Searching for Trade Partners in Developing Countries *

Julia Cajal-Grossi†

June 2022

Abstract

Multinationals' reputation in high-income countries is increasingly tied to the behavior of their foreign suppliers. How do buyers find suitable suppliers in low-income countries? Using customs records from Bangladesh's garment sector, this paper shows that when starting to source a product, buyers experiment with potential suppliers through small-scale, short-lived interactions before forming lasting relationships. While large buyers experiment more than smaller buyers on average, they experiment *less* when quality dispersion among potential suppliers is high—a counterintuitive finding from the perspective of search theory. I rationalize these empirical facts with a model of sequential search with reputational risk, and highlight an important trade-off for international buyers: on the one hand, when quality dispersion is high, buyers experiment more in hopes of finding a high-quality recurrent trade partner; on the other, in doing so they may unknowingly experiment with a low-quality supplier who damages their reputation. The model characterizes the optimal amount of experimentation and the threshold supplier quality at which buyers should settle. It yields two difference-in-differences relationships that I test in the data, exploiting exogenous variation in buyers' reputation concerns after the largest industrial accident in the history of the garment sector. In line with the model, the shock to reputation concerns led to less experimentation and worse matches among large (plausibly more reputation-sensitive) buyers in product categories with high supplier quality dispersion.

Keywords: Buyer-Driven Supply Chains, Search, Experimentation, Reputation.
JEL Codes: L14, D83, F14, F63

*This paper has benefited from valuable feedback from Jean-Louis Arcand, Jie Bai, Richard Baldwin, Greg Crawford, Mirko Draca, Jan Eeckhout, Christian Fischer, Rachel Griffith, Giulia Guipponi, Gordon Hanson, Beata Javorcik, Kala Krishna, Gilat Levy, Rocco Macchiavello, Rahul Mukherjee, Guillermo Noguera, Carlos Noton, Nicola Pavanini, Veronica Rapoport, Ken Teshima, Lore Vandewalle, Chris Woodruff, Daniel Xu and Yuan Zi. This paper has also benefited from helpful discussions at Aix-Marseille, Autonoma de Barcelona, Bayreuth, Cambridge University, DICE, Groningen Univeristy, Harvard Kennedy School, HEC Paris, IHEID, LMU, Mannheim, McGill, Namur, Northwestern (Kellogg), Passau, Penn State, Queen Mary, Rotterdam, UiOslo and Warwick.

†Geneva Graduate Institute (IHEID). E-mail: julia.cajal@graduateinstitute.ch.

1 Introduction

Multinationals' reputation in high-income countries is increasingly bound up with the behavior of their suppliers abroad. Across multiple industries, downstream firms' profitability depends on the timely delivery of parts and components, consistent access to quality inputs, and suppliers' environmental and social compliance, even in arms-length trade. Thus, the identification of suitable, trustworthy suppliers is a key pillar of global supply chain organization. This task is particularly challenging in developing countries, where contracting and information frictions tend to be high (Atkin and Khandelwal, 2020). A large body of literature has shown that in these contexts, long-lasting, recurrent buyer–seller relationships emerge as a solution (for a survey, see Macchiavello, 2021). However, little is known about how firms identify suppliers worthy of these sustained relationships.

This paper studies how firms (*buyers*) choose foreign suppliers in developing countries and how reputation concerns shape this decision process when buyers cannot observe relevant attributes of potential partners ex ante. I refer to these attributes as the supplier's *quality*. In line with the existing literature, quality is a bundle of characteristics that make a supplier, on average, a more appealing trade partner. The empirical context of this paper is Bangladesh's ready-made garment (RMG) sector, where from the perspective of international brands, short lead times, a low incidence of physical product defects and acceptable working conditions in the supply chain are important components of a suppliers' quality.[1] This bundle of characteristics is not fully observable by buyers ex ante, and supplier quality is learned only after an initial trade has taken place. Through an experimental order, the buyer can experience working with the supplier, lab-test products, calculate throughput times, trace the origin of material inputs, and assess supplier adherence to the buyers' code of conduct during production. Poor supplier performance on any of these fronts may severely tarnish the buyer's reputation downstream—hence the relevance of bridging the information gap before forming long-lasting trade partnerships.

This paper highlights a key trade-off in these buyer decisions: on the one hand, buyers need to learn suppliers' quality to find a good trade partner; on the other, because quality is only discernible through trade, buyers may inadvertently engage with suppliers who turn out to be of low quality. Despite an immediate termination, and however small the experimental order, this association can cost the buyer dearly. Stakeholders, consumers and the media

[1] A common assumption—intuitive from an empirical point of view—is that the components that enter the production function for quality are complements (Kremer, 1993; Verhoogen, 2008): a garment manufacturer that uses high-quality fabrics also hires proficient workers who produce garments with low defect rates, has operations managers who make sure deliveries ship on time, and hires social compliance officers who closely monitor working conditions. Thus, a high-quality supplier performs well on all (positively correlated) dimensions of relevant heterogeneity.

may interpret the engagement with the low-quality supplier as a reflection of the buyer's standards - the buyer's quest to find a good supplier carries a reputational risk.

The reputational risks of experimentation with suppliers may vary across buyers and product markets. Some buyers may be more reputation sensitive than others due to either higher exposure to scrutiny or a more valuable brand reputation. Reputation sensitivity is in general hard to measure but positively correlated with firm size (see, for example, Rob and Fishman, 2005). That is, buyers of different sizes may fare differently in the face of uncertainty during experimentation. Similarly, supplier quality may be a more or less critical issue in the sourcing of certain products. In just-in-time production chains, low-quality suppliers can cause severe disruptions in time-sensitive segments (Evans and Harrigan, 2005; Pisch, 2020). As another example, in the presence of O-ring technologies, the assembly of highly differentiated, complex products is particularly susceptible to the risk of sourcing from a bad supplier (a worker, in the original formulation of Kremer, 1993). Finally, in garment production, late deliveries are particularly costly in high-fashion-turnover segments, and social compliance violations are more likely in more differentiated product categories, where there is high dispersion in supplier quality.[2]

The main contribution of the paper is to empirically identify the link between buyers' reputation concerns and the search for trade partners in developing countries. It thus micro-founds the heterogeneous effects of information frictions across different products and buyers. As transitioning to the export of highly differentiated products is one development pathway for low-income countries, the findings here are relevant for the design of development policies aiming to mitigate information frictions. In this regard, the mechanism proposed in this paper is complementary to channels identified in existing literature on why low-income countries specialize in products with low scope for differentiation (Hallak and Schott, 2011; Hummels and Klenow, 2005; Khandelwal, 2010; Schott, 2004).

Using detailed customs records from the RMG sector for the period 2005-2015, I establish that (i) one-off trade interactions between buyers and sellers serve as an experimentation device for buyers searching for a recurrent supplier and (ii) large buyers experiment *less* when potential suppliers' quality dispersion is high. I briefly discuss the evidence supporting each of these two statements in turn. First, I show that a large proportion of interactions between buyers and sellers are one-off and that such interactions are more likely to happen when the buyer is entering a new product category.[3] These one-off interactions, as well as

[2] Section 2 discusses evidence supporting this characterization.

[3] This pattern coincides with existing evidence from other contexts. In terms of volume, most international trade occurs in the context of long-standing relationships; however, most trade relationships involve very short-lived interactions. See the empirical evidence in, for example, Monarch and Schmidt-Eisenlohr (2020) on the United States, Martin et al. (2020) on France and Macchiavello and Morjaria (2015) on Kenya. For

first interactions in recurrent relationships, have significantly smaller volumes than do later trade instances that materialize conditional on the relationship surviving. When buyers start sourcing a new product, those that do not form a recurrent relationship with a supplier are more likely to exit within their first year in the product category. The more one-off interactions that a buyer sustains during this first year in the product category, the larger and longer-lasting the trade under recurrent relationships tends to be if they are formed. Second, I examine the relationship between experimentation (the number of one-off interactions in which the buyer engages), buyer size and quality dispersion across potential partners.[4] I show that (a) large buyers experiment more than small ones, that (b) on average, buyers experiment *less* in product categories with high dispersion in potential supplier quality than in more homogeneous product categories, and that (c) the latter is particularly true of large buyers.

The negative relationship between experimentation and quality dispersion is counterintuitive from a search-theoretic perspective. Other things equal, dispersion in the value or quality of options should induce more search. This is because the presence of outstanding suppliers in the market increases the value of search and makes buyers pickier: they search more intensively in hopes of finding an outstanding option. The presence of low-quality suppliers in the market does not affect the value of search: if drawn, these low-quality suppliers are simply rejected, and no recurrent relationship is formed.[5]

I rationalize the empirical findings described above with a formulation of the search problem in which experimentation is risky for buyers with reputation concerns. There are two key ingredients in this formulation that set this framework apart from standard sequential search models. The first pertains to the *nature of search*. Unlike typical consumer or labor market search problems, the search for partners in international trade involves small-scale trade. The buyer cannot discern all relevant characteristics of the supplier until the trade is realized. In this context, the timing of the resolution of uncertainty becomes relevant. All

a focus on one-off trade, see Geishecker et al. (2019).

[4]This characterization requires a measure of the quality of the individual suppliers available in each product category. I obtain this seller-specific measure as a demand shifter conditional on price in an approach similar to that in Khandelwal (2010), following Berry (1994). I depart from Khandelwal's framework in that instead of a variety being defined in terms of a product–origin combination, the origin in my application is always Bangladesh, and varieties are defined in terms of product–seller pairs observed across multiple export destinations (markets). The outside option in the underlying consumer problem is given by the share of imported garments from other countries to each destination. The instrumentation strategy relies on combining fluctuations in the international price of raw cotton and the cotton content of different product varieties.

[5]For an early formulation of this problem, see Weitzman (1979), who formalizes sequential search for differentiated alternatives as optimization with a reservation utility property. A recent, more general exposition of the problem appears in Armstrong (2017) on ordered search. For a summary of equivalent formulations in the context of labor market search, see Rogerson et al. (2005).

recurrent relationships are formed when uncertainty has been resolved, but all experimental trade—necessary to form recurrent relationships—happens before the buyer learns the supplier's quality. The second ingredient is connected to the *experimentation cost*. In standard search problems, the cost of searching over a low-quality alternative need not differ from that corresponding to a high-quality alternative. This symmetry is unlikely to hold in the context of trade, in particular for reputation-sensitive international buyers, whose supply chain management is subject to public scrutiny. An example illustrates why this is the case. Consider a buyer who allocates a small experimental order with an unknown supplier. After the experimental phase, the buyer discovers that the supplier is of low quality and engages in social compliance violations. The buyer will not form a recurrent relationship with this supplier, but the exploratory interaction cannot be undone and may tarnish the buyer's reputation significantly. To follow Weitzman's metaphor, Pandora's box cannot be closed once it is open (1979). In this way, reputation concerns induce a form of downside risk in the process of experimentation with unknown partners.

Incorporating these two ingredients in a model of sequential search shows that if buyers' reputation concerns are sufficiently strong, they may be willing to settle for lower-quality suppliers and experiment *less* in environments with high quality dispersion. To highlight the role of dispersion in experimentation costs, I discuss comparative statics of the model under mean-preserving spreads of the underlying distribution of suppliers when the buyer is small and when the buyer is large (reputation-sensitive). This delivers a characterization of two key equilibrium outcomes in the buyer's experimentation process: the threshold supplier quality that the buyer is willing to accept and the amount of experimentation needed to find a suitable supplier. Comparing these outcomes for high- and low-dispersion environments and for large and small buyers after a shock to reputational costs, I obtain two difference-in-differences inequalities on these outcomes.

Next, I take these predicted equilibrium relationships to the data. Identification requires an exogenous shift to reputation costs. I leverage an unanticipated event that took place in 2013 and provided a plausibly exogenous increase in the reputational cost of buyers' experimentation with suppliers. In April of that year, the Rana Plaza complex, a multistory building housing garment plants, collapsed, imposing an unprecedented death and injury toll.[6] The episode heightened consumer awareness of buyers' purchasing practices and so-

[6]In a series of studies, Koenig and Poncet (2019, 2020) analyze the effects of the Rana Plaza collapse on the composition of garment imports into France. They show that while imports from Bangladesh continued to increase after the event, consumers penalized brands that were *named and shamed* (i.e., publicly listed as trading with Rana Plaza manufacturers), as evidenced by a drop in imports from the countries in which these brands marketed their products. The authors also show evidence suggesting that these brands turned to substitute, non-Asian sourcing countries.

cial compliance in their supply chains. Boycott campaigns and more stringent scrutiny by conscious consumers increased the potential penalties for buyers found to be dealing with noncompliant suppliers. In line with existing work, I show that trade volumes and the numbers of active buyers and entrants did not change significantly after the event, relative to their trends in previous years. However, the share of trade accounted for by one-off interactions dropped significantly.

The empirical strategy consists of studying buyer experimentation upon entry into product categories before and after the Rana Plaza collapse. To identify the sign of the empirical analogue to the theoretical inequalities in the model, I use a triple-difference equation that examines changes before and after the shock (first difference), for large and small buyers (second difference), and across product categories with varying supplier quality dispersion (third difference).

In line with the proposed reputational mechanism, I show that after the Rana Plaza collapse, relative to small buyers, large buyers reduced their experimentation and the quality of the suppliers with whom they settled in product categories with high supplier quality dispersion. This result holds when I control for the overall volume that buyers import in each category and when I condition on buyer and product–time fixed effects. A large battery of robustness exercises related to the selection of operational definitions, specification of the functional form and inclusion of additional controls supports the baseline findings. Three such exercises are worth highlighting. First, the differential effect of the shock on high-dispersion categories remains unchanged when heterogeneity (in the triple-difference) across other moments in the distribution is allowed for. Of particular interest are measures of central tendency and mass—the first considered to account for first order stochastic dominance and the second to account for market-thickness effects. Second, to relax the assumption that all buyers face the same pool of available suppliers (and, hence, the same dispersion), I construct hypothetical consideration sets for each buyer based on *similarity* sets assembled with reference to the buyer's observed partners. Finally, to assuage concerns over the quality measures being recovered from trade transactions, I redefine the consideration sets by excluding any sellers with whom a buyer has ever traded. As an alternative, I instrument for dispersion across sellers by using exogenous characteristics of the product categories.

The data do not seem to support alternative mechanisms. In particular, examining the evolution of large and small buyers' imported volumes across product categories before and after the shock, I show that the shift in large buyers' experimentation does not arise from a change in preferences. Nor does it follow from large buyers having a different rate of learning about sellers' attributes as they enter subsequent products over time. Finally, I show that the decrease in experimentation does not appear to be accompanied by the deployment of successful substitute screening technologies by large buyers. On the contrary, comparing the average quality of suppliers with whom buyers recurrently traded before and after the collapse, I find that large buyers saw a relative worsening of their matches in high-dispersion categories. This aligns with the model predictions and thus lends support to the reputation mechanism.

The paper contributes to two interconnected bodies of literature. First, it builds on a large empirical literature studying issues of reputation and information frictions between buyers and sellers in developing countries. Examples of these are Macchiavello and Morjaria (2015) in the context of the Kenyan flower market, Antràs and Foley (2015) on the relationship between a US poultry product exporter and its foreign buyers, Macchiavello (2010) on the Chilean wine-for-export industry, McMillan and Woodruff (1999) on trade credit in Vietnam, and Banerjee and Duflo (2000) on software contractors in India.[7] Taken jointly and in brief, these works establish that (i) sustained vertical relationships are valuable for both downstream and upstream firms when market frictions are present and (ii) the (individual or collective) reputation of upstream firms drives whether and how trade occurs between parties. Relative to the existing literature in this space, this paper introduces two novel elements. The first is the study of one-off interactions as experimentation, a precursor to the formation of incentive-compatible, sustained relationships. The second novel element is the transmission of reputation concerns from the downstream product market up the supply chain through buyer sourcing decisions. In contrast, existing work focuses on the role of reputation building within the sourcing relationship.

Second, the role of information frictions between buyers and sellers in shaping global trade flows has been an area of extensive theoretical and empirical research in international trade over the last few decades (see Grossman and Helpman, 2005 for an early contribution). Comprehensive surveys of this literature can be found in Antràs and Chor (2021) and, with a focus on development, Atkin and Khandelwal (2020). Some of the stylized facts in this paper are complementary to findings in Brugués (2021), Heise (2019) and Monarch and Schmidt-Eisenlohr (2020). While these studies focus on within-relationship dynamics, this paper restricts its attention to one-off relationships and conceptualizes experimentation as a buyer–product-level decision.[8] Focusing on homogeneous products, some papers leverage the

[7]Although it has a different goal, the identification strategy in this paper resembles that in studies exploiting media scandals after product recalls and industrial accidents. In particular, spillover effects across firms after reputation-damaging scandals are studied by Bai et al. (Forthcoming) with respect to the Chinese food industry and Freedman et al. (2012) in the context of toy recalls in the US. These works build on a long tradition of business studies focusing on firms' stock prices and sales following consumer reactions to negative news (Davidson III and Worrell, 1992; Dawar and Pillutla, 2000; Heerde et al., 2007).

[8]The results in this paper are also complementary to empirical findings on buyer–seller matching after policy shocks (see, for example, Benguria, 2021 and Sugita et al., Forthcoming). While this stream of

mapping between information frictions and prices to recover information costs (Allen, 2014; Jensen, 2007; Steinwender, 2018). In a differentiated product setting where prices are not fully informative of search frictions, Startz (2018) exploits tailored surveys to observe specific actions that buyers take to address such frictions. In that context, buyers are observed to travel internationally to meet suppliers face to face and find products ahead of the fashion cycle. In line with this approach, by studying one-off interactions as an experimentation device, this paper also relies on observed actions to characterize the nature and role of information frictions.

Section 2 describes the salient features of the empirical context and presents the data on which the analysis of this paper builds. Section 3 presents stylized facts on buyers' experimentation. Section 4 develops the theoretical framework, discusses comparative statics and presents the key theoretical relationships to be taken to the data. Section 5 discusses the identification strategy, presents the main empirical results and robustness, and addresses alternative mechanisms. Final remarks are left for Section 6.

2 The RMG Sector in Bangladesh

This section describes three salient features of the empirical context and presents the data on which the analysis of this paper builds. In brief, in the sourcing of garments from Bangladesh, (i) contracting and information frictions are pervasive, (ii) trade with low-quality or unreliable suppliers may carry large reputational costs for international buyers, and (iii) this reputational risk is larger in more differentiated product categories.

2.1 Background

The fast growth of the Bangladeshi RMG sector over the last few decades has turned the country into the world's second largest garment exporter. With the share of garments in export volumes expanding from 50% in 1990 to 84% in 2015, the sector now accounts for 14% of the country's GDP and 45% of its industrial employment.[9] While European and North American retailers benefit from Bangladesh's expertise and low production costs, challenges

research focuses on relationship churning as the pool of potential partners expands with policy interventions, in this paper, I show results on the quality of matches formed after a shock to reputational costs. This paper is also connected to Albornoz et al. (2012) and Nguyen (2012), who model the behavior of *exporters* facing uncertainty about their profitability in different destination markets. While these papers share the intuition on experimentation acting as an information-gathering strategy for sequential entry into different destinations, the focus here is on buyers (*importers*) experimenting with relationships differently across products due to reputation risk.

[9] Calculations by the author, based on information made available by the Bangladesh Garment Manufacturers and Exporters Association (BGMEA, www.bgmea.bd).

related to finding suitable suppliers and managing the supply chain in the country continue to give rise to bottlenecks (McKinsey Company, 2011; McKinsey Company, 2021). Three factors contribute to these difficulties.

First, contracts between buyers and sellers are hard to enforce, and ensuring the quality and timely delivery of orders produced in compliance with buyers' codes of conduct remains difficult. In this context, trust-based, long-lasting relationships provide incentives to suppliers and assurances to buyers, helping overcome contracting frictions (Cajal-Grossi et al., 2022). However, from buyers' perspective, identifying appropriate suppliers is challenging: despite their investments in screening technologies, international buyers lack relevant information about potential trade partners.[10] To mitigate this information problem, as in many other export-oriented industries in developing countries, buyers scrutinize potential partners in multiple ways. In particular, they commonly place small orders to assess not only garments' physical quality but also suppliers' ability to respond to tight deadlines without compromising social compliance standards. Anecdotal accounts of small-scale trade as an experimentation device in various industries appear in Tewari (1999) and Egan and Mody (1992).[11]

In the garment industry in particular, the allocation of small-scale, experimental orders is typically coupled with social compliance audits, chemical testing (for, among others, allergenic dyes, chlorophenols, and extractable arsenic), analyses of mechanical properties of the garments, and assessments of order traceability and timely delivery.[12] For example, Zara/Inditex performs pre-assessment audits and evaluates orders from potential suppliers through its comprehensive *Picking Program*.[13] Based on company reports, of all the suppli-

[10]A testament to this information problem is the large number of episodes in which industrial accidents reveal questionable production practices by suppliers of well-known brands (which were allegedly unaware of the violations). Table E2 presents a non-exhaustive list of media controversies of this nature, overlapping our sample period. Section 5 discusses in detail the last episode included in Table E2. In this vein, and in the context of the apparel sector in Cambodia, Brown et al. (2004) underscore the importance of social compliance indicators for reputation-sensitive international buyers.

[11]In the latter, this process is described as follows: *No matter how careful the selection process, the real test of a buyer's decision comes when the buyer and supplier are working together. For this reason, buyers tend to remain cautious after the final selection. For example, buyers often begin with small order.* (pp. 327, Egan and Mody, 1992). Formalizations of such experimentation appear in Watson (1999), Rauch and Watson (2003) and companion papers by the same authors with a game-theoretic approach and a specific focus on relationships between buyers and developing-country suppliers.

[12]It is well documented that observational information gathering and audits alone are not fully effective for screening. Studying 32,000 garment orders shipped to 30 large international buyers, Caro et al. (2021) find that a model with a comprehensive set of observable firm characteristics improves prediction of illegal subcontracting by only 33% over random guessing. Similarly to Plambeck and Taylor (2016), Caro et al. argue that stringent auditing gives rise to more refined cheating strategies.

[13]The various documents detailing the assessment process for suppliers joining the Zara/Inditex's supply chain are available online; of particular relevance is the documentation on physical testing requirements (Inditex, 2021a), the code of conduct (social compliance, including subcontracting) (Inditex, 2021b), safety

ers to whom Zara/Inditex allocated orders in 2015, 65 were rejected after order assessment; in 2014, 56 plants were rejected after assessment.[14] Naturally, not all elements of potential supplier screening are observable at the industry scale. In Section 3, I show evidence that one-off, small-scale orders are likely a good proxy for buyer experimentation with potential suppliers.

The second factor contributing to the aforementioned difficulties is that buyers' reputations are tightly linked to the behavior of upstream partners. Fast-fashion brands, operating with high-turnover collections and minimal in-house quality control, rely directly on high supplier performance. Consumer reports of systematic quality defects or brands missing a collection deadline by a few days can lead to significant mark-downs.[15] In addition, consumers with social and environmental concerns heavily penalize brands sourcing from plants that do not meet safety, sustainability and ethical standards. Harrison and Scorse (2010) study consumer boycotts against Nike, Adidas and Reebok in the context of their sourcing in Indonesia, and Koenig and Poncet (2020) focus on the Bangladeshi case. Any association, however brief and indirect, with an offending plant can impose large reputational costs. An article published in *The Guardian* soon after the Rana Plaza accident illustrates this point:

> *US giant Walmart, for example, is not involved in helping victims despite documentary evidence that its products were made in the building just a year ago. The retailer says that the work was unauthorised and no production was being carried out at Rana Plaza at the time of the accident.* (Guardian, 2013)

The quote refers to unauthorized subcontracting, a practice that buyers' codes of conduct typically ban yet one that is often at the heart of compliance violations.[16] Analyzing over 32,000 garment orders delivered from various developing countries to 30 large international buyers, Caro et al. (2021) find that unauthorized subcontracting was detected in 36.4% of the orders.

Third, a buyer's risk of unknowingly engaging with a low-quality supplier is higher in more differentiated product categories. This is not only because quality defects are more frequent in complex products but also because production of fashion items is prone to subcontracting

analyses (Inditex, 2021c) and chemical testing (Inditex, 2021d).

[14] Available online, accessed on 10 September 2021: `http://static.inditex.com/annual_report_2015/en/sustainability-balance/sustainability-balance/`.

[15] For instance, Zara/Inditex's 12 to 16 collections per year are maintained with minimal inventories and command their full price for an average of three weeks. This system has allowed the brand to reduce mark-downs to 15-20% of its stock; for comparison, H&M's mark-downs are closer to 45% of their inventories (Financial Times, 2014; Forbes, 2015; see also Carugati et al., 2008).

[16] For a sample of over 30 listed firms, Jacobs and Singhal (2017) show that the impact of the Rana Plaza collapse on firms' stock values was large and negative but short-lived. This stems from the sample being formed by early signatories of two leading remediation schemes. Studying the evolution of import flows into France, Koenig and Poncet (2020) show that offending brands were penalized with a decrease in demand.

and social compliance breaches. The Better Buying Initiative collects information on the purchasing practices of large international buyers, surveying over 750 garment suppliers across various countries, including 58 Bangladeshi plants (Better Buying Institute, 2020). Among other information, the survey measures the riskiness of sourcing practices based on order size volatility throughout the year. This metric aims to characterize *"the challenges suppliers face in adjusting to dramatic peaks and troughs of orders,"* and it is higher for fashion or differentiated products than for basic garments.[17] In line with common industry accounts and summarizing the responses from the survey, Better Buying's 2018 report states the following:

> *What happens as a result of inconsistency in monthly volume? Some suppliers indicated they were reluctant to be completely honest about how monthly order fluctuation impacted working conditions. ...Over 42% indicated that the month-to-month variability they experienced in buyer orders did not impact working conditions. Others reported a range of impacts, including: overtime within the law or code requirements, overtime in excess of law or code requirements, hiring of temporary workers, unauthorized subcontracting, reduced hours/ underemployment, and layoffs/ retrenchment of workers.* (pp 20, Better Buying Institute, 2018)

Furthermore, the 2020 edition of the report recounts the following response:

> *One supplier described the challenges saying,* "To balance out wrong forecasts and delays in materials, development, etc., we double-book our capacity." (pp 10, Better Buying Institute, 2020)

Lending further support to this characterization, Table E1 shows that plants specialized in products with high scope for differentiation are more likely to misreport the work hours of their employees, to impose overtime hours exceeding the legal limit on workers, to have (permanently or temporarily) overcrowded facilities and to fail on social compliance requirements for subcontracted work. To establish these facts, I exploit social compliance surveys conducted by the Better Work program of the International Labour Organization with 209 garment exporters in Bangladesh.[18] Combining these surveys with data (to be described in the next subsection) on these exporters' trade flows, I recover the main exporting product of each plant. I categorize these according to their scope for quality differentiation based on the measures constructed in Khandelwal (2010) using US data. Table E1 shows that relative to plants producing homogeneous products, those that specialize in products with

[17] As stated directly in the report, *"[O]rders that are primarily basic products are different from those that are primarily fashion products because suppliers acknowledge that there are no fashion products without order volatility"* (pp. 18, Better Buying Institute, 2018).

[18] The outcomes studied here are only a subset of the dimensions collected in the Better Work surveys and have been selected to match as closely as possible the statements in the anecdotal evidence above. This is at the expense of balance in the panel. In Section 3.2, I exploit the entire, balanced panel of plants, using an aggregated social compliance score constructed by the Better Work program.

a high scope for differentiation are 15% more likely to engage in misreporting and illegal overtime, 7% more likely to produce in overcrowded facilities and 17% more likely to violate social compliance requirements for subcontracted work (although the latter is not precisely estimated due to a smaller sample size).

The three features described above point to a delicate balance in the process of finding a trade partner. On the one hand, buyers would like to experiment with potential suppliers to gather necessary information and eventually settle with a high-quality, reliable partner. On the other hand, buyers need to minimize association in any capacity with potentially low-quality suppliers to reduce reputational risks. This trade-off appears particularly acute in product categories with high differentiation.

2.2 Data

The empirical analysis in this paper exploits a rich dataset recording all export transactions between Bangladeshi RMG manufacturers and buyers in the rest of the world. The primary source of this dataset is the compilation of bills of entry (exit) by the main custom stations in Bangladesh from January 2005 through to September 2015. Each record corresponds to a (six-digit HS-coded) product within a shipment from a seller to a buyer and occurring on a given date. The data include details on the statistical values, quantities, destinations and terms of trade. Importantly, they include identifiers for all buyers and sellers.

The almost 5 million export transactions in the data correspond to over ten thousand buyers and almost eight thousand sellers. Table 1 reports key summary statistics for sellers (Panel A), buyers (Panel B) and relationships (Panel C). Buyers import an average of 9.2 different products and trade with almost 15 sellers throughout their life span in the data (see Panel B). Within a product, buyers trade with just over three sellers on average, with a median of one seller per product. The median buyer–seller pair trades one product only (see Panel C). In terms of buyer–seller–product triplets (a *relationship*), the majority last for just one calendar quarter, and only about a third of uncensored relationships last two quarters or more, with the average duration being almost 4 quarters.[19] Table E3 presents further statistics on relationships' survival profile and shows that the probability of survival, while very low after the first interaction (quarter), improves significantly from the second interaction onward. I label relationships that last at most one calendar quarter as *one-off* and refer to all other relationships as *recurrent*.[20]

[19] Censoring is corrected by dropping all cells whose first observation falls in the first year or whose last observation falls in the last year of the data.

[20] Woven garments are typically produced under a *utilization declaration procedure*. For these exports, using identifiers recorded in the customs data, we can group different shipments into the orders (or contracts) that buyers place. The procedure and aggregation of the data to the order level are described in Cajal-Grossi et

3 Experimentation with Partners

This section exploits the observed trade between buyers and sellers to establish two empirical regularities. First, one-off trade between buyers and sellers appears as an experimentation device for buyers searching for a recurrent supplier. Second, larger buyers—which are plausibly more reputation sensitive—engage in *less* experimentation of this type when potential suppliers' quality dispersion is high.

3.1 One-off Trade as Experimentation

The analysis that follows shows that there is a large proportion of one-off interactions between buyers and sellers and that these are more likely to occur when the buyer is entering a new product category. When entering product categories, most buyers have some one-off interactions, and a small proportion of them engage in recurrent trade with at least one partner. Buyers who do not form a recurrent relationship are more likely to exit within their first year in the product category. I show that one-off interactions, as well as first interactions in recurrent relationships, are characterized by significantly smaller trade volumes than those arising under later trade instances conditional on the relationship surviving. Finally, buyers with more one-off interactions during their first year in a product category have longer-lasting recurrent relationships and larger trade volumes further down the line. The patterns presented in this section are descriptive, conditional correlations. Considered jointly, they suggest that one-off interactions function, at least partly, as an experimentation device whereby buyers test suppliers before settling into a recurrent relationship.

One-off interactions and entry. One-off interactions serve as an experimentation device when buyers start sourcing a product. In line with the summary statistics in the previous subsection, Figure 1 shows that approximately 60% of relationships are one off. This is the case across both uncensored and all buyer–seller–product triplets. This high number of one-off interactions is not accounted for by short-lived buyers. When I focus on a subsample of buyer–product combinations active for at least two years in the panel, the incidence of one-off relationships remains as high.

In what follows, I show that buyers are more likely to engage in one-off interactions when they enter new product categories. I follow the specification

$$one - off_{bt} = \delta_b + \delta_t + \alpha Entry^j_{bt} + \epsilon_{bt},$$

al. (2022). The median duration of woven orders over the period 2005-2012 is just over three months or one quarter (see Panel A of Table 1 in Cajal-Grossi et al., 2022).

where bt indexes a buyer–quarter duplet and δ_b and δ_t stand for buyer- and quarter-specific intercepts. The outcome, $one-off_{bt}$, takes value one whenever the buyer has at least one one-off interaction in the corresponding quarter and zero otherwise. The regressor of interest, $Entry^j_{bt}$, is a dummy variable indicating whether the buyer is entering at least one product category j, corresponding to an HS6 code, for the first time in the panel. The baseline specification, whose results are presented in column (3) of Table 2, conditions on buyer and quarter fixed effects, δ_b and δ_t. For reference, columns (1) and (2) discard the fixed effects, and column (1) presents a nonlinear estimation alternative. Across all specifications, the likelihood of one-off interactions significantly increases (by 35% to 39%) in quarters in which the buyer starts sourcing at least one product category.

The remaining columns in Table 2 show that this pattern holds when I take into account two plausible motives for short-lived interactions. First, buyers might adopt a sourcing strategy whereby they channel their steady demand through recurrent relationships and resort to unknown suppliers when core partners face capacity constraints. Under this setup, one-off interactions are more likely when the buyer faces unusually high demand or when its existing suppliers are producing at full capacity. Second, following termination of a recurrent relationship, the buyer might need to allocate volumes with unknown suppliers to either meet its demand or find a long-term replacement for the severed tie.

Column (4) of Table 2 addresses the capacity motive. For each of the buyer's recurrent suppliers, capacity is estimated as the largest quarterly volume observed throughout the panel. The capacity utilization of a seller in a given quarter is defined as the ratio between the volume shipped in that quarter and the seller's maximum capacity. The measure of interest, $\overline{Capacity}^s_{bt}$, is the average capacity utilization of all the recurrent suppliers with whom the buyer trades in the relevant quarter, and it is, naturally, bounded between zero and one. As expected, the probability of observing one-off interactions for a given buyer increases with the capacity utilization of existing suppliers. Columns (5) and (6) focus on the hypothesis of one-off interactions resulting after—or concomitantly with—termination of recurrent relationships. $Any\ Breakup_{bt}$ takes value one if the buyer sees the end of a recurrent relationship in quarter $t-1$ or t. The span over the previous quarter captures delays incurred between the processes of terminating trade partnerships and engaging with unknown suppliers. The table also presents an alternative specification leveraging the breakup count ($\#Breakups_{bt}$). The results show a positive, monotonic, significant association between the termination of recurrent relationships and the probability of engaging in one-off trade.

Across all specifications, when buyer-specific unobservables, quarter fixed effects, capacity constraints and relationship terminations are accounted for, buyers' likelihood of displaying one-off interactions is higher in periods of product entry. Table E4 presents further evidence supporting this characterization. With a more disaggregated specification, the table shows that the probability of observing one-off interactions decreases significantly (by between 15% and 36%) after the buyer forms its first recurrent relationship, conditional on the buyer staying in the market.[21] To complement this evidence, Appendix A characterizes the entry instances observed in the data and studies buyers' survival after product market entry. I show that when no recurrent relationships are formed, buyers are more likely to exit the product category within a year of entering. In other words, surviving buyers typically have at least one recurrent partner.

Starting small. Experimentation occurs via small-scale trade. The first interaction in any relationship (including one-off interactions) tends to be small. In the words of Rauch and Watson, buyers from developed countries *start small in unfamiliar environments* (2003).[22] In what follows, I show that this pattern can be verified empirically: relationships have a first trade interaction that is on average smaller than subsequent trade instances, conditional on the relationship surviving.

I study this pattern by means of two specifications, both decomposing the (log) volume traded by the buyer–seller–product–quarter combination, q_{sbjt}. The first specification exploits variation across relationships. It conditions on all buyer–product (δ_{bj}) and seller–product (δ_{sj}) characteristics and also controls for time-varying factors specific to the destination (δ_{dt}) or the product category (δ_{jt}):

$$q_{sbjt} = \alpha one - off_{sbj} + \delta_{bj} + \delta_{sj} + \delta_{dt} + \delta_{jt} + \epsilon_{sbjt},$$

where $one - off_{sbj}$ indicates whether the buyer–seller–product combination has a unique interaction. α denotes the difference in volume in quarterly interactions that are one-off relative to those that are recurrent. The second specification conditions on buyer–seller

[21] The specification is $One - off_{bjt} = \delta_b + \delta_{jt} + \alpha After\ Recurrent_{bjt} + \epsilon_{bjt}$, where the bjt subscript refers to buyer–product–quarter combinations and the outcome $One - off_{bjt}$ is an indicator for buyer–product–quarter triplets in which at least one one-off interaction with a seller is registered. The interest lies in α, which reflects the change in the probability of observing such one-off interactions in any quarter occurring after the buyer forms the first recurrent relationship in the product category. This event is indicated by the dummy $After\ Recurrent_{bjt}$. Buyer and product–quarter fixed effects remove any variation in the incidence of short-lived relationships that arises from time-varying factors common to all buyers in a product category or from buyer-specific characteristics. The results stated in the main text hold for all (uncensored) observations and in subsamples of buyer–product combinations that feature long survival in the panel.

[22] The authors model the buyer's decision on whether to start relationships with a testing project to gather information on the supplier's ability to fill large orders in due time and form or to allocate full-sized export orders from the start.

effects and therefore uses the variation in traded volumes within a relationship over time:

$$q_{sbjt} = \beta \mathbf{I}\{i_{sbjt} = 1^{st}\} + \delta_{sb} + \delta_{dt} + \delta_{jt} + \nu_{sbjt}.$$

In this case, the regressor of interest is $\mathbf{I}\{i_{sbjt} = 1^{st}\}$ and indicates the quarter in which the buyer–seller–product combination starts trading and takes value zero in any other quarter.

Table 3 collects the results of these exercises. Column (1) shows that one-off interactions are on average 46.9% smaller than those that take place in the context of recurrent relationships. Column (2) indicates that the volume of the first interaction in a relationship is on average 39.1% smaller than that under subsequent instances of trade. Columns (3) and (4) repeat the exercise from column (2) in two relevant samples. First, note that the base category in column (2) includes relationships of any duration, conditional on survival after the first interaction. Column (3) shows that in a sample restricted to relationships that last at least a year, first interactions are 34.1% smaller than subsequent ones. Second, note that the sample studied in columns (1) to (3) includes buyers of any size. The distribution of buyers' volumes is highly skewed, with a large mass of very small buyers and a top tail of large buyers who account for most of the country's exports. Small buyers, whose orders are (in equilibrium) small, might not be able to cut volumes significantly in first interactions due to both fixed costs of trade and the minimum-volume restrictions that many suppliers impose. Column (4) restricts attention to relationships involving the largest 200 buyers, who account for approximately 70% of the traded volume in the sample. In this subsample, the first interaction of surviving relationships is on average 51.3% smaller in volume than trade interactions in further quarters.[23]

Experimentation and relationship success. How does experimentation relate to success after entry? Buyers who experiment more (have more one-off relationships) during the first year in the product category have longer recurrent relationships with larger trade volumes. To establish these correlations, I focus on the following specification:

$$y_{sbj} = \alpha Experiment_{bj}^{entry} + \delta_j + \delta_s + \delta_b + \delta_{c(sbj)} + \delta_{c(bj)} + \epsilon_{sbj},$$

[23]Partial year effects as discussed in Bernard et al. (2017) are less of a concern in the specifications presented here because of the shorter horizon of aggregation over time. Table E5 reproduces the results here, correcting first-quarter volumes to aggregate over 92-day windows (rather than calendar quarters) and finding significant and qualitatively similar, albeit attenuated, results.

which studies outcomes y_{sbj} of recurrent relationships, i.e., buyer–seller–product triplets active for more than one quarter.[24] The outcomes reflect the performance of recurrent relationships along three dimensions: duration, overall volume and intensity (number of shipments per kilo of trade).[25] All regressions condition on product (δ_j), buyer (δ_b) and seller (δ_s) fixed effects and intercepts for the quarter in which the buyer–seller–product relationship starts and that in which the buyer enters the product (the *cohort*, $c(sbj)$ and $c(bj)$). These account for determinants of relationship performance such as seasonality of traded products, seller production capacity or buyer size. $Experiment_{bj}^{entry}$ collects the (log) count of one-off interactions by buyer b within the first year after it enters product category j.

Table 4 presents the results. Columns (1) and (2) show that in product categories where the buyer engages in more one-off trade shortly after entering the market, relationships tend to last longer once they become recurrent. Column (1) uses the count of effective quarters of interaction as a measure for duration, while column (2) uses the time span between the first and last interactions. Column (3) shows that the recurrent relationships formed after high experimentation upon entry tend to have larger trade volumes. Conditional on volumes, these relationships also involve a higher number of shipments (column (5); see column (4) for reference). This result is consistent with close relationships between buyers and sellers featuring a high number of small shipments, as opposed to large, sporadic interactions.

Multiple mechanisms may account for these patterns. On the one hand, if experimentation improves selection into relationships, conditional on survival, relationships that follow higher experimentation should ex post be more successful. On the other hand, buyers who expect to import large volumes or anticipate the need for sustained supply may engage ex ante in higher experimentation upon entry.

To sum up, this section showed evidence of buyer experimentation via one-off interactions. All relationships start small and have a high probability of failure. Buyers' survival in a market is closely connected to the formation of recurrent relationships, whose success in turn tends to be higher when buyers engage in more experimentation upon entry.

[24] The sample restricts attention to the outcomes of recurrent relationships of buyers who do some experimentation upon entry. This means that buyers who form no recurrent relationship are not included in this analysis, nor are those without at least one one-off interaction within the first year in the category.

[25] See Taylor and Wiggins (1997) for a model in which relational sourcing induces more frequent shipments, conditional on traded volumes.

3.2 Experimentation and Differentiation

I next study how experimentation varies across buyers and product categories. To characterize the latter, I construct a measure of a product's *dispersion in the quality of available sellers*. The first subsection below describes this measure. The second subsection shows that (i) on average, experimentation is *lower* in products with higher dispersion across sellers and (ii) large buyers tend to experiment more than smaller buyers but do so less markedly in markets with high dispersion.

Measuring quality. I recover the seller-specific quality, θ_s, using a demand model that defines the seller's quality as a market-share shifter conditional on prices, where markets are destination–product–time triplets. The approach is akin to the recovery of variety-specific quality in Khandelwal (2010), based on Berry (1994). A detailed discussion of the estimation procedure is presented in Appendix C.

By construction, the measure reflects any vertical attribute that conditionally increases demand for the seller on average across all markets. As such, $\hat{\theta}_s$ bundles aspects of physical quality, nonphysical quality such as delivery predictability, and social compliance performance, among other characteristics. To illustrate this, using data from social compliance assessments conducted by the Better Work program of the International Labour Organization, Figure C1 in the appendix shows that among 193 plants, those with an estimated high $\hat{\theta}_s$ exhibit significantly better compliance performance.[26] Albeit based on a restricted sample, this result suggests that the estimated quality shifters at least partially capture heterogeneity across suppliers on the social compliance front.

To characterize how heterogeneous the sellers active in a product category are, I measure the dispersion across $\hat{\theta}_s$ within the product category by computing the standard deviation across $\hat{\theta}_s$ for all sellers ever trading product j, which I label $Dispersion_j$.[27] The metric of

[26] Specifically, in a balanced panel of five assessment cycles for each of the 193 plants, Better Work constructs a *noncompliance score* capturing the number of social compliance violations over a large number of dimensions. These 193 plants are active in the customs records and have a $\hat{\theta}_s$ estimated by means of the procedure described here. Partitioning these 193 estimates into low, medium and high groups, Figure C1 shows that sellers with a high $\hat{\theta}_s$ have 10% lower noncompliance scores than those with low $\hat{\theta}_s$. Those with a medium $\hat{\theta}_s$ show 4% lower noncompliance than plants with low estimated quality, but this difference is not significantly different from zero.

[27] This approach is analogous to that in Khandelwal (2010), where the (log) difference between the minimum and the maximum demand shifter in a category is defined as the scope for product differentiation or the length of the *ladder*. Here, dispersion is measured with the standard deviation, which not only exploits the extrema but all seller-specific estimates. The metric is thus constructed as the square root form $\sqrt{\sum_{s\in j}(\hat{\theta}_s - \overline{\hat{\theta}_s})^2/\#\{s \in j\}}$, where $s \in j$ denotes all sellers who sell product j at least once, $\#\{s \in j\}$ is the cardinality of this set and $\overline{\hat{\theta}_s}$ corresponds to the average across all sellers in seller set j. For ease of interpretation, in regressions, I use the log of this dispersion metric. The robustness of the main results in

seller dispersion across the product market correlates well with observable product characteristics normally associated with horizontal and vertical differentiation. Table C3 shows that product categories with higher seller dispersion more often correspond to categories of garments made with fabrics other than cotton—the least differentiated material—and garments produced for women—items typically characterized by quicker fashion turnover. Products with higher dispersion are also shown to be more complex, in that they require the combination of a higher number of inputs and feature a high scope for product differentiation, according to Khandelwal's (2010) quality ladders for the US. Organizing product categories into quartiles according to their quality dispersion, such that products in the first quartile exhibit low dispersion across sellers (a low standard deviation in $\hat{\theta}_s$) and those in the fourth quartile feature high dispersion (a high standard deviation), Table C4 lists the largest product categories in each of these quartiles. Visual inspection of the goods' descriptions confirms the correlations offered in Table C3.

Experimentation, large buyers and dispersion. I investigate the relationship between experimentation, buyer size and quality dispersion across the sellers available in a product category. A priori, various mechanisms may mediate these relationships. For example, large buyers trading high volumes may find it more valuable to secure a good match, leading to them engaging in more experimentation than smaller buyers. In addition, buyers of different sizes may incur different experimentation costs. Across products, quality dispersion across potential partners in highly differentiated categories is high. This may induce high experimentation both because a thick bottom tail implies that buyers reject a large mass of suppliers and because a thick upper tail promises an outstanding match.[28]

To study the relationships described above, I measure the amount of experimentation that a buyer b engages in during calendar quarter t in product category j as the (log) count of one-off interactions in the buyer–product–time triplet. I label this outcome as $Experiment_{bjt}$ and decompose it by means of

$$Experiment_{bjt} = \Delta + \gamma_1 Large_b + \gamma_2 Dispersion_j + \gamma_3 Dispersion_j \times Large_b + \epsilon_{bjt},$$

where Δ denotes a set of fixed effects that vary across alternative specifications. $Dispersion_j$ corresponds to (the log of) the measure of heterogeneity across seller quality in product cat-

this paper to alternative constructions of this variable is left for consideration in Appendix D.1.

[28]For a formalization of this positive relationship between dispersion and search for the best alternative, see Weitzman's (1979) formulation of the search problem as a reservation utility one. For an analogous formulation in the labor literature, see Rogerson et al. (2005) (equation (5), pp. 962). I revisit these formalizations in Section 4.

egory j, as defined above.[29] $Large_b$ is a dummy that takes value one if the buyer is among the top 200 garment importers in Bangladesh. To focus on specific sources of variability in the data, I include product–time fixed effects ($\Delta = \delta_{jt}$) in some specifications and exploit variation across buyers to identify γ_1. Similarly, I present exercises including buyer–time fixed effects ($\Delta = \delta_{bt}$) to leverage variation across products for the estimation of γ_2. The richest specifications include buyer and product–quarter effects ($\Delta = \delta_b + \delta_{jt}$) and identify γ_3 only.

Table 5 shows that on average, large buyers experiment more than their smaller counterparts and experimentation is lower in more dispersed environments. This is more markedly the case for large buyers. Columns (1) and (2) show that large buyers have 15% more one-off interactions than small ones.[30] Column (3) of Table 5 shows that on average, buyers experiment *less* in products that exhibit high dispersion in available suppliers' quality. This difference across products is particularly stark for large buyers (columns (4) to (9)). When I condition flexibly on buyer- and product–time-specific determinants of one-off interactions, such as demand shocks common to all buyers or buyer-specific screening technologies, large buyers experiment approximately 5% less in product categories with the highest dispersion across sellers (see column (5) of Panel B). This pattern is also observed when attention is restricted to quarters in which the buyer's motives to search for a supplier are heightened; this is in line with the earlier discussion in Section 3 on buyers entering a product category or experiencing the end of a recurrent relationship.

One possible explanation for large buyers experimenting less in high-dispersion environments is that these product categories may have lower-quality sellers on average or may have fewer sellers available overall. To test these ideas, I take into account other moments in the quality distribution. In particular and as expected, large buyers tend to experiment more in thicker markets and in categories where the median supplier is of high quality. When I control for these factors, the results on dispersion and experimentation remain unchanged.[31]

[29] For ease of interpretation, I study an alternative specification where the dispersion metric is replaced by the quartile categories described in preceding paragraphs: $Experiment_{bjt} = \Delta + \gamma_1 Large_b + \sum_{i=2}^{4} \gamma_2^i \{Quartile\ Dispersion_j = i\} + \sum_{i=2}^{4} \gamma_3^i \{Quartile\ Dispersion_j = i\} \times Large_b + \epsilon_{bjt}$. $Quartile\ Dispersion_j$ is the quartile of product j in the distribution of seller heterogeneity, such that products in the first quartile (i.e., $Quartile\ Dispersion_j = 1$) feature low dispersion across sellers while the fourth quartile represents large dispersion.

[30] This follows partly from volumes being positively correlated with the number of relationships (of any duration). Still, when I control for the volume imported by the buyer in the product–year combination, buyers classified as large have 13% higher experimentation (not reported in the table, for brevity).

[31] It is possible that large buyers have lower demand for highly differentiated products. If this is the case, costly experimentation may appear less appealing. Columns (8) and (9) of Table 5 control for the volume that the buyer imports in the product category and for the number of recurrent relationships that it forms to show that these factors do not drive the dispersion and experimentation patterns.

This section showed that large buyers experiment less when quality dispersion is high. Having been shown not to be driven by buyer-specific characteristics, imported volumes or other moments in the quality distribution, the pattern appears counterintuitive. In models of search for differentiated alternatives, dispersion in the underlying quality of alternatives induces *more* experimentation, not *less*. The following section presents a model that rationalizes the counterintuitive relationship between dispersion and experimentation discussed here.

4 Theoretical Framework

This section presents a simple model of a buyer's sequential experimentation with potential suppliers. The framework characterizes the amount of experimentation that a buyer engages in as a function of the cost of testing suppliers and the expected value of forming a relationship with sellers who differ on some ex ante unobservable, which I refer to broadly as quality. The structure of the model builds on the discussions in Sections 2.1 and 3. First, the evidence on buyer' experimentation upon entry into product categories lends the model its structure as a search-theoretic framework. Second, the relatively small trade volume under first interactions leads to the simplification that experimentation yields no revenue to buyers. Third, the fact that large buyers tend to experiment more than their smaller counterparts is reflected by the value of relationships being directly proportional to buyer scale. Fourth, the fact that large buyers experiment less in high-dispersion categories is reflected in the form of a convex experimentation cost function: inadvertently experimenting with a low-quality supplier can have disproportionately large reputational costs for large buyers.

The model delivers a characterization of the two equilibrium outcomes in the buyer's experimentation process: the threshold seller quality that the buyer is willing to accept and the amount of experimentation in which the buyer needs to engage to find a suitable supplier. The framework highlights the countervailing forces behind buyer size and supplier dispersion. On the former, large buyers benefit from positive scale effects but are more exposed to reputational risks than their smaller counterparts. On the latter, environments with high dispersion across suppliers encourage experimentation in hopes of securing high-value relationships but impose higher downside risk than low-dispersion markets. Section 4.1 presents the model, and Section 4.2 discusses a set of comparative statics that highlight the forces described here. This section ends with the derivation of two triple-difference inequalities on the equilibrium outcomes of the model. These compare the experimentation responses of small and large buyers in high- and low-dispersion environments after a shock to experimentation costs. These inequalities are empirically assessed in Section 5.

4.1 Formalization

Quality and relationship value. A buyer of size q enters a product market populated by a continuum of sellers. Manufacturers are heterogeneous in a vertical attribute (quality θ) drawn independently from a distribution $F(\theta;\rho)$ over $[\underline{\theta},\overline{\theta}]$, $\underline{\theta}>0$, with a differentiable PDF $f(\cdot)$; $(1-F)$ is assumed to be log concave and $F(\cdot)$ to be twice continuously differentiable in θ and parameter ρ. When trading recurrently, the pair (q,θ) produces value for the buyer according to a supermodular function $v(q,\theta)$, with strictly positive first partial derivatives.

Experimentation. When starting to source a product, the buyer sequentially searches for a suitable trade partner. To do so, she draws a potential supplier from $F(\cdot)$ and allocates an experimental order. This testing stage yields no profits to the buyer but informs her of the supplier's quality. The cost of experimenting with a seller of quality θ is given by a function $r(\theta,q;\alpha)$, twice continuously differentiable and parametrized by α, to be described momentarily.

Choice of partner. For simplicity, I set the value function to feature multiplicative complementarity in the buyer and seller characteristics, $v(q,\theta)=q\theta$. The search behavior of a buyer whose current best alternative is of standard Θ follows the optimal stopping rule.[32]

$$g(\Theta)=\int_{\Theta}^{\overline{\theta}}(\theta-\Theta)f(\theta;\rho)d\theta=\frac{1}{q}\int_{\underline{\theta}}^{\overline{\theta}}r(\theta,q;\alpha)f(\theta;\rho)d\theta. \tag{1}$$

Using c to denote the unit experimentation cost (i.e., the right-hand side of equation (1)), the $\hat{\Theta}(c)$ that solves the above equation represents the threshold seller quality for a buyer facing search cost c, such that if the current match is of $\theta>\hat{\Theta}(c)$, the buyer does not search for another supplier. Note that the decision to continue experimenting follows from a comparison between the buyer's current best option and the benefits, net of experimentation costs, of drawing an alternative. If such an alternative gives higher payoffs, it is preferred to the existing best option. Otherwise, free recall guarantees that the buyer keeps her existing partner.

Standard derivations show that the $g(\Theta)$ function above is monotonically decreasing and has a unique solution over the interval $[\underline{c},min\{\overline{c},E[\theta]\}]$, where $\overline{c}$ and $\underline{c}$ are the upper and lower bounds of the unit experimentation cost. Taking the first and second derivatives of equation (1) shows that $\hat{\Theta}(c)$ is decreasing and convex in c over the relevant interval and

[32] Θ corresponds to the best option that the buyer has encountered so far. As is standard in formulations of this class of sequential search optimization, the choice problem follows a stationary reservation utility strategy. See Kohn and Shavell (1974) for a derivation and Weitzman (1979) for a more general presentation of the problem.

$\hat{\Theta}(E[\theta]) = \underline{\theta}$ and $\hat{\Theta}(\underline{c}) = \overline{\theta}$.[33] In this context, when the buyer meets supplier s, the probability that she retains s as a supplier is given by $Prob[\theta_s > \hat{\Theta}(c)] = 1 - F(\hat{\Theta}(c))$. Analogously, her probability of returning to the pool is $\Sigma \equiv F(\hat{\Theta}(c))$.

Reputation. The buyer's reputation concerns are captured in reduced form by the shift parameter α and the shape of function $r(\cdot)$. In general terms, for any given α, the cost function is decreasing and convex in θ, is increasing in q and has a negative cross-partial derivative. Thus, drawing a low-quality supplier is particularly costly for large buyers. For intuition, a parametrization of $r(\cdot)$ has $\alpha \in (0, 1]$ as the probability of the experimentation process being scrutinized and $r(\cdot) = \alpha(q/\theta)^2$. Reputational costs are high when the probability of scrutiny is high, when the buyer is large (and either is more visible to consumers or has a reputation to protect) and when the supplier is of low quality.[34]

4.2 Comparative Statics

This section illustrates, by means of two comparative statics exercises, the relevant model mechanics. First, I show how the threshold quality that a buyer is willing to accept varies across markets with different quality dispersion. As more dispersed environments carry higher reputational risk for buyers, this may lower the quality of sellers with whom buyers are willing to settle. This reputation mechanism reverses the otherwise-positive relationship between search and dispersion found in standard formulations. Second, I describe the effects of an increase in the experimentation cost shifter across markets with varying dispersion.

Dispersion in quality. Let ρ denote the parameter that governs the dispersion across θ under $F(\cdot)$, such that $F(\theta; \rho)$ characterizes an environment with lower quality dispersion relative to $F(\theta; \rho')$ whenever $\rho < \rho'$.[35] In Appendix B, I show that the threshold quality at which a buyer is willing to settle decreases with dispersion ρ whenever

$$- \int_{\hat{\Theta}}^{\overline{\theta}} F_\rho(\Theta; \rho) < \int_{\underline{\theta}(\rho)}^{\overline{\theta}(\rho)} \frac{r(\theta, q; \alpha)}{q} f_\rho(\theta; \rho) d\theta, \tag{2}$$

[33] For a given F, $\hat{\Theta}'(c) = \frac{-1}{1-F(\hat{\Theta}(c))} < 0$ and $\hat{\Theta}''(c) = \frac{f(\hat{\Theta}(c))[\hat{\Theta}'(c)]^2}{1-F(\hat{\Theta}(c))} > 0$ (see Moraga-González et al., 2017).

[34] In general, the restrictions are that $r(\theta, q; \alpha)$ satisfies $\partial r(\cdot)/\partial\theta < 0$, $\partial^2 r(\cdot)/\partial\theta^2 < 0$, $\partial r(\cdot)/\partial q > 0$, $\partial^2 r(\cdot)/\partial\theta\partial q > 0$, $\partial^2 r(\cdot)/\partial\theta\partial\alpha > 0$, $\partial^2 r(\cdot)/\partial q\partial\alpha > 0$, $\partial^3 r(\cdot)/\partial\theta\partial q\partial\alpha > 0$.

[35] The formulation here corresponds to mean-preserving spreads of $F(\cdot)$. I focus on comparative statics around changes in dispersion with the mean held constant because the implications for search behavior of differences in average quality across environments are immediate. In Appendix B, ρ is introduced as the parameter of increase in risk from Rothschild and Stiglitz (1970) and Diamond and Stiglitz (1974).

where the subindexes in F_ρ and f_ρ denote the derivatives with respect to the dispersion parameter. The acceptable quality threshold increases otherwise.

The sign rule in equation (2) is intuitive and highlights the departure of this model from standard formulations of sequential search for differentiated alternatives with or without heterogeneous costs (Weitzman, 1979). The left-hand side of the inequality shows a *revenue effect* of quality dispersion on experimentation. As dispersion increases, with the mean held constant, the CDF of the spread-out distribution, $F(\theta;\rho')$, may lie above or below the CDF of the original distribution, $F(\theta;\rho)$, depending on the location Θ.[36] This is represented by the sign of $F_\rho(\Theta;\rho)$. It is clear that $\int_{\hat{\Theta}}^{\overline{\theta}} F_\rho(\Theta;\rho)$ is always smaller than or equal to zero, and so the left-hand side of the inequality is always positive.

In standard search models, the experimentation cost does not depend on the dispersion parameter ρ, so the right-hand side of the inequality is equal to zero. As a result, increases in quality dispersion do not induce decreases in the acceptable supplier threshold due to higher search costs. Whether increases in dispersion lead to more or less experimentation depends on the sign of $F_\rho(\Theta;\rho)$. Buyers whose threshold acceptable partner is located in a region of θ where $F(\theta;\rho')$ runs above $F(\theta;\rho)$ unequivocally increase their search intensity. This is because the probability of drawing low-quality partners is higher. In the example of the normal, this would be the case for any buyer whose threshold supplier is below the mean. On the other hand, buyers for whom $F(\theta;\rho')$ falls below $F(\theta;\rho)$ face a smaller mass of below-standard suppliers. In this case, the total amount of experimentation in an environment characterized by ρ' might be higher or lower than that under ρ.

The right-hand side of (2) shows the *cost effect* of quality dispersion on experimentation. As the cost function $r(\cdot)$ is convex in θ, this expression is positive: increases in dispersion increase the experimentation cost. In turn, the threshold supplier with whom buyers are willing to settle decreases. This effect unequivocally lowers experimentation whenever $F(\theta;\rho')$ falls below $F(\theta;\rho)$.[37]

To summarize, the revenue effect, which is present in standard search models, implies that the threshold partner is weakly higher for all buyers (i.e., all buyers are pickier) in more dispersed than in less dispersed environments. In the presence of nonzero cost effects of dispersion, the acceptable quality threshold may instead drop with quality dispersion. To illustrate the mechanism, Figure 2 exploits a parametrization of $F(\cdot)$ and $r(\cdot)$ to show the relationship between the threshold acceptable supplier, $\hat{\Theta}(c)$, and quality dispersion, ρ,

[36] For instance, in the case of the normal with mean μ, the CDF after a mean-preserving spread lies above the original curve for all $\Theta < \mu$ ($F_\rho > 0$) and below for all $\Theta > \mu$ ($F_\rho < 0$).

[37] Note that in the model, scale favors large buyers: the larger the buyer, the higher are the marginal gains from a better match with a supplier. Equivalently, the larger the buyer, the lower is the per-unit cost of an instance of experimentation. Other things equal, larger buyers have higher acceptance thresholds $\hat{\Theta}(c)$.

when only revenue effects are present and when both revenue and cost effects are at play. Naturally, under the standard model (revenue effects only), as dispersion increases, the buyer sets a higher acceptance threshold. In other words, the partner with whom she is willing to settle is of higher quality. As discussed above, when the cost effect is allowed, it drives the acceptance threshold down. Under the parametrization of Figure 2, the cost effect overrides any revenue effect, and $\hat{\Theta}(c)$ is decreasing over the entire range of ρ.

A shock to reputation costs. I consider an increase in the reputational risk of experimentation. For clarity, I abstract from the role of buyer size in the convexity of reputation costs, i.e., $r(\theta, q; \alpha) = r(\theta; \alpha)$, and examine a shift in α. This simplification allows a thought experiment in which large and small buyers are equally exposed to reputational costs, providing a useful benchmark. From the optimal stopping rule in equation (1) and the definition of α, it is apparent that an increase in this parameter (i) lowers the threshold acceptable supplier, $\hat{\Theta}(c)$, and (ii) does so more markedly for buyers with low scale. Based on the discussion on quality dispersion, these effects ought to be more pronounced in environments where quality dispersion is high.

The derivations presented so far, like those in Appendix B, hold generally for any $F(\cdot)$ and $r(\cdot)$ satisfying the functional restrictions imposed above. I return to a specific parametrization to graphically represent the reputation shock: Figure 3 shows the effect of an increase in the cost shifter α from α_1 to α_2 for a small buyer (characterized by q_1) and a large buyer (characterized by q_2) in a low-dispersion environment (ρ_1, top subfigure) and a high-dispersion environment (ρ_2, bottom subfigure).

Three observations are in order. First, the threshold supplier quality and amount of experimentation drop for small and large buyers in environments with both low and high dispersion. Using $\Delta_\alpha \Theta_{b,m}$ to indicate the change in acceptable supplier quality for a buyer of size $b \in \{q_1, q_2\}$ with $q_1 < q_2$ in an environment of dispersion $m \in \{\rho_1, \rho_2\}$ with $\rho_1 < \rho_2$ after an increase in α, the above can be stated as $\Delta_\alpha \Theta_{b,m} < 0, \forall b, m$. Second, for a buyer of any size, the threshold supplier quality at which the buyer is willing to settle reacts more markedly in high-dispersion environments. More compactly, $|\Delta_\alpha \Theta_{q_1,\rho_1}| < |\Delta_\alpha \Theta_{q_1,\rho_2}|$ and $|\Delta_\alpha \Theta_{q_2,\rho_1}| < |\Delta_\alpha \Theta_{q_2,\rho_2}|$. Third, in any environment, the mitigating effects of buyer scale imply that small buyers see a larger downward shift in their acceptable quality threshold.[38] That is, $|\Delta_\alpha \Theta_{q_1,\rho_1}| > |\Delta_\alpha \Theta_{q_2,\rho_1}|$ and $|\Delta_\alpha \Theta_{q_1,\rho_2}| > |\Delta_\alpha \Theta_{q_2,\rho_2}|$. Taken together, these observations lead

[38] Note that scale enters the problem in direct and indirect ways. In equation (1), any cost shock is scaled by q (the direct effect). In addition, Θ is decreasing and convex in the experimentation cost, and so small buyers, who have high unit experimentation costs, have more sensitive thresholds (the indirect effect).

to the triple-difference statement

$$[\Delta_\alpha \Theta_{q_1,\rho_1} - \Delta_\alpha \Theta_{q_2,\rho_1}] - [\Delta_\alpha \Theta_{q_1,\rho_2} - \Delta_\alpha \Theta_{q_2,\rho_2}] \geq 0.$$

In the absence of reputational costs disproportionately affecting large buyers, $r(\theta, q; \alpha) = r(\theta; \alpha)$, the triple-difference above is positive. This means that the threshold supplier with whom buyers are willing to settle drops more markedly for small buyers in high-dispersion than in low-dispersion environments.

Conversely, if large buyers are particularly exposed to reputational costs, the inequality reverses to

$$[\Delta_\alpha \Theta_{q_1,\rho_1} - \Delta_\alpha \Theta_{q_2,\rho_1}] - [\Delta_\alpha \Theta_{q_1,\rho_2} - \Delta_\alpha \Theta_{q_2,\rho_2}] < 0. \tag{3}$$

In prose, a shock to reputational costs decreases the threshold supplier with whom large buyers, relative to small buyers, are willing to settle more in high-dispersion environments than in homogeneous ones. Using Σ to denote the amount of experimentation a buyer engages in, if such costs are sufficiently high,

$$[\Delta_\alpha \Sigma_{q_1,\rho_1} - \Delta_\alpha \Sigma_{q_2,\rho_1}] - [\Delta_\alpha \Sigma_{q_1,\rho_2} - \Delta_\alpha \Sigma_{q_2,\rho_2}] < 0. \tag{4}$$

Compared to small buyers, large buyers are less likely to reject a draw in high-dispersion environments (i.e., they experiment *less*) than in less-dispersed ones.

In the next section, I bring the triple-difference relationships in inequalities (3) and (4) to the data.

5 Empirical Analysis

In this section, I exploit quasi-experimental variation in trade patterns following the 2013 Rana Plaza (RP) collapse to empirically assess the triple-difference relationships derived in Section 4. Section 5.1 starts by describing the shock exploited for identification. In brief, the collapse directed international scrutiny toward the sourcing decisions of international brands. The empirical strategy leverages this unanticipated shock to identify changes in experimentation after the incident among large relative to small buyers in high- in comparison to low-dispersion product categories. The results align with the model predictions: when reputational costs increase, large buyers reduce their experimentation, particularly in product categories with high dispersion. The section concludes with a brief discussion of key robustness exercises, presented in full in Appendix D. Finally, Section 5.2 discusses alternative mechanisms that could drive the results.

5.1 A Shock to the Reputational Cost of Experimentation

The RP collapse on 24 April 2013 lends itself to interpretation through the lenses of the model in Section 4. The episode exacted a toll of over 1,200 garment worker fatalities and some 4,000 reported injuries. Intense media coverage ensued, and several activist campaigns of international reach put the names of brands at the center of the debate on responsible global supply chains. Consumer boycotts followed, targeting specific brands believed to have purchased garments, either directly or indirectly, from RP factories. The aftermath of the event saw a rise in public awareness of social compliance issues and a renewed consumer willingness to penalize potentially offending brands by moving away from their products. The potential cost to a buyer of being discovered trading with a noncompliant supplier increased dramatically.

Before turning to the main triple-difference exercise, it is informative to explore key industry trends before and after the shock. Table E6 studies the evolution of product–year and buyer–product–year outcomes through the sample period, showing that the volume of garments exported from Bangladesh continued to grow after RP. The total number of available suppliers and active buyers continued to increase, following the precollapse trend. Within buyer–product combinations, the total traded volumes and number of active relationships followed the positive trend of previous years. However, the share of traded volumes that buyers allocated to one-off interactions decreased markedly after RP. Altogether, these patterns suggest that while RP did not appear to shift general industry trends, it was accompanied by a drop in one-off trade interactions.

Empirical strategy. I study whether large buyers' increased reputational costs and exposure after RP led to observable changes in their experimentation. I focus on instances of buyer entry into sourcing product markets, where one-off interactions are most likely related to experimentation. By definition, buyers enter a particular market only once. Thus, the exercises here compare buyer experimentation upon market entry prior to RP with entry instances into other products after RP. The baseline specification is as follows:

$$\begin{aligned} Experiment_{bj}^{entry} = \delta_b + \delta_{jc(bj)} + \beta_1 Post_{c(bj)} \times Large_b + \beta_2 Dispersion_j \times Large_b \\ + \beta_3 Dispersion_j \times Post_{c(bj)} \times Large_b + \epsilon_{bj}, \quad (5) \end{aligned}$$

where $Experiment_{bj}^{entry}$ corresponds to the (log) count of buyer b's one-off interactions in the first year after entering product category j. Two sets of fixed effects, following the richest specification discussed in Section 3.2, condition on important drivers of this outcome. δ_b is a buyer-specific fixed effect capturing relevant buyer characteristics common to all products that the buyer enters. These include any centralized screening technology to which the buyer may have access, buyer visibility to consumers and idiosyncratic preferences over trade partners. $\delta_{jc(bj)}$ is a product–time fixed effect, where time corresponds to the quarter of buyer b's entry into product j (again, $c(\cdot)$ stands for cohort). This term controls for any seasonal and product-specific demand or supply shifters common to all buyers. In particular, it captures both the overall effect of RP across all products categories and product-specific departures from these effects. This term would account, for instance, for all buyers experimenting less (or more) aggressively in product categories with high supplier heterogeneity after the shock.[39] $Post_{c(bj)}$ is an indicator of calendar quarters after RP, and $Large_b$ and $Dispersion_j$ are defined as discussed above.[40] In this specification, β_3 is the empirical counterpart of the triple-difference relationship in equation (4).

Results. Table 6 presents the estimation results for equation (5) and shows that relative to small buyers, large buyers decreased their experimentation after RP when entering high-dispersion environments ($\hat{\beta}_3 < 0$). In particular, a large buyer entering the most dispersed categories after the collapse would see a drop in its experimentation activity between 11% and 15.5% greater than that of smaller buyers. These results hold when I consider both the discretized (columns (1) and (2)) and the continuous (columns (3) and (4)) metrics of dispersion and control for the total volume that the buyer purchases in its first year in the product category (columns (2) and (4)).[41] On average, prior to RP, the amount of experimentation upon entry among large buyers relative to that among small buyers appears not to differ with the underlying heterogeneity across suppliers (i.e., $\hat{\beta}_2 \approx 0$). Large buyers' response after RP does not seem to have differed from that of small buyers in low-dispersion environments, but it did in categories with high heterogeneity, where large buyers reduced their experimentation activity significantly.

Table 7 and Figure 4 present a set of exercises showing that the estimation relies on

[39] A saturated specification with buyer–product and buyer–time fixed effects is not possible. By definition, buyers enter product categories only once, such that for each buyer–product combination there is only one observation in the restricted sample used to estimate equation (5). In addition, the time dimension used in this exercise $c(bj)$ is very disaggregated to account for rich forms of seasonality. Buyers typically do not enter more than one product category in a calendar quarter, and so including buyer–time effects is not possible.

[40] Once more, for ease of interpretation, an alternative recovers vectors $\{\boldsymbol{\beta}_2, \boldsymbol{\beta}_3\}$, using the quartiles of dispersion $\{Quartile\ Dispersion_j = i\}$ for $i = 2, 3, 4$ and $i = 1$ excluded as the base category.

[41] Note that the first row of Table 6, showing coefficient $\hat{\beta}_1$, is fully consistent across specifications. The interpretation of this coefficient in columns (1) and (2) is that large buyers seem not to have experimented more (or less) than small buyers after RP on average in low-dispersion environments (first quartile). The quantitative interpretation corresponds to percentage changes in the outcome after unit changes in the interactions. In columns (3) and (4), the baseline corresponds to environments whose continuous metric of dispersion is equal to zero, and β_3 is interpreted as an elasticity.

suitable identifying variation: (i) the results are not driven by unbalancedness in the data, (ii) the main effects do not follow from differential pretrends, and (iii) placebo shocks produce no significant effect. I discuss each of these tests in turn. First, I identify the coefficient of interest, β_3, by comparing buyer entries in products with different dispersion across sellers before and after RP. Column (1) of Table 7 reproduces the baseline result (of column (3) of Table 6) but restricts the sample to only buyers classified as *incumbents*, i.e., those with at least one entry before and at least one entry after RP.[42] Columns (2) to (5) further restrict the sample to incumbent buyers with at least x entry instances after RP, with $x = \{4, 5, 6, 7\}$.[43] Across all of these exercises, the coefficient of interest remains stable between -0.163 and -0.250. Second, columns (6) and (7) augment the baseline specification to allow for pretrends in the time-varying interactions. Reassuringly, there appears to be no significant trend in large buyers' experimentation in markets with different dispersion before RP. Third, and to further assuage any concerns over pretrends, Figure 4 presents the estimates of the main coefficient of interest from 400 placebo regressions, drawing random dates to partition the pre- and postperiods, using the data from before 2013. The average across these estimates is -0.003; all estimates fall within $[-0.054, 0.068]$ and cannot reject the null of nonsignificance at 10%.

Robustness. Appendix D.1 details several robustness exercises. In particular, qualitatively and quantitatively equivalent effects of RP are obtained when a nonlinear count data model is used, when the dispersion measure is constructed with metrics other than the standard deviation, when the buyer size cutoff is shifted, when an alternative definition for the outcome variable is used, and when other moments of the distribution of seller heterogeneity are controlled for. The estimated coefficients on the triple-interaction term across 78 robustness specifications are presented in Figure 5. In the case of linear models, all estimated coefficients fall in the interval $[-0.196, -0.072]$, comparable with the findings in non-linear alternatives. Importantly, 72 out of the 78 specifications (95% of the estimates) give a coefficient on the triple interaction that is significantly different from zero (and negative).

In addition to the sensitivity analysis in Figure 5, Appendix D presents alternative estimations addressing two important conceptual issues, discussed below.

First, transitioning from the model to the empirical analysis, it is necessary to define the pool of sellers from which the buyer draws potential partners. The analysis so far assumed

[42] The breakdown of all entry instances in the data by buyer status is summarized in Table E7. Note that nonincumbent buyers do not contribute directly to identification of β_3, but they provide variation for estimation of the fixed effects and β_2.

[43] As there are effectively only fewer than two years of usable data after RP, restricting the sample further on survival reduces the sample size and power significantly. The point estimate remains very similar.

that all buyers face the same environment at all times when entering a given product category. This formulation may represent the buyer's search problem poorly if dispersion varies over time and across buyers. Such a setting may be relevant if search is directed or if a buyer-specific screening technology reduces the consideration set for certain buyers. To address this concern, I construct hypothetical consideration sets for each buyer, based on the information available in the data. Appendix D.2 explains the construction of these sets. Table D4 reproduces the main results from this exercise (see column (3) of Table 6) and compares them to alternative dispersion measures, using iteratively more restrictive consideration sets. Navigating Table D4 from columns (1) to (5) shows that the key results remain qualitatively unchanged across specifications. The point estimate on the triple interaction is still negative, large and significant. Naturally, its magnitude drops under more stringent consideration sets, as dispersion (the right-hand-side variable) mechanically decreases. Still, always significantly negative, β_3 is in the range of $[-0.09, -0.067]$ when estimated on all buyers and within the interval $[-0.104, -0.069]$ when I focus on incumbent buyers.

The second conceptual issue pertains to the use of active or observed sellers to measure dispersion. The dispersion that the buyer faces is a function of seller-specific shifters, computed using market shares that sum over all the buyers with whom a given seller trades (in a product category, time and destination combination). A buyer may be particularly influential in determining a seller's shifter, and this influence, in turn, may shape the observed dispersion in the market. To mitigate this concern, I recompute the dispersion that each buyer faces when entering a product category, excluding from the computation all sellers with whom a buyer has ever traded. The results based on this approach are presented in column (6) in Table D4. Reassuringly, the coefficient of interest remains quantitatively large, negative and significant.[44]

Despite the conservative strategy above, a remaining concern is that low-quality sellers may refrain from serving product markets in which frictions are low and competition is high. They may, instead, trade in niches in which search frictions are high, reinforcing the thickness of the observed lower tail of seller quality distribution. In the structure of equation (5), this mechanism would bias the results discussed so far as long as dispersion (a function of seller shifter estimates) is correlated with unobservable drivers of a buyer's experimentation in a product category beyond buyer and product–time fixed effects. To address this concern, in Appendix D.1, I instrument for dispersion in the environment that the buyer faces by using exogenous characteristics varying across product categories. I leverage the correlations

[44] The point estimate is slightly lower in magnitude than the baseline estimate (-0.088 and -0.096 for all buyers and incumbents only, respectively). This follows from the fact that in general, the larger the buyer, the more trade partners it has and the more pronounced is the trimming of the choice set that this approach imposes.

presented in Table C3, showing that product categories with higher quality dispersion across sellers typically correspond to products that are made of materials other than cotton, that are produced for women, that require the combination of more inputs and that correlate with quality ladder length in a developed downstream market. Column (4) of Table D1 presents the results of this IV estimation of the baseline equation, showing a quantitatively stronger result, with $\hat{\beta}_3 = -0.229$.

5.2 Discussion

The previous subsection shows that after RP, large buyers disproportionately reduced their experimentation upon entry in high-dispersion environments. This is interpreted to follow from increases in their reputational costs of experimenting with low-quality suppliers. These costs are relevant only when buyers are exposed to public scrutiny (i.e., when they are large and visible) and when there is sufficient heterogeneity across suppliers (i.e., in high-dispersion environments). This subsection discusses alternative mechanisms and presents complementary results. In brief, it shows that the relative decrease in large buyers' experimentation in high-dispersion environments after RP (i) did not follow from shifts in preferences, (ii) cannot be accounted for by differential learning profiles of large buyers, and (iii) was not accompanied through deployment of successful substitute screening technologies. This section concludes by showing that on the contrary, the observed decline in large buyers' experimentation was accompanied by a lowering of their supplier quality threshold.

A first alternative explanation for the patterns found so far is that the RP shock changed buyers' tastes for or perceptions of the value of relationships in high-dispersion categories. Under this hypothesis, the RP collapse should be observed to have affected not only experimentation but also buyer trade volumes. Using the structure of previous specifications, I decompose the quarterly (and alternatively, yearly) volumes following

$$
\begin{aligned}
q_{bjt} = \delta_b + \delta_{jt} + \pi_1 Post_t \times Large_b + \pi_2 Dispersion_j \times Large_b \\
+ \pi_3 Dispersion_j \times Post_t \times Large_b + \epsilon_{bjt}, \quad (6)
\end{aligned}
$$

where q_{bjt} is the (log) volume that buyer b imports of product j at time t. All other regressors are as described in earlier sections, and $Post_t$ is a dummy taking value one for quarters after RP. I estimate equation (6) in the entire sample of buyer–product–quarter combinations and a restricted sample of triplets falling within the first two years of buyer entry into the product category. In alternative specifications, I aggregate volumes at the year level (q_{bjy}) and consider only volumes channeled through recurrent relationships.[45] The results are collected in Table 8. Large buyers tend to have smaller trade volumes in high-dispersion categories than other buyers ($\pi_2 < 0$). This is consistent with the fact that some of these buyers are key importers of fairly homogeneous, basic products, such as men's plain polo shirts. Relative to other buyers, large buyers reduced their volumes after RP ($\pi_1 < 0$) but did so to a lesser extent in high-dispersion categories ($\pi_3 > 0$). This evidence rejects the hypothesis that the lower experimentation of large buyers in high-dispersion categories after RP followed from a decrease in their demand for these products. The results of different specifications with flexible pretrends, different sample trimmings and different levels of aggregation of traded volumes point in the same direction.

A second competing mechanism involves buyers learning how to choose better trade partners as they enter more product categories. In particular, entering the first market may require significant buyer experimentation, and the knowledge gained from this experience may reduce the need for experimentation in subsequent entries.[46] To study whether different learning profiles across buyers and markets induce the results discussed in the previous subsection, I study how experimentation evolves along a buyer's sequence of entry instances. The specification follows the structure of equation (5) and includes interactions with $Age\ Trend_{n(bj)}$, a linear "age" variable that counts the order of a buyer's entries into different markets. $Age\ Trend_{n(bj)}$=1 if j is the first product that the buyer is observed entering, =2 if it is the second, etc. The results are presented in Table 9. Column (1) shows that experimentation, on average, tends to increase slightly as the buyer ages (i.e., enters subsequent product markets).[47] Large buyers do not exhibit a different trend in environments with any level of dispersion. Moreover, when I condition on the learning effects discussed here (see columns (2) and (3) of Table 9), the effects of the RP shock remain negative, significant and quantitatively similar to those in the baseline specification.

Finally, it is plausible that large buyers have easier access to screening technologies that are suitable substitutes for experimentation. In this case, a decrease in experimentation may not lead to changes in the formation of recurrent relationships. To explore this idea, I use the

[45] At the quarterly level, the incidence of observations with nonzero trade but zero trade with recurrent partners is high. Yearly aggregation mitigates this issue.

[46] The analysis of experimentation upon entry, discussed in the previous subsection, always conditions on unobserved buyer and product–time heterogeneity. This specification implicitly controls for the buyer's cohort in the industry but does not account for the learning or age effect discussed here.

[47] This is consistent with buyers entering higher-dispersion categories later in their tenure in the industry. Table E8 regresses the dispersion in the market across entry instances on the age trend and buyer and time fixed effects. It shows that all buyers tend to enter more dispersed environments at later stages of their tenure and that this pattern is slightly more pronounced for large buyers. When interacted with an indicator for the RP collapse, these trends do not add up to differential selection into products by large buyers after the event.

estimated quality measures for each seller to study whether the average quality of recurrent partners remained stable for large and small buyers after RP. Specifically, I estimate

$$\begin{aligned} Average\ Quality_{bjy} = \delta_b + \delta_{jy} + \phi_1 Post_y \times Large_b + \phi_2 Dispersion_j \times Large_b \\ + \phi_3 Dispersion_j \times Post_y \times Large_b + \epsilon_{bjy}. \end{aligned} \quad (7)$$

The left-hand side, $Average\ Quality_{bjy}$, is the average quality across the sellers with whom buyer b trades recurrently in product j and year y. The right-hand side closely follows the structure of the estimated equations already discussed.[48] Table 10 reports the results of this estimation. While match quality tends to improve over time and after RP for large buyers ($\phi_1 > 0$) and large buyers seek high-quality partners in differentiated categories ($\phi_2 > 0$), the triple interaction shows a relative worsening of large buyers' match quality relative to that of small buyers in high-dispersion environments after RP ($\phi_3 < 0$). This result suggests that large buyers' disproportionate decrease in experimentation documented in the previous subsection does not seem attributable to their deployment of alternative screening technologies, particularly in high-dispersion categories.

An important final observation on the last exercise presented here is that ϕ_3 is a close empirical counterpart of the triple-difference relationship in equation (3). The theoretical relationship pertains to buyers' endogenously determined threshold supplier quality. While this object is not observed in the data, under the model, the quality of recurrent suppliers must be above this threshold, as must the average across them. The result here lends evidence to the conclusion that the relative decline in large buyers' experimentation in high-dispersion environments was accompanied by a lowering of their quality threshold.

6 Conclusion

This paper shows that one-off trade relationships occur when buyers cannot discern supplier quality until trade actually takes place. Low levels of experimentation in highly differentiated product categories arise as large buyers, exposed to public scrutiny, are discouraged from searching by the downside risk of experimentation. As recurrent relationships are formed only after the experimentation phase, buyers make decisions on whom to trade recurrently with under full information. However, before forming such relationships, they must decide on how much to experiment and what their acceptable quality threshold is before uncertainty can be resolved. This leads to under-experimentation and weakly worse matches between

[48] Note here that as the outcome is aggregated to year level, the postshock variable is set to match this aggregation: $Post_y$ is defined to take value one for years from 2013 onward.

sellers and reputation-sensitive buyers in categories with high dispersion in supplier quality. The Rana Plaza collapse in 2013, interpreted as an exogenous shock to buyers' reputation concerns, is leveraged to identify the mechanism proposed here. After the collapse, large buyers experimented less when entering new product categories—and more markedly so in highly differentiated categories.

Through its empirical application, this paper speaks to a fast-growing literature connecting consumer social and environmental awareness to the organization of global supply chains (Alfaro-Ureña et al., 2021; Dragusanu and Nunn, 2020; Hainmueller et al., 2015; Hart and Zingales, 2017). Harrison and Scorse (2010) study a wage increase in Indonesia in the wake of antisweatshop boycotts targeting multinationals sourcing apparel and footwear from the country. Koenig and Poncet (2020) analyze the trade responses to the Rana Plaza collapse, looking at garment imports into France. Also close to the empirical context of this paper is the setting of Boudreau (2020), who designs an intervention to introduce workers' safety committees in a sample of Bangladeshi garment firms, all of them suppliers of large multinationals. Amengual and Distelhorst (2019) study the penalties and provisions in the code of conduct of one such multinational. While these works focus on the role of social responsibility provisions in relationship- or firm-level outcomes with existing partners, this paper offers a simple framework to link reputation concerns over social compliance to the process of forming new relationships. Thus, this paper's findings may contribute to ongoing policy debates on multinational-driven initiatives aimed at fostering sustainable private sector development in low-income countries.

References

Albornoz, Facundo, Héctor F Calvo Pardo, Gregory Corcos, and Emanuel Ornelas, "Sequential Exporting," *Journal of International Economics*, August 2012, *88* (1), 17–31.

Alfaro-Ureña, Alonso, Benjamin Faber, Cecile Gaubert, Isabela Manelici, and José Vasquez, "Responsible Sourcing? Theory and Evidence from Costa Rica," *mimeo*, July 2021.

Allen, Treb, "Information Frictions in Trade," *Econometrica*, November 2014, *82* (6), 2041–2083.

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

Antràs, Pol and Davin Chor, "Global Value Chains," *mimeo*, 2021.

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

Armstrong, Mark, "Ordered Consumer Search," *Journal of the European Economic Association*, 06 2017, *15* (5), 989–1024.

Atkin, David and Amit K. Khandelwal, "How Distortions Alter the Impacts of International Trade in Developing Countries," *Annual Review of Economics*, 2020, *12* (1), 213–238.

Bai, Jie, Ludovica Gazze Gazze, and Yukun Wang, "Collective Reputation in Trade: Evidence from the Chinese Dairy Industry," *Review of Economics and Statistics*, Forthcoming.

Banerjee, Abhijit and Esther Duflo, "Reputation effects and the limits of contracting: a study of the Indian software industry," *Quarterly Journal of Economics*, 2000, *115*, 989 to 1017.

Benguria, Felipe, "The matching and sorting of exporting and importing firms: Theory and evidence," *Journal of International Economics*, 2021, *131*, 103430.

Bernard, Andrew B., Esther Ann Boler, Renzo Massari, Jose-Daniel Reyes, and Daria Taglioni, "Exporter Dynamics and Partial-Year Effects," *American Economic Review*, October 2017, *107* (10), 3211–28.

Berry, Steven T, "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 1994, *25* (2), 242–262.

Better Buying Institute, "2020 Better Buying Index Report," 2018.

— , "2020 Better Buying Index Report," 2020.

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

Brown, Drusilla K., Alan Deardorff, and Robert Stern, *The Effects of Multinational Production on Wages and Working Conditions in Developing Countries*, University of Chicago Press, February

Brugués, Felipe, " Take the Goods and Run: Contracting Frictions and Market Power in Supply Chains," *mimeo*, 2021.

Cajal-Grossi, J., R. Macchiavello, and G. Noguera, "International Buyers' Sourcing and Suppliers' Markups in Bangladeshi Garments," *mimeo*, January 2022.

Caro, Felipe, Leonard Lane, and Anna Saenz de Tejada Cuenca, "Can Brands Claim Ignorance? Unauthorized Subcontracting in Apparel Supply Chains," *Management Science*, 2021, *67* (4), 2010–2028.

Carugati, A., R. Liao, and P. Smith, "Speed-to-fashion: Managing global supply chain in Zara," *2008 4th IEEE International Conference on Management of Innovation and Technology*, 2008, pp. 1494–1499.

Correia, Sergio, Paulo Guimarães, and Thomas Zylkin, "Verifying the existence of maximum likelihood estimates for generalized linear models," 2019.

_ , _ , and _ , "Fast Poisson estimation with high-dimensional fixed effects," 2020.

Dawar, Niraj and Madan M. Pillutla, "Impact of Product-Harm Crises on Brand Equity: The Moderating Role of Consumer Expectations," *Journal of Marketing Research*, 2000, *37* (2), 215–226.

Diamond, Peter A and Joseph E Stiglitz, "Increases in risk and in risk aversion," *Journal of Economic Theory*, 1974, *8* (3), 337 – 360.

Dragusanu, Raluca and Nathan Nunn, "The Effects of Fair Trade Certification: Evidence From Coffee Producers in Costa Rica," 2020.

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

Evans, Carolyn L and James Harrigan, "Distance, Time, and Specialization: Lean Retailing in General Equilibrium," *The American economic review*, 2005, *95* (1), 292–313.

Financial Times, "Fashion: A better business model," 2014.

Forbes, "Zara leads in fast fashion," 2015.

Freedman, Seth, Melissa Kearney, and Mara Lederman, "Product Recalls, Imperfect Information, and Spillover Effects: Lessons from the Consumer Response to the 2007 Toy Recalls," *The Review of Economics and Statistics*, May 2012, *94* (2), 499–516.

Geishecker, Ingo, Philipp J. H. Schröder, and Allan Srensen, "One-off export events," *Canadian Journal of Economics*, February 2019, *52* (1), 93–131.

Grossman, Gene M. and Elhanan Helpman, "Outsourcing in a Global Economy," *Review of Economic Studies*, 2005, *72* (1), 135–159.

Guardian, The, "Bangladeshi factory deaths spark action among high-street clothing chains," 2013.

Hainmueller, Jens, Michael J. Hiscox, and Sandra Sequeira, "Consumer Demand for Fair Trade: Evidence from a Multistore Field Experiment," *The Review of Economics and Statistics*, 2015, *97* (2), 242–256.

Hallak, Juan Carlos and Peter K. Schott, "Estimating Cross-Country Differences in Product Quality," *The Quarterly Journal of Economics*, 2011, *126* (1), 417–474.

Harrison, Ann and Jason Scorse, "Multinationals and Anti-sweatshop Activism," *American Economic Review*, March 2010, *100* (1), 247–73.

Hart, Oliver and Luigi Zingales, "Companies Should Maximize Shareholder Welfare Not Market Value," *Journal of Law, Finance, and Accounting*, 2017, *2* (2), 247–275.

Heerde, Harald Van, Kristiaan Helsen, and Marnik G. Dekimpe, "The Impact of a Product-Harm Crisis on Marketing Effectiveness," *Marketing Science*, 2007, *26* (2), 230–245.

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

Hummels, David and Peter J. Klenow, "The Variety and Quality of a Nation's Exports," *American Economic Review*, June 2005, *95* (3), 704–723.

III, Wallace N. Davidson and Dan L. Worrell, "Research notes and communications: The effect of product recall announcements on shareholder wealth," *Strategic Management Journal*, 1992, *13* (6), 467–473.

Inditex, "Online Document 1," https://www.inditex.com/documents/10279/241088/PTR-2021.pdf/7a827080-2d64-2305-72ef-c2e5835e630e 2021. Online; accessed 10 September 2021.

_ , "Online Document 2," https://www.inditex.com/documents/10279/241035/Inditex+Code+of+Conduct+for+Manufacturers+and+Suppliers/e23dde6a-4b0e-4e16-a2aa-68911d3032e7 2021. Online; accessed 10 September 2021.

_ , "Online Document 3," https://www.inditex.com/documents/10279/241148/Safety+Product+Policy_Inditex/e5ddb37f-ca6e-42e5-1d70-579bd87913b4 2021. Online; accessed 10 September 2021.

_ , "Online Document 4," https://www.inditex.com/documents/10279/241097/Health+Product+Policy_Inditex/c2a62984-0d6b-3e74-dcb8-8f6b274ea2c5 2021. Online; accessed 10 September 2021.

Jacobs, Brian W. and Vinod R. Singhal, "The effect of the Rana Plaza disaster on shareholder wealth of retailers: Implications for sourcing strategies and supply chain governance," *Journal of Operations Management*, 2017, *49-51*, 52–66. Competitive Manufacturing in a High-Cost Environment.

Jensen, Robert, "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector," *The Quarterly Journal of Economics*, 2007, *122* (3), 879–924.

Khandelwal, Amit, "The Long and Short (of) Quality Ladders," *Review of Economic Studies*, August 2010, *77* (4), 1450–1476.

Koenig, Pamina and Sandra Poncet, "Social responsibility scandals and trade," *World Development*, 2019, *124*, 104640.

_ **and** _ , "The effects of the Rana Plaza collapse on the sourcing choices of French importers," *mimeo*, 2020.

Kohn, Meir and Steven Shavell, "The theory of search," *Journal of Economic Theory*, 1974, *9*, 93–123.

Kremer, Michael, "The O-Ring Theory of Economic Development," *The Quarterly Journal of Economics*, 1993, *108* (3), 551–575.

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

_ , "Relational Contracts and Development," *mimeo*, August 2021.

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

Martin, Julien, Isabelle Mejean, and Mathieu Parenti, "Relationship stickiness, international trade, and economic uncertainty," *CEPR Discussion Papers No. 15609*, 2020.

McKinsey Company, "Apparel, Fashion & Luxury Practice: Bangladesh's Ready-Made Garments Landscape; The Challenge of Growth," Technical Report, McKinsey and Company, https://www.mckinsey.com/ November 2011.

_ , "What's next for Bangladesh's garment industry, after a decade of growth?," Technical Report, McKinsey and Company, https://www.mckinsey.com/ March 2021.

McMillan, J. and C. Woodruff, "Interfirm relations and informal credit in Vietnam," *The Quarterly Journal of Economics*, 1999, *114* (4), 1285 to 1320.

Monarch, Ryan and Tim Schmidt-Eisenlohr, "Longevity and the Value of Trade Relationships," *mimeo*, 2020.

Moraga-González, José Luis, Zsolt Sándor, and Matthijs R Wildenbeest, "Prices and heterogeneous search costs," *The RAND Journal of Economics*, 2017, *48* (1), 125–146.

Nguyen, Daniel X., "Demand uncertainty: Exporting delays and exporting failures," *Journal of International Economics*, 2012, *86* (2), 336–344.

Pisch, Frank, "Managing Global Production: Theory and Evidence from Just-in-Time Supply Chains," *mimeo*, November 2020.

Plambeck, Erica L. and Terry Taylor, "Supplier evasion of a buyer's audit: implications for motivating supplier social and environmental responsibility," *Manufacturing & Service Operations Management*, 2016, *18* (2), 184+.

Rauch, James E. and Joel Watson, "Starting small in an unfamiliar environment," *International Journal of Industrial Organization*, September 2003, *21* (7), 1021–1042.

Rob, Rafael and Arthur Fishman, "Is Bigger Better? Customer Base Expansion through Word-of-Mouth Reputation," *Journal of Political Economy*, 2005, *113* (5), 1146–1162.

Rogerson, Richard, Robert Shimer, and Randall Wright, "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, December 2005, *43* (4), 959–988.

Rothschild, Michael and Joseph E Stiglitz, "Increasing risk: I. A definition," *Journal of Economic Theory*, 1970, *2* (3), 225 – 243.

Schott, Peter K., "Across-Product Versus Within-Product Specialization in International Trade," *The Quarterly Journal of Economics*, 2004, *119* (2), 647–678.

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

Steinwender, Claudia, "Real Effects of Information Frictions: When the States and the Kingdom Became United," *American Economic Review*, March 2018, *108* (3), 657–96.

Sugita, Yoichi, Kensuke Teshima, and Enrique Seira, "Assortative Matching of Exporters and Importers," *Review of Economics and Statistics*, Forthcoming.

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.

Tewari, Meenu, "Successful Adjustment in Indian Industry: the Case of Ludhiana's Woolen Knitwear Cluster," *World Development*, September 1999, *27* (9), 1651–1671.

Verhoogen, Eric, "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector," *The Quarterly Journal of Economics*, 2008, *123* (2), 489–530.

Watson, Joel, "Starting Small and Renegotiation," *Journal of Economic Theory*, 1999, *85* (1), 52–90.

Weitzman, Martin L., "Optimal Search for the Best Alternative," *Econometrica*, 1979, *47* (3), 641–654.

Tables and Figures

Table 1: Summary Statistics

	Count of Cells	Mean	Std. Dev.	P10	P25	P50	P75	P90
Panel A: Sellers (s), seller-years (sy), seller-products (sj), seller-year-products (sjy)								
$Count_s^j$	7,925	10.5	12.2	1	2	6	15	27
$Count_s^b$	7,925	18.8	25.3	1	2	8	27	51
$Count_s^t$	7,925	14.4	13.2	1	3	9	24	37
$Count_{sy}^j$	37,308	5.09	5.06	1	2	3	7	11
$Count_{sy}^b$	37,308	7.02	7.19	1	2	5	10	16
$Count_{sj}^b$	83,053	3.67	8.01	1	1	1	3	7
$Count_{sj}^t$	83,053	4.29	6.60	1	1	2	4	11
$Count_{sjy}^b$	190,047	2.48	3.40	1	1	1	2	5
Panel B: Buyers (b), buyer-years (by), buyer-products (bj), buyer-year-products (bjy)								
$Count_b^j$	10,149	9.21	14.5	1	2	4	10	22
$Count_b^s$	10,149	14.7	36.6	1	2	4	12	33
$Count_b^t$	10,149	11.6	11.4	1	3	7	17	30
$Count_b^t$ (uncensored)	4,856	5.61	5.25	1	2	4	8	13
$Count_{by}^j$	46,270	4.91	7.68	1	1	2	5	11
$Count_{by}^j$	46,270	5.66	11.6	1	1	2	5	12
$Count_{bj}^s$	93,465	3.26	8.02	1	1	1	3	6
$Count_{bj}^s$ (one off)	60,275	1.92	3.53	1	1	1	2	3
$Count_{bj}^t$	93,465	4.39	6.74	1	1	2	4	11
$Count_{bj}^t$ (uncensored)	63,659	2.40	2.82	1	1	1	3	5
$Count_{bjy}^s$	227,241	2.07	3.22	1	1	1	2	4
Panel C: Seller-buyers (sb), seller-buyer-products (sbj)								
$Count_{sb}^j$	149,298	2.04	2.62	1	1	1	2	4
$Count_{sb}^t$	149,298	3.11	4.37	1	1	2	3	7
$Count_{sbj}^t$	304,723	2.39	3.35	1	1	1	2	5
$Count_{sbj}^t$ (uncensored)	240,293	1.94	2.18	1	1	1	2	4
$Count_{sbj}^t$ (uncensored, count > 1)	84,279	3.69	2.98	2	2	3	4	7

Super- and subscripts are as follows: b corresponds to buyers, s to sellers, j to HS6 product categories, y to years and t to quarters. $Count_y^x$ is the number of x per y. For example, $Count_{sjy}^b$ is the number of buyers per seller–product–year combination. The note "uncensored" corresponds to a trimming of the sample that drops cells whose first and/or last instance is censored, i.e., observed in the first (2005) or last (2015) year of the data. Thus, $Count_{sbj}^t$ (uncensored) corresponds to the count of quarters in the sample of buyer–seller–product triplets whose first trade interaction was in or after 2006 and whose last interaction was in or before 2014. The note "one off" refers to one-off relationships: $Count_{bj}^s$ is the count of sellers with whom the buyer trades as a one-off interaction (an interaction lasting at most one calendar quarter).

Table 2: Probability of One-Off Interactions

	Probability of $one-off_{bt}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$Entry^{j}_{bt}$	0.356***	0.398***	0.395***	0.394***	0.398***	0.386***
	(0.005)	(0.007)	(0.008)	(0.008)	(0.008)	(0.008)
$\overline{Capacity}^{s}_{bt}$				0.043***	0.046***	0.049***
				(0.015)	(0.015)	(0.015)
$Any\ Breakup_{bt}$=1					0.087***	
					(0.008)	
# $Breakups_{bt}$=1						0.060***
						(0.008)
# $Breakups_{bt}$=2						0.189***
						(0.015)
# $Breakups_{bt}$=3						0.243***
						(0.022)
# $Breakups_{bt}$=4						0.301***
						(0.031)
# $Breakups_{bt}$=5+						0.298***
						(0.032)
Model	Probit	Linear	Linear	Linear	Linear	Linear
Fixed Effects	.	.	b, t	b, t	b, t	b, t
R^2	.	0.18	0.44	0.44	0.44	0.45
Obs.	17,592	17,592	17,592	17,592	17,592	17,592

Standard errors in parentheses, clustered at the buyer level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcome in all specifications is a dummy that takes value one if the buyer has at least one *one-off* interaction in a given quarter. The outcome is thus defined at the level of a buyer–time tuple, and in the data, it equals one for 49% of the bt combinations. By construction, the regression sample retains only buyer–quarter combinations with nonzero trade. In addition, I restrict attention to buyers active for at least a year. This is necessary to allow for sufficient variation in the regressors, conditional on the relevant fixed effects. The regressors are as follows. $Entry^{j}_{bt}$ takes value one if the buyer enters at least one product j in quarter t or $t+1$ and zero otherwise. $\overline{Capacity}^{s}_{bt}$ is a measure of capacity utilization of the buyer's recurrent sellers. We define the seller's capacity as the maximum volume that it trades across all quarters in the data and the utilization rate as the ratio between the volume of the seller–quarter and seller capacity. $\overline{Capacity}^{s}_{bt}$ is the average of the utilization rate across all recurrent sellers with whom the buyer trades in the quarter. $Any\ Breakup_{bt}$ is an indicator that takes value one if at least one of the buyer's relationships appears to end (i.e., trade ceases) in quarter t or $t-1$. The companion set of indicators # $Breakups_{bt} = n$ for $n \in \{1, 2, 3, 4, 5+\}$ reflects the breakup count, with $n = 0$ being the default. As an example, # $Breakups_{bt} = 3$ is an indicator that takes value one whenever buyer b ceases to trade with three sellers in t or $t-1$. Column (1) shows probit marginal effects, and all other columns report results of linear probability models. Columns (1) and (2) include no fixed effects, while columns (3) to (6) include buyer and quarter fixed effects.

Table 3: Traded Volumes in One-off and First Interactions

	q_{sbjt}			
	(1)	(2)	(3)	(4)
$one-off_{sbj}$	-0.469***			
	(0.012)			
$\mathbf{I}\{i_{sbjt}=1^{st}\}$		-0.391***	-0.349***	-0.513***
		(0.023)	(0.022)	(0.037)
Fixed Effects	bj, sj, jt	sb, jt	sb, jt	sb, jt
Duration	Any	Any	1y+	1y+
Buyers	All	All	All	Large
R^2	0.55	0.54	0.53	0.41
Obs.	526,163	532,864	294,939	144,154

Standard errors in parentheses, clustered at the buyer level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcome in all specifications is the log volume traded by the seller–buyer–product–quarter tuple, q_{sbjt}. Column (1) studies the correlation between the outcome and an indicator that takes value one if the tuple corresponds to a one-off interaction, i.e., a buyer–seller–product triplet with interactions for one quarter only ($one-off_{sbj}$). The specification includes buyer–product, seller–product and product–quarter fixed effects. Columns (2)-(4) study the first interaction in a relationship by means of an indicator that takes value one if a given quarter corresponds to the first interaction of the buyer–seller–product triplet ($\mathbf{I}\{i_{sbjt}=1^{st}\}$). In each of columns (2)-(4), seller–buyer and product–quarter fixed effects are included. Different columns study different samples. Columns (1) and (2) include all tuples not affected by censoring. Column (3) restricts attention to buyer–seller–product triplets active for at least one year. Column (4) further restricts the sample to consider only the 200 largest buyers, who account for 70% of the volume traded in the industry throughout the sample period.

Table 4: Performance of Recurrent Relationships and Experimentation upon Entry

	(1) $Duration_{sbj}$	(2) $Span_{sbj}$	(3) q_{sbj}	(4) $Transactions_{sbj}$	(5)
$Experiment_{bj}^{entry}$	0.021***	0.039**	0.059***	0.059***	0.021**
	(0.007)	(0.017)	(0.019)	(0.014)	(0.008)
q_{sbj}					0.644***
					(0.008)
FEs	$j, b, s, c(bj), c(sbj)$				
R^2	0.39	0.43	0.56	0.43	0.82
Obs.	27,179	27,179	27,179	27,179	27,179

Standard errors in parentheses, clustered at the product level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The unit of observation is a buyer–seller–product triplet, which is defined as a recurrent relationship (a triplet active for more than one quarter). In addition, only buyer–product entries observed to have at least one one-off interaction are included in the sample. $Experiment_{bj}^{entry}$ collects the (log) count of one-off interactions that buyer b engages in within the first year after entering product category j. All columns report OLS regressions on this variable and fixed effects for the buyer (b), seller (s) and product category (j). In addition, the quarter in which the buyer–seller–product starts trading and the quarter in which the buyer enters the product category are conditioned upon $c(sbj)$ and $c(bj)$, respectively (where c stands for the cohort). The outcomes correspond to $Duration_{sbj}$, the log count of quarters of trade by the triplet, in column (1); $Span_{sbj}$, the log time span between the first and last observed shipment for the triplet, in column (2); q_{sbj}, the log volume traded in the relationship, in column (3); and $Transactions_{sbj}$, the log count of shipments in the relationship, in columns (4) and (5). Column (5) differs from column (4) only in that it controls for traded volumes, q_{sbj}.

Table 5: Experimentation, Dispersion and Large Buyers

Panel A: Continuous measure of dispersion (baseline)

	$Experiment_{bjt}$								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
q_{bjy}	0.028*** (0.001)							-0.002*** (0.001)	
$Large_b$=1		0.156*** (0.004)							
$Dispersion_j$			-0.141*** (0.003)	-0.085*** (0.005)					
$Large_b$=1 × $Dispersion_j$				-0.117*** (0.007)	-0.059*** (0.006)	-0.041*** (0.006)	-0.073*** (0.009)	-0.074*** (0.009)	-0.083*** (0.009)
$Large_b$=1 × $Count^s_j$						0.003*** (0.000)	0.004*** (0.000)	0.004*** (0.000)	0.005*** (0.000)
$Large_b$=1 × $Med\ Quality_j$						0.018*** (0.002)	0.010*** (0.003)	0.010*** (0.003)	0.012*** (0.003)
$Count^{recurrent}_{bj}$									-0.023*** (0.001)
Fixed Effects	jt	jt	bt	bt	b, jt	b, jt	b, jt	b, jt	b, jt
Sample	All	All	All	All	All	All	Search	Search	Search
R^2	0.16	0.17	0.36	0.36	0.45	0.45	0.38	0.38	0.38
Obs.	409,373	409,373	358,195	358,195	408,201	408,201	182,841	182,841	182,841

Panel B: Dispersion measured in quartiles

	$Experiment_{bjt}$						
	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$Quartile\ Dispersion_j$=2	0.097*** (0.001)	0.107*** (0.002)					
$Quartile\ Dispersion_j$=3	-0.053*** (0.002)	-0.038*** (0.003)					
$Quartile\ Dispersion_j$=4	-0.136*** (0.003)	-0.107*** (0.005)					
$Large_b$=1 × $Quartile\ Disp._j$=2		-0.022*** (0.003)	0.031*** (0.002)	-0.006** (0.003)	0.001 (0.004)	0.000 (0.004)	-0.003 (0.004)
$Large_b$=1 × $Quartile\ Disp._j$=3		-0.032*** (0.004)	-0.018*** (0.004)	-0.011*** (0.004)	-0.014** (0.006)	-0.015** (0.006)	-0.020*** (0.006)
$Large_b$=1 × $Quartile\ Disp._j$=4		-0.059*** (0.006)	-0.051*** (0.006)	-0.027*** (0.007)	-0.052*** (0.009)	-0.053*** (0.009)	-0.060*** (0.010)
$Large_b$=1 × $Count^s_j$				0.003*** (0.000)	0.004*** (0.000)	0.004*** (0.000)	0.005*** (0.000)
$Large_b$=1 × $Med\ Quality_j$				0.017*** (0.002)	0.006* (0.003)	0.006* (0.004)	0.008** (0.003)
q_{bjy}						-0.002*** (0.000)	
$Count^{recurrent}_{bj}$							-0.023*** (0.001)
Fixed Effects	bt	bt	b, jt	b, jt	b, jt	b, jt	b, jt
Sample	All	All	All	All	Search	Search	Search
R^2	0.39	0.39	0.45	0.45	0.38	0.38	0.38
Obs.	358,195	358,195	408,201	408,201	182,841	182,841	182,841

Standard errors in parentheses, bootstrapped from 400 samples of entire vectors of products (HS6) with replacement. Columns (1) and (2) of Panel A cluster the standard errors at the buyer level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcome of these regressions is the log count of one-off interactions (+1) in a buyer–product–quarter combination. By construction, the regression sample retains only buyer–product–time triplets with nonzero trade. All columns correspond to OLS regressions. $Large_b$ is an indicator that takes value one if the buyer is among the top 200 buyers in the industry. $Dispersion_j$ measures the dispersion (log of the standard deviation) in seller quality within a product j. Results using this measure are collected in Panel A. Under the alternative measure, $Quartile\ Dispersion_j$, products are arranged in quartiles, with quartile 1 corresponding to the js with the lowest heterogeneity, which is the base category. The results using these quartiles are presented in Panel B. Columns (1) to (6) in both panels exploit all buyer–product–time triplets in the sample, while columns (7) to (9) consider buyer–product–time combinations where search for trade partners is more likely to occur. This restricted sample includes observations bjt in which t or $t-1$ is the first quarter in which bj is active (entry) or the buyer has experienced at least one relationship termination. Fixed effects are added sequentially. Columns (1) and (2) include product–quarter fixed effects, columns (3) and (4) have buyer–quarter fixed effects, and the rest of the columns use the baseline set of fixed effects: buyer and product–quarter intercepts. Additional controls are as follows: q_{bjy} is the log volume imported by the buyer in the product-year combination; $Count_{bj}^{recurrent}$ is the (log) count of relationships that the buyer forms in the product category; $Count_j^s$ is the log count of sellers available in product market j; and $Med.\ Quality_j$ is the median quality across the estimated $\hat{\theta}_s$ of sellers available in j.

Table 6: A Shock to the Cost of Experimentation

	$Experiment_{bj}$			
	(1)	(2)	(3)	(4)
$Large_b$=1 × $Post_{c(bj)}$=1	-0.004	-0.008	0.212**	0.154**
	(0.032)	(0.031)	(0.083)	(0.076)
$Quartile\ Dispersion_j$=2 × $Large_b$=1	0.009	0.009		
	(0.014)	(0.014)		
$Quartile\ Dispersion_j$=3 × $Large_b$=1	-0.010	-0.012		
	(0.015)	(0.015)		
$Quartile\ Dispersion_j$=4 × $Large_b$=1	0.018	-0.002		
	(0.023)	(0.022)		
$Quartile\ Dispersion_j$=2 × $Large_b$=1 × $Post_{c(bj)}$=1	0.033	0.033		
	(0.041)	(0.040)		
$Quartile\ Dispersion_j$=3 × $Large_b$=1 × $Post_{c(bj)}$=1	-0.041	-0.018		
	(0.043)	(0.040)		
$Quartile\ Dispersion_j$=4 × $Large_b$=1 × $Post_{c(bj)}$=1	-0.154***	-0.112**		
	(0.052)	(0.050)		
$Dispersion_j$ × $Large_b$=1			0.004	-0.012
			(0.022)	(0.021)
$Dispersion_j$ × $Large_b$=1 × $Post_{c(bj)}$=1			-0.161***	-0.116**
			(0.054)	(0.051)
$q_{bj}^{year=1}$		-0.062***		-0.062***
		(0.001)		(0.001)
Fixed Effects	b, jt	b, jt	b, jt	b, jt
R^2	0.25	0.29	0.25	0.29
Obs.	60,297	60,297	60,297	60,297

Standard errors in parentheses, bootstrapped 400 times stratifying by product. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcome $Experiment_{bj}^{entry}$ is the log count of one-off interactions that buyer b has in product j within the first year after entering the sourcing market. $Large_b$ is an indicator that takes value one if the buyer is among the top 200 buyers in the industry. The postshock variable, $Post_{c(bj)}$, is a dummy that takes value one if the entry of buyer b in product j occurs in a quarter after the RP event. $Quartile\ Dispersion_j$ measures the dispersion in seller quality within a product j, and products are arranged in quartiles, with quartile 1 corresponding to js with the lowest heterogeneity, which is the base category. $Dispersion_j$ is the continuous metric of dispersion across seller quality in each product category j. Columns (1) and (2) use the quartile-based measure of dispersion, while columns (3) and (4), instead, use the continuous measure of dispersion (see specification (5)). All columns include buyer and product–time fixed effects, where time is the quarter of entry, $t = c(bj)$. In addition, columns (2) and (4) control for the (log) volume that buyer b imports of product j throughout the first year of trade in the product category, $q_{bj}^{year=1}$.

Table 7: Trimmed Samples and Pretrends

	$Experiment^{entry}_{bj}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Large_b$=1 × $Post_{c(bj)}$=1	0.223***	0.212	0.247	0.294*	0.344*	0.245**	0.279**
	(0.085)	(0.141)	(0.154)	(0.157)	(0.182)	(0.109)	(0.125)
$Large_b$=1 × $Dispersion_j$	0.008	0.032	0.043	0.098*	0.089	0.021	0.045
	(0.027)	(0.045)	(0.052)	(0.056)	(0.070)	(0.055)	(0.059)
$Post_{c(bj)}$=1 × $Large_b$=1 × $Dispersion_j$	-0.163***	-0.144	-0.181*	-0.216**	-0.250**	-0.177**	-0.201**
	(0.056)	(0.095)	(0.104)	(0.107)	(0.121)	(0.072)	(0.082)
$Year$ =2011 × $Large_b$=1						0.003	0.101
						(0.124)	(0.144)
$Year$ =2010 × $Large_b$=1						0.091	0.146
						(0.113)	(0.133)
$Year$ =2009 × $Large_b$=1						-0.211*	-0.152
						(0.117)	(0.140)
$Year$ =2008 × $Large_b$=1						0.140	0.153
						(0.113)	(0.144)
$Year$ =2007 × $Large_b$=1						0.113	0.088
						(0.123)	(0.139)
$Year$ =2006 × $Large_b$=1						0.042	0.076
						(0.115)	(0.143)
$Year$ =2011 × $Large_b$=1 × $Dispersion_j$						-0.009	-0.072
						(0.085)	(0.098)
$Year$ =2010 × $Large_b$=1 × $Dispersion_j$						-0.040	-0.077
						(0.077)	(0.090)
$Year$ =2009 × $Large_b$=1 × $Dispersion_j$						0.148*	0.110
						(0.079)	(0.096)
$Year$ =2008 × $Large_b$=1 × $Dispersion_j$						-0.109	-0.134
						(0.077)	(0.098)
$Year$ =2007 × $Large_b$=1 × $Dispersion_j$						-0.069	-0.066
						(0.084)	(0.097)
$Year$ =2006 × $Large_b$=1 × $Dispersion_j$						-0.004	-0.046
						(0.079)	(0.100)
Fixed Effects	b, jt	b, jt	b, jt	b, jt	b, jt	b, jt	b, jt
Sample	Incumbent	First 4	First 5	First 6	First 7	All	Incumbent
R^2	0.26	0.35	0.37	0.39	0.41	0.25	0.26
Obs.	32,462	13,072	10,696	8,692	6,971	60,297	32,462

Standard errors in parentheses, bootstrapped 400 times stratifying by product. $*(p < 0.10), **(p < 0.05), * * *(p < 0.01)$. The table complements the results presented in Table 6. Table 7 reproduces the baseline of column (3) of Table 6 under different sample trimmings and conditioning on pretrends. The outcome $Experiment^{entry}_{bj}$ is the log count of one-off interactions that buyer b has in product j within the first year after entering the sourcing market. $Large_b$ is an indicator that takes value one if the buyer is among the top 200 buyers in the industry. Here the postshock variable $Post_{c(bj)}$ is a dummy that takes value one if the entry of buyer b in product j occurs in a quarter after the RP event. $Dispersion_j$ is the continuous metric of dispersion across sellers in market j. All specifications include buyer and product–time fixed effects, where time is the quarter in which the entry takes place $t = c(bj)$. Column (1) here restricts the sample of column (3) from Table 6 to only buyers classified as *incumbents*, i.e., buyers with at least one entry before and at least one entry after RP. The breakdown of all entries in the data by buyer status is summarized in Table E7. Columns (2) to (5) further restrict the sample to keep the first x entry instances after RP by buyers surviving at least x instances, with $x = \{4, 5, 6, 7\}$. Restricting the sample further reduces power significantly. Columns (6) and (7) are based on the sample of all entries and entries by all incumbent buyers, respectively. In addition to the baseline regressors, they include pretrends for the time-varying interactions. These are constructed for all years prior to RP (2012, 2011, 2010 . . .), with 2012, the year immediately prior to the shock, excluded as the base category. Entries in 2013 prior to the collapse (before April 2013) are classified as part of the baseline together with any entries in 2012.

Table 8: Traded Volumes in All and Recurrent Relationships

	q_{bjt}			q_{bjy}		$q_{bjy}^{recurrent}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Large_b$=1 × $Post$=1	-0.601***	-0.553**	-0.907***	-0.429**	-0.659**	-0.412	-0.116
	(0.169)	(0.229)	(0.214)	(0.187)	(0.304)	(0.512)	(0.766)
$Large_b$=1 × $Dispersion_j$	-0.280***	-0.235*	-0.266***	-0.449***	-0.612***	-1.251***	-1.040**
	(0.060)	(0.124)	(0.065)	(0.067)	(0.179)	(0.176)	(0.441)
$Large_b$=1 × $Dispersion_j$ × $Post$=1	0.446***	0.397**	0.651**	0.358***	0.513**	0.348	0.115
	(0.120)	(0.164)	(0.147)	(0.132)	(0.212)	(0.353)	(0.521)
Fixed Effects	b, jt	b, jt	b, jt	b, jy	b, jy	b, jy	b, jy
Pretrends	No	Yes	No	No	Yes	No	Yes
Time	Quarter	Quarter	Quarter	Year	Year	Year	Year
Partners	All	All	All	All	All	Recurrent	Recurrent
R^2	0.46	0.46	0.47	0.43	0.43	0.27	0.27
Obs.	150,761	150,761	125,953	106,402	106,402	106,402	106,402

Standard errors in parentheses, bootstrapped 400 times stratifying by product. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table studies the evolution of buyer–product–time volumes before and after RP, with the time dimension defined in terms of quarters (columns (1) to (3)) and years (columns (4) to (7)): q_{bjt} is the volume imported by buyer b in product t during calendar quarter t; q_{bjy} is the log volume imported by the buyer in the product in year y; $q_{bjy}^{recurrent}$ is the log volume (+1) imported by the buyer–product–year triplet, sourced from recurrent partners (i.e., volumes purchased from one-off partners are excluded from the volume aggregation). $Large_b$ is an indicator that takes value one if the buyer is among the top 200 buyers in the industry. $Dispersion_j$ measures the dispersion in seller quality within a product j. $Post$ is an indicator taking value one if the time period falls after the RP event. It is defined at the level of quarters in columns (1) to (3) and years in columns (4) to (7). All specifications include buyer and product–time fixed effects. In addition, columns (2), (5) and (7) include a complete set of pretrends for the time-varying interactions, ($Large_b \times Post$) and ($Large_b \times Dispersion_j \times Post$). These are interactions of time dummies and $Large_b$ and $Large_b$ x $Dispersion_j$ for all periods prior to RP (see columns (6) and (7) of Table 7, which show the structure of these pretrends). Column (3) reproduces the exercise in column (1) on the sample of quarters that fall within the first two calendar years of the buyer entering the product.

Table 9: Learning Across Entry Instances

	$Experiment_{bj}$		
	(1)	(2)	(3)
$Age\ Trend_{n(bj)}$	0.001***	0.001***	
	(0.000)	(0.000)	
$Large_b$=1 × $Age\ Trend_{n(bj)}$	-0.000	-0.001	
	(0.001)	(0.001)	
$Large_b$=1 × $Dispersion_j$	0.036	0.042	0.045
	(0.040)	(0.041)	(0.039)
$Large_b$=1 × $Dispersion_j$ × $Age\ Trend_{n(bj)}$	-0.001	-0.001	
	(0.001)	(0.001)	
$Large_b$=1 × $Post_{c(bj)}$=1		0.215**	0.259***
		(0.088)	(0.097)
$Large_b$=1 × $Dispersion_j$ × $Post_{c(bj)}$=1		-0.148**	-0.198***
		(0.059)	(0.066)
$Age\ Pretrend_{n(bj)}$			0.002***
			(0.000)
$Large_b$=1 × $Age\ Pretrend_{n(bj)}$			-0.001
			(0.001)
$Large_b$=1 × $Dispersion_j$ × $Age\ Pretrend_{n(bj)}$			-0.001
			(0.001)
Fixed Effects	b, jt	b, jt	b, jt
R^2	0.25	0.25	0.25
Obs.	60,297	60,297	60,297

Standard errors in parentheses, bootstrapped 400 times stratifying by product. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcome $Experiment_{bj}^{entry}$ is the log count of one-off interactions that buyer b has in product j within the first year after entering the sourcing market. $Age\ Trend_{n(bj)}$ is a linear trend for the buyer's sequence of entries. It takes value one if j is the first product that the buyer enters, two if it is the second product that the buyer enters, etc. $Age\ Pretrend_{n(bj)}$ is an analogous count that keeps track of the order of entries prior to the RP event. $Large_b$ is an indicator that takes value one if the buyer is among the top 200 buyers in the industry. $Dispersion_j$ measures the dispersion in seller quality within a product j. $Post_{c(bj)}$ is an indicator taking value one if the entry of buyer b in product j occurs in a quarter after the RP event, $t = c(bj)$. All specifications include buyer and product–time fixed effects, where time is a quarter.

Table 10: Average Quality of Recurrent Partners

	Average Quality$_{bjy}$			
	(1)	(2)	(3)	(4)
$Large_b$=1 × $Post_y$=1	2.873***	2.747***	2.424*	3.167**
	(0.956)	(0.999)	(1.292)	(1.283)
$Large_b$=1 × $Dispersion_j$	1.640***	1.539***	1.681***	1.477***
	(0.431)	(0.462)	(0.461)	(0.432)
$Large_b$=1 × $Dispersion_j$ × $Post_y$=1	-2.033***	-1.931***	-1.916**	-2.110***
	(0.712)	(0.742)	(0.815)	(0.773)
q_{bjy}		-0.191***		
		(0.015)		
$Large_b$=1 × $Med\ Quality_j$			-0.147	
			(0.125)	
$Large_b$=1 × $Med\ Quality_j$ × $Post_y$=1			-0.182	
			(0.241)	
$Large_b$=1 × $Count^s_{jy}$				-0.088
				(0.055)
$Large_b$=1 × $Count^s_{jy}$ × $Post_y$=1				-0.039
				(0.114)
Fixed Effects	b, jy	b, jy	b, jy	b, jy
R^2	0.79	0.80	0.79	0.79
Obs.	45,849	45,849	45,849	45,849

Standard errors in parentheses, bootstrapped 400 times stratifying by product. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. A unit of observation is a buyer–product–year combination. The outcome is the (simple) average of the estimated quality of all recurrent partners with whom the buyer trades in the product–year combination, $Average\ Quality_{bjy}$. Buyer–product–year triplets featuring no trade with recurrent partners are excluded. $Large_b$ is an indicator that takes value one if the buyer is among the top 200 buyers in the industry. $Dispersion_j$ measures the dispersion in seller quality within a product j. $Post_y$ is an indicator taking value for 2013, when RP took place, and the years after. All specifications include buyer and product–year fixed effects. Column (1) presents the baseline specification using the double and triple interactions of interest. The rest of the columns add various controls. Column (2) controls for the (log) volume imported by the buyer in the product and year, q_{bjy}. Column (3) controls for double and triple interactions with the median type of seller in market j, $Med\ Quality_j$. Column (4) adds double and triple interactions with the count of active sellers in the product–year combination.

Figure 1: Relationship Duration in Terms of Quarters

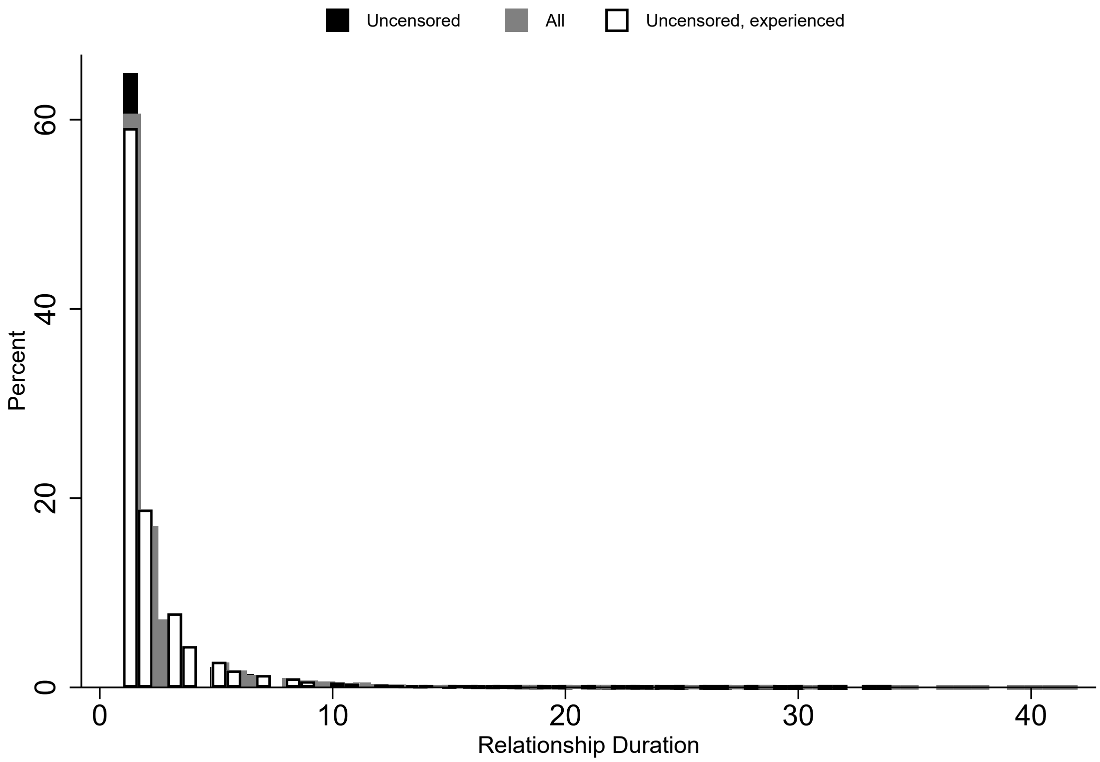

The figure shows the histogram of relationship duration in terms of quarters of effective trade. A relationship is defined as a buyer–seller–product triplet. The vertical axis reports percentages. The solid black bars correspond to uncensored relationships (i.e., those starting in or after 2006 and ending in or before 2014). The solid gray bars show all relationships. The white bars show uncensored relationships of buyer–product combinations active in the data for at least two years.

Figure 2: Threshold Supplier and Quality Dispersion

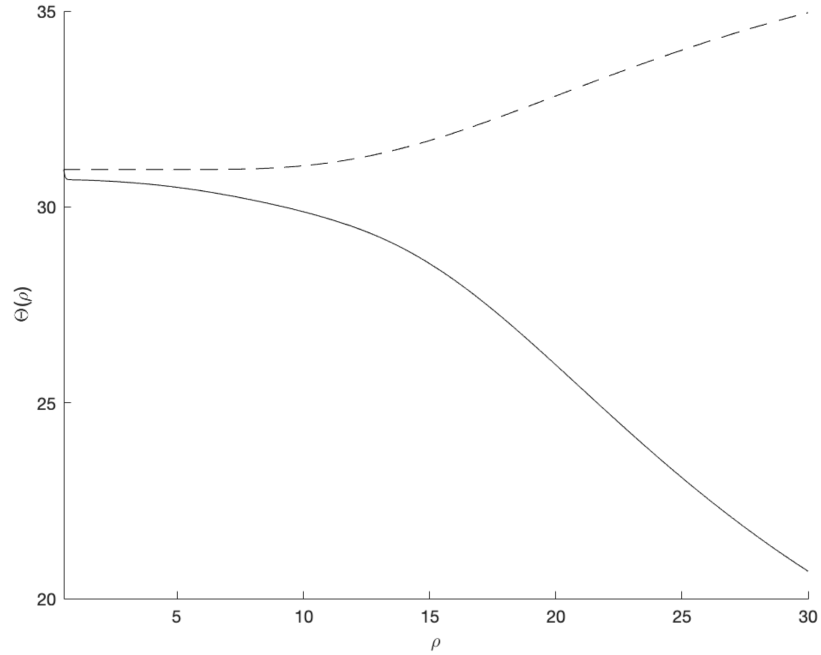

The figure illustrates the relationship between the threshold supplier quality, $\hat{\Theta}(c)$, and quality dispersion, ρ, when only revenue effects are present (dashed line) and when both revenue and cost effects are at play. The parametrization for this figure is as follows: $F(\theta;\rho)$ is a truncated normal over $[\underline{\theta},\overline{\theta}]$ with $\underline{\theta}=1$ and $\overline{\theta}=100$; ρ varies on $[0.5, 30]$; $r(\theta,q;\alpha)=\alpha q/\theta$, with $\alpha=1000$; and $q=1$. For the case with zero cost effects, the cost is set to be that corresponding to $\rho=0.5$, while revenues are allowed to vary freely with ρ.

Figure 3: An Increase in α in Low- and High-Dispersion Environments

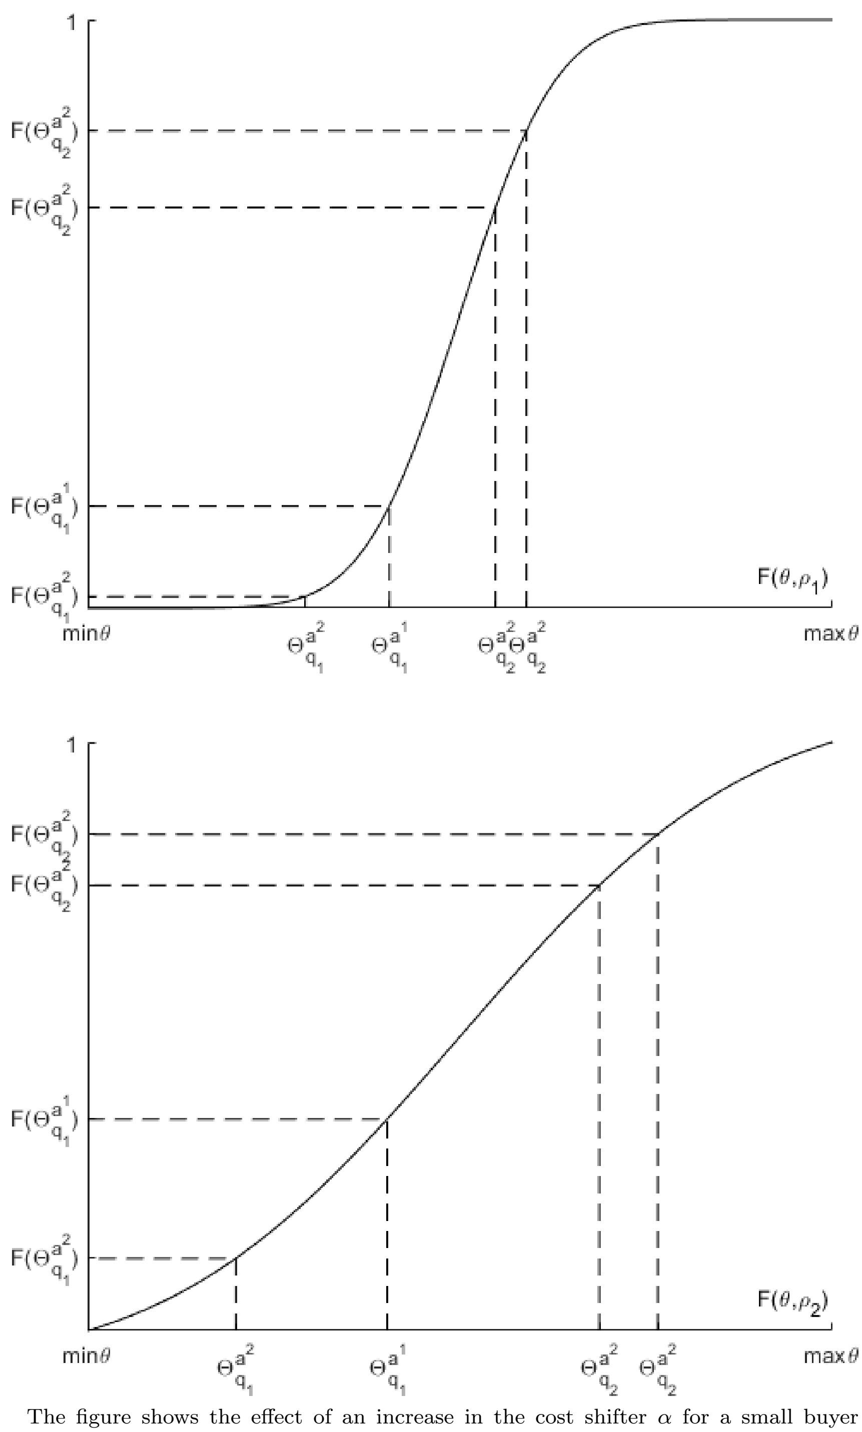

The figure shows the effect of an increase in the cost shifter α for a small buyer (characterized by q_1) and a large buyer (characterized by q_2) in a low-dispersion (ρ_1, top subfigure) and a high-dispersion (ρ_2, bottom subfigure) environment. In both subfigures, the horizontal axis depicts θ, and $F(\theta,\rho)$ is plotted against the vertical axis. The parametrization for these figures is as follows: $F(\theta;\rho)$ is a truncated normal over $[\underline{\theta},\overline{\theta}]$ with $\underline{\theta} = 1$ and $\overline{\theta} = 100$; $\rho_1 = 10$ and $\rho_2 = 30$; $r(\theta;\alpha) = \alpha/\theta$, with $\alpha_1 = 1000$ and $\alpha_2 = 2000$; and $q_1 = 2$ and $q_2 = 20$. Note that to simplify the comparative statics over α, $r(\cdot)$ does not depend on q, and so buyer size drives the scale mechanics only. For printing purposes, in the axis, α is replaced by a.

Figure 4: Placebo Regressions on 400 Randomly Drawn Dates

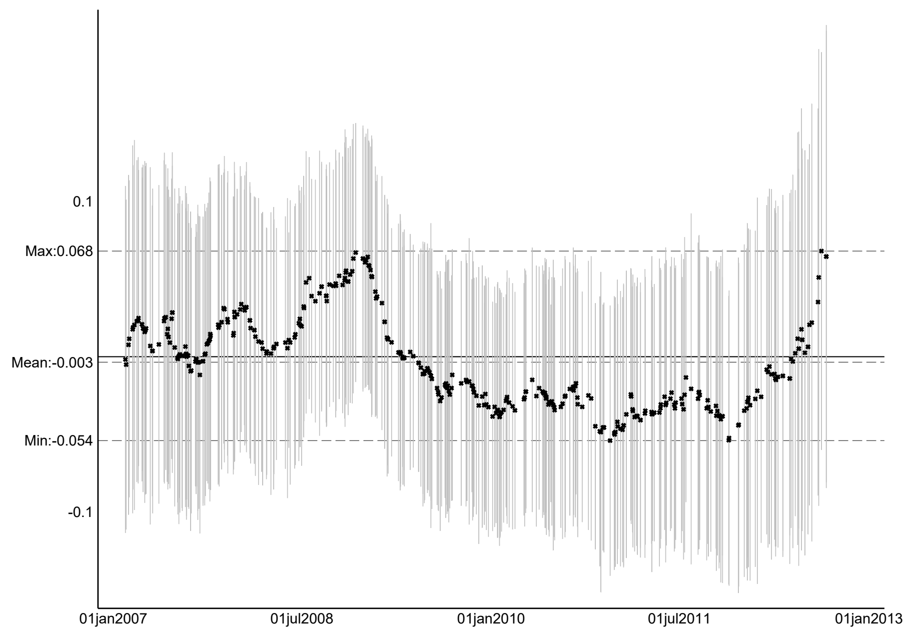

The figure shows the estimated coefficient $\hat{\beta}_3$ of equation (5) using 400 randomly selected placebo dates for the RP shock. Data for the period after the RP incident are discarded. Dates between January 2007 and October 2012 are drawn at random to serve as the cutoff date for the variable $Post_{c(bj)}$. For example, if the date drawn is 1 June 2009, $Post_{c(bj)}$ takes value one for all entry instances bj that take place after 1 June 2009 and zero otherwise. The bounds for these draws (January 2007 and October 2012) are set so that there is sufficient data in both the pre- and postperiods for all dates. Thus, the latest possible draw, from October 2012, has entries between this date and the RP incident as the postperiod (six months of calendar time). For each draw, the regression following (5) runs bootstrapping of the standard errors, with 100 bootstrap samples over the product strata. Point estimates are depicted with black markers, and 95% confidence intervals are shown in gray. The average estimated $\hat{\beta}_3$ is -0.003, the lowest is -0.054, and the highest is 0.068. None of the coefficients are significant at 10%.

Figure 5: Robustness to Alternative Specifications and Definitions

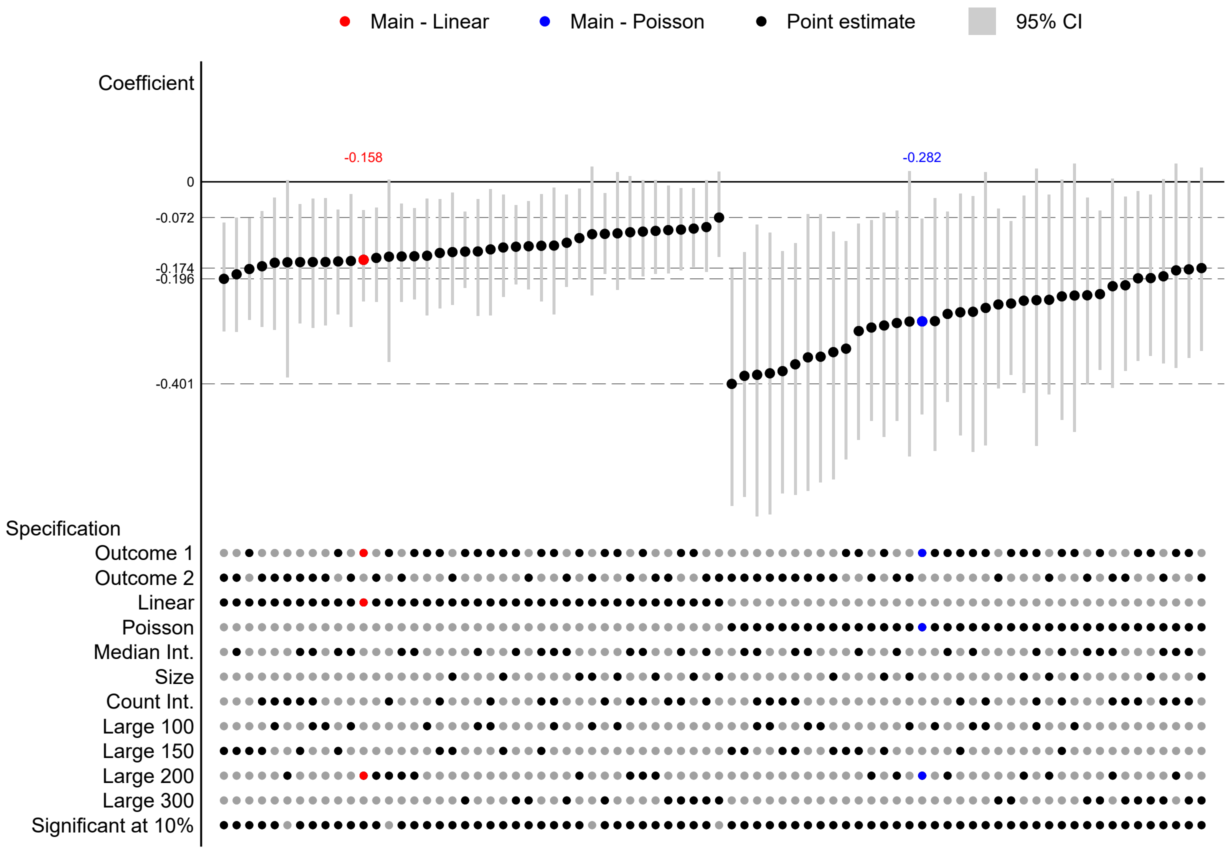

The figure presents the point estimate of $\hat{\beta}_3$ under the alternative definitions and specifications discussed in Appendix D.1, alongside confidence intervals constructed with bootstrapped standard errors, drawing entire vectors of products (400 bootstrap samples). The left-hand side of the figure shows estimates of β_3 in the linear specification of equations (5) and (D1), while the right-hand side of the figure shows the results of Poisson pseudo-maximum likelihood (PPML) estimations with the count outcome in levels, following equation (D2). The point estimate of the baseline linear specification is highlighted in red, while the baseline for the PPML estimation appears in blue. All point estimates of the linear specifications fall within the interval $[-0.196, -0.072]$, while those of the PPML fall within the interval $[-0.401, -0.174]$. Of the 78 different specifications presented here, 74 (95%) produce statistically significant coefficients at 10%. Outcome 1 refers to the baseline outcome, based on the count of one-off interactions upon the buyer's entry into the product category within the first calendar year. Outcome 2 corresponds to the count of one-off interactions until the first recurrent interaction. The labels "Median Int." and "Count Int." indicate the inclusion of interactions with the median quality and the count of available suppliers in j, following equation (D1). Specifications with the "Size" label control for the total volume imported by the buyer within the first year in the product category. The labels "Large x" for $x = \{100, 150, 200, 300\}$ correspond to different cutoffs for the dummy variable for large buyers. The baseline throughout the paper is $x = 200$. Markers shaded in black reflect the specification for each point estimate. For example, the left-most point estimate corresponds to Outcome 2, is linear, and does not include the size control or interactions with the median or count of sellers, and the buyer cutoff is set to 150. This figure was produced with code adapted from Hans Sievertsen's version of *speccurve*, accessed here: `https://github.com/hhsievertsen/speccurve`.

A Recurrent Relationships and Exit

The evidence presented in Section 3 suggests that when entering a new product category, lacking information on relevant characteristics of potential partners, buyers experiment with relationships that may not develop beyond the first interaction. In the presentation that follows, I show that when no recurrent relationships are formed, buyers are more likely to exit the product category within a year of entering. In other words, surviving buyers typically have at least one recurrent partner.

To establish this, I start by discussing descriptive patterns of the observed entry instances, defined as buyer–product combinations whose first shipment is not censored.[49] There are over sixty thousand such instances in the data, and of these, 79.5% involve at least one one-off relationship during the first year in which the buyer is active in the product category (see Panel A of Table A1). Only the remaining 20.5% of the entry instances are ones in which the buyer continues to trade with its first partner. Of the entry instances that feature some experimentation in the first year (i.e., at least one one-off interaction), almost 15% see at least one recurrent relationship at some point.[50] In this case, there are on average 1.44 one-off interactions in the first year. With a median of one, the very skewed distribution has three and five one-off interactions at the 95^{th} and 99^{th} percentiles, respectively. Figure A1 shows that the high incidence of one-off interactions upon entry is positively correlated with buyer size and the number of recurrent relationships that the buyer establishes in the product category. In sum, (i) most entry instances feature one-off interactions, (ii) only a small share of entries eventually lead to a recurrent relationship, and (iii) when this happens, it takes between one and two one-off interactions to get there, with this number of attempts being greater for larger buyers and those forming multiple recurrent relationships.

I turn to studying the link between buyers' exits from the product categories that they enter and the formation of recurrent relationships. I do so by exploiting the specification

$$Exit\ 1^{st}\ year_{bj} = \alpha Forms\ Recurrent_{bj} + \delta_j + \delta_b + \delta_{c(bj)} + \epsilon_{bj}.$$

The unit of observation described by this specification is an entry instance by buyer b into product category j. The outcome of interest is an indicator taking value one if the buyer exits (stops importing) the product category within a year of the entry date. $Forms\ Recurrent_{bj}$ indicates whether the firm forms at least one recurrent relationship in the product category within a year of entering. In this context, α reflects the change in the probability of the firm exiting the market within a year if it forms a recurrent relationship, relative to the probability when it has one-off interactions only. Product fixed effects δ_j absorb differences in average buyers' survival rates across products. Buyer fixed effects δ_b and buyer–product cohort effects (the buyer's quarter of entry in the product category) $\delta_{c(bj)}$ account for idiosyncratic drivers of buyer exit common to all the product categories that they enter as well as time trends

[49] The censoring rule implies the discarding of the first shipments observed in the first and last years of the data. Studying the time elapsed between shipments within uncensored buyer–product combinations, I find that more than 95% of the shipments happen within 365 days of each other.

[50] For the more than seven thousand entries in sample (c) in Table A1, the time span between the buyer's entry date and the date at which a relationship becomes recurrent (i.e., when the buyer trades for a second quarter with some seller) has a median of just over a year (366 days).

common to all entries. An augmented specification also controls for the volume that the buyer imports while active during the first year.

Table A2 shows that the likelihood of a buyer exiting the product category within a year of entering drops by more than a third when the buyer forms a recurrent relationship. In the most demanding specification, with product, buyer and cohort fixed effects and a control for traded volumes, buyers that form a recurrent relationship are 33% less likely to exit within their first year than they are in cases with only one-off interactions (column (3)). The result does not arise mechanically: buyers could stay in the market and source products via multiple short-lived relationships. In contrast, it appears that in the study context, buyers that remain active also trade recurrently with at least one partner.

Table A1: Entry Instances: one-off and recurrent relationships

Panel A: Buyer-product entry instances, by type of relationships formed

One-off in first year	Recurrent Relationships throughout the sample		
	None	At least one	Total
None	0	13,024 (b)	13,024
At least one	43,490 (a)	7,145 (c)	50,635
Total	43,490	20,169	63,659

Panel B: Count of relationships in buyer-product entry instances

Count of relationships	Sample	N	Mean	Std. Dev.	P10	P25	P50	P75	P90
One-off in first year	(a)	43,490	1.11	0.38	1	1	1	1	1
Recurrent in sample	(b)	13,024	1.30	0.89	1	1	1	1	2
One-off in first year	(c)	7,145	1.44	0.97	1	1	1	2	2
Recurrent in sample	(c)	7,145	1.78	1.87	1	1	1	2	3

The table reports descriptives of buyer–product entry instances. An entry instance is defined as a buyer–product duple observed for the first time in the data. All cells whose first and/or last instance are censored, i.e., observed in the first (2005) or last (2015) year of the data, are dropped. Panel A classifies all the uncensored entry instances in the data by their relationship formation status. The rows in this panel classify all entries according to whether they feature at least one one-off interaction in the first calendar year after entry. A one-off interaction is a relationship that spans no more than one calendar quarter. The columns in this panel classify entries according to whether they feature, at any point in the data, at least one recurrent relationship. Recurrent relationships are those that span more than one quarter. The combination of the row and column criteria lead to three subsamples, labeled (a), (b) and (c). For example, sample (c) corresponds to all uncensored buyer–product entry instances in which there is at least one one-off interaction in the first year and at least one recurrent relationship at some point in the data. Panel B reports the count of relationships of each type (one-off in first year or recurrent at any point) that buyer–product entry instances exhibit across the different samples. For example, focusing on entry instances that have no one-off interactions in the first year and that form at least one recurrent relationship, I find that the average number of recurrent relationships formed is 1.3, with a median of 1 and a 90th percentile of 2.

Table A2: Probability of Exit and Recurrent Relationships

	Probability of exit in 1^{st} $year_{bj}$			
	(1)	(2)	(3)	(4)
$Forms\ Recurrent_{bj}$	-0.350***	-0.433***	-0.332***	-0.311***
	(0.008)	(0.008)	(0.006)	(0.008)
Model	Probit	Linear	Linear	Linear
Fixed Effects	.	.	j, b, c	j, b, c
R^2	.	0.29	0.35	0.35
Obs.	63,659	63,659	61,473	61,473

Standard errors in parentheses, clustered at the product level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcome in all specifications is a dummy that takes value one if the buyer exits the product category within one year of entering (based on the date of its first transaction). In this context, an observation in these regressions is a buyer–product duple. The main regressor across all specifications is a dummy that takes value one if the buyer forms at least one recurrent relationship in the product category within the first year. Column (1) reports the marginal effect in a linear probit model of forming a recurrent relationship on the probability of exit within the first year. Column (2) re-estimates the effect using a linear probability model. Columns (3) and (4) augment the LPM with product, buyer and cohort (quarter) fixed effects. Column (4) further conditions on the total volume that the buyer trades in the product category while still active.

Figure A1: Entry Instances: One-Off Interactions, Buyer Size and Recurrent Relationships

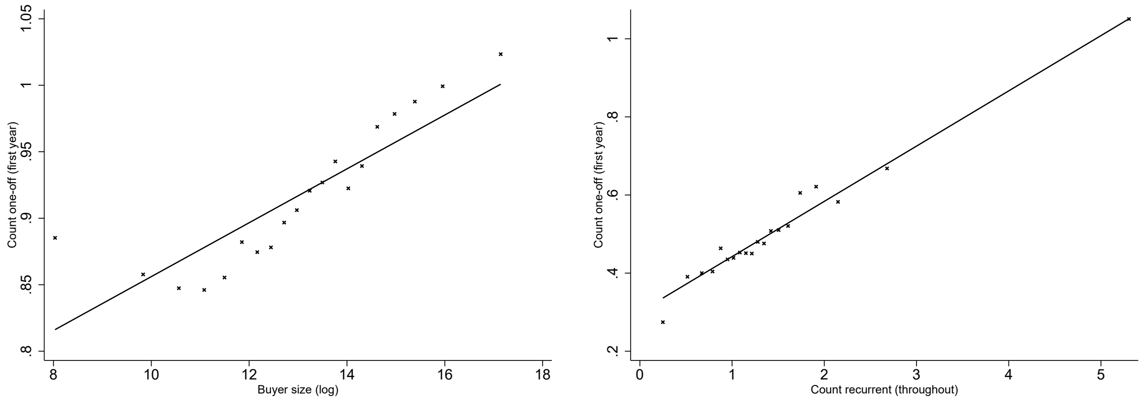

The figures show the relationship between the count of one-off interactions observed in each buyer–product entry instance and buyer size (left) and the number of recurrent relationships (right). Only uncensored entry instances are considered. In both graphs, the scatter markers correspond to averages of the underlying data partitioned into equally sized bins. The solid line shows the linear fit in the data, conditional on a set of controls. The left-hand-side graph shows the log buyer size on the horizontal axis. This is computed by aggregating the full volume imported by the buyer across all product categories in the data. The vertical axis collects the number of one-off interactions within the buyer's first year in a product category. The graph is constructed on 63,659 uncensored entry instances (see Table A1). The linear fit conditions on product fixed effects and on the number of recurrent relationships eventually formed in the buyer–product combination. The right-hand-side figure restricts the sample to entry instances in which a recurrent relationship is eventually formed. It includes a total of 20,169 data points (see Table A1). The horizontal axis shows the count of recurrent relationships formed in the buyer–product combination. The vertical axis, as before, shows the count of one-off interactions within the buyer's first year in a product category. The regression conditions on product fixed effects and buyer size.

B Model: Dispersion and Experimentation

In Section 4, I study changes in search outcomes when the dispersion in unobserved supplier quality changes. For specific parametrizations of $F(.)$, a suitable definition of dispersion might arise naturally. In what follows, I propose a general proof of the results in the body of the paper. To do so, I adopt the definition of variability in Rothschild and Stiglitz (1970) and Diamond and Stiglitz (1974). Accordingly, a definition of *greater riskiness* is compatible with three common ways of understanding dispersion: (i) the addition of mean-zero noise to an original distribution, (ii) the preference on the part of an individual with concave utility for an original distribution over an alternative with equal mean, and (iii) the shifting of mass from the center of an original distribution toward the tails, leaving the mean unchanged. These define a mean-preserving spread and, as such, capture the intuition on dispersion that the application in the paper is concerned with.

Let ρ be the parameter of increasing risk of Diamond and Stiglitz (1974) defined by the integral conditions (i) $\int_{\underline{\theta}}^{\overline{\theta}} F_\rho(\theta;\rho)d\theta = 0$ and (ii) $\int_{\underline{\theta}}^{x} F_\rho(\theta;\rho)d\theta \geqslant 0, \forall\theta \in [\underline{\theta},\overline{\theta}]$. Equation (1) in the main text yields the implicit characterization of $\hat{\Theta}$,

$$\int_{\hat{\Theta}}^{\overline{\theta}} (\theta - \hat{\Theta}) f(\theta;\rho)d\theta - \int_{\underline{\theta}}^{\overline{\theta}} c(\theta) f(\theta;\rho)d\theta = 0. \tag{B1}$$

Integration by parts on the first term gives

$$\int_{\hat{\Theta}}^{\overline{\theta}} (\theta - \hat{\Theta}) f(\theta;\rho)d\theta = \Big[(\theta - \hat{\Theta})F(\theta;\rho)\Big]_{\hat{\Theta}}^{\overline{\theta}} - \int_{\hat{\Theta}}^{\overline{\theta}} F(\theta;\rho)d\theta = \overline{\theta} - \hat{\Theta} - \int_{\hat{\Theta}}^{\overline{\theta}} F(\theta;\rho)d\theta.$$

Using this expression to rewrite (B1),

$$\overline{\theta} - \hat{\Theta} - \int_{\hat{\Theta}}^{\overline{\theta}} F(\theta;\rho)d\theta - \int_{\underline{\theta}}^{\overline{\theta}} c(\theta) f(\theta;\rho)d\theta = 0.$$

Applying Leibniz's rule in the implicit differentiation of $\hat{\Theta}$ with respect to the dispersion parameter gives

$$\frac{d\hat{\Theta}}{d\rho} = \frac{1}{1 - F(\hat{\Theta};\rho)} \times \left[- \int_{\hat{\Theta}}^{\overline{\theta}} F_\rho(\theta;\rho)d\theta - \int_{\underline{\theta}}^{\overline{\theta}} c(\theta) f_\rho(\theta;\rho)d\theta \right]. \tag{B2}$$

The leading factor is always positive, and so the sign of the expression for $\frac{d\hat{\Theta}}{d\rho}$ depends on the sign of the terms in the square brackets. By the integral conditions in the definition of ρ, $\int_{\hat{\Theta}}^{\overline{\theta}} F_\rho(\theta;\rho)d\theta \leq 0$, and so -1 times that expression is positive. By the fundamental theorem of risk, the sign of the second term depends on the sign of the second derivative of $c(\theta)$: whenever $c''(\theta) > 0$, the integral expression is positive, making the whole second term negative.[51] The threshold $\hat{\Theta}$ can increase or decrease with dispersion, depending on which term dominates.

[51] Strictly speaking, if the upper and lower bounds of the integral are functions of ρ, handling this case requires a small departure from the standard derivation. Two additional terms carry over, collecting the cost function and PDF evaluated at the bounds and the derivative of the bounds with respect to ρ (these are of the form $c(x)f(x)dx/d\rho$ with $x \in \{\overline{\theta}, \underline{\theta}\}$). Naturally, these additional terms disappear when the bounds are not a function of ρ. Noting that the PDF is always positive and that $c \to 0$ as $\overline{\theta} \to \infty$, the sign result when $c(\cdot)$ is decreasing and convex holds whenever $d\underline{\theta}/d\rho \leq 0$.

C Seller Heterogeneity

This section presents the demand model and econometric approach for recovering an estimate of the overall quality of each supplier. The developments here follow the standard framework in Berry (1994) for estimating consumer choice models in the presence of differentiation. Widely used in the empirical industrial organization literature, this approach has also been adapted to data and discrete choice problems typical of the international trade domain. The seminal application in trade is Khandelwal's (2010) estimation of quality ladders.

I make a relevant departure from previous work. In existing applications, a variety is defined as a product or a product–origin combination. In the setup of this paper, the origin is Bangladesh in all cases, and varieties are defined as product–seller combinations. Observing market shares across different destinations offers an additional source of variation for identification of the average appeal of a particular seller, which is the main object of interest.

C.1 Demand Model

I model the demand for different varieties of Bangladeshi garments across different countries. Countries indexed with $d \in \{1, 2, \ldots, D\}$, for destination, are populated by consumers $n \in \{1, 2, \ldots, N_d\}$, with preferences defined over products $j \in \{1, 2, \ldots, J\}$.[52] In the data, these are products at the HS6-code level: *woven men's shirts made of cotton* or *women's blouses made of synthetic woven fabric*, for example. Each of these products is available in distinct varieties and supplied by sellers $s \in \{1, 2, \ldots, S\}$. At time t, a consumer chooses the variety within a product category (i.e., a nest) that grants her the highest indirect utility

$$v_{ndtjs} = \delta_{dt} + \delta_{tjs} + \alpha p_{dtjs} + \sum_{j}^{J} \mu_{ndtj} i_{js} + (1 - \sigma)\epsilon_{ndtjs}. \tag{C1}$$

The value that the consumer derives from the best alternative is separable into four components. First, δ_{dt} denotes a shifter that captures differences in taste for destination–time pairs that are common across all consumers and varieties. This captures general demand shocks in the destination. Second, δ_{tjs} is the average taste across all consumers for variety js at time t. I decompose this term further as $\delta_{tjs} = \theta_s + \xi_{tjs}$, where θ_s is the average, time-invariant attractiveness of the seller's products to all consumers and ξ_{tjs} is a mean-zero deviation reflecting product–time-specific departures from the average attractiveness to the seller. Third, p_{dtjs} is the unit value of the variety in the destination and time and renders (dis)utility α to the consumer. The fourth and final component of (C1) accounts for horizontal differentiation: i_{js} is an indicator that takes value one whenever the variety of product j offered by seller s is available, and μ_{ndtj} is the consumer's taste for all varieties of product j. As standard in nested logit models, ϵ_{ndtjs} is Type I extreme-value distributed, and $\sigma \in (0, 1]$ captures the within-nest (i.e., within j) correlation, with the case of $\sigma = 0$ returning simple logit substitution.

[52]Naturally, a given consumer n is active in only one destination d. Thus, the destination is a consumer characteristic that is invariant across the choice instances with which she is presented.

The outside option for consumer n is to purchase imported garments from origins other than Bangladesh, rendering utility

$$v_{ndt\varnothing} = \delta_{dt} + \alpha p_{dt\varnothing} + \mu_{ndt\varnothing} + (1-\sigma)\epsilon_{ndt\varnothing}. \tag{C2}$$

The mean utility of not purchasing from Bangladesh is normalized to zero. The volume of non-Bangladeshi garments imported in each country, $q_{dt\varnothing}$, is observed, and the total size of market dt is $q_{dt} = q_{dt\varnothing} + \sum_{js} q_{dtjs}$, where q_{dtjs} is the total volume exported in the destination–time–seller–product tuple. The market share of each inside variety is computed as the ratio $\mathcal{S}_{dtjs} = q_{dtjs}/q_{dt}$, and demand for the variety can be expressed as (Berry, 1994)

$$ln(\mathcal{S}_{dtjs}) - ln(\mathcal{S}_{dt\varnothing}) = \delta_{dt} + \theta_s + \alpha p_{dtjs} + \sigma ln(\mathcal{NS}_{dtjs}) + \xi_{tjs}, \tag{C3}$$

where $\mathcal{NS}_{dtjs} = q_{dtjs}/q_{dtj}$ is the share of the seller's variety in the nest.

C.2 Estimation and Results

The object of interest for the purpose of this paper is the vector of seller-specific quality shifters, θ_s. These are recovered as fixed effects in the least squares estimation of the following equation:

$$ln(\mathcal{S}_{dtjs}) - ln(\mathcal{S}_{dt\varnothing}) = \delta_{dt} + \theta_s + \alpha p_{dtjs} + \sigma ln(\mathcal{NS}_{dtjs}) + \varepsilon_{dtjs}. \tag{C4}$$

While I use OLS as a benchmark, both the nest shares and the prices are likely correlated with the quality deviations ξ_{tjs}, captured in the econometric error ε_{dtjs}. I instrument for prices using the daily international cotton price as an exogenous shifter. The instrument is constructed as follows. Denote with p^c_τ the international price of raw cotton at market closing time on day τ. Let $q_{dtjs\tau}$ be the size of an export transaction to destination d in year t by seller s in product category j on date τ. The weighted average cotton price for the $dtjs$ tuple is $p^c_{dtjs} = \sum_\tau (q_{dtjs\tau}/q_{dtjs}) p^c_\tau$. To account for the potentially weaker pass-through from cotton price shocks to garment prices when manufacturers use synthetic, man-made or other fibers, I include an interaction of the cotton price with the share of cotton-made products in the total volume exported by the seller in the HS4 category in that year, sh^c_{Jst}. To instrument for the share of the variety in the nest, I exploit the number of sellers who export the product to the destination and the number of products that the seller exports to the destination: N^s_{dtj} and N^j_{dts}, respectively.

The data contain 489,451 observations aggregated at the destination–year–product–seller level. From these data, I take 1,000 bootstrap samples with replacement of entire vectors of observations by seller. Each $\hat{\theta}_s$ is obtained as the average estimate of the corresponding parameter across all bootstrap iterations.

Table C1 reports the results of the OLS and IV estimation procedures on equation (C4). Diagnostics and the first stage corresponding to the IV approach are presented in Table C2.

The instruments are strong and correlated with the endogenous regressors with the expected sign: increases in the average international cotton price shift the price of garments upwards, and less so if the garment has a low-cotton composition; the share of any seller in the nest is negatively related to the number of competitors and the presence of the seller in other products (nests). The instrumentation strategy corrects the OLS coefficients in the intuitive direction. In Table C1, the coefficient on the price appears significantly upwardly biased under the OLS benchmark. The IV estimation produces a large, negative and strongly significant coefficient, compatible with manufacturers facing highly elastic demand. Similarly, the OLS estimation suggests excessively high within-nest correlation in preferences for varieties, which the IV estimation corrects downwards. In the instrumented specification, the nest coefficient is smaller in magnitude and not statistically different from zero.

Of relevance for the developments in this paper, the IV estimation recovers 8,644 seller-specific demand shifters, conditional on price. Precisely estimated, 95% of the estimates are statistically different from zero.

Table C1: Estimation of Demand Equation

	OLS	IV
	(1)	(2)
Price coefficient α:		
Coefficient	-0.059	-0.850
Standard Error	0.004	0.143
CI Lower Bound	-0.066	-1.182
CI Upper Bound	-0.052	-0.602
Nest Share coefficient σ:		
Coefficient	0.330	0.094
Standard Error	0.004	0.069
CI Lower Bound	0.322	-0.023
CI Upper Bound	0.338	0.253
Quality estimates θ_s:		
Average t-statistic	23.964	7.881
Median t-statistic	17.410	7.385
25^{th} percentile t-statistc	7.810	5.553
75^{th} percentile t-statistc	31.952	9.543
Estimated coefficients	8,644	8,644
Statistically non-zero	8,159	8,185
Number of Observations	489,451	
Number of Destination-Year	1,063	

The table reports the results of the OLS and IV estimations of equation (C4) in columns (1) and (2), respectively. The estimation procedure is bootstrapped 1,000 times, drawing with replacement all observations ($dtjs$ tuples) for each seller. Coefficients, standard errors and confidence intervals are constructed using the means and standard deviations of the bootstrap-based distributions of the coefficient estimates. The top panel reports information on α, the coefficient on prices, p_{dtjs}, followed by estimates of σ, the coefficient on the nest share, $ln(\mathbb{NS}_{dtjs})$. The estimation routines recover $8,664$ seller-specific intercepts in the demand equation. The table shows summary statistics of the t-statistics of these intercepts, constructed as ratios between the bootstrapped means and standard deviations. The critical value used for the count of seller shifters statistically different from zero is 1.96.

Table C2: First Stage of Demand Estimation

Equation:	(1) p_{dtjs}		(2) $ln(\mathbb{NS}_{dtjs})$	
	Coeff	S.E.	Coeff	S.E.
Share of cotton: sh^c_{Jst}	0.445	0.111	0.021	0.021
Raw cotton price: p^c_{dtjs}	0.097	0.213	-0.800	0.044
Interaction: $p^c_{dtjs} \times sh^c_{Jst}$	-0.217	0.117	-0.070	0.021
Nr. of Sellers: N^s_{dtj}	-0.004	0.000	-0.007	0.000
Nr. Of Products: N^j_{dts}	-0.029	0.019	0.071	0.005
Fixed effects	dt,s			
Observations	489,451			
Kleibergen-Paap rk (F weak)	21.08			
Kleibergen-Paap rk (LM underid)	104.27			

The table reports results of the first stage in the IV strategy used for the estimation of equation (C4). This corresponds to the second stage outcomes reported in column (2) of Table C1. In all cases, the first-stage equations include seller–product–year fixed effects. The block labeled (1) corresponds to the price equation, where p_{dtjs} is instrumented, while the block labeled (2) presents the equation for the share of variety in the nest, $ln(\mathbb{NS}_{dtjs})$. In all cases, the equations include destination–year and seller fixed effects. The excluded instruments are the weighted average international raw cotton price, p^c_{dtjs} (see the text for the construction of the average); the share of cotton exports in the seller's trade, sh^c_{Jst}; their interaction; and the numbers of sellers who export the product to the destination and of products that the seller exports to the destination, N^s_{dtj} and N^j_{dts}, respectively. Standard errors are clustered at the variety level. The bottom panel of the table reports test statistics for underidentification (LM) and weak instruments (F). The LM test corresponds to the Kleibergen–Paap rank test, and in all cases, all exogenous regressors (including the fixed effects) are partialed out (χ^2 distributed).

Table C3: Measures of Sellers' Dispersion in Product Markets

	$Dispersion_j = St.Dev._j(\hat{\theta}_s)$			
	(1)	(2)	(3)	(4)
$Not\ Cotton_j$=1	0.244***			
	(0.063)			
$Female_j$=1		0.304***		
		(0.053)		
$Complexity_j$			0.232**	
			(0.113)	
$Quality\ Ladder_j$				0.077**
				(0.031)
R^2	0.05	0.07	0.03	0.03
Obs.	241	241	218	222

Heteroskedasticity-robust standard errors in parentheses. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. All columns correspond to OLS regressions whose outcome is our measure of supplier dispersion. This is defined as the standard deviation across all the estimated $\hat{\theta}_s$ for sellers active in product category j. See Appendix C for details on the estimation of $\hat{\theta}_s$. The regressor in column (1), $Not\ Cotton_j$, is a dummy that takes value one if product category j corresponds to garments made of fabrics other than cotton. These include synthetic fibers, man-made fibers, furs, wools and mixed fabrics. The regressor used in column (2), $Female_j$, is a dummy that takes value one if product category j corresponds to garments that are for women, for girls or unisex (this includes categories without gender classification). In column (3), $Complexity_j$ is a measure of the complexity of product j. For some product categories, the customs data allow matching of export orders with the imported inputs used to produce them. Details on this matching can be found in Cajal-Grossi et al. (2022). The number of different inputs (typically fabrics) combined in the production of a garment is used in the industry as a measure of garment complexity. $Complexity_j$ is the (log) average count of inputs used to produce orders in product category j. Finally, column (4) uses the measure of the length of the quality ladder in product j as estimated in Khandelwal (2010) (available for download from the author's website). The unit of observation in all regressions is a product category disaggregated at the HS6 level. The number of observations varies across columns because not all regressors are defined for all product categories. In particular, the measure of input complexity is not available for product categories for which inputs and outputs cannot be matched at the order level. This is a feature of the institutional environment and the use of export declaration procedures. In addition, the quality ladders constructed in Khandelwal (2010) for product codes in the US are not available for 19 of the 241 product codes in the Bangladeshi data.

Table C4: Quartiles of Sellers' Dispersion: Main Product Codes

Code	Type	Product Description
Quartile # 1 (Low Dispersion)		
610510	Knitted	Shirts; men's or boys', of cotton, knitted or crocheted
611020	Knitted	Jerseys, pullovers, cardigans, waistcoats and similar articles; of cotton, knitted or crocheted
610462	Knitted	Trousers, bib and brace overalls, breeches and shorts; women's or girls', of cotton, knitted or crocheted
610342	Knitted	Trousers, bib and brace overalls, breeches and shorts; men's or boys', of cotton, knitted or crocheted
610610	Knitted	Blouses, shirts and shirt-blouses; women's or girls', of cotton, knitted or crocheted
Quartile # 2		
610910	Knitted	T-shirts, singlets and other vests; of cotton, knitted or crocheted
620342	Woven	Trousers, bib and brace overalls, breeches and shorts; men's or boys', of cotton (not knitted or crocheted)
611090	Knitted	Jerseys, pullovers, cardigans, waistcoats and similar articles; of textile materials (other than wool or ...
620462	Woven	Trousers, bib and brace overalls, breeches and shorts; women's or girls', of cotton (not knitted or crocheted)
620520	Woven	Shirts; men's or boys', of cotton (not knitted or crocheted)
Quartile # 3		
620333	Woven	Jackets and blazers; men's or boys', of synthetic fibres (not knitted or crocheted)
610821	Knitted	Briefs and panties; women's or girls', of cotton, knitted or crocheted
620630	Woven	Blouses, shirts and shirt-blouses; women's or girls', of cotton (not knitted or crocheted)
620433	Woven	Jackets and blazers; women's or girls', of synthetic fibres (not knitted or crocheted)
620463	Woven	Trousers, bib and brace overalls, breeches and shorts; women's or girls', of synthetic fibres ...
Quartile # 4 (High Dispersion)		
621710	Woven	Clothing accessories; other than those of heading no. 6212 ...
621210	Woven	Brassieres; whether or not knitted or crocheted
620193	Woven	Anoraks (including ski-jackets), wind-cheaters, wind-jackets and similar articles; men's or boys', of man-made ...
620112	Woven	Coats; men's or boys', overcoats, raincoats, car-coats, capes, cloaks and similar articles, of cotton, other than ...
620293	Woven	Anoraks (including ski-jackets), wind-cheaters, wind-jackets and similar articles; women's or girls', of man-made ...

The table lists the five largest products (according to exported volumes) within each of the four quartiles of product-level seller dispersion. As described in the main text, seller dispersion is measured with the standard deviation of seller-specific shifters, $\hat{\theta}_s$, for each product j. The metric is thus constructed as $\sqrt{\sum_{s\in j}(\hat{\theta}_s - \overline{\hat{\theta}_s})^2/\#\{s \in j\}}$, where $s \in j$ denotes all sellers who sell product j at least once, $\#\{s \in j\}$ is the cardinality of this set, and $\overline{\hat{\theta}_s}$ corresponds to the average across all sellers in seller set j. Product markets are organized in quartiles, such that those in quartile 1 feature low dispersion across sellers and those in quartile 4 feature high dispersion. The first column reports the HS6 code corresponding to j, the second column classifies the garment according to its type (knitted or woven), and the final column shows the product description of the HS6 classification of j.

Figure C1: Noncompliance and Estimated Thetas

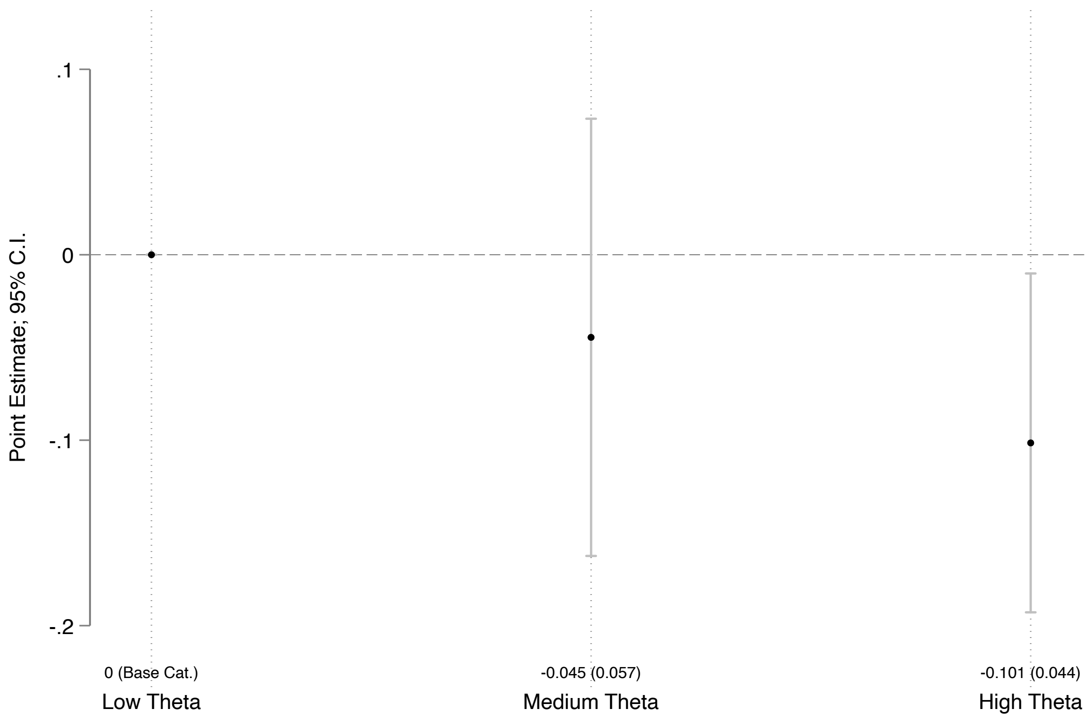

The figure plots estimated coefficients of a regression of plant-level noncompliance scores and the estimated θ_s, grouped in three equally sized bins: low, medium and high. The underlying regression runs on a panel of 193 plants, assessed over five cycles of compliance evaluations. These evaluations are performed by the Better Work program of the International Labour Organization and produce a *noncompliance score*. The score, which takes nonnegative integer values, ranges from 6 to 38 in the data. High scores reflect a high number of social compliance violations. The left-hand-side variable in the underlying regression is the log of the noncompliance score. The regression runs over 965 observations (i.e., plant–cycle pairs). The specification includes fixed effects for the cycle and the seller's cohort (the year that the seller is first observed in the customs data) and largest product (the HS6 code accounting for the largest share in the seller's exports). The regression controls for the number of employees in the plant and the log exported volumes of the seller. The regressor of interest arranges the estimated θ of the 193 plants in thirds, where *Low Theta* corresponds to the third of sellers with the lowest estimated θ. This is the excluded base category, and we identify the coefficients for the medium and high categories. For the latter, the point estimate reads as follows: sellers with a high estimated θ have 10% lower noncompliance scores than those of sellers with a low θ. This difference is significant at the 5% level, with standard errors clustered at the product level.

D Robustness

D.1 Alternative Definitions and Specifications

This section shows that the main results presented in Table 6 are robust to alternative specifications and variations in operational definitions of key variables.

Dispersion measure. I show that the results are robust to alternative treatments of the quality dispersion measure used in Table 6. As discussed in the main text (see Section 3.2), the use of a discrete measure of dispersion simplifies interpretation of the coefficient on the triple-interaction variable. In practice, this approach is implemented by arranging products j in ascending order according to the standard deviation across seller characteristics in j. The first robustness exercise here reconstructs these quantiles based on the interquartile range as the underlying dispersion measure. For each product category, the dispersion in seller quality is measured as $IqR_j(\hat{\theta}_s) = \hat{\theta}_s^{75^{th}} - \hat{\theta}_s^{25^{th}}$. Products are arranged in ascending order according to this dispersion measure, and the support is partitioned into quartiles.

The rest of the robustness exercises in this appendix are carried out with the continuous dispersion measure, $Dispersion_j$, used in the main text. Using the quartiles would require keeping track of three coefficients to characterize the triple-interaction terms (corresponding to a vector $\boldsymbol{\beta}_3$). The continuous measure, instead, involves a scalar β_3. Given the large number of robustness combinations discussed in this appendix (78 specifications), this continuous metric proves useful as the benchmark.

The third robustness exercise on the dispersion measure instruments for $Dispersion_j$ by using exogenous characteristics varying across product categories. For the instrumentation strategy, I leverage the correlations presented in Table C3, showing that product categories with higher quality dispersion across sellers typically correspond to products that are made of materials other than cotton, are produced for women, require the combination of more inputs and correlate with quality ladder length in a developed downstream market. Using the labels of Table C3, the exogenous shifters of $Dispersion_j$ are the dummies $Not\ Cotton_j$ and $Female_j$ and the continuous variables $Complexity_j$ and $Quality\ Ladder_j$. The four shifters are interacted with $Large_b$ and $Post_{c(bj)}$ to instrument for $Dispersion_j \times Large_b$ and $Dispersion_j \times Post_{c(bj)} \times Large_b$ (eight exogenous instruments for two potentially endogenous variables).

The IV approach mitigates any potential concerns over the use of a function of estimated demand shifters as a right-hand-side variable. One such concern arises if unaccounted drivers of buyer experimentation, such as buyer taste for a particular type of seller in a specific category, shapes the market composition (i.e., which sellers enter product category j) or the structure of the demand schedule (in particular the setting of prices).

Table D1 presents the results of the various robustness tests discussed here. Columns (1) and (3) of Table D1 simply reproduce columns (1) and (3) of Table 6 for ease of comparison. Column (2) presents the estimation of equation (5) using the quartiles based on the IqR measure. Column (4) presents the results of estimating equation (5) while instrumenting for the continuous seller measure. The first-stage coefficients and diagnostics are collected in Table D2. The results presented here are consistent both qualitatively and quantitatively with those of the main exercise in the body of the paper.

Buyer size cutoff. The regressions in the body of the paper focus on interactions of the cost shock, dispersion in the environment and a buyer-specific size measure. In the baseline specification, size is captured by a dummy that takes value one if the buyer is among the top 200 importers, labeled $Large_b$. Alternatively, this cutoff is set at 100, 150 and 300. The results of estimating equation (5) and its variations with these different operational definitions of $Large_b$ are included in the specification curve of Figure 5. The coefficient of interest remains robust to changes in the buyer size cutoff. Moreover, the small differences in point estimates across these alternatives are intuitive: cutoffs capturing smaller buyers have a slightly smaller coefficient estimate, and the smaller the buyer set, the noisier is the point estimate.

Outcome variable. The outcome studied in equation (5) is constructed to collect the amount of experimentation that buyers engage in when they first enter a product category. In practice, this is the log count of one-off interactions observed in the data during the buyer's first year in the product category ($Experiment_{bj}^{entry}$). An alternative definition uses the count of one-off interactions until the first recurrent relationship is observed—this is, until the first time that the buyer is observed to have a trade relationship exceeding one quarter in total length. The results of estimating equation (5) under the alternative definitions of the outcome variable are all collected in Figure 5, where Outcome 1 corresponds to the baseline metric (aggregating over the whole first year) and Outcome 2 corresponds to the alternative explored here. The figure shows that when a linear model is used, the coefficient of interest using Outcome 2 is not systematically above or below the one obtained when Outcome 1 is used. Under a nonlinear count model, discussed below, the coefficient obtained when Outcome 2 is used on the left-hand side of the estimating equation tends to be lower than the one obtained with Outcome 1. This is consistent with Outcome 2 being more skewed toward low counts than Outcome 1. In all cases, with linear and nonlinear models, the coefficients of interest based on the alternative outcome variable remain very close to those in the baseline.

Other moments. A potential concern in the interpretation of $\hat{\beta}_3$ in equation (5) is that it might capture variation responding to other moments of the seller heterogeneity distribution that are correlated with the dispersion measure. Two particularly important characteristics are the median seller quality and the number of available suppliers. In principle, high-dispersion markets may coincide with those in which sellers are, in general, of low quality. If this is the case, a negative $\hat{\beta}_3$ may simply respond to the gains from experimentation being low rather than to risk in the experimentation process. Similarly, high dispersion may arise in thin markets, i.e., those with few potential suppliers. To study this interaction between the dispersion mechanism and other market characteristics, I augment specification (5) to include appropriate interactions of the shock:

$$\begin{aligned} Experiment_{bj}^{entry} = \delta_b + \delta_{jc(bj)} + \beta_1 Post_{c(bj)} \times Large_b + \beta_2 Dispersion_j \times Large_b \\ + \beta_3 Dispersion_j \times Post_{c(bj)} \times Large_b + \gamma_2 Moment_j \times Large_b \\ + \gamma_3 Moment_j \times Post_{c(bj)} \times Large_b + \epsilon_{bj}, \quad \text{(D1)} \end{aligned}$$

where $Moment_j$ is either the median quality across available suppliers (i.e., observed trading) in market j or is the number of available sellers. Richer specifications in which both sets of interactions are included are also explored. The β_3 estimates from these augmented specifications are presented under the "Median Interaction" and "Count Interaction" labels in Figure 5. Across all augmented specifications, $\hat{\beta}_3$ remains negative, significant and close in magnitude to the baseline point estimate.

Nonlinear count model. A specific characteristic of the econometric specification in equation (5) is that the outcome is a count variable. The results discussed in the main text estimate the parameters of interest by means of OLS with the outcome (+1) in logs. I discuss an alternative that accounts for the existence of zeros in the data (i.e., entrants with no one-off interactions upon entry) and for the count nature of the outcome variable. I use a PPML procedure with high-dimensional fixed effects, following Correia et al. (2020) and Correia et al. (2019), under the mean specification

$$E[Experiment_{bj}^{entry}|\boldsymbol{X}] = exp(\delta_b + \delta_{jc(bj)} + \beta_1 Post_{c(bj)} \times Large_b + \beta_2 Dispersion_j \times Large_b \\ + \beta_3 Dispersion_j \times Post_{c(bj)} \times Large_b), \quad \text{(D2)}$$

where $Experiment_{bj}^{entry}$ is the count outcome in levels and $\boldsymbol{X}$ is shorthand notation for all fixed effects and regressors. All the specifications run in linear form are repeated with this nonlinear approach. The results are presented on the right-hand side of Figure 5, which shows the coefficients of the Poisson procedure. Focusing on the specification with the baseline structure, highlighted in blue in the figure, I estimate a coefficient of -0.282.

D.2 Buyer-Specific Choice Sets

This section explains the construction of buyer-specific choice sets. These are used as an alternative to the baseline definition of the environment that a buyer faces when experimenting with suppliers. The intuition for the construction of these hypothetical choice sets is that the buyer's environment includes sellers whom the buyer knows (via past trade) and sellers unknown to the buyer but similar on observable characteristics to the buyer's trade partners.

I consider a *trade instance* to be a buyer–product–seller–quarter combination in the data ($bjst$, in the notation of the paper). In other words, an instance corresponds to realized trade between parties in a given quarter and product combination. I define the potential set of sellers with whom the buyer can trade in each instance to include sellers s' satisfying the following conditions:

1. The buyer trades with the seller, i.e. $s' = s$.
2. Seller s' is in the choice set if it is active and known to the buyer:
   (a) Buyer b has traded with s' before t, perhaps in another product category.
   (b) Seller s' has traded product j with some buyer before quarter t.
   (c) Seller s' has not exited the product group (HS4) by date t.

3. Seller s' is in the choice set if it is active, not known to the buyer and similar to the buyer's chosen seller s in prices and quantities:

   (a) Buyer b has not traded with s' before t in any product category.

   (b) Seller s' has traded product j with some buyer before date t.

   (c) Seller s' has not exited the product group (HS4) by date t.

   (d) For each seller satisfying conditions (a)-(c), compute the squared Mahalanobis distance between s' and s in two dimensions:[53]

      - The average price across all transactions that the seller has in product category j with any buyer.
      - The average size (volume) across all transactions that the seller has in product category j with any buyer.

   (e) Based on a preset cutoff, retain sellers s' who are *sufficiently similar* to s based on the Mahalanobis distance over prices and quantities. Four different cutoffs are considered: (i) all sellers, irrespective of distance; (ii) the closest x sellers where x coincides with the 75^{th} percentile on the size distribution of choice sets; (iii) the closest x sellers where x coincides with the 50^{th} percentile on the size distribution of choice sets; and (iv) the closest x sellers where x coincides with the 25^{th} percentile on the size distribution of choice sets.

Based on the criteria above, for every trade instance of buyer b in product category j, there is a set of sellers with whom the buyer could have traded, i.e., a hypothetical choice set for each trade instance.

To create a measure of the dispersion that the buyer faces when entering product category j, I consider all the sellers s' in all hypothetical choice sets of the buyer in product–year jy. This leads to four possible constructions of buyer–product–year-specific choice sets, combining the four different cutoff points on the similarity metric (see point (e) above). Mimicking the baseline measure of dispersion in seller characteristics, the metrics are constructed as the standard deviation across the seller-specific shifters in the demand equation of Appendix C.

I consider a final alternative to construct the consideration sets that takes into account all sellers in the product category, excluding those who have ever traded with the buyer. This construction leads to buyer–product-specific measures of dispersion but leverages only sellers not linked to the buyer.

Table D3 shows the size distribution of the different consideration sets across all buyer–product entry instances. The baseline refers to product-specific consideration sets (the level of aggregation is the product category j). Under this definition, the consideration sets are the largest: across all entry instances in the data, the average number of sellers on which the quality dispersion is computed is $1,441.67$. The alternative metric that leaves out current or past trade partners of the buyer produces sightly smaller choice sets, suggesting that the majority of the buyers trade with a small number of sellers. The largest buyers, who trade with many sellers, are those for whom the dispersion under the baseline construction and the

[53]The Mahalanobis distance between vector $\boldsymbol{x}_s$ and $\boldsymbol{x}_{s'}$ is given by $\sqrt{(\boldsymbol{x}_s - \boldsymbol{x}_{s'})' V^{-1} (\boldsymbol{x}_s - \boldsymbol{x}_{s'})}$, where V is the covariance matrix of $\boldsymbol{x}$.

alternative option differ most. The retrospective set includes all of the buyer's past sellers, following the criterion in point 2 above. The count of entries that have a retrospective set is smaller than the count with the baseline measure, reflecting that some entries are the buyer's first in the sector or in the HS4 product. The choice sets based on seller similarity with different cutoffs follow the construction described in point 3 above. These constructions require computation of an inverse covariance matrix, which is not possible in six product categories in which the matrix is singular or near singular.[54]

[54]The products for which the Mahalanobis distance scores cannot be computed are 610792, 611512, 611522, 621141, 621310 and 621510.

Table D1: Robustness: Measures of Dispersion

	$Experiment^{entry}_{bj}$			
	(1)	(2)	(3)	(4)
$Quartile\ Dispersion_j$=2 × $Large_b$=1	0.009	-0.014		
	(0.014)	(0.014)		
$Quartile\ Dispersion_j$=3 × $Large_b$=1	-0.010	-0.002		
	(0.015)	(0.015)		
$Quartile\ Dispersion_j$=4 × $Large_b$=1	0.018	0.017		
	(0.021)	(0.022)		
$Quartile\ Dispersion_j$=2 × $Post_{c(bj)}$=1 × $Large_b$=1	0.033	0.037		
	(0.040)	(0.039)		
$Quartile\ Dispersion_j$=3 × $Post_{c(bj)}$=1 × $Large_b$=1	-0.041	0.032		
	(0.043)	(0.041)		
$Quartile\ Dispersion_j$=4 × $Post_{c(bj)}$=1 × $Large_b$=1	-0.154***	-0.087*		
	(0.053)	(0.049)		
$Dispersion_j$ × $Large_b$=1			0.004	-0.062
			(0.023)	(0.058)
$Dispersion_j$ × $Post_{c(bj)}$=1 × $Large_b$=1			-0.161***	-0.229**
			(0.055)	(0.112)
Fixed Effects	b, jt	b, jt	b, jt	b, jt
Measure	St.Dev.	IqR	St.Dev.	St.Dev.
Estimation	OLS	OLS	OLS	IV
R^2	0.25	0.25	0.25	
Obs.	60,297	60,297	60,297	58,265

Standard errors in parentheses, bootstrapped 400 times stratifying by product. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The regressions follow the specification in (5). Columns (1) and (3) simply reproduce the baseline estimates presented in columns (1) and (3) of Table 6 for ease of comparison. Column (2) replaces the quartile-based measure of dispersion with an alternative constructed using quartiles based on the IqR measure. Column (4) estimates equation (5) instrumenting the dispersion measure by means of the IV approach described in this appendix.

Table D2: First Stage Regressions of IV on Dispersion Interactions

	$Dispersion_j \times Large_b$=1	$Dispersion_j \times Large_b$=1 $\times Post_{c(bj)}$=1
	(1)	(2)
$Post_{c(bj)}$=1 $\times$ $Large_b$=1	-0.071	1.087***
	(0.068)	(0.075)
$Not\ Cotton_j$=1 $\times$ $Large_b$=1	0.094**	-0.002***
	(0.037)	(0.001)
$Female_j$=1 $\times$ $Large_b$=1	0.131***	-0.002***
	(0.032)	(0.001)
$Complexity_j$ $\times$ $Large_b$=1	0.221*	-0.002
	(0.116)	(0.002)
$Quality\ Ladder_j$ $\times$ $Large_b$=1	-0.014	-0.001**
	(0.019)	(0.000)
$Post_{c(bj)}$=1 $\times$ $Not\ Cotton_j$=1 $\times$ $Large_b$=1	0.035	0.132***
	(0.027)	(0.046)
$Post_{c(bj)}$=1 $\times$ $Female_j$=1 $\times$ $Large_b$=1	0.035	0.169***
	(0.029)	(0.041)
$Post_{c(bj)}$=1 $\times$ $Complexity_j$ $\times$ $Large_b$=1	0.091	0.303***
	(0.114)	(0.113)
$Post_{c(bj)}$=1 $\times$ $Quality\ Ladder_j$ $\times$ $Large_b$=1	0.028**	0.018
	(0.014)	(0.025)
Fixed Effects	b, jt	b, jt
F-Test	7.408	651.41
Obs.	58,265	58,265
SH J Test of OverId. (χ^2)	7.323	
KP LM Test of UnderId. (χ^2)	23.335	

Standard errors in parentheses, clustered by product. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table presents the first-stage regressions of the two instrumented interactions in the IV regression of column (4) in Table D1. The bottom of the table presents statistics for Fisher tests for each individual regression. The χ^2 statistic for the Sargan–Hansen test of overidentifying restrictions is also included and does not reject the null hypothesis of valid instruments (not rejected with a p-value of 0.2920). The Kleibergen–Paap rk statistic for the null hypothesis of underidentification is also included (rejected with a p-value of 0.0015).

Table D3: Size Distribution of Consideration Sets under Alternative Constructions

Description	Aggr.	Count	Mean	St. Dev.	p10	p25	p50	p75	p90
Baseline: All sellers in product	j	63659	1441.67	1416.16	133	321	1015	1984	3701
Leave-Out: All sellers in product, excluding trade partners	bj	63659	1439.93	1415.61	132	320	1014	1983	3696
Retrospective: All past sellers of the buyer	bjy	44942	5.18	5.88	1	2	3	6	12
Similar: All sellers similar to the buyer's partners	bjy	59821	410.32	382.77	25	90	306	628	1016
Retrospective + Similar	bjy	63659	387.88	383.42	14	67	272	605	1000
Similar sellers within large radius	bjy	59821	309.13	219.22	25	90	306	537	537
Similar sellers within medium radius	bjy	59821	195.88	113.19	25	90	261	261	276
Similar sellers within small radius	bjy	59821	77.97	35.89	25	83	83	83	93

The table presents the size distribution of consideration sets in the buyer–product entry instances in the data. The baseline refers to product-specific consideration sets (the level of aggregation is the product category j). Under this definition, across all entry instances in the data, the average number of sellers on which the quality dispersion is computed is $1,441.67$. The leave-out alternative excludes all sellers who have ever traded with the buyer in the product, such that the dispersion measure is buyer–product specific (bj). The retrospective set includes all past sellers of the buyer and is constructed as described in point 2 of Appendix D.2. The similar set includes all sellers similar to the buyer's trade partners, as described in point 3 of Appendix D.2. The set labeled "Retrospective + Similar" combines the previous two. The bottom three sets use only similar sellers within different distance score radii. These correspond to the trimmings defined in point 3, item (e), subitems (ii) to (iv) in Appendix D.2.

Table D4: Consideration Sets: Alternative Constructions

Panel A: All buyers

	$Experiment_{bj}^{entry}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$Large_b$=1 × $Post_{c(bj)}$=1	0.212***	0.098	0.114**	0.106**	0.066	0.113*
	(0.082)	(0.061)	(0.056)	(0.051)	(0.042)	(0.063)
$Large_b$=1 × *Dispersion*	0.004	-0.003	0.002	0.000	0.001	-0.001
	(0.021)	(0.016)	(0.015)	(0.013)	(0.012)	(0.016)
$Post_{c(bj)}$=1 × $Large_b$=1 × *Dispersion*	-0.161***	-0.079**	-0.090**	-0.088***	-0.067**	-0.088**
	(0.054)	(0.037)	(0.035)	(0.034)	(0.029)	(0.039)
Fixed Effects	b, jt	b, jt	b, jt	b, jt	b, jt	b, jt
Construction	Baseline	Retro+Sim	Retro+Sim	Retro+Sim	Retro+Sim	Leave-Out
Similarity Cut-Off	.	All	75^{th} *pctile*	50^{th} *pctile*	25^{th} *pctile*	.
Aggregation	j	bjy	bjy	bjy	bjy	bj
R^2	0.25	0.25	0.25	0.25	0.25	0.25
Obs.	60,297	58,649	58,649	58,649	58,649	60,297

Panel B: Incumbent buyers

	$Experiment_{bj}^{entry}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$Large_b$=1 × $Post_{c(bj)}$=1	0.223**	0.126**	0.140**	0.123**	0.075*	0.133**
	(0.088)	(0.060)	(0.060)	(0.054)	(0.043)	(0.068)
$Large_b$=1 × *Dispersion*	0.008	0.016	0.019	0.010	0.005	0.009
	(0.027)	(0.018)	(0.018)	(0.017)	(0.015)	(0.018)
$Post_{c(bj)}$=1 × $Large_b$=1 × *Dispersion*	-0.163***	-0.094**	-0.104***	-0.096***	-0.069**	-0.096**
	(0.058)	(0.038)	(0.038)	(0.034)	(0.030)	(0.042)
Fixed Effects	b, jt	b, jt	b, jt	b, jt	b, jt	b, jt
Construction	Baseline	Retro+Sim	Retro+Sim	Retro+Sim	Retro+Sim	Leave-Out
Similarity Cut-Off	.	All	75^{th} *pctile*	50^{th} *pctile*	25^{th} *pctile*	.
Aggregation	j	bjy	bjy	bjy	bjy	bj
R^2	0.26	0.26	0.26	0.26	0.26	0.26
Obs.	32,462	32,099	32,099	32,099	32,099	32,462

Standard errors in parentheses, bootstrapped 400 times stratifying by product. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table re-estimates the baseline regression of experimentation upon buyers' entries into product markets, using different conceptualizations of dispersion in the relevant environment. Panel A exploits all buyer–product entry instances, while Panel B uses only entry instances by *incumbent* buyers, i.e., buyers with at least one entry prior to and at least one entry after RP. In all cases, the outcome $Experiment_{bj}^{entry}$ is the log count of one-off interactions that buyer b has in product j within the first year after entering the sourcing market. $Large_b$ is an indicator that takes value one if the buyer is among the top 200 buyers in the industry. The postshock variable $Post_{c(bj)}$ is a dummy that takes value one if the entry of buyer b in product j occurs in a quarter t after the RP event. *Dispersion* is a variable measuring seller quality dispersion in the environment that the buyer faces when entering a product category, and different columns adopt different constructions for this variable. Columns labeled (1) reproduce the baseline results for ease of comparison. They use the baseline definition of dispersion, namely, the standard deviation in quality across all sellers ever active in product category j. In this case, the consideration sets are product specific (Aggregation: j). Columns (2) to (5) compute dispersion within buyer–product–year-specific choice sets (Aggregation: bjy), following the routines described in Appendix D.2. In all cases, all past known sellers are included in the consideration set (Retrospective). In addition, all sellers similar to effective trade partners (Similar) are also included in column (2). Columns (3) to (5) subsequently reduce the buyer's consideration sets to keep the closest (most similar) sellers, with cutoffs set on the size of the choice set coinciding with the 75^{th}, 50^{th} and 25^{th} percentiles, respectively. Columns labeled (6) recompute the dispersion across all sellers ever active in the product category but exclude (or *leave out*) the buyer's trade partners. In this case, the consideration sets are buyer–product specific (Aggregation: bj). Refer to Table D3 for the size distribution of choice sets across the different columns here.

E Appendix Tables and Figures

Table E1: Performance on Social Compliance and Product Specialization

Panel A: Linear Specifications				
	(1) $Time\ Records_{sn}$	(2) $Overtime_{sn}$	(3) $Overcrowd_{sn}$	(4) $Subcontract_{sn}$
$Quality\ Ladder_{\bar{j}(s)}$=Medium	0.066	0.068*	0.003	0.111*
	(0.044)	(0.036)	(0.015)	(0.059)
$Quality\ Ladder_{\bar{j}(s)}$=High	0.152**	0.159***	0.063**	0.180
	(0.073)	(0.050)	(0.025)	(0.118)
Fixed Effects	$n, c(s)$	$n, c(s)$	$n, c(s)$	$n, c(s)$
Count of Plants	209	209	209	113
Controls	Yes	Yes	Yes	Yes
R^2	0.08	0.09	0.05	0.13
Obs.	431	431	431	179
Panel B: Nonlinear Specifications				
	(1) $Time\ Records_{sn}$	(2) $Overtime_{sn}$	(3) $Overcrowding_{sn}$	(4) $Subcontract_{sn}$
$Quality\ Ladder_{\bar{j}(s)}$=Medium	0.066*	0.072**	0.001	0.105**
	(0.039)	(0.033)	(0.021)	(0.050)
$Quality\ Ladder_{\bar{j}(s)}$=High	0.159**	0.153***	0.073**	0.172
	(0.068)	(0.040)	(0.052)	(0.111)
Fixed Effects	$n, c(s)$	$n, c(s)$	$n, c(s)$	$n, c(s)$
Count of Plants	209	209	167	113
Controls	Yes	Yes	Yes	Yes
Obs.	431	431	316	179

Standard errors in parentheses, clustered at the product level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The regressions are based on an unbalanced sample of 209 plants assessed between one and four times over consecutive years starting in 2014. The underlying data are collected by the Better Work Program of the International Labour Organization as part of its periodic social compliance evaluations of enrolled plants. Each of these plants is observed exporting during the period 2005-2015 in the customs records. From the trade data, I obtain the main product exported by the plant, denoted $\bar{j}$ (i.e., the HS6 code with the highest share in the exporter's sales). I assign each of these main products the measure of scope for quality differentiation constructed in Khandelwal (2010) using US data. I organize these into three equally sized categories reflecting low, medium and high scope for quality differentiation. The low category is excluded in all regressions, and $Quality\ Ladder_{\bar{j}(s)}$=Medium or =High denotes the other two categories. The outcomes are dummy variables reflecting whether the firm has been found to have engaged in social compliance breaches of different types. $Time\ Records_{sn}$ takes value one if plant s assessed in evaluation cycle n does not produce working time records that reflect hours actually worked. $Overtime_{sn}$ takes value one if the employer does not comply with the legal limits on overtime work. $Overcrowd_{sn}$ takes value one if the evaluator found that the production process is overcrowded (i.e., at least part of the work space has more workers per unit of space than recommended). $Subcontract_{sn}$ takes value one if the plant does not comply with requirements connected to subcontracted work. All specifications include cycle fixed effects (n) and fixed effects for the cohort (year) of the plant, as measured by the first export transaction in the customs data ($c(n)$). In addition, all specifications control for the total export volumes of the plant and the number of employees. Panel A reports OLS regressions, and Panel B shows marginal effects after probit for each of the outcomes. In both panels, column (4) is based on a smaller number of firms because the corresponding question on subcontracting was included only in some cycles. The discrepancy across panels in the number of firms in column (3) reflects the fact that the outcome is perfectly predicted in the nonlinear model by being in one cycle. Those observations are therefore dropped.

Table E2: Industrial Accidents Associated to Brands' Names - Selected Episodes

Date	Episode	Brands (allegedly involved)
April 2005	Spectrum Factory collapse during night of forced work: 64 deaths + 75 injuries	Inditex, Carrefour, among others
February 2006	Fire at KTS Textile, child labor, locked exits: 61 deaths + 100 injuries	ATT, VIDA Andrew Scott
February 2006	Collapse Phoenix Building, unauthorized plant: 22 deaths + 50 injuries	Unreported. Destinations: Germany, Switzerland and Denmark
February 2006	Imam Group explosion and blocked exits: 57 injuries	Kmart, Folsom Corporation.
March 2006	Fire Sayem Fashions: 3 deaths + 50 injuries	Inditex (various brands).
February 2010	Fire in Garib and Garib, no ventilation, suffocation, blocked exits: 21 deaths + 57 injuries	H&M, El Corte Ingles.
December 2010	Fire That's It Sportswear (Hameem Group), no exits, no drills, illegal inaccessible top floors: 29 deaths + 11 injuries	Gap, PVH Corp., VFCorporation.
December 2011	Boiler Explosion Eurotex, stampede, collapse of stairs with exits blocked: 2 deaths + 64 injuries	Tommy Hilfiger (PVH Corp.), Inditex, Gap, C&A.
November 2012	Fire in Tazreen Fashions, captive workers, child labor: 115 deaths + 200 injuries	C&A, Walmart, others (unauthorized production)
April 2013	Building Collapse in Rana Plaza, forced labor, blocked exits: 1,132 deaths + 1,800+ injuries	Benetton, Kik, Mango, Primark, Walmart

The table presents a nonexhaustive compilation of industrial accidents involving well-known brands over the sample period, according to media accounts. The list is nonexhaustive in that (i) it covers only episodes reported in international media outlets and (ii) it refers to brands that these outlets "named and shamed", which are often believed to be a small subset of all brands sourcing from offering plants. Episodes listed from 2005 to 2011 are taken from reports by the Clean Clothes Campaign. Details on recent episodes are obtained from international news outlets.

Table E3: Relationship Survival

Interaction i (1)	Share of relationships ending in i (2)	Probability of continuing to $i+1$ (3)	Improvement in survival in i relative to $i-1$ (4)
1	0.619	0.381	-
2	0.172	0.548	0.440
3	0.071	0.661	0.205
4	0.040	0.713	0.079
5	0.025	0.743	0.042
6	0.017	0.769	0.035
7	0.012	0.781	0.016
8	0.009	0.797	0.020
9	0.006	0.816	0.025
10	0.005	0.811	-0.006
11	0.004	0.814	0.004
12	0.003	0.835	0.025
13	0.003	0.834	-0.002
14	0.002	0.842	0.010
15	0.002	0.841	-0.001
16	0.001	0.846	0.006
17	0.001	0.841	-0.005
18	0.001	0.844	0.004
19	0.001	0.846	0.001
20	0.001	0.835	-0.012
21	0.001	0.813	-0.026
22	0.001	0.828	0.018
23	0.000	0.830	0.002
24	0.000	0.829	-0.001
25	0.000	0.858	0.036
26	0.000	0.833	-0.030
27	0.000	0.754	-0.095
28	0.000	0.823	0.093
29	0.000	0.817	-0.007
30	0.000	0.777	-0.050

The table studies the survival patterns of (uncensored) relationships in the data. A relationship is a buyer–seller–product triplet, where a product corresponds to an HS6 code. There are over 270 thousand such triplets in the data. Column (1) labels the first thirty interactions within relationships, $i = 1, \ldots 30$. An interaction is a quarter of trade. Column (2) reports the share of relationships that do not survive after interaction i. Column (3) shows the probability of the relationship surviving to interaction $i+1$ which, by construction, conditions on survival up until i. Column (4) measures the improvement in the probability of survival relative to the probability in the previous period. It is calculated as the entry in column (3) for i minus that of $i-1$ divided by that of $i-1$. Naturally, the improvement in survival is not defined for $i = 1$. Given these definitions, as an example, the third row of the table would read as follows: 7.1% of the relationships in the data (column (2), 0.071) end on the third interaction (column (1), 3); conditional on reaching a third interaction, 66.1% of relationships continue onto a fourth interaction (column (3), 0.661); the probability of survival after the third interaction is 20.5% higher than that after the second interaction (column (4), 0.205).

Table E4: Probability of One-off Interactions after the Formation of a Recurrent Relationship

	Probability of $one-off_{bjt}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$After\ Recurrent_{bjt}$	-0.364***	-0.420***	-0.351***	-0.117***	-0.138***	-0.155***	-0.345***
	(0.002)	(0.004)	(0.004)	(0.004)	(0.005)	(0.006)	(0.006)
Model	Probit	Linear	Linear	Linear	Linear	Linear	Linear
Fixed Effects	.	.	b, jt	bj, jt	bj, jt	bj, jt	b, jt
Duration	1y+	1y+	1y+	1y+	2y+	3y+	3y+
R^2	.	0.17	0.36	0.52	0.48	0.46	0.35
Obs.	94,170	94,170	94,170	93,917	75,081	58,802	58,802

Standard errors in parentheses, clustered at the buyer level. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The outcome in all specifications is a dummy that takes value one if the buyer has at least one *one-off* interaction in a given product–quarter combination. The outcome is thus defined at the level of a buyer–product–time tuple, and in the data, it equals one for 33% of the bjt combinations. By construction, the regression sample retains only buyer–product–quarter combinations with nonzero trade. In addition, we restrict attention to buyer–products active for at least a year. The regressor of interest, $After\ Recurrent_{bjt}$, is an absorbent indicator that takes value one for all buyer–product–time triplets that take place after the buyer trades *recurrently* (i.e., for at least a second time) with at least one seller in the product category. Column (1) shows probit marginal effects, and all other columns report results of linear probability models. Columns (1) and (2) include no fixed effects. Column (3) augments the specification in column (2) with buyer and product–time fixed effects. Column (4) retains the sample from column (3) but exploits within-buyer–product variation by including buyer–product and product–time fixed effects. Columns (5) and (6) repeat the exercise from column (4) while trimming the sample to consider buyer–product pairs active for two or more and three or more calendar years, respectively. Column (7) reproduces the regression from column (3) on the restricted sample from column (6).

Table E5: Traded Volumes in One-off and First Interactions: Correcting for Partial First Quarter

	q_{sbjt}			
	(1)	(2)	(3)	(4)
$One-off_{sbjt}$	-0.590***			
	(0.012)			
$\mathbf{I}\{i_{sbjt}=1^{st}\}$		-0.357***	-0.093***	-0.231***
		(0.027)	(0.018)	(0.027)
Fixed Effects	bj,sj,jt	sb,jt	sb,jt	sb,jt
Duration	Any	Any	1y+	1y+
Buyers	All	All	All	Large
R^2	0.56	0.54	0.54	0.41
Obs.	526,163	414,829	243,704	122,610

Standard errors in parentheses, clustered at the buyer level. $*(p<0.10), **(p<0.05), ***(p<0.01)$. The table reproduces the results of Table 3, correcting the volume traded in the first calendar quarter of a relationship, for entries that occur at different points in the quarter. To this end, the volume in the first quarter of all relationships is aggregated over a 92-day window from the first observed instance of trade. The outcome in all specifications is this corrected log volume traded by the seller–buyer–product–quarter tuple, q_{sbjt}. Column (1) studies the correlation between the outcome and an indicator that takes value one if the tuple corresponds to a one-off interaction, i.e., a buyer–seller–product triplet interacting for one quarter only ($one-off_{sbj}$). The specification includes buyer–product, seller–product and product–quarter fixed effects. Columns (2)-(4) study the first interaction in a relationship by means of an indicator that takes value one if a given quarter corresponds to the first interaction of the buyer–seller triplet ($\mathbf{I}\{i_{sbjt}=1^{st}\}$). In each of columns (2)-(4), seller–buyer and product–quarter fixed effects are included. Different columns study different samples. Columns (1) and (2) include all tuples not affected by censoring. Column (3) restricts attention to buyer–seller–product triplets active for at least one year. Column (4) further restricts the sample to only the 200 largest buyers, who account for 70% of the volumes traded in the industry throughout the sample period.

Table E6: Industry Patterns: Before and After Rana Plaza

Panel B: Product-Year Outcomes

	(1) q_{jy}	(2) $Count^{s}_{jy}$	(3) $Count^{b}_{jy}$	(4) $Count^{e}_{jy}$	(5) $Count^{sb}_{jy}$
year=2014	0.370***	0.085***	0.085**	-0.071*	0.125***
	(0.091)	(0.029)	(0.035)	(0.041)	(0.037)
year=2013	0.267***	0.069***	0.073***	-0.027	0.105***
	(0.071)	(0.021)	(0.024)	(0.033)	(0.024)
year=2011	-0.121*	-0.028	-0.051	0.005	-0.072**
	(0.071)	(0.028)	(0.031)	(0.037)	(0.032)
year=2010	-0.060	-0.044	-0.083**	0.023	-0.103**
	(0.092)	(0.035)	(0.041)	(0.046)	(0.042)
year=2009	-0.296***	-0.072**	-0.130***	-0.023	-0.143***
	(0.092)	(0.035)	(0.043)	(0.051)	(0.044)
year=2008	-0.244**	-0.062	-0.098**	0.038	-0.118**
	(0.103)	(0.042)	(0.047)	(0.054)	(0.050)
year=2007	-0.258**	-0.101**	-0.142***	0.047	-0.172***
	(0.104)	(0.043)	(0.050)	(0.053)	(0.052)
Fixed Effects	j	j	j	j	j
R^2	0.89	0.95	0.93	0.88	0.94
Obs.	1,684	1,684	1,684	1,684	1,684

Panel B: Buyer-Product-Year Outcomes

	(1) q_{bjy}	(2) $Count^{sb}_{bjy}$	(3) $Count^{OneOff}_{bjy}$	(4) $Share^{OneOff}_{bjy}$
year=2014	0.148***	0.033***	0.012*	-1.397***
	(0.022)	(0.007)	(0.007)	(0.425)
year=2013	0.062***	0.033***	0.014**	-0.871***
	(0.016)	(0.004)	(0.005)	(0.337)
year=2011	-0.051***	-0.002	0.021***	2.180***
	(0.016)	(0.004)	(0.005)	(0.345)
year=2010	-0.089***	-0.007	0.029***	2.395***
	(0.020)	(0.006)	(0.007)	(0.394)
year=2009	-0.139***	-0.023***	0.015**	2.538***
	(0.023)	(0.008)	(0.007)	(0.439)
year=2008	-0.100***	-0.017**	0.032***	2.890***
	(0.024)	(0.008)	(0.008)	(0.480)
year=2007	-0.099***	-0.026***	0.015*	2.422***
	(0.026)	(0.009)	(0.008)	(0.522)
Fixed Effects	bj	bj	bj	bj
R^2	0.73	0.75	0.53	0.60
Obs.	131,916	131,916	131,916	131,916

Standard errors in parentheses, clustered at the product level in Panel A and at the buyer level in Panel B. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The table shows results of linear regressions of yearly dummies on outcomes defined at the product-year level (jy in Panel A) or buyer–product–year level (bjy in Panel B). The estimating equation for all columns in Panel A is $y_{jy} = \delta_j + \sum_y \beta_y \mathbb{I}\{year = y\} + \epsilon_{jy}$ and that of Panel B $y_{bjy} = \delta_{bj} + \sum_y \beta_y \mathbb{I}\{year = y\} + \epsilon_{bjy}$. δ_j and δ_{bj} correspond to product and buyer–product fixed effects, and in all cases, the coefficients of interest are the set of β_ys. The year 2012 (just before RP) is excluded as the base category, and all observations for 2005 and 2015 are dropped from the sample. The outcome variables are as follows: q_{jy} corresponds to the log exported volume in jy; $Count^{s}_{jy}$ corresponds to the log count of uniquely identified sellers exporting in jy; $Count^{b}_{jy}$ corresponds to the log count of uniquely identified buyers importing in jy; $Count^{e}_{jy}$ is the log count of entrants in jy, that is, the number of uniquely identified buyers observed for the first time in jy; $Count^{sb}_{jy}$ is the log count of unique buyer–seller pairs trading in the jy combination; q_{bjy} is the log volume imported by buyer b in the product–year combination jy; $Count^{sb}_{bjy}$ is the log count of buyer–seller relationships trading in bjy; $Count^{OneOff}_{bjy}$ is the log count of these relationships that correspond to one-off interactions; and $Share^{OneOff}_{bjy}$ is defined as the percentage of all volume (q_{bjy}) traded via one-off interactions in the bjy combination.

Table E7: Buyer Status and Buyer-Product Entry Counts

	Incumbents	Entrants	Exiters	All Buyers
$Post_{c(bj)}$=0	26,740	0	28,742	55,482
$Post_{c(bj)}$=1	7,170	1,007	0	8,177
Total	33,910	1,007	28,742	63,659

The table divides all buyer–product entry instances observed in the data by buyer status and the pre-/post-RP collapse indicator. Incumbent buyers are defined as those with at least one product entry prior to and at least one product entry after RP. Entrants are buyers who do not feature in the data prior to RP and have at least one entry after RP. Exiters are buyers with at least one entry prior to RP and no further entries after RP. Note that exiters might still trade in the data, i.e., continue to purchase products in categories that they entered prior to the collapse.

Table E8: Selection into Product Categories

	$Dispersion_j$		
	(1)	(2)	(3)
$Age\ Trend_{n(bj)}$	0.003***	0.002***	0.001***
	(0.000)	(0.000)	(0.000)
$Large_b$=1 × $Age\ Trend_{n(bj)}$		0.001***	0.002***
		(0.000)	(0.000)
$Post_{c(bj)}$=1 × $Age\ Trend_{n(bj)}$			0.001***
			(0.000)
$Large_b$=1 × $Post_{c(bj)}$=1 ×$Age\ Trend_{n(bj)}$			-0.001***
			(0.000)
Fixed Effects	b, t	b, t	b, t
R^2	0.19	0.19	0.19
Obs.	61,476	61,476	61,476

Standard errors in parentheses, bootstrapped 400 times stratifying by product. $*(p < 0.10), **(p < 0.05), ***(p < 0.01)$. The unit of observation is an entry instance of a buyer into a product category at a point in time. The outcome is defined at the level of the product only: $Dispersion_j$ measures seller quality dispersion within a product j. All specifications include buyer and time fixed effects, where the time corresponds to the calendar quarter in which the buyer entered product j. $Age\ Trend_{n(bj)}$ is a linear trend for the sequence of entries of the buyer. It takes value one if j is the first product that the buyer enters, value two if it is the second product that the buyer enters, etc. $Large_b$ is an indicator that takes value one if the buyer is among the top 200 buyers in the industry. $Post_{c(bj)}$ is an indicator taking value one if the entry of buyer b in product j occurs in a quarter t taking place after the RP event.