2)) with the relationship level metric, centered around the buyer's mean (see definition of $\widetilde{Relational}_{sb}$ following after equation (8)). Column (5) also studies $\widetilde{Relational}_{sb}$ and includes buyer fixed effects, absorbing all the variation in $\widetilde{Relational}_b$.

Table D13: Sourcing and Reliability in Delivery

	(1) $Lead_{sbjo}$	(2) $Lead_{sbjo}$	(3) μ_{sbjo}	(4) μ_{sbjo}
$Relational_b$	-0.047***		0.025***	
	(0.015)		(0.006)	
$Relational^D_b$=1		-0.157***		0.061***
		(0.039)		(0.016)
$Lead_{sbjo}$			-0.037***	-0.036***
			(0.003)	(0.003)
FEs	sjt,d	sjt,d	sjt,d	sjt,d
Controls	B,R,O	B,R,O	B,R,O	B,R,O
R^2	0.36	0.36	0.42	0.42
Obs.	15,647	15,647	15,647	15,647

Standard errors in parentheses, clustered at the level of the buyer. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcome in the first two columns is the 'lead time' of the order (technically the throughput time), $Lead_{sbjo}$. $Lead_{sbjo}$ is constructed as the log number of days elapsed between the shipment into the country containing imported fabric for the order and the first shipment out of the country containing garment fulfilling the order. This variable is used as a regressor in other specifications. In columns (1) and (2) the regressor of interest is the measure of sourcing of the buyer - continuous, standardized, increasing in the relational characteristic of the buyer (column (1)) or its discrete alternative, picking up the top 10% of the sourcing distribution (column (2)). We present both measures for ease of interpretation. The outcome in columns (3) and (4) is the log markup factor, μ_{sbjo}. All columns in the table include seller-product-time (sjt) fixed effects, as well as destination fixed effects (d) and buyer-, relationship- and order-level covariates, as follows. Buyer controls (B): fixed effect for the main destination of the buyer, cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship controls (R): Cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order controls (O): size of order (log volume), log price of fabric of the order.

Table D14: The Downstream Market

	(1)	(2)	(3)	(4)	(5)	(6)
	μ_{sbjo}	μ_{sbjo}	μ_{sbjo}	μ_{sbjo}	μ_{sbjo}	μ_{sbjo}
$Relational_b$	0.026***	0.025***	0.024**	0.023***	0.101**	0.099**
	(0.007)	(0.007)	(0.011)	(0.007)	(0.041)	(0.041)
$Downstream_{bt}$					-0.008	
					(0.019)	
FEs	sjt,d	sjt,djt	sjt,sd	sjt,cjt	sjt,d	sjt,d
Controls	B,R,O	B,R,O	B,R,O	B,R,O	B,R,O	B,R,O
R^2	0.41	0.47	0.50	0.44	0.44	0.44
Obs.	15,647	15,186	14,856	15,030	5,471	5,471

Standard errors in parentheses, clustered at the buyer level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. In all columns the outcome is the log markup factor, μ_{sbjo}. The main regressor in all cases is the baseline, buyer-specific metric of relational sourcing and it is standardized. All columns in the table include buyer-, relationship- and order-level covariates, as follows. Buyer controls (B): fixed effect for the main destination of the buyer, cohort of the buyer (year first observed in the data), size (log volume imported by the buyer throughout our data and across all woven products), age of the buyer at the time of the order (log number of months elapsed since first observed in the data), a dummy indicating whether the buyer is a signatory of the Accord as of 2019. Relationship controls (R): Cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order controls (O): size of order (log volume), log price of fabric of the order. All columns include seller-product-year fixed effects. Column (1) also includes destination fixed effects and as such, simply reproduces the results of column (3) in Table 6. Column (2), instead, allows for destination-product-year fixed effects. Column (3) allows for seller-destination effects. Column (4) instead of the destination of the buyer uses the country to which the order is shipped, here denoted with c and allows for country-product-year fixed effects. Columns (5) and (6) use the specification of column (1) on a sub-sample of orders for which we have data on the total sales in clothing of the buyer in the year. These data were obtained from Euromonitor and covers all years in our sample, 2005-2012. Column (5) includes as a regressor $Downstream_{bt}$, the log annual sales of clothing of the buyer in its main country of operations, defined by the size of its retail sales. Column (6) reproduces the baseline regression of column (1) in the restricted sample of column (5). We note that the relational metric is re-standardized in columns (5) and (6) over the buyers in the smaller sample of these columns. While the sample size is considerably reduced in columns (5) and (6) we note that the raw correlation between the variable $Downstream_{bt}$ and the buyer size control that we use as baseline in the paper (i.e. the total volume of garments imported by the buyer) is positive and high (0.52). Moreover, conditional on destination and year, that correlation is 0.79.

Table D15: A Change in Sourcing Strategy - VF's Case

	(1) p_{sbjo}	(2) $(F/Q)_b$	(3) mc_{sbjo}	(4) μ_{sbjo}
VF_o=1 × $Post$=1	0.180***	-0.098	-0.019	0.200*
	(0.032)	(0.087)	(0.113)	(0.105)
FEs	sjt,b	sjt,b	sjt,b	sjt,b
Controls	R,O	R,O	R,O	R,O
R^2	0.78	0.51	0.68	0.45
Obs.	1,388	1,388	1,388	1,388

Standard errors in parentheses, clustered at the level of the buyer and year. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. We focus on export orders in the products of interest, manufactured by sellers that traded at some point with VF. Among those orders, we consider the orders placed by VF or by a main buyer of the seller. A main buyer is either the largest buyer (in volumes) of the supplier over the entirety of the sample period, before 2010 or after 2010. The estimated equation is in all cases $y_{sbjo} = \delta_{sjt} + \delta_b + \beta VF_o \times I_{r>2010} + \gamma Z_{sbjo} + \varepsilon_{sbjo}$. Outcomes are the log price of the order, p_{sbjo}, in column (1), the buy-to-ship ratio, $(F/Q)_{sbjo}$, in column (2), the log marginal cost, mc_{sbjo}, in column (3) and, in column (4), the log markup factor, μ_{sbjo}. Across all specifications, the regressor of interest is a treatment variable that takes value one for all orders placed by VF after its 2010 change in sourcing strategy. All specifications include seller-product-year fixed effects, buyer fixed effects and relationship- and order-level covariates. These are defined as follows. Relationship controls (R): Cohort of the relationship (year first observed in the data), size (log volume traded by the buyer and seller throughout our data and across all woven products), age of the relationship at the time of the order (log number of months elapsed since first observed in the data), share of the seller in all of buyer's trade, share of the buyer in all of seller's trade. Order controls (O): size of order (log volume), log price of fabric of the order.

The table complements the results presented in Figure 5. Column (1) of Table D15 shows that, relative to other buyers, the prices paid by VF increased after the transition. Column (2) detects a reduction in the buy-to-ship ratio which results in a decrease in marginal costs, as presented in column (3). These estimates are, however, noisy and we cannot reject a zero effect. Consequently, column (4) shows that the markup factor increases following VF's transition to relational sourcing. The difference between the price increase ($\approx$ 18%) and the markup increase ($\approx$ 20%) arises from the (imprecisely estimated) increase in efficiency through a reduction of the buy-to-ship ratio.

Figure D1: Cross-sectional Variation in the Relational Sourcing Metric

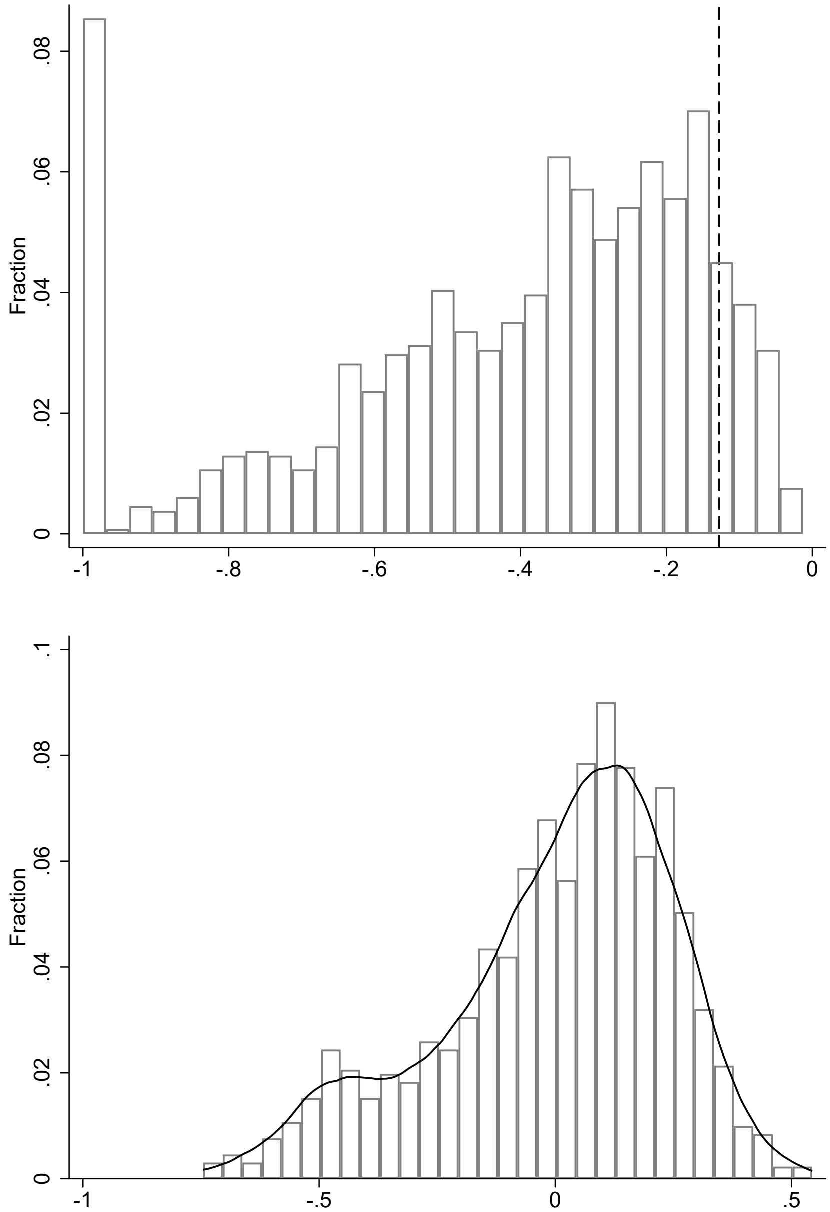

The histograms show the cross-sectional variation in the baseline metric of relational sourcing used in the paper. This is defined at the level of the buyer, as (minus) the ratio of sellers to shipments, as a weighted average in products excluded from the analysis. By definition, the measure ranges the interval $[-1, 0)$, where -1 corresponds to the most extreme spot sourcing and $\rightarrow 0$ is the most relational extreme. The top panel presents a histogram of this metric over the 1,311 buyers that ever trade orders in the analysis sample. The minimum (maximum) value that the metric takes is -1 (-0.012); the mean (median) is -0.413 (-0.344) and the standard deviation is 0.265. The vertical dashed line marks the top 10^{th} percentile (-0.127). The bottom panel presents the histogram (and kernel approximated density) of a residualized version of the sourcing metric, having regressed it on the overall buyer's traded volumes and the cohort of the buyer.

Figure D2: Buyers' Sourcing Strategies in Different Countries

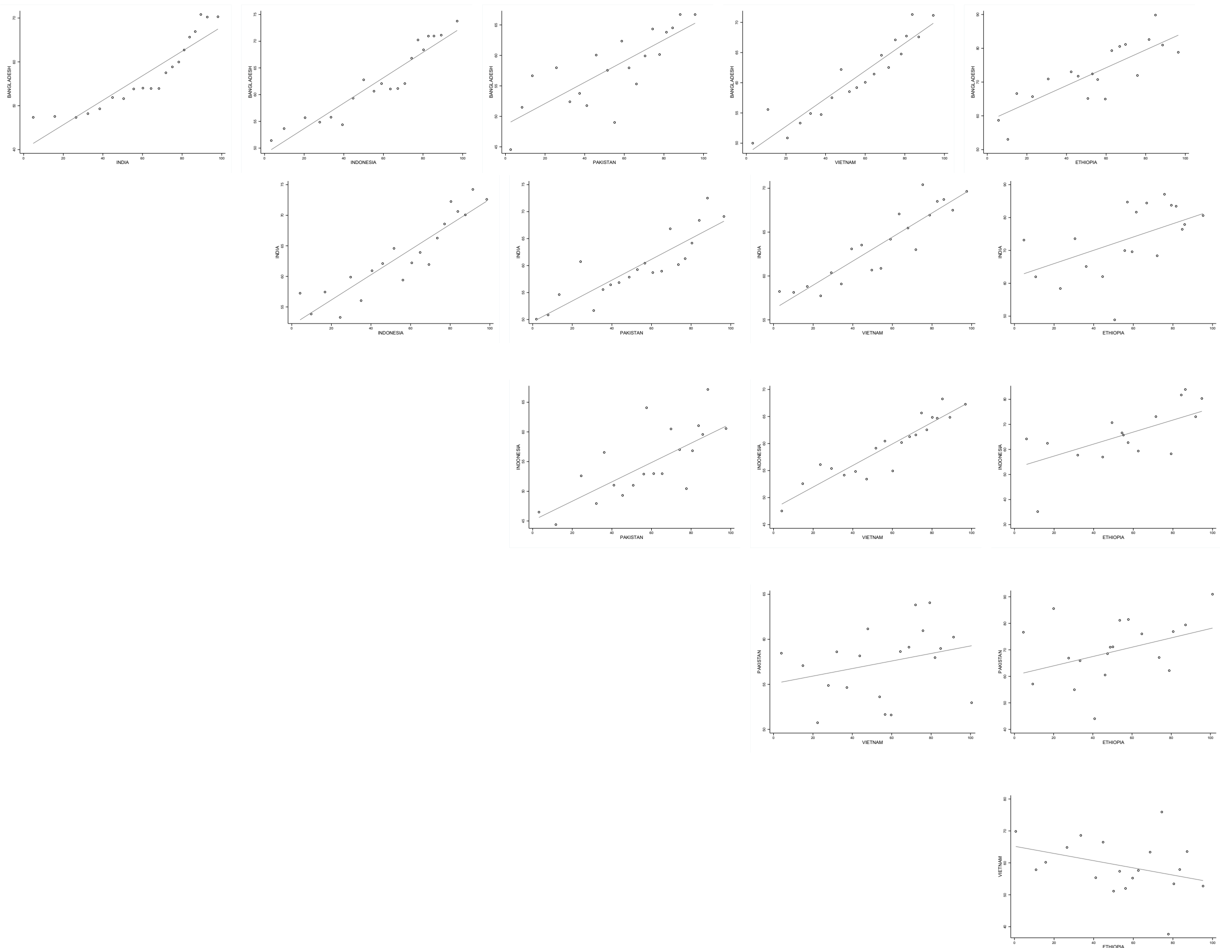

The graphs show two-way comparisons of buyers' sourcing strategies in different countries. A datapoint used for the construction of these graphs is a buyer-product-country combination, where the buyer-product is active in the two countries in the corresponding plot. The variable being plotted measures the sourcing strategy of the buyer in the product-country, $Relational_{bjc}$, and it is measured as (minus) the ratio between the number of sellers and the number of shipments. These measures are standardized within product-country pairs, and arranged in 100 quantiles in each country. The scatter markers correspond to averages in a partition over 20 bins. The solid line depicts the linear fit after a regression of the sourcing metric of buyers in the country indicated on the vertical axis, over the sourcing metric of these buyers in the country of the horizontal axis, conditional on product fixed effects. Each graph is produced on a different number of observations (buyer-product combinations present in both the horizontal axis and vertical axis countries). We report here the number of observations, point estimate of the slope coefficient and standard errors (clustered by product), corresponding to each graph. Bangladesh-India (top-left) : N= 12842 , Coeff= 0.291 , SE= 0.000. Bangladesh-Indonesia: N= 4196 , Coeff= 0.231 , SE= 0.000. Bangladesh-Pakistan: N= 3136 , Coeff= 0.163 , SE= 0.001. Bangladesh-Vietnam: N= 5131 , Coeff= 0.233 , SE= 0.001. Bangladesh-Ethiopia: N= 181 , Coeff= 0.314 , SE= 0.004. India-Indonesia: N= 5114 , Coeff= 0.217 , SE= 0.000. India-Pakistan: N= 3701 , Coeff= 0.192 , SE= 0.000. India-Vietnam: N= 5674 , Coeff= 0.170 , SE= 0.000. India-Ethiopia: N= 192 , Coeff= 0.208 , SE= 0.004. Indonesia-Pakistan: N= 876 , Coeff= 0.167 , SE= 0.002. Indonesia-Vietnam: N= 5248 , Coeff= 0.206 , SE= 0.000. Indonesia-Ethiopia: N= 100 , Coeff= 0.168 , SE= 0.021. Pakistan-Vietnam: N= 990 , Coeff= 0.041 , SE= 0.002. Pakistan-Ethiopia: N= 66 , Coeff= 0.273 , SE= 0.013. Vietnam-Ethiopia (bottom-right): N= 153 , Coeff= -0.112 , SE= 0.03.