Vertical Integration & Relational Contracts: The Threat Point Effect

Russell Morton*

January 10, 2024

Click Here for Latest Version

Abstract

Firms in low-income countries often face frictions sourcing inputs and rely on vertical integration or relational contracts rather than spot markets. Empirical evidence highlights that vertical integration and relational contracts co-exist both within industries and individual firms. What are the economic and policy implications of this co-existence? This paper identifies and quantifies a new mechanism which I call the *threat point effect*, which is the contract change from firms improving their bargaining position due to partial vertical integration. I build and estimate a structural model to quantify the threat point effect in the context of a large Indian garment manufacturer that adds integrated capacity bargaining with its relational fabric suppliers. Model estimation uses the universe of the manufacturer's fabric purchase transaction data. The threat point effect reduces input prices by 6.7% for small constrained suppliers that highly value the reduced exposure to demand shocks in the relational contract. I analyze counterfactuals that *i)* increase downstream buyer competition and *ii)* create the missing market, specifically insurance against demand shocks, that leads small firms to offer discounts in the relational contract. Only the latter shifts surplus to small firms.

*University of Michigan, rpmorton@umich.edu. I am extremely grateful to Lauren Falcao Bergquist, Zach Brown, Ying Fan, and Dean Yang for invaluable their guidance and support. I also thank Achyuta Adhvaryu, Jie Bai, Hoyt Bleakley, Max Huppertz, Chuqing Jin, Francine Lafontaine, David A. Miller, Ameet Morjaria, Emir Murathanoglu, and Anant Nyshadham for their excellent feedback and comments, as well as seminar participants at EARIE 2023 and the University of Michigan. I appreciate the team at Shahi Exports and Good Business Lab for teaching me about fabric procurement and providing the data. All errors are my own.

1 Introduction

In many low-income countries, institutions to support interfirm trade are limited. For example, firms might not trust courts to enforce contracts effectively (Banerjee and Duflo, 2000). As a response to these contracting and other frictions,[1] firms often rely on either vertical integration or relational contracts to circumvent input markets. When firms vertically integrate, they bypass input markets by sourcing in-house. Alternatively, firms can use relational contracts to facilitate cooperative, rather than opportunistic, trading that creates surplus for both parties. These contracts prevent opportunistic behavior by punishing it through ending the relationship, blocking access to future surplus (Baker et al., 2002). Empirical evidence highlights that vertical integration and relational contracts often co-exist within an industry, and even within individual firms.[2]

What are the economic and policy implications of the co-existence of vertical integration and relational contracts within a firm? I propose that co-existence improves the integrated firm's bargaining position within its relational contracts. Consider a buyer bargaining with external suppliers, although similar effects apply when suppliers integrate. When the buyer integrates, its threat point (outside option) changes from other external suppliers to sourcing internally. This shift in the buyer's threat point alters its bargaining position with external suppliers, empowering the buyer to appropriate more of the relational contract surplus.[3]

This contract change caused by vertical integration altering bargaining positions has policy implications for policy makers with distributional preferences. Bargaining positions shape how buyers and suppliers share gains from trade, but we know little about how this surplus is shared in relational contracts, despite their ubiquity in low-income countries (Macchiavello, 2022; Macchiavello and Morjaria, 2023). Furthermore, understanding the distribution of relational contract surplus speaks to concerns that global value chains (GVC), which often rely on relational contracts, are inequitable with low prices for low-income country producers (Boudreau et al., 2023)—the World Bank notes that "gains from GVC participation are not

[1]Other frictions include hold-up (which often result from contracting frictions) (Lafontaine and Slade, 2007), input market power (Hortaçsu and Syverson, 2007), and quality non-contractability (Hansman et al., 2020).

[2]Many firms in low-income countries are partially integrated (*i.e.*, they source inputs both internally and externally). For example, 62% of integrated firms in Karnataka source some inputs externally (Garg et al., 2023) and Breza and Liberman (2017) study a partially integrated Chilean retailer with relational suppliers. As integration suggests the presence of frictions, many external suppliers for partially integrated firms are likely relational suppliers.

[3]When a supplier partially integrates, its threat point changes to selling to the integrated buyer rather than market buyer, improving contracts with external buyers.

distributed equally."

This paper studies how vertical integration affects bargaining in the context of relational contracts between a large Indian garment manufacturer (hereafter, buyer) and its external fabric suppliers. In their relational contract, the buyer receives a discount relative to market prices in exchange for reducing suppliers' quantity variability (*i.e.*, demand assurance); many relational contracts include quantity assurance and a price wedge relative to market prices. Suppliers benefit as reduced quantity variability helps them control costs and stabilize profits, which can be especially important as risk aversion can shape firm behavior in low-income countries.[4] I leverage a combination of rich transaction-level data on the universe of the garment manufacturer's fabric purchases and an increase in vertical integration, as the buyer partially vertically integrates by constructing a new in-house fabric mill, to analyze how integration affects contracts with 471 suppliers, both relational and market.

I define the threat point effect as the change in contracts with external firms caused by the change in outside option (threat point) due to vertical integration. As the buyer benefits from the relational contract through discounts, the threat point effect in my setting is the additional decrease in relational prices after the buyer adds integrated capacity (*i.e.*, the extra discount the buyer receives after integration).

Then, I build and estimate a structural model to quantify the threat point effect. The model illustrates how vertical integration changes contracts between a buyer and supplier in a relational contract, showing that threat point effects are heterogeneous with respect to buyer and supplier characteristics. Motivated by both reduced-form evidence of the relational contract and institutional details of the empirical setting, the model features a buyer and a supplier bargaining over a relational contract pricing function that determines surplus. Pre-integration, trading with the market is the threat point (outside option) for both the buyer and supplier; after integration, the buyer's threat point becomes its integrated supplier. Vertical integration allows the buyer to credibly threaten to use its integrated supplier to displace the external supplier.[5] The model emphasizes heterogeneity of the threat point effect, which varies based on the pre-integration sharing rule for and level of relational contract surplus. The threat point effect is large when pre-integration surplus is high, as occurs for small constrained suppliers; then, the buyer can leverage its threat point change

[4]For example, risk aversion influences productive behavior and contracting of low- and middle-income country coffee mills (Blouin and Macchiavello, 2019) and Ghanaian agriculture producers (Karlan et al., 2014).

[5]I support this assumption in a companion paper (Morton, 2023) that empirically documents substitutability between vertical integration and relational contracts leveraging exogenous variation in relational contracting costs at the product level.

to reduce relational prices.

I relate the sharing rule for surplus and level of surplus to buyer and supplier characteristics through two parameters: a buyer bargaining parameter and supplier risk aversion. The buyer bargaining parameter determines the buyer's share of relational contract surplus. Supplier risk aversion shapes the level of relational surplus as risk averse suppliers highly value profit smoothing from demand assurance in the relational contract. I model the buyer, a large sophisticated firm purchasing nearly 230 million USD of fabric yearly, as risk neutral.

I estimate the model to quantify the threat point using unusually rich transaction data on the universe of fabric purchases by the buyer, where I observe the price, quantity, detailed product information, date, and supplier identity (including purchases from integrated suppliers). The model quantifies the threat point effect for varied buyer and supplier types, which is key as the threat point effect is especially large for risk-averse suppliers. This heterogeneity is empirically relevant as a wide range of firms use relational contracts—from major U.S. airlines (Gil et al., 2022) to small-holder Rwandan coffee farmers (Macchiavello and Morjaria, 2020). I provide a novel approach to separately identify the buyer bargaining parameter from supplier risk aversion based on the shape of the relational contract pricing function. Risk-averse suppliers prefer contracts that equalize profits; therefore, relational contracts with more risk-averse suppliers provide more profit stability against demand shocks by adjusting prices. Conversely, the buyer bargaining parameter shifts prices regardless of demand shock.

I find that the decrease in relational prices caused by vertical integration (*i.e.*, the threat point effect) can be as large as 6.7%, depending on supplier risk aversion and the buyer bargaining parameter. At 6.7%, the threat point effect is economically meaningful–larger than Indian VAT for apparel at the time.[6] It is largest for small constrained suppliers that highly value demand assurance in the relational contract because these suppliers receive relational surplus even when prices are low. I validate the model by showing it produces both a close match to untargeted moments and superior out-of-sample fit than two alternative estimators. I also validate the model using the buyer's construction of a new mill as a natural experiment. I compare contracts before and after integration with external suppliers, leveraging heterogeneity in supplier exposure to integration.[7]

This finding has important distributional implications, as vertical integration decreases relational contract prices most for small constrained suppliers that highly value demand assurance in the relational contract. Because these suppliers receive relational surplus even

[6]GST rates for textiles and apparel in India.

[7]Supplier exposure to integration derives from the supplier's product mix, with suppliers producing fabric that can be made at integrated supplier classified as exposed to integration.

when prices are low, the buyer can leverage its change in threat point to negotiate lower prices without threatening the relational contract's existence. Small firms (suppliers in this setting) in low-income countries are also highly policy relevant. They comprise the majority of firms in low- and middle-income countries (Hsieh and Olken, 2014) and receive specific attention from policy makers, including the IMF and World Bank.

I evaluate two counterfactual policies: *i)* increasing downstream buyer competition, inspired by the focus on market power in industrial organization; and *ii)* creating the missing market for suppliers to insure profit against demand shocks given evidence of the deleterious effects of missing markets from development economics. Increasing downstream buyer competition has a trivial effect on prices paid to small risk-averse suppliers. The buyer responds to enhanced buyer competition by increasing the share of surplus that suppliers receive—raising prices—to prevent them from switching to another relational buyer in the long-run (as relational contracts require time and effort to establish, suppliers cannot switch in the short-run). However, once the buyer's threat point is its integrated supplier, the surplus in the relational contract is so small that even allocating all of it to the supplier does not meaningfully shift prices. It follows that omitting the threat point effect results in large overestimates of the benefit of this policy—when the buyer's threat point is market supply, there is sufficient surplus that reallocation changes prices.

Second, creating the missing market for profit insurance against demand shocks increases relational prices and effectively eliminates the threat point effect. Suppliers value insurance because it reduces their reliance on the relational contract to smooth profits, thereby mitigating the buyer's ability to use its threat point change to reduce relational prices. This result suggests policies to improve small firms' ability to smooth profit, such as increasing access to savings accounts to facilitate self-insurance or financial products to hedge exposure (*e.g.*, derivatives), can generate large benefits through both facilitating intertemporal smoothing and shifting bargaining positions and, thereby, profits.

I contribute to a rich literature studying how relational contracts react to shocks to economic conditions, such as supply or demand, or market structure and competition–specifically, whether the relational contracts survive or breakdown as a result (Macchiavello and Morjaria, 2015, 2020, 2023; Ghani and Reed, 2022; Gil et al., 2022). I focus on surplus sharing as the primary outcome of interest and analyze how buyer and supplier characteristics shape the level of and sharing rule for relational contract surplus, as well as how firms can behave strategically to increase their share of the surplus by integrating. Additionally, I study the effects of a different shock–*vertical* organization of production–rather than shocks to market structure, which is *horizontal* organization of production.

The contracting literature (*e.g.,* Harris and Nguyen (2023)) also studies how transaction governance forms (*i.e.*, whether transaction terms are determined by spot markets, relational contracts, or an integrated firm) interact. Specifically, prior work shows that effective spot markets can, on the one hand, make relational contracts less desirable but, on the other hand, support relational contracts (*e.g.*, by providing transparent information about the value of the outside option). This paper also studies interactions between transaction governance forms, but analyzes the interaction between vertical integration and relational contracts and considers interaction within a firm rather than across firms,[8] illustrating that vertical integration can make relational contracts more desirable by improving contract terms.

This research also complements prior work on the motivations for and effects of vertical integration. Both theoretical and empirical studies illustrate that frictions which limit the effectiveness of market transactions result in vertical integration (Williamson, 1979; Lafontaine and Slade, 2007; Acemoglu et al., 2009; Forbes and Lederman, 2009; Macchiavello, 2012; Breza and Liberman, 2017; Hansman et al., 2020; Garg et al., 2023).[9] The literature on the effects of vertical integration focus on quantities, such as increasing input purchases due to reducing double marginalization (Hortaçsu and Syverson, 2007), or on other downstream buyers, such as foreclosure to or increased input prices for downstream rivals to weaken downstream competition (Asker, 2016). I contribute to these literatures by proposing a novel effect of vertical integration–changing bargaining position and contracts with upstream firms–that creates an additional rationale for integration. This mechanism can affect either upstream or downstream firms, depending on which firm integrates, although it applies only to partially integrated firms.

Understanding the motivations for and effects of vertical integration influences efficiency through competition policy, which can have large welfare effects in low-income countries where many markets are fragmented and uncompetitive (Bergquist and Dinerstein, 2020; Banerjee et al., 2022; Bergquist et al., 2023). Because the threat point effect can shift input market prices, it can have either pro- or anti-competitive effects downstream depending on the direction of the input market price change–the threat point effect decreases input prices when a buyer partially integrates and vice-versa when a supplier partially integrates.

[8]My companion paper (Morton, 2023) discusses interaction effects across firms, highlighting possible anticompetitive effects of relational contracts.

[9]Other studies provide evidence that firms integrate to transfer and spread intangible inputs, such as management practices and organizational capabilities (Atalay et al., 2014), across the firm. These two motivations are similar as, arguably, using vertical integration for intangible assets reflects the difficulties in contracting for intangible inputs given that they are typically firm-specific and hard to measure.

However, in settings where relational contract quantities are fixed and are inframarginal, as occurs when spot market input transactions persist after integration, the threat point effect has neglibile direct effects on competition because it does not change marginal input costs–only inframarginal input prices change.

Because the trade-off between incentives and insurance (hold-up and demand assurance in my setting) is central to many vertical relationships (Lafontaine and Slade, 2007), my empirical model incorporates insurance and takes seriously both estimating risk aversion and separately identifying it from bargaining parameters. Standard models of vertical relationships (*e.g.*, Cuesta et al. (2019)) omit risk aversion, resulting in underestimates of the bargaining parameter for the more risk-averse party. And, even with correct parameter estimates, models generate incorrect prices if one firm provides insurance (*e.g.*, quantity assurance) but it is outside the model. In my empirical setting, omitting the demand assurance the supplier receives in the relational contract (a vertical relationship) would lead to underestimated supplier surplus and overestimated prices.[10] Additionally, I develop a novel strategy to separately identify bargaining parameters from risk aversion by leveraging variation in both price levels and the convexity of prices with respect to capacity.

Last, I contribute to the literature estimating structural models of relational contracts by developing an approach that does not require observing relational contracts breaking down (Galenianos and Gavazza, 2017; Startz, 2021; Harris and Nguyen, 2023). I directly estimate the value of the outside option leveraging data on costs and commonly available spot market prices. As relational contracts should generally not break down in equilibrium (and breakdowns may suggest the presence of changes in the economic environment), this approach can have broad applicability and generalize to other settings where either no relational contract breakdowns occur or breakdowns are not clearly observable.

The remainder of the paper is organized as follows. Section 2 describes the data and context. Section 3 presents the conceptual model of the relational contract pre- and post-vertical integration. In section 4, I estimate and validate the conceptual model. Then, I use it to quantify the threat point effect for different buyer and supplier types. Section 5 discusses the policy implications of the threat point effect and analyzes policies to reallocate surplus to small firms. Section 6 concludes and considers broader implications.

[10]Fortunately, the biases from omitting risk aversion and from omitting insurance offset each other. However, they are unlikely to counteract each other perfectly and could be different orders of magnitude, biasing price estimates.

2 Data and Context

My empirical setting focuses on an Indian garment manufacturer purchasing fabric from internal and external suppliers, where external suppliers include both relational contract suppliers and market suppliers not in a long-term relationship with the buyer. I first describe the data before providing context on the firm and the increase in integrated capacity at the buyer. Last, I describe the hold-up problem in this setting that motivates the relational contract, including documenting empirical patterns in the data consistent with my characterization of the relational contract.

2.1 Data

I use two data sources: transaction data and cost data. The rich transaction data provide information about the universe of fabric purchases completed by the buyer between September 1, 2016 and December 31, 2019, including the transaction price, the date, the supplier which provided the fabric (including for internal suppliers), the quantity, and detailed information about the specific fabric. The 34,681 transactions represent over one billion square meters of fabric purchased from 622 distinct suppliers (471 pre-integration). There are two integrated suppliers, with one built during the timeframe of the study and the other joining the firm many years prior to the start of my data. Additionally, the transaction data include estimated production costs for fabrics which are produced at the internal mills; these estimates are for a mill that is not capacity constrained and roughly represent average variable costs.

The cost data from the internal mills provide data on capacity utilization and costs at the two internal mills operated by the buyer. These data are at a monthly level from March 2019 to April 2020 and include a broad support of levels of capacity utilization from 41% to 102%.

2.2 Context: Vertical Integration at a Large Buyer

Figure 1: Integrated Mill Constructed in June 2018

Prior to building its integrated mill, the buyer sources $\sim$ 250 million square meters of fabric per year (worth $\sim$ 228 million USD over 10,000 distinct orders) both at a pre-existing internal mill ($\sim$ 21% on average) and from external suppliers, both relational contract suppliers and market suppliers. Figure 2 illustrates three key facts about this integration that make this case study well-suited to analyze the threat point effect. First, the integration is large enough for the buyer to credibly threaten to bring production in-house. After integration, the percent produced internally increased greatly from $\sim$ 21% to $\sim$ 44%, with roughly half produced at the newly built internal mill. For comparison, the largest external supplier produced only 7%. Second, sourcing from external suppliers remains important after integration, as external suppliers still produce the majority of fabric. It follows that prices from external suppliers are still economically meaningful for the buyer, such that the buyer still cares about bargaining with them. Last, downstream demand from end clients (*e.g.,* garment retailers like Old Navy) does not change much (decreasing slightly to $\sim$ 215 million square meters of fabric per year), such that the buyer must make strategic decisions to reduce some combination of the number of external suppliers and the average volume per external supplier.

Figure 2: Fabric Production Over Time by Supplier Type

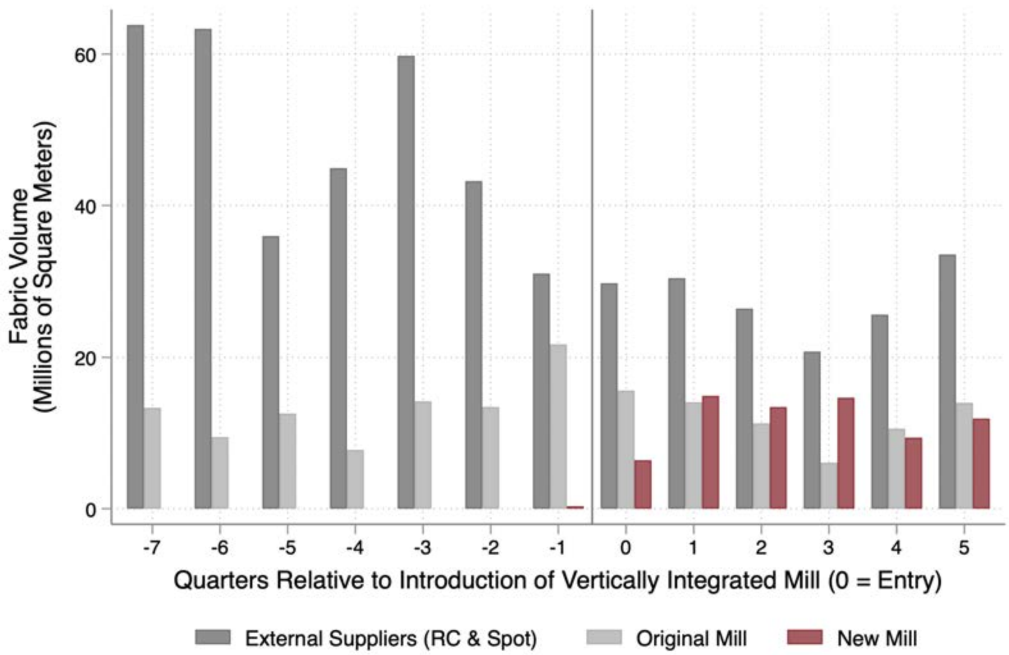

Note: Data from universe of fabric transactions by buyer. Bars represent the quantity of fabric purchased during the quarter from each supplier type.

However, the integration is not so large that the buyer must reduce relational quantities– the buyer could reduce volumes from only market suppliers. Figure 3 illustrates this point, as the quantity produced internally at the new mill is much less than the pre-integration quantities sourced from market suppliers. It follows that the buyer can (and, as I later show, does) displace market rather than relational suppliers with integrated supply.

Figure 3: Fabric Production Over Time for Relational Contract and Market Suppliers

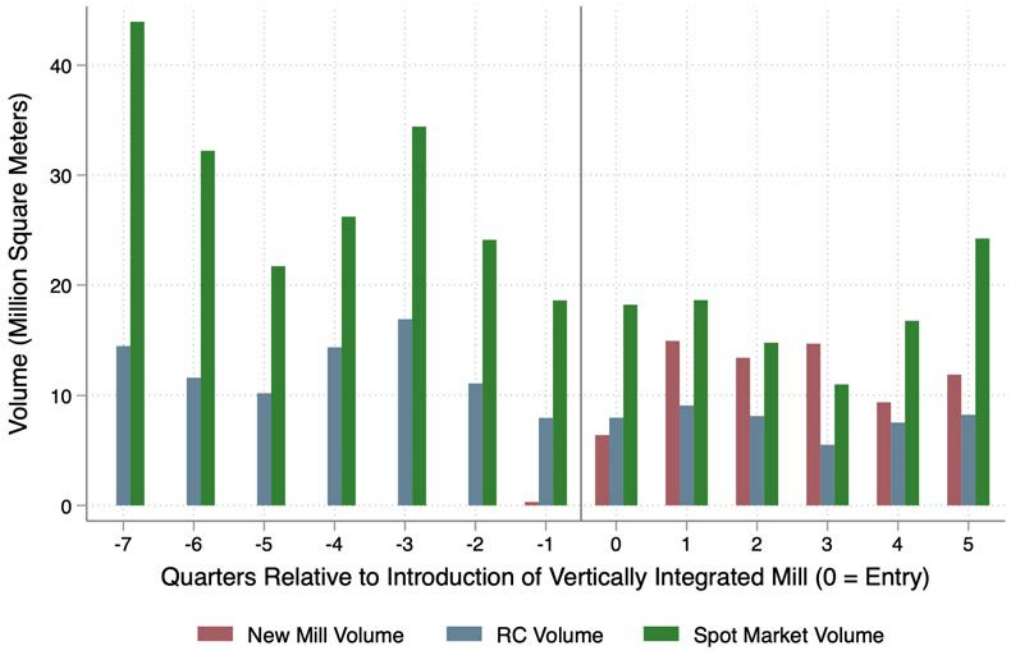

Note: Data from universe of fabric transactions by buyer. Bars represent the quantity of fabric purchased during the quarter from each supplier type.

2.3 Context: Fabric Procurement and Garment Manufacturing

Garment manufacturing, in this context, begins with designers who work for the garment manufacturer (the buyer) conceptualizing and planning the garment, sketching out the shape and color, and choosing the broad type of fabric (shown in Figure 4). The buyer then selects a supplier (a mill) to provide fabric for a small number of examples of the garment, called "samples"; although the buyer does pay for this fabric, this cost is largely negligible given the small quantity of fabric ordered and does not constitute a commitment from the supplier to the same price for future orders even for the same fabric. Fabric procurement costs for sample garments are so trivial that the buyer does not even record them in a centralized database. The samples are subsequently shown to an end client (*e.g.*, Old Navy), who then decides which samples to order and the quantities for any order it places. Only once the order is placed by the end client does the buyer return to the initial supplier to purchase the fabric necessary for the garment. Finalized designs are confirmed just-in-time (*i.e*, as close to shipment and sell date as possible) to ensure that they reflect recent trends. In general, there are only roughly 30 days between the initial fabric purchase for the sample

Figure 4: Garment Manufacturing Process

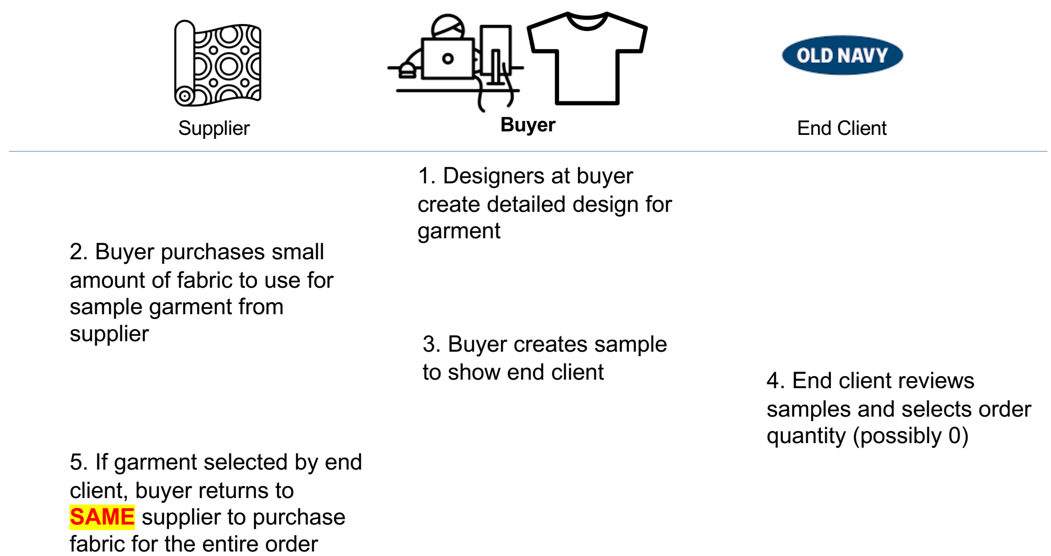

and sourcing fabric for approved orders, in between which the sample is shown to the end client and the quantity is confirmed.

A key feature in this process is that the supplier which provided fabric for the sample almost always receives the order for the fabric because the buyer does not want to change the fabric after the end client has seen it.[11] This supplier rigidity reflects that many subtle features of fabric are difficult to specify and copy, leading the buyer to believe that the end client might not accept a garment with a similar, but not identical, fabric made by another supplier. Additionally, timelines are short, preventing switching to a cheaper supplier and engaging in trial-and-error over time to match the fabric. The buyer's emphasis on ensuring that the garment produced for the end client is the same as the sample emphasizes the importance of the end client to the buyer. The importance of the end client is consistent with other research about the same buyer demonstrating that the buyer's production process prioritizes serving end clients over maximizing productive efficiency (Adhvaryu et al., 2019).

[11] Based on discussions with fabric procurement managers at the buyer, the probability of a supplier receiving a fabric order given that it provided the sample and is a large enough mill is at least 80%.

Therefore, when the buyer selects a supplier to provide a sample, it is selecting the supplier for the entire order despite the fact that neither the buyer nor the supplier know if the garment will be selected nor what the quantity would be (although both the buyer and the supplier likely have some beliefs over purchase probabilities and volumes from the history of transactions with the end client). Unsurprisingly, in the rare cases where the supplier is changed, production is typically brought in-house where quality can presumably be controlled intensely and for orders with long time before delivery to ensure that there is sufficient time to match the fabric.

The rigidity in supplier selection once the supplier has provided fabric for the sample creates a classic hold-up problem, allowing the supplier to charge a price that extracts at least the static surplus (*i.e.,* profit) the buyer receives from the order. The supplier could likely even extract some of the dynamic surplus from the relationship between the buyer and the end client given the importance of the relationship with the end client to the buyer.[12]

2.4 Context: Relational Contracts in Fabric Procurement

The hypothesized relational contract between the buyer (the garment manufacturer) and its suppliers (the external mills) centers around the buyer receiving a discount from the supplier, relative to the high price in the input market where the supplier holds up the buyer, in exchange for demand assurance in the form of reduced capacity variance as compared to transactions in the spot market. This is a relational contract as the cooperative equilibrium occurs, rather than the static equilibrium where the buyer is held up, and it is enforced by the continuation value of future transactions rather than an explicit contract that governs any individual transaction.

Even risk-neutral suppliers value demand assurance due to the connection between capacity and costs. Empirical evidence of the importance of capacity for costs includes that the buyer purchases a non-trivial amount of fabric at prices below estimated costs for a mill

[12] It is unlikely in this setting for buyer to be able to hold up the seller given the timing of the garment production process. Specifically, the buyer and supplier negotiate the transaction price once the end client has approved the garment and specified the quantity. At this point, while the specific fabric is very valuable to the buyer, the supplier could still sell the relevant capacity to other buyers, such as garment manufacturing firms targeting domestic markets and even as stuffing for pillows, etc. The buyer's main opportunity to hold up the supplier would be after the supplier has paid the sunk cost to actually produce the fabric, at which point the buyer could try to renegotiate prices. However, the buyer is unlikely to do so as the explicit contract for the specific transaction is likely sufficiently simple and straightforward to be easily enforceable in court (and, even if not, the buyer may have concerns about developing a reputation for not upholding prior price agreements which could make future contracting difficult; as a large firm, the buyer has strong incentives to maintain its reputation).

operating at reasonable capacity (19% of transactions and 26% of volume), using average variable cost estimates from the buyer's own internal mills. This result is consistent with suppliers charging lower prices when capacity utilization is low, and are not driven by new suppliers trying to signal low prices before increasing prices afterwards–the frequency of fabric purchased below cost for suppliers who have long histories of transactions with the manufacturer is similar to the overall frequency (19% of transactions and 23% of volume).

The importance of capacity for supplier costs likely reflects the environment having demand heterogeneity over time, due to the seasonality of the fashion industry, and input adjustment costs associated with regulations and financial constraints.[13] It follows that suppliers which can predict future quantities accurately can use that information to accept or reject various sampling (and subsequent trading) opportunities with other buyers to manage capacity, and thereby costs. Discussions about fabric procurement with the buyer are consistent with this story. For example, a fabric procurement manager at the buyer who works with H&M states that he even provides quarterly estimates of future transaction volume to the most important suppliers and that past volume "is the most important factor whenever we address and negotiate the prices and lead time."[14]

The central features of the relational contract between the buyer and the supplier, specifically the discount provided to the buyer in exchange for a quantity guarantee to the supplier, have testable empirical implications. First, the discount in the relational contract suggests that suppliers which charge lower average markups should be more likely to be in a relational contract compared to other suppliers. Second, the quantity guarantee in the relational contract suggests that volumes for relational contract suppliers should be more stable than volumes for other suppliers, as long as the quantity guarantee does not change meaningfully over the course of my sample.

As suppliers' relational contract status, long-run markups, and implicit quantity guarantees are not directly observable in the data, the econometrician must define these objects, or proxies for them, to bring these hypotheses to the data. Therefore, I define relational

[13] For example, the Factories Act of 1948 specifices that employers must pay double the standard wage to overtime workers.

[14] Similarly, the fabric procurement manager for Target highlights the relevance of capacity, noting that he tracks "in a spreadsheet style wise/Mill wise/ season wise how much business a mill can handle considering the mill capacity" based on previous transactions with the supplier.

contract suppliers as suppliers which transact every month.[15,16] Economically, focusing on transaction frequency reflects that relational contracts enforce themselves through future transactions; therefore, parties in relational contracts must transact frequently. I classify 19 of 471 external suppliers as relational contract suppliers. These 19 suppliers account for 25% of fabric volume prior to integration.

Because the relational contract includes a discount, suppliers who provide a discount should be much more likely to be in relational contracts. There is little reason for other suppliers to offer a discount in equilibrium, although some suppliers may do so when costs are low due to idiosyncratic capacity shocks and the supplier underestimates industry-wide capacity utilization and, therefore, market prices.[17] In this setting, as production technology is anecdotally similar across suppliers (at least for the same fabric and among the set of suppliers that have passed the buyer's inspections and review process to be in the set of suppliers considered), there is little scope for suppliers to have persistently different costs, *except* through targeting capacity.

To test this hypothesis empirically, I estimate the relationship between supplier prices and relational contracting status. I measure supplier prices as the average standardized price per supplier. I standardize prices both at the fabric level, as costs differ across fabrics due to varying production times and input prices, and at the monthly level, as industry-wide capacity changes over time (importantly, while some seasons generally have higher capacity utilization, such as the early fall for December holiday shopping, individual fabrics experience

[15] Defining relational contract status based on transaction frequency is common in the empirical contracting literature. For example, Macchiavello and Miquel-Florensa (2017) similarly defines relational contracts based on consistent trading every period ("at least three consecutive seasons") and Macchiavello and Morjaria (2015) also use transaction frequency. Cajal-Grossi et al. (2023) assign a buyer to a contracting strategy, relational versus spot, using the ratio of shipments to suppliers as a measure of supplier concentration. This approach is quite similar to using transaction frequency in practice as there should be a strong correlation between transaction number and transaction frequency, especially as it seems unlikely that buyers in their study would concentrate purchases in time given that end clients purchase garments throughout the year and their demand is based on somewhat unpredictable trends, except for the case where the buyer is purchasing an uncommon fabric that is "on-trend", which is not likely to be part of a relational contract. Additionally, both approaches reflect a similar economic logic: relational contracts require future transactions for enforcement and should, therefore, have many transactions.

[16] I only consider months prior to the construction of the new internal supplier for defining relational contract status. This time restriction ensures that subsequent analyses of the effects of increased vertical integration on relational contracts do not allow treatment (*i.e.,* the creation of the internal mill) to also impact supplier classification as a relational contract supplier. Appendix Figure A.2 shows the distribution of the percentage of months with transactions across suppliers.

[17] It is possible that suppliers trying to become relational contract suppliers might offer discounts to try to signal their interest in forming a relational contract. This concern is unlikely in this setting given that the buyer is a large established firm that presumably cannot add more relational contracts, otherwise it would. I discuss this point in more depth in A.5.

large unpredictable heterogeneity, as shown in Appendix Figure A.3).

This measurement approach aims to capture the extent to which suppliers hold up the buyer. For example, consider two fabric purchases from different suppliers that occur at the same time for different fabrics with different production costs due to different material costs. If the two suppliers charge the same price, the supplier of the high cost fabric is not holding up the buyer to the same extent. It follows that using unadjusted prices as a measure of supplier prices would not account for cost heterogeneity across fabrics and over time. An alternative approach would be to measure markups directly using the buyer's cost data for fabrics they produce. However, this approach would both limit the sample, recalling that the buyer's cost data are only available for the subset of fabrics they produce, and fail to adjust for industry-wide seasonality in capacity usage.

Figure 5 illustrates the connection between long-run markups (relative to other suppliers, rather than marginal cost), proxied as the volume-weighted average standardized markup over the same nine month window,[18] and relational contracting. It shows a spline estimate of the association between long-run markups and whether a supplier is a relational contract or market supplier. Consistent with the connection between relational contract suppliers and discounts, the analysis illustrates that suppliers that do not provide discounts relative to other suppliers are very likely to be market suppliers (at least 90%). Additionally, the probability of being a relational contract supplier decreases in long-run markups, while the probability of being a market supplier increases in long-run markups. Similar patterns emerge when using one year rather than nine months to calculate long-run markups (see Appendix Figure A.4) and using a bin-scatter rather than a spline (see Appendix Figures A.5, A.6).

[18]Specifically, the data are from the nine months prior to the introduction of the new vertically integrated mill.

Figure 5: Discounts and Supplier Type

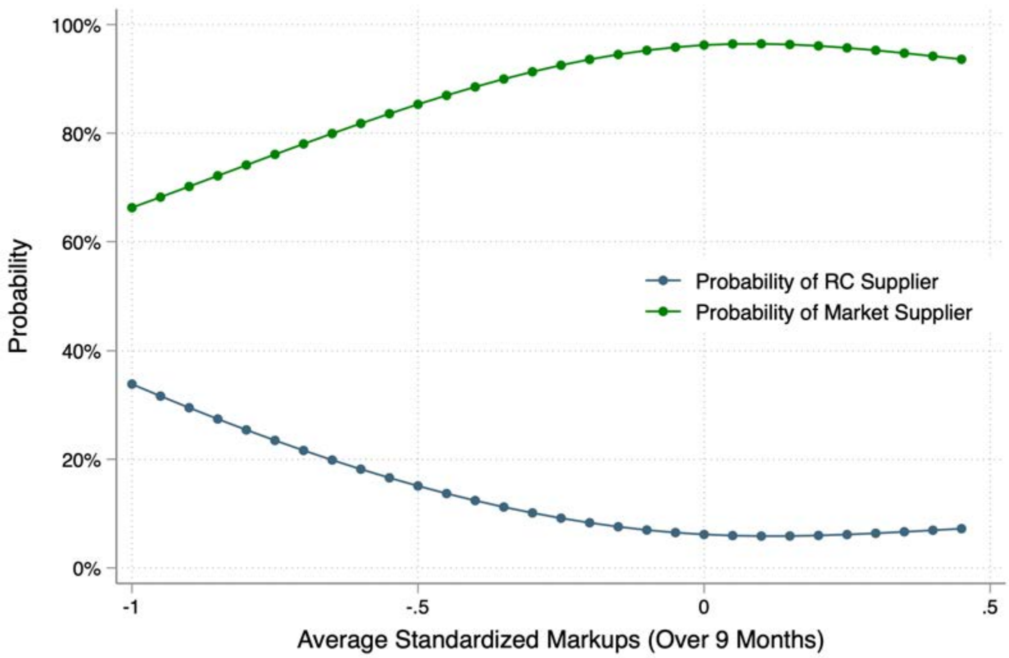

Note: Data from universe of fabric transactions by buyer. Probability estimates represent marginal effects from a non-parameteric series estimator calculated using transactions from the three quarters prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Suppliers benefit from the relational contract through the demand assurance that facilitates suppliers both controlling capacity and, therefore, costs, and smoothing profits. Demand assurance means relational contract suppliers receive relatively stable volumes over time, conditional on the quantity level promised in the relational contract not changing in an economically meaningful way over the course of the sample. To bring this hypothesis to the data, I first compare the coefficient of variation for relational contract suppliers versus market suppliers. As the quantity guarantee is designed to stabilize capacity, relational contract suppliers with quantity guarantees should experience more stable capacity utilization than other suppliers. Figure 6 illustrates that suppliers in relational contracts have less volatile demand than market suppliers. The difference in average coefficient of variation by supplier type is statistically significant ($p < .001$). Importantly, this result is not entirely mechanical, as the criterion to define relational contracting status only incorporates transaction frequency and not volume. It follows that there could be suppliers identified as relational contract suppliers with extremely volatile demand, as long as those suppliers receive at least one order per month. In Appendix A.3, I show that quantity patterns for relational contract suppliers are consistent with the buyer directing volumes to fulfill quantity guarantees.

Figure 6: Coefficient of Variation and Supplier Type

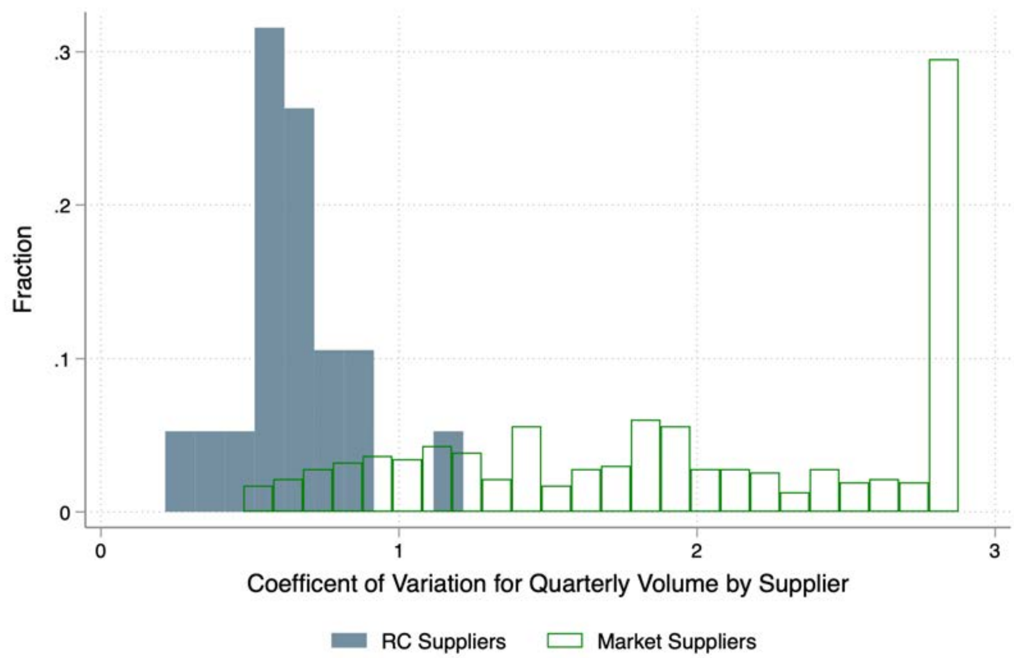

Note: Data from universe of fabric transactions by buyer. Coefficients of variation are calculated using transactions from all quarters prior to integration.

An explicit contract at the time of sampling that specifies a price for the fabric should the end client order the garment seems, intuitively, like an attractive alternative approach to resolve the hold-up problem. In Appendix A.4, I discuss why the buyer and seller rely on a relational rather than explicit contract. As the buyer prefers sourcing fabrics from relational contract suppliers to market suppliers, in principle, the buyer should only use relational contract suppliers. Appendix A.5 clarifies why the buyer is unable to source exclusively from relational contract suppliers in this setting.

2.5 Alternative Possible Explanations for Observed Patterns: Supplier Heterogeneity and Analysis of Placebo Suppliers

Although the evidence is consistent with the hypothesized relational contract, I conduct additional analysis to mitigate concerns that unobserved supplier heterogeneity explains the patterns in the data rather than relational contracts. For example, some suppliers could be more efficient, perhaps due to better management given that technology is fairly standardized (at least among the suppliers that produce similar fabrics and meet the buyer's screening criteria). More efficient suppliers might pass through some of their cost savings to the buyer, leading to lower prices, and might receive more (although not necessarily more stable) volume. That said, this scenario seems unlikely as there are no incentives for efficient suppliers to pass on any cost savings due to efficiency outside a relationship–statically, even very efficient firms would optimally charge the high hold-up price.

To mitigate concerns that unobserved supplier heterogeneity, rather than relational contracting, explains observed patterns in the data, I reproduce the analysis but compare relational contract suppliers to a subset of market suppliers that receive large orders but do not transact as frequently with the buyer ("placebo suppliers"). As the relational contract both depends on transaction frequency for the value of the future and provides demand assurance, placebo suppliers are plausibly not in a relational contract. Importantly, many stories of unobserved heterogeneity suggest that supplier volume, rather than transaction frequency, should differentiate contracting, such as the example of heterogeneity in supplier efficiency. In other words, plausible unobserved supplier heterogeneity is likely highly correlated with supplier volume. It follows that finding different patterns in the data for placebo and relational contract suppliers suggests that unobserved heterogeneity is not a valid alternative explanation for the empirical patterns. Note that this analysis alone does not suggest that relational contract suppliers are more or less efficient than other suppliers, but rather that the discounts they offer are due to the relational contract rather than any possible efficiency advantages they might have.

I construct a set of placebo suppliers by finding the suppliers with the largest volumes pre-integration that are not relational contract suppliers; to keep the set of suppliers as directly comparable to the set of relational contract suppliers, I select the same number of placebo suppliers as relational contract suppliers. I then verify that volumes are indeed similar for relational contract and placebo suppliers, as this approach to defining placebo suppliers does not mechanically ensure that placebo suppliers receive similar volumes to relational contract suppliers.

Figure 7: Volume Comparison: Placebo and Relational Contract Suppliers

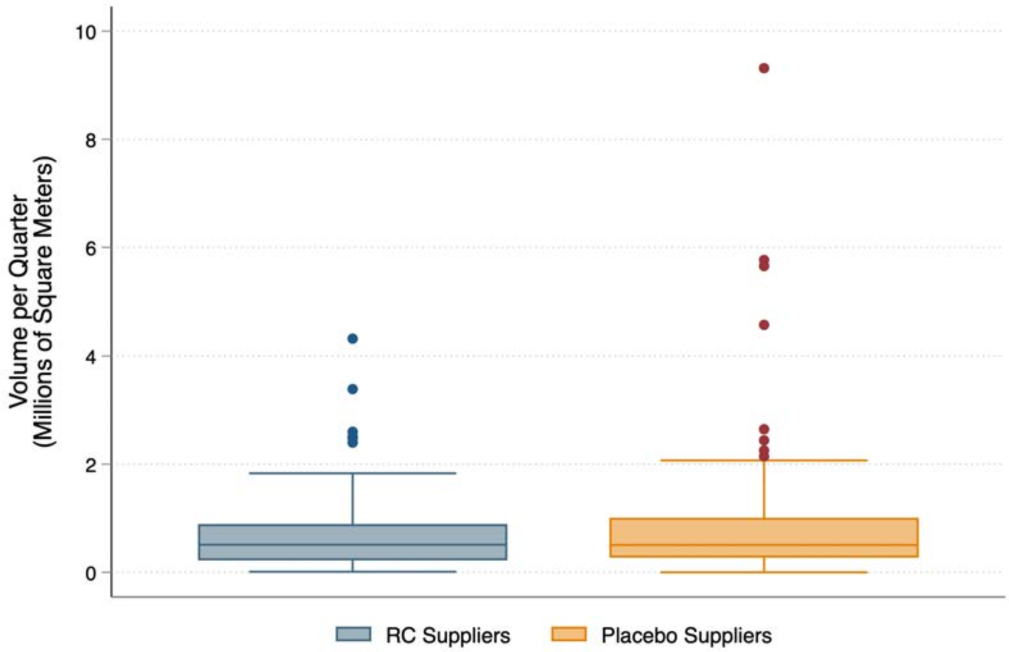

Note: Data from universe of fabric transactions by buyer. Boxplots show volume per supplier per quarter using all quarter prior to integration. Placebo suppliers are defined as the suppliers with the largest average quarterly volume that are not relational contract suppliers regardless of transaction frequency chosen such that the number of placebo suppliers matches the number of relational contract suppliers.

Figure 7 illustrates that relational contract and placebo suppliers have similar quarterly volumes and that the distributions of quarterly volumes have fairly similar support. If anything, the placebo suppliers have slightly *larger* volumes than the relational contract suppliers, with higher 25th and 75th percentiles, higher mean, and almost identical medians ($\approx 1\%$ different).

Appendix A.6 documents findings that placebo suppliers both charge higher prices and have less stable volumes than relational contract suppliers.

3 Conceptual Model of the Threat Point Effect

This section presents a conceptual model of the relational contract between a buyer and supplier to illustrate how buyer and supplier characteristics shape the magnitude of the threat point effect, which is the difference between the relational contracting equilibrium outcomes (quantity and prices) before and after integration. I first describe the model, which is motivated by the features of the empirical setting and microfounds relational contract surplus through supplier characteristics (risk aversion), demand assurance, and the production technology. I then study threat point effect heterogeneity through the lens of the model with respect to buyer and supplier characteristics (*i.e.,* relationship type).

When the vertical integration is sufficiently large to more than crowd out all market suppliers, quantities in the relational contract must change (holding downstream demand faced by the buyer fixed). In this setting, the prices in the relational contract will also likely change as the new equilibrium relational contract quantities decrease, which would presumably increase relational prices by reducing the value of the demand assurance provided. However, my empirical setting allows me to isolate the change in relational contract *prices* due to the change in threat point as the buyer still purchased over half of its fabric inputs from external suppliers after integration, which is more than pre-integration relational contract quantities ($\sim 25\%$). The effect of the change in threat point on prices applies broadly regardless of the level of integration and should impact relational contract prices even in settings where the buyer would need to reduce relational contract quantities.[19]

3.1 Model Description and Set-Up

I model the prices in the relational contract as solving the Nash bargaining problem between the buyer and supplier, where the bargain is over a relational contract pricing function that maps quantities to prices. For simplicity and to focus on the key economics of this problem, the main version of the model presented does not include any dynamic incentives to govern the relational contract and make it self-enforcing. Extending the model in Appendix A.7.1 to incorporate standard relational contract dynamics (specifically, the dynamic incentive compatibility constraint as described in Macchiavello and Morjaria (2023)) does not change any key results. Rather, relational contract dynamics serve as the unmodeled mechanism that facilitates the buyer and supplier reaching a cooperative equilibrium; without dynamics, the supplier's optimal strategy would be to hold up the buyer in every period. Furthermore, the buyer and supplier bargain over a function that maps quantities to prices. As relational contracts require future trade, such an agreement is especially important in the context of a relational contract.

[19] Excepting complete integration, which renders external suppliers irrelevant.

$$\underset{\substack{p^{RC}(q^{RC}) \\ \in\Gamma:\mathbb{R}^+\to\mathbb{R}^+}}{\operatorname{argmax}} \left(\underbrace{-\mathbb{E}_{q^{RC}}[p^{RC}(q^{RC})q^{RC}] - O_B}_{\text{Buyer Surplus}}\right)^{\alpha_B} \left(\underbrace{\mathbb{E}_{q^{RC}}[U(p^{RC}(q^{RC})q^{RC} - C(q^{RC}))] - O_S}_{\text{Supplier Surplus}}\right)^{1-\alpha_B} \tag{1}$$

The relational contract solves the Nash bargaining problem in (1), specifying the price schedule in the relational contract for each observed quantity (capacity). The buyer bargaining parameter is α_B; higher values of α_B result in relational contract pricing functions that allocate more surplus to the buyer as its surplus factors more heavily in the objective function. As the buyer and supplier split the surplus, their bargaining parameters must add to one. Therefore, the supplier's bargaining parameter is $1 - \alpha_B$.

The buyer's surplus is the difference in profit when buying fabric in the relational contract–purchasing quantity q^{RC} at price $p^{RC}(q^{RC})$–versus buying fabric from the outside option–with value O_B. As revenue from the buyer's sales to the end client do not depend upon fabric prices since the negotiation over production prices and quantities precedes bargaining over fabric purchasing, buyer revenue is the same in both the relational contract and with the outside option. Therefore, I can express the buyer surplus from the relational contract as the difference in costs from purchasing the fabric in the relational contract as compared to the outside option. I assume the buyer is risk neutral as the buyer is a large firm. Additionally, the buyer requires suppliers to purchase some insurance for the main risk of production delays, and each supplier's quantity is small relative to the buyer's overall quantity, such that the buyer is reasonably risk neutral, especially for any individual relational contract.

The supplier's surplus is the utility of profit in the relational contract–with CARA utility function $U(\cdot)$ and production costs given capacity $C(q)$–less the value of the outside option O_S. I assume that suppliers are risk averse, as in Blouin and Macchiavello (2019) and Karlan et al. (2014). Supplier risk aversion could derive from, among other possibilities, owners' preferences over profits mirroring individuals' preferences over consumption as small supplier profits are likely the sole income source for owners. Alternatively, firms may need to make at least some payments every period (*e.g.,* payroll) and financial market frictions make it challenging or expensive to borrow or save to pay such expenses rather than using contemporaneous cash flows. I assume the supplier cost function is increasing and has increasing marginal costs (*i.e.,* is convex). This assumption reflects that marginal costs can increase due to machine inputs being fixed in the short-run; therefore, increasing production requires more congestion at machines. Additionally, under Indian labor laws, labor costs could increase dramatically as quantity increases–the Factories Act of 1948 requires employers to pay double wages for overtime work.[20] Using data from the integrated supplier, I also empirically validate cost function convexity (see Figure 10).

I now describe the outside option for both the buyer and the supplier, starting with unit prices and costs and then describing quantities in the model. Before vertical integration, the outside option for both the buyer and the supplier is to transact with the market. Even though the buyer is a large firm, the buyer (and suppliers) are small relative to Indian fabric markets (let alone global fabric markets).[21] Therefore, I model the market as perfectly elastic with a constant price p^M, which includes hold-up, faced by both the buyer and the supplier.[22]

After integration, the outside option for the buyer changes to use the same cost function as the supplier, reflecting that the buyer and its relational contract suppliers presumably use similar technology–anecdotally, these firms are all users of (near-)frontier technology rather than firms with large research and development departments trying to innovate and compete by advancing fabric production technology. Allowing for some cost heterogeneity between the buyer and supplier does not change results meaningfully (see Appendix A.7.2); intuitively, the lack of change reflects that shifting costs impacts both the inside and outside options in a fairly similar way such that the overall change is minimal.[23] This change in the outside option assumes that the buyer treats their integrated supplier and their relational contract suppliers as alternatives (*i.e.,* substitutes). Morton (2023) shows strong evidence of substitutability in this setting, leveraging exogenous variation at the fabric level. It

[20]Furthermore, labor market frictions limit firms' ability to simply hire more workers, leading to firms running expensive extra shifts at night or on Sunday to meet demand shocks.

[21]Recall the buyer uses fewer than 300 million square meters of fabric per year, which is less than 0.5% of the 71 billion square meters of fabric produced in India in 2019, and India is the sixth largest producer by country globally (India Brand Equity Foundation).

[22]Note that incorporating market power for the buyer when purchasing inputs in the spot market should not affect the threat point effect in the context of the bargain between the buyer and a single relational contract supplier. The buyer cannot both threaten to use the integrated supplier to displace the relational contract supplier and use it to decrease the quantity purchased in the spot market simultaneously; therefore, changing the outside option does not affect market prices as it does not change market quantities. Overall, aggregating across suppliers, if the buyer has market power in the spot market, then prices in the spot market would decrease as the buyer replaces market supply with integrated supply, but it would not impact the threat point effect directly as only market prices would change. That said, incorporating indirect effects through the outside option changing in the Nash bargain would lead to underestimates of the threat point effect, as the outside option improves for the buyer and worsens for the supplier as lower overall demand reduces market prices.

[23]Furthermore, large heterogeneity in costs functions seems implausible–if the buyer were much more efficient, it should completely integrated and if the supplier were much more efficient the buyer should never integrate.

follows that the integrated supplier is a credible threat to displace external relational contract suppliers, which corresponds with the investment in integrated capacity as a meaningful and visible sunk cost.

Quantities in the model reflect the stochastic nature of demand in the industry, with quantities determined by the end client before bargaining for fabric prices rather than by the buyer or supplier as part of the relational contract. As the relational contract is designed to provide demand assurance rather than large demand volume, I model market quantities as a mean-preserving spread of demand quantities. I set mean quantities (for both the relational contract and market) to 100% capacity (and ensure that, given market prices and productive technology, the supplier's optimal capacity utilization is 100%). This assumption does not imply that the supplier only works with the buyer. Instead, it could be interpreted as the capacity utilization relative to the capacity saved by the supplier for the buyer, such as effectively reserving floor space or time for the buyer. In this sense, how important the relational contract is to the supplier relative to their contracts with other buyers is incorporated through the risk aversion parameter (*i.e.,* if the supplier has other income streams, then they might be less risk averse over production for this one buyer in the model.)[24] Formally, the outside options are:

$$\begin{aligned}\text{Buyer before integration: } & O_B = -p^M \mathbb{E}_{q^{RC}}[q^{RC}] \\ \text{Buyer after integration: } & O_B = -\mathbb{E}_{q^{RC}}[C(q^{RC})] \\ \text{Supplier before and after integration: } & O_S = \mathbb{E}_{q^M}[U(p^M q^M - C(q^M))]\end{aligned}$$

For the exact parameterization of the model and a discussion of how equilibrium relational contract prices are computed, please see Appendix A.7.3; parameters and functional forms are chosen to illustrate the economics of the threat point effect rather than calibrated to match data or parameter estimates from other studies.

[24] In theory, the supplier working with other buyers could also affect the cost convexity for the supplier. For example, if the supplier has a diversified buyer pool, then the supplier might be able to balance out demand shocks from different buyers; although, seasonality in industry demand reduces any individual suppliers' ability to smooth demand by contracting with many buyers (*i.e.*, demand shocks across buyers are correlated). Recall, however, that changes in costs do not affect results for the magnitude of the threat point effect, as shown in Appendix A.7.2.

3.2 Threat Point Effect and Buyer Bargaining Parameter

Figure 8: Relational Contract Prices and the Buyer Bargaining Parameter

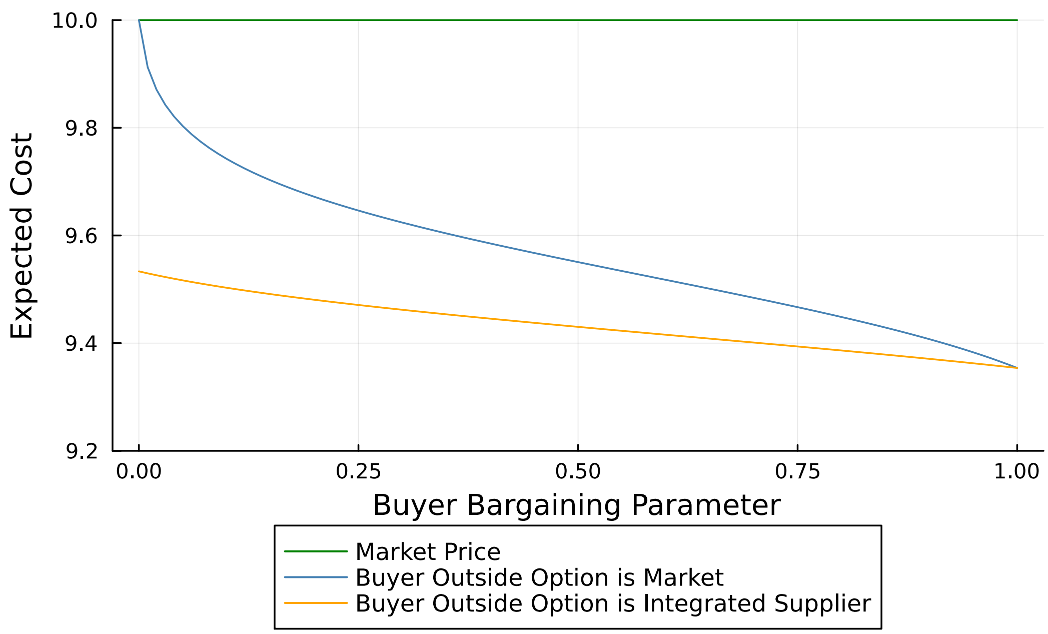

Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.7.3.

Figure 8 shows how the buyer's expected fabric procurement cost in the relational contract varies with respect to both the buyer bargaining parameter and the buyer's outside option, holding supplier risk aversion constant. When the buyer bargaining parameter approaches zero and the buyer's outside option is the market, the expected cost to the buyer approaches the market price. The economic intuition for this result is that the buyer's outside option determines the relational contract price when the buyer bargaining parameter is low–the supplier extracts all surplus so relational contract prices increase until expected buyer surplus is zero. As the buyer bargaining parameter increases, the buyer's share of the surplus increases through lower relational contract prices, decreasing expected cost. When the buyer's outside option is the integrated supplier, as the buyer's threat point improves (*i.e.,* is lower cost), relational contract prices decrease. As the buyer bargaining parameter increases, the supplier's outside option plays an increasingly important role in determining relational contract prices. Therefore, when the buyer bargaining parameter is one, the supplier's outside option determines relational contract prices and the buyer's outside option is effectively irrelevant, resulting in the same relational contract prices (and expected cost) regardless of buyer outside option.

The threat point effect is visible in the graph as the difference between the relational contract price when the outside option is the market (the blue line) and the relational contract price when the outside option is the integrated supplier (the orange line). The threat point effect can be quantitatively important (over 5% of market prices) when the buyer's outside option influences relational contract prices. However, when the buyer has a high bargaining parameter, changing the buyer's outside option has a negligible effect on relational contract prices since prices are effectively determined by the supplier's outside option (*i.e.*, the supplier's individual rationality/participation constraint), which does not change when the buyer adds integrated capacity. Equivalently, the threat point effect is small when the buyer bargaining parameter is large because the buyer already extracts all the surplus pre-integration–there is no scope for the buyer to extract more surplus such that improvements in the buyer's bargaining position do not affect relational contract outcomes.

3.3 Threat Point Effect and Supplier Risk Aversion

Figure 9: Relational Contract Prices and Supplier Risk Aversion

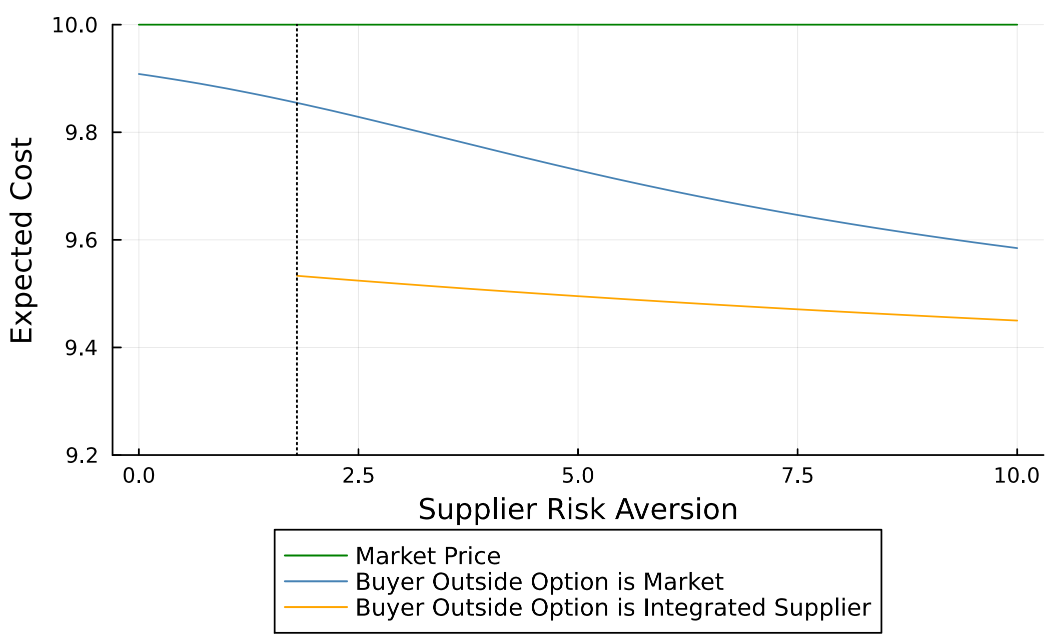

Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.7.3.

Figure 9 illustrates how the expected cost to the buyer varies with respect to both supplier risk aversion and the buyer's outside option, holding the buyer bargaining parameter constant. First, the expected cost in the relational contract when the buyer's outside option is the market and supplier risk aversion is zero highlights that risk aversion is not the only feature of the model that makes demand assurance valuable for suppliers–the combination of cost convexity and lower quantity variance in the relational contract compared to the market also creates surplus. As supplier risk aversion increases, the value of the outside option, with high quantity, cost, and profit variance, decreases relative to the inside option, creating higher surplus in the relational contract (holding prices fixed). As the buyer bargaining parameter is not zero in Figure 9, some of this surplus is shared with the buyer through lower relational contract prices. When the buyer outside option is the integrated supplier, surplus in the relational contract decreases dramatically because the buyer's outside option improves.

Given the lower quantity variance in the relational contract and cost convexity, there is a cutoff level of risk-aversion below which the relational contract does not create any surplus and breaks down, recalling that the supplier does have some surplus when selling to the market as market prices are greater than costs due to high hold-up prices. In this scenario, assuming the buyer has the ability to threaten the outside option of its choice after vertical integration (either market or integrated supplier),[25] the buyer would prefer to convince the supplier that the relevant outside option is the market to keep the relational contract. In the empirical context studied, for example, the buyer could argue that the integrated supplier is irrelevant for bargaining by clarifying that the integrated supplier displaces volume from other suppliers or even claiming that the integrated supplier produces different fabrics. If the supplier is sufficiently risk averse for a relational contract to exist when the buyer outside option is the integrated supplier, then the buyer prefers to use the integrated supplier as the outside option to reduce prices.

[25] A model where the buyer selects which outside option to use in bargaining, even if buyer can only threaten to use the integrated supplier as an outside option if there is capacity at the integrated supplier, would not change this result as long as the buyer bargains first with its relational contract suppliers and strategically chooses which outside option to use in the negotiation. With this model timing structure, the buyer would only threaten to use the integrated supplier as the threat point when doing so would not break the relational contract (assuming the buyer cannot simply create new relational contracts). Equivalently, if the buyer can choose the bargaining order, the same equilibrium would hold, as bargaining with market suppliers is trivial–the price is always the market price–so the buyer would choose to bargain with relational contract suppliers first.

As in Figure 8, the threat point effect is the difference between the relational contract price when the outside option is the market (the blue line) and the relational contract price when the outside option is the integrated supplier (the orange line). First, there is no threat point effect for low levels of supplier risk aversion, as the buyer strategically keeps the market as the outside option and does not try to leverage an improved bargaining position due to vertical integration because doing so would break, rather than improve, the relational contract. Second, when supplier risk aversion is above the cutoff for the relational contract to exist with the integrated buyer as the outside option, the threat point effect decreases slightly as risk aversion increases. This result reflects that total surplus differentially increases more in risk aversion when the buyer's threat point is the market due to the algebra of the Nash product.[26] As some of this surplus accrues to the buyer in the form of lower prices, relational contract prices decrease more quickly with respect to risk aversion when the buyer threat point is the market.

4 Structural Model

4.1 Why a Structural Model

I estimate the conceptual model structurally, recovering the price schedule that solves the Nash bargaining problem specified in equation 1. This structural modelling approach is useful for two key reasons: *i)* quantifying the threat point effect for heterogeneous buyer and supplier types, reflecting the importance of heterogeneity in the conceptual model, and *ii)* evaluating policies to reallocate surplus to small constrained suppliers.

Although reduced-form methods can measure the threat point effect in a case study (and I later validate my structural estimates using a difference-in-difference approach that directly estimates the threat point effect in this setting), they do not facilitate quantifying the threat point effect for heterogeneous buyer and supplier types. Specifically, conducting a rich analysis that incorporates heterogeneity by relationship type with only reduced-form methods similar to my approach requires data and a setting with not only enough relationships to estimate the effect for relationships of that type but also a variety of different relationship types (and knowledge of what type each relationship is). Finding data and a setting for such analysis would be challenging, if not practically impossible–it would require finding case

[26]Specifically, because total surplus is a *product* rather than a *sum*, total surplus is much higher when the buyer's threat point is the market. Formally, a marginal increase in supplier surplus increases total surplus more when the level of buyer surplus is higher.

studies where vertical integration occurred (and without an increase in downstream demand for the buyer), with buyers sufficiently large that it purchases externally from both market and relational contract suppliers after integration, with the desired heterogeneity in buyer bargaining parameters and supplier risk aversion, and with meaningful heterogeneity in supplier exposure to integration. As large buyers presumably intentionally select their relational contract suppliers based on specific characteristics, their relational contract suppliers likely resemble each other, making it largely implausible to find sufficient supplier heterogeneity from any reduced-form analysis.

My structural approach enables quantifying the threat point effect for out-of-sample relationships and conducting counterfactual policy analysis through, although requires imposing additional assumptions. As other settings demonstrate that even small firms, such as smallholder Rwandan farmers in Macchiavello and Morjaria (2020) and small fishing enterprises in Sierra Leone (Ghani and Reed, 2022), use relational contracts, it is important to consider the quantitative importance of the threat point effect for these presumably out-of-sample firms. Specifically, as buyer and supplier characteristics determine the threat point effect, threat point effects for these firms could be very different than what is observed for in-sample fabric suppliers.

Not only does the structural approach facilitate quantifying the threat point effect for small constrained suppliers, it also provides a framework for evaluating policies to reallocate surplus to them. Specifically, I consider two policies to improve the relational contract for small firms: *i)* increasing downstream buyer competition; and *ii)* creating the missing market that generates demand for relational contracts from the small suppliers by allowing them to purchase insurance. I validate the model three ways to ensure that model captures key patterns in the data and produces reasonable price estimates.

I begin by discussing estimation and identification of the model before describing the parameter estimates and model validation against untargeted moments and out-of-sample fit. I then quantify the threat point effect for heterogeneous buyer and supplier (*i.e.*, relationship) types. Last, I validate the model using a difference-in-differences approach that compares the model-implied threat point effect, which depends upon parameter estimates for in-sample suppliers, with difference-in-difference estimates of the threat point effect. Additionally, I show that event study coefficients using model-implied prices match untargeted event study coefficients using actual prices.

4.2 Estimation Approach

4.2.1 Buyer and Supplier Primitives

I use the structural model to estimate the buyer bargaining parameter and supplier risk aversion for each supplier relationship governed by the relational contract. Additionally, I estimate a tuning parameter to penalize deviations from risk neutrality, motivated both by economic intuition that firms should not have extreme deviations from risk neutrality and econometric concerns of overfitting given that there are relatively few observations per supplier. Overall, my approach reflects that quantifying the threat point effect for heterogeneous types is largely an out-of-sample exercise; therefore, I develop an estimation strategy that naturally allows for testing out-of-sample fit. Specifically, I primarily estimate parameters using pre-integration data and test out-of-sample fit using post-integration data. Although this approach is admittedly out-of-sample in a different sense, namely quarters for in-sample suppliers rather than for out-of-sample relationship types, it provides at least supportive evidence that the model performs reasonably well. I first describe the "inner" optimization used to estimate the buyer bargaining parameter and supplier risk aversion parameter given the tuning parameter and then describe the "outer" optimization to estimate the tuning parameter given the previously estimated buyer bargaining parameter and supplier risk aversion.

The "inner" optimization has three steps. First, I derive the model-implied relational contract pricing function $p^{RC}(q^{RC};\boldsymbol{\Omega}_s)$ for each relationship by finding the function that optimizes the Nash product as in (1), where $\boldsymbol{\Omega_s}$ represents the parameter vector for each relationship: the buyer bargaining parameter and supplier risk aversion parameter. The relational contract pricing function depends upon the buyer bargaining parameter and the supplier risk aversion parameter through the Nash product, as shifting these parameters shifts the value of the Nash product and, therefore, the optimal relational contract prices.[27] Second, I calculate the price estimates for the quantities observed in the relational contract

[27] I use a second-order Taylor approximation, meaning $p^{RC}(q^{RC};\boldsymbol{\Omega}_s) = \beta_{0,s} + \beta_{1,s}q^{RC} + \beta_{2,s}(q^{RC})^2$ to find the $\beta_{0,s}, \beta_{1,s}, \beta_{2,s}$ that maximize the Nash product. As suppliers typically operate near full capacity, I estimate the second-order Taylor approximation around a capacity of one. Identification of $\beta_{0,s}, \beta_{1,s}, \beta_{2,s}$ follows the logic of identification of buyer bargaining parameter and supplier risk aversion in 4.3. This logic is clearest for $\beta_{0,s}$, which shifts the level of relational contract prices, and $\beta_{2,s}$ which shapes the convexity of relational contract prices. Including $\beta_{1,s}$ incorporates additional flexibility for the relational contract to respond to increasing marginal costs. For example, average cost per unit is larger for capacity utilization of 1.2 than for capacity utilization of .8, but the relational contract price would be the same without including $\beta_{1,s}$. It follows that supplier risk aversion, along with convexity in costs, together lead to larger $\beta_{1,s}, \beta_{2,s}$ as these parameters help smooth profit, with $\beta_{1,s}$ and $\beta_{2,s}$ makes the relational contract stabilize supplier profits across different capacity realizations.

for each supplier implied by the model-derived relational contracting pricing function. Last, I compute the objective function per supplier, which is the root mean squared error of the price estimates from the relational contracting pricing function using *only the pre-integration data.*

I penalize extreme estimates of the risk aversion parameter by adding the absolute value of the risk aversion parameter, weighted by the tuning parameter, to the objective function. As I use CARA utility, a risk aversion parameter of zero is risk neutral, meaning that deviations from risk neutrality in either direction are equally penalized. This approach is similar to the set-up of a LASSO regression by constraining extreme parameter values. Incorporating the tuning parameter to push parameter estimates towards risk neutrality reflects both economic motives, as firms are unlikely to have extreme deviations from risk neutrality, as well as econometric concerns around overfitting given the relatively few observations per supplier.

Formally, the "inner" optimization finds the buyer bargaining parameter and supplier risk aversion parameter, given the tuning parameter, to minimize the sum of the per supplier objective functions:

$$\min_{\{\alpha_{B,s}\},\{\theta_s\}} \underbrace{\sum_{s\in S}}_{\substack{\text{Sum over}\\ \text{suppliers}}} \Bigg\{ \Bigg[\overbrace{\sum_{\{q^{RC}_{s,t}\}}}^{\substack{\text{Sum over}\\ \text{pre-integration}\\ \text{quarters}}} \underbrace{\sqrt{\frac{1}{T}\left(p_s^{RC}(q^{RC}_{s,t};\boldsymbol{\Omega}_s) - p^{RC}_{s,t}\right)^2}}_{\text{RMSE: Model \& Actual Prices}} \Bigg] + \overbrace{\lambda}^{\substack{\text{Tuning}\\ \text{Parameter}}} \Big| \underbrace{\theta_s}_{\substack{\text{Supplier}\\ \text{Risk Aversion}\\ \text{Parameter}}} \Big| \Bigg\}$$

In this optimization problem, S represents the set of suppliers, $\{q^{RC}_{s,t}\}$ represents the set of observed quantities for relational contract supplier s across pre-integration periods,[28] T is the number of periods, $p_s^{RC}(q^{RC}_{s,t};\boldsymbol{\Omega}_s)$ is the price estimated by the relational contract pricing equilibrium for supplier s for quantity $q^{RC}_{s,t}$, $p_s^{RC}(q^{RC}_{s,t})$ is that actual relational contract price for supplier s for quantity $q^{RC}_{s,t}$, λ is the tuning parameter, and θ_s is the supplier risk aversion parameter.

The "outer" optimization finds the tuning parameter than maximizes fit in post-integration quarters. Specifically, the optimization routine finds the tuning parameter that leads to the

[28] While it would be appealing to use monthly rather than quarterly data for this analysis to increase the sample size, the relational contracts appear to operate at the quarterly level rather than the monthly level. Not only are the reduced form patterns consistent with the relational contract clearest at the quarterly level, but anecdotally staff procuring fabric at the buyer emphasize that meetings with key suppliers about capacity projects are once per quarter (and not once per month).

relational contract pricing function that best fits the post-integration data, recalling that the relational contract pricing function is determined using only the pre-integration data. The tuning parameter impacts post-integration fit by influencing the estimated relational contract pricing function: a larger tuning parameter pushes estimated supplier risk aversion parameters towards risk neutrality. Therefore, the “outer” optimization solves:

$$\min_{\lambda} \underbrace{\sum_{s \in S}}_{\substack{\text{Sum over}\\\text{suppliers}}} \left[\overbrace{\sum_{\{q_{s,t}^{RC}\}}}^{\substack{\text{Sum over}\\\text{post-integration}\\\text{quarters}}} \underbrace{\sqrt{\frac{1}{T}\left(p_s^{RC}(q_{s,t}^{RC}; \boldsymbol{\Omega}_s) - p_{s,t}^{RC}\right)^2}}_{\text{RMSE: Model \& Actual Prices}} \right]$$

4.2.2 Outside Estimation

To estimate the model, it is necessary to find additional moments from the data. Table 1 describes the additional moments used:

Table 1: Additional Moments for Estimating the Structural Model

Parameter(s) in Model	Moment(s) to Match	Data
Cost Function	Actual Costs and Capacity	Cost and Capacity Data from Integrated Supplier
Mean Capacity	Average Log Quantity	Transaction Data
$\frac{\mathbb{V}[q^M]}{\mathbb{V}[q^{RC}]}$	Variance in Quantities for non-RC Suppliers Compared to Variance in Quantities for Relational Contract Suppliers	Transaction Data
p^M	Ratio of Prices from Market Suppliers to Estimates from Integrated Supplier	Transaction Data

I use data on capacity and costs at the integrated supplier to estimate the convexity of the cost function. As, anecdotally, relational contract suppliers use similar technology to the integrated suppliers, their cost functions should be similar to the cost function at the integrated supplier (nevertheless, I include a robustness check that incorporates supplier cost heterogeneity). Therefore, I estimate the total cost function for suppliers as $c(q) = \phi_0 + \phi_1 q + \phi_2 q^2$. I include q^2 given the importance of convexity, although noting that I impose no structure to force the model to yield $\phi_2 > 0$, which would be consistent with convexity. Estimates are equivalent to a second-order Taylor approximation of the true cost function.

The data used are the cost data from the integrated supplier, which are a panel at the division (*i.e.*, fabric type, such as knits and striped fabrics) and month level from April 2019 to March 2020.[29] Additionally, note that, especially at the month level, it is reasonable to treat these estimates as recovering the true causal effect of capacity on costs because quantities are, at least locally, exogenous. Quantity exogeneity derives from the institutional features of how quantities are determined in garment manufacturing. Recall that the fabric supplier is locked in at the time of making the sample garment, which is *before* the end quantity for the order is determined. It follows that the exact quantity used at the integrated supplier is determined by stochastic downstream demand rather than choices by the buyer.[30] Therefore, as quantities are (at least locally) exogenous, OLS regressions yield causal estimates of the effect of capacity utilization on costs.

[29] For convexity, it is important to use the most disaggregated data temporally rather than combine. To fix ideas, consider a supplier that operates for two months at 100% capacity and another supplier than operates first at 80% capacity and another month at 120%. For both suppliers, the average capacity is 100%. Therefore, a regression of costs on capacity at an aggregated level would show different costs for the same capacity level that do not correspond to actual capacity differences apparent in the disaggregated data.

[30] Based on previous transactions with end clients, the buyer presumably has some sense of how common certain volumes are from various end clients. The buyer likely uses this knowledge when sampling to try to ensure that capacity utilization at the integrated supplier is not too extreme. However, even with this knowledge, for short time periods, it is impossible to achieve exactly targeted capacity usage–consistent with the heterogeneity in capacity usage estimated.

Figure 10: Cost Convexity at the Integrated Supplier

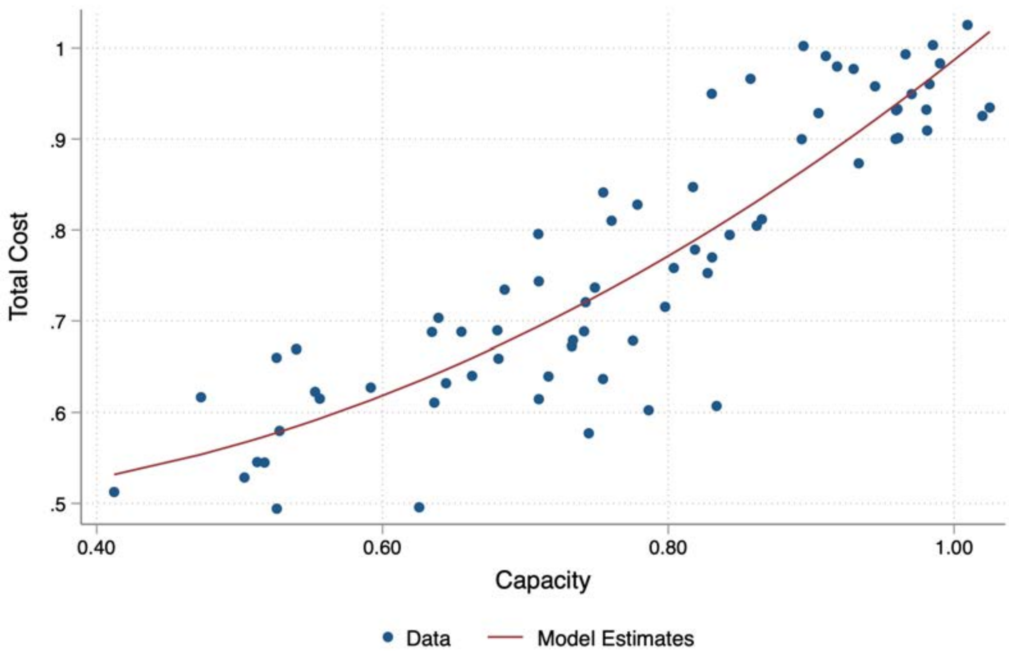

Note: Data from integrated suppliers at division by month level. Costs used are the volume-weighted average cost, expressed as a percent of average volume-weighted cost for the divsion. Capacity is reported directly in the data.

The estimates show that the model fits the data well. The R^2 is .795 and the RMSE is .07, which is small relative to the variance in costs overall in the data. Additionally, the estimate finds strong convexity, with the coefficient on q^2 at .7815 and a p-value of .02. I transform costs so units are in percentage terms relative to the mean costs for the fabric group, as costs are reported by fabric group, facilitating analysis combining data from all fabrics. The strong fit of the model suggests that heterogeneity by fabric type is not likely a central feature of the cost function.

As the cost function is measured in capacity units, transaction quantities need to be converted into capacity units to estimate the structural model. Two assumptions underlie my conversion: first, I assume that log quantities are relevant for capacity usage. This assumption reflects that orders with larger quantities typically also have longer time between order date and delivery, such that capacity usage does not increase linearly with quantity (see Appendix Table A.1). Therefore, increases in quantity result in smaller increases in capacity usage. Expressing quantities in logs captures this pattern in a reduced-form manner without requiring estimating any additional parameters that govern how quantities are spread over time to map to capacity. Additionally, the log specification is also reasonable for revenue, as larger orders are associated with lower prices (see Appendix Table A.2), which may reflect lower costs from learning-by-doing for the specific fabric on the supplier side, consistent with evidence in the industry in Adhvaryu et al. (2023).

The second assumption is that the supplier's total capacity for the relational contract is the mean capacity in the pre-period. This assumption reflects that the relational contract is conceptualized as providing consistent capacity usage around an agreed upon capacity. To estimate the pricing function in the relational contract, I assume that relational contract quantities are drawn from the gamma distribution based on the method of moments estimator for the observed relational contract quantities. I chose a gamma distribution because it is a simple two parameter distribution with positive support and fits the data well.[31] Furthermore, using a gamma distribution nests many possible distributions, including other commonly used distributions (*e.g.,* chi-squared and exponential distributions are special cases of a gamma distribution). Overall, this approach incorporates scale economies, as suppliers with larger capacity experience the same deviation in actual quantity as a smaller change in capacity. It follows that large suppliers who receive many orders can likely reduce their capacity variance as they receive more stochastic end client demand shocks that, on average, tend to offset each other over many orders.

Additionally, it is necessary to estimate the counterfactual capacity variance that relational contract suppliers would experience if they were not in the relational contract. As the relational contract is designed to reduce capacity variance, I assume that relational contract suppliers' counterfactual market capacity utilization distribution would have the same mean as the relational contract but higher variance, as in the conceptual model. Because I do not have data from suppliers about total capacity utilization by buyer type (*i.e.,* relational contract and market), I assume the ratio of capacity variances by buyer type matches the empirical ratio of average within-supplier capacity variance for non-relational contract suppliers relative to relational contract suppliers.

Using all non-relational contract suppliers would likely lead to overestimates of the variance that relational contract suppliers would face outside the relational contract given that non-relational contract suppliers work with other buyers and may even be in relational con-

[31]The smallest p-value from a Kolmogorov–Smirnov test for each supplier is .55 for a test where the null is that the empirical distribution is different from the gamma distribution. Given concerns that the Kolmogorov–Smirnov can be too conservative in small samples, we can also evaluate the p-values for whether the empirical distribution is larger or smaller than the theoretical gamma distribution. Even using one-sided tests, the smallest p-value is still .28.

tracts with other buyers. Therefore, I restrict the set of non-relational contract suppliers for this calculation to the placebo suppliers.[32] Focusing on placebo suppliers to define the relevant set reflects that placebo suppliers are arguably the closest counterfactual for relational contract suppliers.[33]The estimates support this interpretation, as the counterfactual variance when working with the market is estimated to be 8.8 times larger due to demand assurance in the relational contract reducing average standard deviation of within-supplier capacity utilization to only .047. For comparison, using all market suppliers to estimate the counterfactual variance would result in an unreasonably large estimate of the variance that the relational contract suppliers would face when selling to the market–971 times the variance that relational contract suppliers experience in the relational contract.

Last, I estimate the market price relative to costs using transaction data. Specifically, for each transaction, I compute the margin charged by the supplier as the ratio of the transaction price over the estimated production costs for the desired quantity for a mill at 100% capacity utilization, where estimates use cost data from the integrated supplier.[34] The market price, with incorporates hold-up, has an estimated margin of 7 percentage points over costs, which seems intuitively reasonable: it is a large enough to create incentives for strategic sourcing through vertical integration and relational contracts but not so large as to be implausible.

4.3 Separate Identification of Buyer Bargaining Parameter and Supplier Risk Aversion

It is not immediately apparent that separate identification of the buyer bargaining parameter and supplier risk aversion is possible as increasing either buyer bargaining parameter alone or increasing supplier risk aversion alone both reduce relational contract prices. With a high buyer bargaining parameter, the buyer receives a larger share of the surplus generated by the relational contract through lower prices. Similarly, when supplier risk aversion increases, holding the buyer bargaining parameter fixed, relational contract surplus increases. As some of this surplus is shared with the buyer (as long as the buyer bargaining parameter is not zero), higher risk aversion also decreases prices.

My identification strategy leverages two key sets of variation in the data. First, there

[32]As $log(0)$ is undefined, I interpret periods with no orders as having zero capacity utilization.

[33]Even using placebo suppliers, data limitations prevent using the obvious empirical analogue: capacity variance at the supplier level for placebo suppliers. Given possible errors in this moment, I take model fit seriously in section 4.5 to mitigate concerns that errors in outside moments lead to poor model performance.

[34]Using the internal cost estimates means that market prices are constructed using the fabrics that are produced internally as the buyer does not estimate costs for fabrics not produced internally.

is variation in how prices change with respect to capacity. The covariance of prices and capacity helps identify supplier risk aversion, as shifts to supplier risk aversion alter the *convexity* of the prices in the relational contract. More convex contracts are associated with more risk averse suppliers because convex prices equalize profit across fluctuations in capacity. This logic reflects that profit is highest when capacity is near one. Specifically, when capacity is low, even though cost per unit is low, small quantities require high prices to equalize profits. When capacity is high, cost per unit is high due to increasing marginal costs; therefore, prices need to be high to equalize profits. Second, there is variation in the level of relational contract supplier prices relative to market prices. The variation in the level of relational contract prices is key for identifying the buyer bargaining parameter, as shifting this parameter reduces prices for all capacity realizations.

If only one price capacity tuple is observed, it is impossible to separately identify buyer bargaining parameter and risk aversion. Specifically, for any price capacity tuple, given a buyer bargaining parameter, shifting risk aversion changes prices (with higher risk aversion leading to a lower price and vice-versa). Analogously, given any risk aversion, changing the buyer bargaining parameter can shift the price for a given quantity to reach the same price. As shown in Figure 11, both a buyer bargaining parameter of .25 with a supplier risk aversion of 5 and a buyer bargaining parameter of .72 with a supplier risk aversion of .5 result in the same price when capacity is one.

Figure 11: Relational Contract Prices for Different Parameters

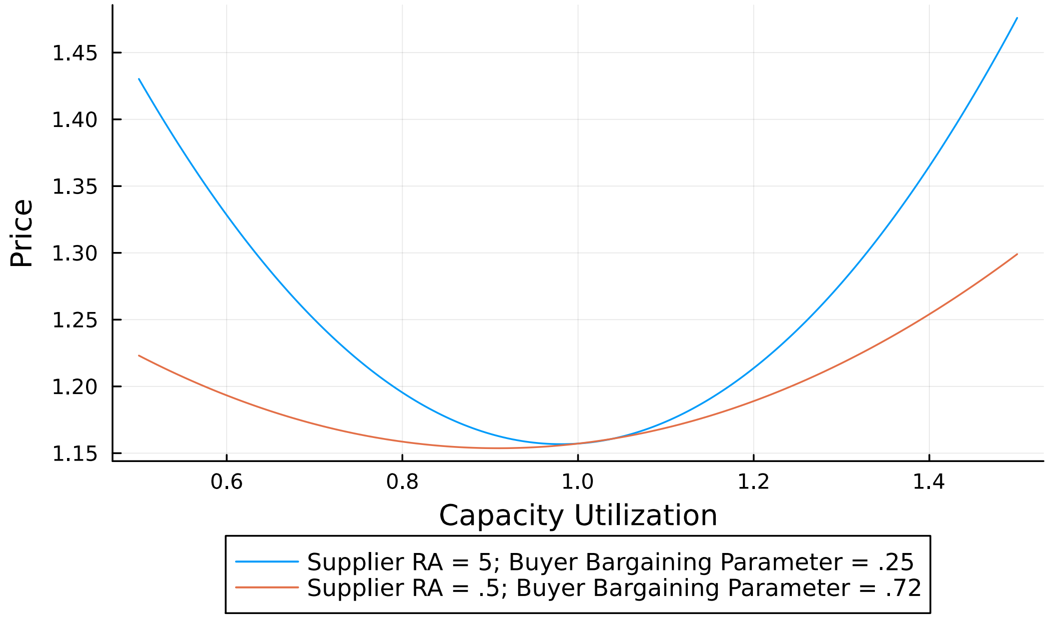

However, once a *second* price capacity tuple is observed for a different capacity, it becomes possible to separately identify buyer bargaining parameter and supplier risk aversion. Figure 11 highlights how lower supplier risk aversion results in a less convex contract. Therefore, as the second price capacity tuple pins down the convexity of the contract, both parameters are separately identified.

To further visualize this identification strategy, figures 12 and 13 show the effects of individually shifting supplier risk aversion and buyer bargaining parameter, respectively. They highlight that shifting the risk aversion parameter mostly shifts convexity while shifting the buyer bargaining parameter mostly shifts levels.

Figure 12: Shift Supplier Risk Aversion

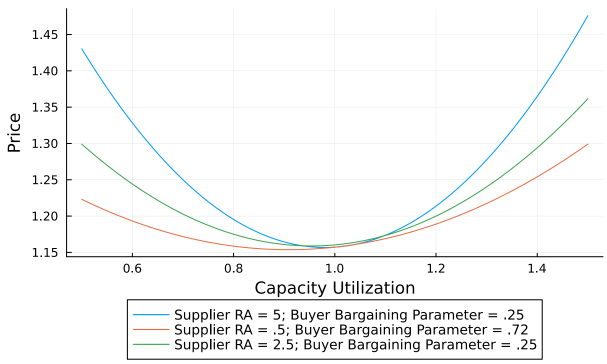

Figure 13: Shift Buyer Bargaining Parameter

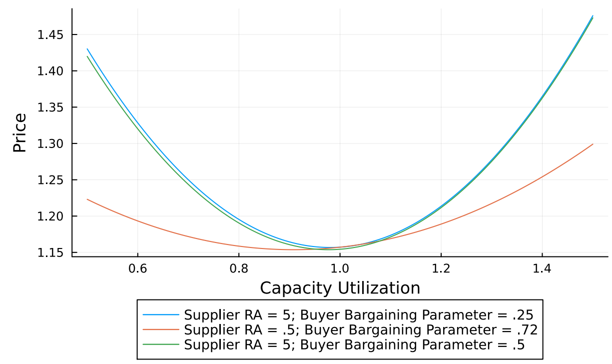

To bring this identification argument to the data, I leverage supplier-specific capacity shocks that reflect stochastic end client demand. Specifically, although the relational contract has a target capacity, institutional features of the garment production process–specifically, that the buyer knows neither which samples will be ordered by the client nor what quantities would be conditional on an order being placed–mean that the relational contract reduces capacity variance but is unable to hit the capacity target every period. In Appendix A.8.1, I discuss this variation in more detail and provide additional evidence that these capacity shocks serve as exogenous variation.

4.4 Buyer Bargaining Parameter and Supplier Risk Aversion Estimates

Table 2: Parameter Estimates from Structural Model

	Buyer Bargaining Parameter	Supplier Risk Aversion
Minimum	.0094	0
25th Percentile	.0301	.0026
Median	.6214	.0085
75th Percentile	.9155	.0223
Max	.9985	14.6811
Mean	.5113	.8282
Standard Deviation	.4286	3.4573
Correlation Between Supplier Risk Aversion & Buyer Bargaining Parameter	.2847	

Note: Estimates from model.

Parameter estimates in Table 2 (confidence intervals in Appendix Table A.10) illustrate that the buyer typically receives slightly more than half of the relational contract surplus,[35] It follows that the selected relational contract suppliers are also likely reasonably large firms, consistent with the low estimates of supplier risk aversion (with the exception of one outlier supplier with high risk aversion and a high buyer bargaining parameter). Together, these results suggest that relational contract suppliers are at least fairly large sophisticated firms, consistent with positive selection into relational contracting. This positive selection matches the intuition that creating and maintaining relational contracts may require organizational capabilities within the firm, at least for more complex relational contracts as in this setting (Macchiavello and Morjaria, 2023). The estimate of the tuning parameter λ is small at .002, suggesting that a large tuning parameter is not the reason why estimated risk aversion is low.

I compare estimated risk aversion with estimates from other studies in Appendix Table A.5.[36] Estimated risk aversion for the relational contract fabric suppliers is close to that of exporting coffee mills studied in Blouin and Macchiavello (2019) (.0085 for fabric suppliers and .0068 for coffee mills), who describe the mills as "large firms by developing country standards" which average over \$3.5 million a year in sales and about \$2 million in total assets. As these firms are likely the most similar in size, sophistication, and setting both with respect to market institutions (as firms in low- and middle-income countries) and stakes of risk (contracts with buyers) to the fabric suppliers studied in this paper, it is promising that risk aversion estimates are quite similar. Table A.5 also highlights that risk aversion estimates in the literature vary widely based on the methodology, the setting, the economic agent, and stakes of the risky activity. Specifically, as the stakes increase, estimated risk aversion increases widely, reaching all the way up to 12.90.

4.5 Model Validation and Fit

Given the importance of quantifying the threat point effect for heterogeneous buyer and supplier types, including out-of-sample types, I prioritize model validation. First, I illustrate that the model matches an untargeted moment. Second, I show that the model has better performance in the post-integration period than two alternative estimators, recalling that post-integration observations are largely out-of-sample as model estimates primarily use pre-

[35] I explore one possible microfoundation for heterogeneity in buyer bargaining parameter in A.16.

[36] As many studies do not report CARA coefficients, I convert the reported risk aversion to a CARA coefficient to be comparable given the estimate and utility function used.

integration data.[37] Validation using difference-in-difference estimates follows the discussion of the threat point effect magnitude as a function of buyer and supplier characteristics, as it is necessary to quantify the threat point effect for the estimated in-sample buyer and supplier types before comparing it to the difference-in-difference estimates.

4.5.1 Model Validation: Untargeted Moment

As the model is non-linear and estimated parameters minimize root mean squared error of model-implied prices, estimated prices could be biased. Specifically, minimizing root mean squared error minimizes a combination of variance and bias. For example, an outlier high price can lead the model to increase all price estimates because root mean squared error penalizes large deviations heavily. In that case, the mean (and median) residual price would be biased, as the model overestimates other prices to reduce the residual for the outlier high price, and the residuals would not be symmetrically distributed around zero.[38] Therefore, I study the distribution of residuals to validate model fit. Figure 14 highlights that the distribution of residuals is reasonably symmetric and the mean residual is nearly 0 at $-.007$. I conclude that the model performs well for an untargeted moment: the mean residual.[39]

[37]Specifically, supplier risk aversion and buyer bargaining parameter estimates use only pre-integration data, conditional on the tuning parameter.

[38]Simulations of this exact scenario verify that indeed the model does result in biased average prices. The bias can be positive or negative depending on what quantities have outlier prices, as outlier prices also influence the convexity of the relational contract. If outlier high prices are in the middle of the quantity distribution, then the model reduces risk aversion to reduce convexity and, concordantly, reduces the buyer bargaining parameter to prevent price levels from decreasing too much. As price levels shift up on average due to the buyer bargaining parameter change, the mean bias is positive. On the other hand, when the outliers occur for extreme capacities, then risk aversion tends to increase as convexity increases, leading to lower buyer bargaining parameters. As price levels then tend to decrease, mean bias is negative.

[39]The median residual, which is also untargeted, is even closer to 0 at .002.

Figure 14: Untargeted Moments: Bias

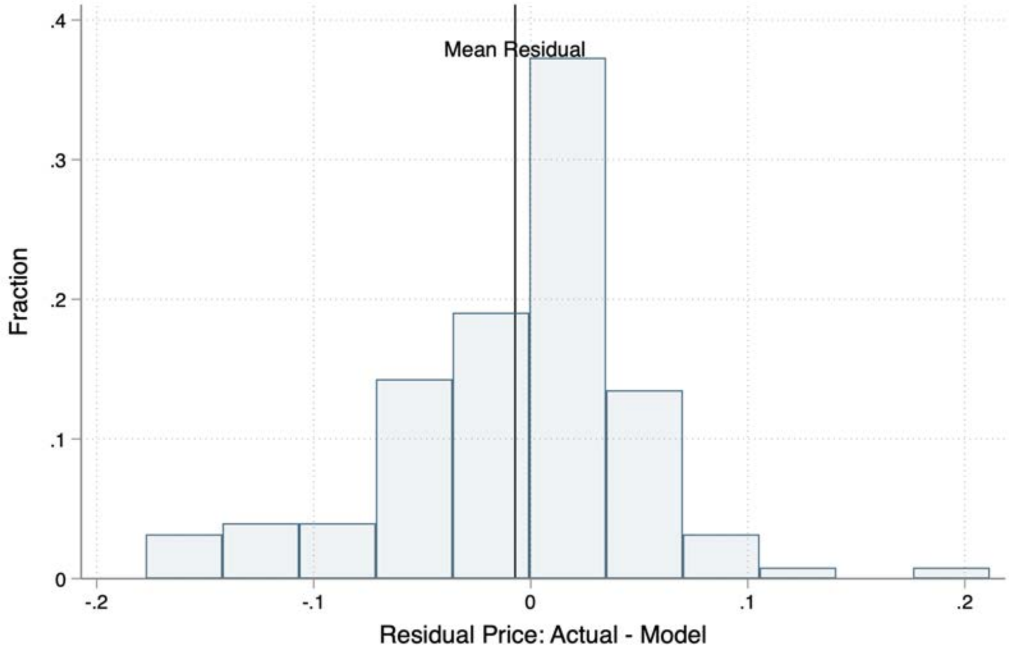

Note: Comparison of in-sample actual and model prices. Mean residual is $-.007$.

4.5.2 Out-of-Sample Fit

I also analyze out-of-sample fit, comparing the distribution of supplier root mean squared error in the post period from the model with the distribution of supplier root mean squared error in the post period using two alternative estimators in Table 3. First, I compare model estimates to the mean price for the supplier from the pre-integration quarters. My relational contract bargaining model performs better than the mean pre-integration price broadly throughout the distribution of supplier root mean squared errors, including at the minimum, the max, and all quartiles. Specifically, mean and median supplier root mean squared error are seven percent higher using the mean pre-integration price than using the model.

I also estimate out-of-sample prices using an OLS regression of prices on capacity and a constant in the pre-period. This estimation routine has two parameters per supplier, namely the coefficient on capacity and the constant, meaning that it uses the same number of parameters per supplier as the structural model (buyer bargaining parameter and supplier risk aversion). It follows that gains relative to these OLS estimates emphasize the benefit of the structural model. Median and mean root squared error using the OLS estimates are, respectively, 23 and 12 percent higher than using the model. Appendix Figure A.22 shows the distribution of root mean squared error across suppliers for each estimator.

Table 3: Out-of-Sample Fit Comparison

	Structural Model	Supplier Mean	OLS
Minimum	.0146	.0182	.0183
25th Percentile	.0482	.0551	.0683
Median	.0741	.0790	.0913
75th Percentile	.1003	.1004	.1091
Max	.1818	.1900	.1701
Mean	.0787	.0844	.0883
Standard Deviation	.0441	.0417	.0366

Note: Comparisons are for the distribution of supplier root-mean squared error in the post-integration error by estimator.

As the OLS model does not have a natural economic motivation for the tuning parameter, whereas economic intuition suggests extreme values for firm risk aversion are unlikely, OLS results do not incorporate a tuning parameter. However, as there is an econometric motive for including the tuning parameter, namely to reduce overfitting, Appendix A.12 compares model fit to OLS estimates with a tuning parameter. The fit of the OLS estimates with a tuning parameter is between the pre-integration supplier mean and standard OLS estimates without the tuning parameter.

I show robustness to *i)* incorporating supplier heterogeneity in marginal costs and *ii)* estimating the entire model, including the cost convexity, simultaneously using GMM with the outside parameters as additional moments to match.[40]

4.6 Quantifying the Threat Point Effect by Buyer and Supplier Characteristics

I then study threat point effect heterogeneity with respect to buyer and supplier characteristics for different buyer and supplier types, both in-sample and out-of-sample. Figure 15 shows the estimated threat point effect, expressed as a percent of average pre-integration relational contract price. Estimates show that the minimum risk aversion such that the threat point effect influences relational contract prices in this setting is 16.6, which is slightly larger than the upper bound of estimates in the literature (12.9 in Gandelman and Hernandez-Murillo

[40] Email author results for these robustness checks, which are all works in progress.

(2014)) but the same order of magnitude. However, this threshold level would decrease in settings with more cost convexity or where the relational contract provides more risk reduction relative to transacting in the market, as both of these factors decrease the cut-off level of risk aversion below which there is no relational contract. Relational contracts can provide more risk reduction in environments where they mitigate risk on both the demand and supply side of the market, such as for agricultural workers facing productivity risk as in Jayachandran (2006). This example is relevant as agricultural settings in low- and middle-income countries can feature relational contracts even for smallholder farmers as in Macchiavello and Morjaria (2020); these farmers are also likely to have highly convex production costs given that a key input is mostly fixed (land) as well as higher levels of risk aversion as low-income individuals in settings without strong risk-coping institutions. Additionally, CARA risk aversion parameters also tend to increase for larger stakes. Estimates of supplier risk aversion in this study are for one quarter's worth of orders presumably for a large supplier; the same volume would presumably be a larger stake for a smaller firm with lower volumes, magnifying the potential for small firms to be more risk averse in general.

Figure 15: Threat Point Effect Magnitude (% of pre-integration relational price)

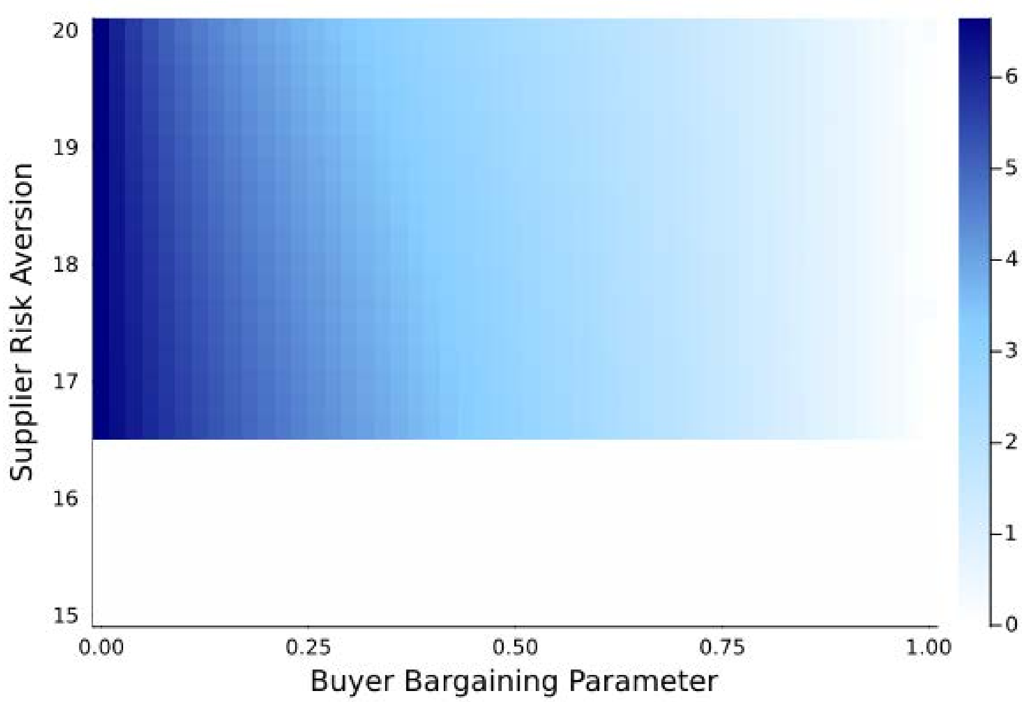

Note: Model estimates for the change in relational contract price as a percent of pre-integration relational contract price. Heatmap legend is to the right of the heatmap. Model estimates use the average capacity and demand assurance across in-sample relational contract suppliers.

As illustrated in the conceptual model, the threat point effect is largest when the buyer bargaining parameter is small, as then the supplier's outside option, which is unaffected by vertical integration, determines prices. The threat point effect also decreases slightly as risk aversion increases due to the algebra of the Nash product.

The results highlight that the threat point effect can be quantitatively important, reducing prices by up to 6.7% of pre-integration relational contract when the market is the buyer's threat point. For comparison, 6.7% is larger than the VAT in India during the study time period for apparel, which was only 5%.[41] As another comparison, when suppliers operate at the profit maximizing capacity, market prices are 7.36% higher than unit costs at 100% capacity, reflecting suppliers' ability to hold-up the buyer. Therefore, the threat point effect is also large relative to markups charged by suppliers holding up the buyer.

We can also analyze the threat point effect magnitude by examining the gains from relational contracts, from the buyer's perspective, relative to market prices in figure 16, which shows the discount relative to market prices expressed as a percent of market prices. Overall, the discount increases in both the buyer bargaining parameter, as the buyer's larger share of the surplus comes through lower prices, and in supplier risk aversion, as the surplus from the relational contract is larger for the same prices when supplier risk aversion is higher, some of which is distributed to the buyer for any positive bargaining parameter in the form of lower prices.

[41] GST rates for textiles and apparel in India.

Figure 16: Buyer Gains from Relational Contracts (% of market price)

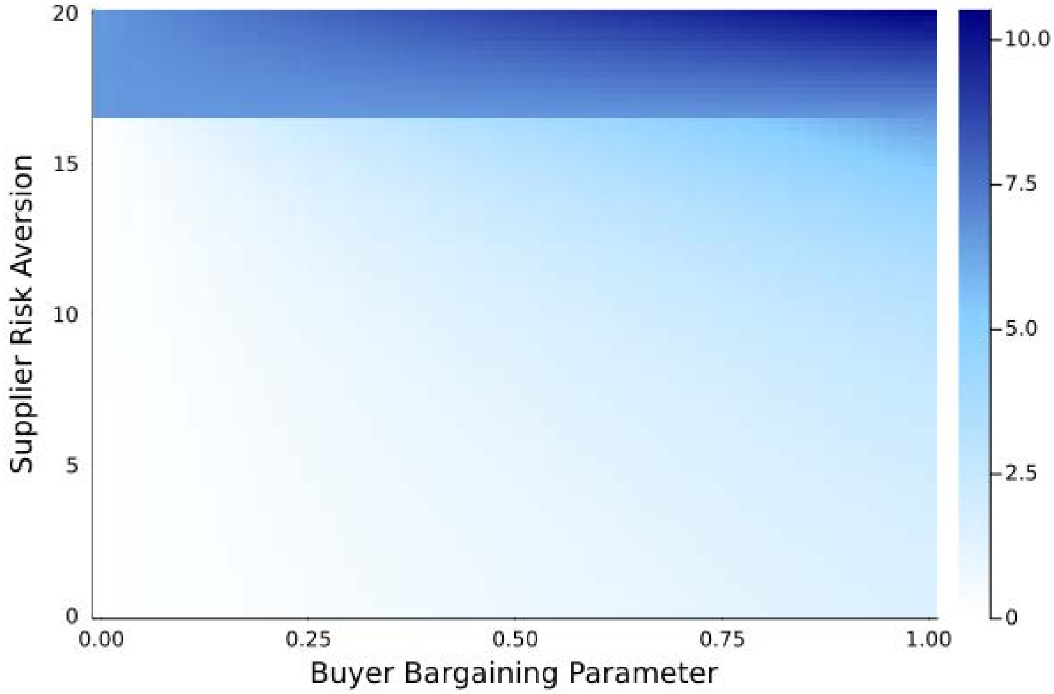

Note: Model estimates for the discount in the relational contract relative to the market as a percent of the market price. Heatmap legend is to the right of the heatmap. Model estimates use the average capacity and demand assurance across in-sample relational contract suppliers.

The heatmap emphasizes the quantitative importance of the threat point effect, with a large discontinuity visible when the threat point effect reduces relational contract prices. This heatmap explicitly highlights how relational contracts, even when effective, magnify incentives to vertically integrate. For example, when surplus is split equally between the buyer and the supplier, the gain from the relational contract increases by 76% from 3.8% to 6.7% relative to market prices when the threat point effect starts to reduce prices. This difference is economically meaningful–relational contract prices for in-sample relational contract suppliers pre-integration are on average only 2.2% lower than market prices and this relatively small gain is still sufficient for the buyer to find it worthwhile to establish and maintain relational contracts.

These results highlight how that the threat point effect can be quantitatively important, creating incentives for buyers to vertically integrate even when relational contracts resolve the hold-up problem.

4.7 Model Validation using Difference-in-Difference Estimates

In addition to the validation and fit exercises in 4.5, I also validate the structural model using a difference-in-difference design. The difference-in-difference design directly estimates the threat point effect by comparing suppliers exposed to vertical integration based on the fabrics they produce with suppliers unexposed to vertical integration before and after the additional integrated capacity is added. Specifically, suppliers are considered exposed to vertical integration (*i.e.,* "treatment" suppliers) if the fabrics they produce pre-integration can also be made at the integrated supplier, such that the buyer switching to the integrated supplier is a credible threat to displace these suppliers. Alternatively, suppliers whose pre-integration fabrics do not overlap with the integrated supplier are not exposed to integration (*i.e.,* "control" suppliers). Importantly, all comparisons are within supplier type, meaning that relational contract suppliers are compared to other relational contract suppliers.

I then validate the model by showing that the direct estimate of the threat point effect using the difference-in-ddifference model coincides with the model-implied threat point effect for in-sample suppliers. In addition to validating the magnitude of the threat point effect, the difference-in-difference event study coefficients also serve as additional untargeted moments to validate the model as the structural model only targets the difference between model-implied and actual prices without incorporating any comparison across suppliers based on exposure to integration. As the difference-in-difference estimates are for both the pre-integration and post-integration quarters while the structural estimates use only pre-integration data, validating the model with these untargeted moments includes both in-sample and out-of-sample validation.

4.7.1 Difference-in-Difference Estimation Approach

In a standard difference-in-differences design in this setting, the identifying assumption would be that counterfactual potential outcomes evolve in a parallel pattern for treatment and control suppliers. This assumption is *not* that assignment of suppliers to treatment status is random, which could be problematic as the analysis in Morton (2023) demonstrates that fabrics were not randomly chosen to be integrated: fabrics with more relational contract suppliers are less likely to be brought in-house because vertical integration and relational contracts are substitutes. However, as suppliers are not selected into relational contracts based on individual fabrics (as discussed in Morton (2023)), suppliers are unlikely to be intentionally selected to be exposed to integration.

Furthermore, the long history included in the data facilitate comparing trends in outcomes for many pre-treatment periods to support the assumption of parallel counterfactual outcomes. However, the data suggest that fabrics were brought in-house based on trends in fabric demand over time, creating non-random assignment in exposure to vertical integration. Therefore, I use a doubly-robust difference-in-difference design, described in detail in Appendix A.13.2. This design weakens the parallel assumptions trend to conditional parallel trends; in my empirical setting, comparisons are conditional on supplier exposure to trends in fabric demand, where trends are for fabric demand across all suppliers.

4.7.2 Validation: Threat Point Effect Magnitude Comparison

The model makes sharp predictions for the effects of vertical integration on both prices and quantities. As the largest estimated supplier risk aversion is below the threshold for the threat point effect to exist (see Appendix Figure A.24), difference-in-difference estimates should not show any meaningful change in prices for relational contract suppliers. Additionally, the model predicts that there should be no change in quantities for relational contract suppliers, with volume reduced from market suppliers rather than relational contract suppliers given that downstream demand from end clients does not increase. I find reduced-form estimates of the threat point effect consistent with the model.

Figure 17: Difference-in-Difference Estimates for Effect of Vertical Integration on Standardized Prices

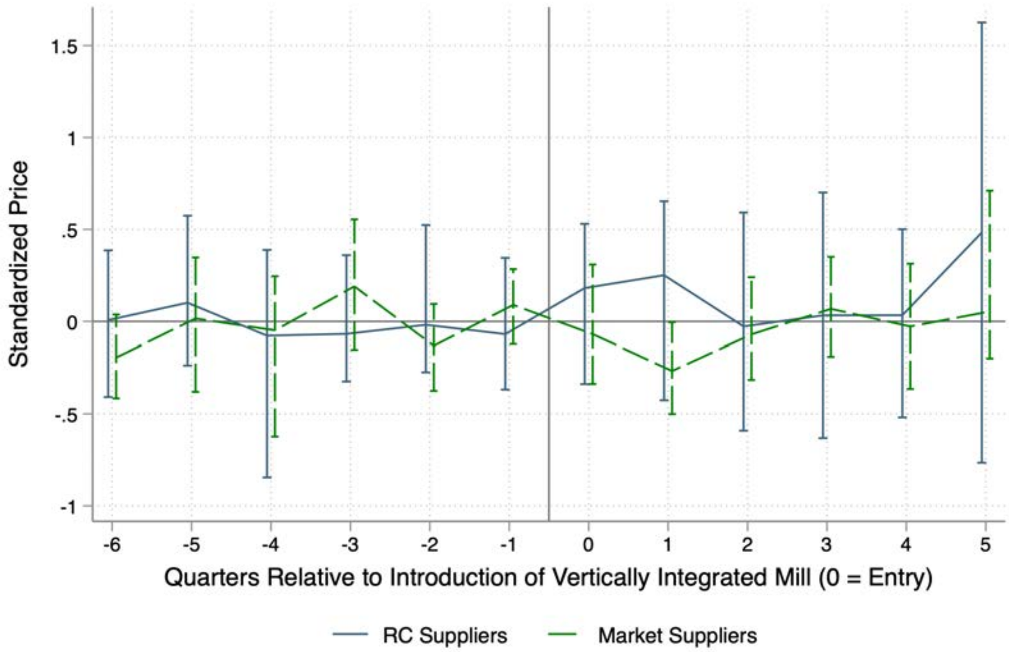

Note: Data from universe of fabric transactions by buyer. Doubly-robust difference-in-differences event study estimates with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

Figure 18: Difference-in-Differences Estimates for Effect of Vertical Integration on Quarterly Volume

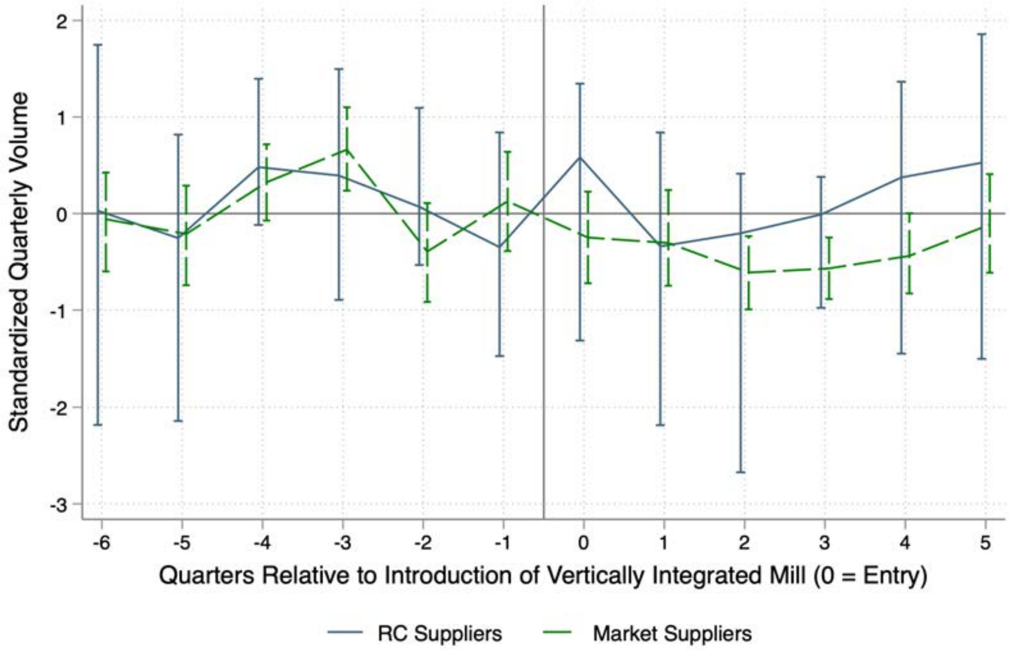

Note: Data from universe of fabric transactions by buyer. Doubly-robust difference-in-differences event study estimates with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

The difference-in-differences results in Figure 17 provide evidence that vertical integration did not affect prices for either supplier type, matching model predictions given estimated buyer bargaining parameters and supplier risk aversion parameters. Additionally, Figure 18 highlights that quantities (standardized by supplier) for relational contract suppliers did not decrease while quantities for non-relational contract suppliers did decrease. For both outcome variables and both supplier types, there are no obviously apparent pre-trends–the most visible pre-trend is that volumes for relational contract suppliers were *decreasing* pre-integration. Although this pre-trend is fairly weak, and there is no clear reason why it would occur, the reversal in sign after integration at least suggests that relational contract volumes did not decrease due to vertical integration.

As the few relational contract suppliers in the data limits statistical power, I also analyze the average treatment effect across all pre and post periods to increase power in Table 4.[42]

[42]Bootstrapped confidence intervals account for clustering at the supplier level, as highlighted in Bertrand et al. (2004), as well as uncertainty in the propensity score and regression outcome models.

The pooled results illustrate that non-relational contract suppliers experience a statistically significant, and economically meaningful at .4 standard deviations (on average $\sim$ 120,000 square meters of fabric per quarter), decrease in quantities. However, relational contract suppliers do not have any statistically significant treatment effect for quantities or prices; although confidence intervals are not always small, point estimates are economically close to zero. Additionally, I show that results are robust to using an alternative measure of volume that has many desirable properties, including scale-invariant t-statistics: the quartic root (Thakral and Toh, 2023).[43]

I also verify that vertical integration does not have unanticipated effects on other outcome variables by analyzing effects on reliability, measured as the percent of on-time deliveries. I also analyze effects on transaction counts, as the decrease in market supplier quantity can reflect a combination of fewer transactions or reallocating sampling opportunities with smaller volumes to market suppliers. I do not find statistically significant results for any pooled coefficient for these outcome variables. I include the full event study plots for these outcome variables in Appendix A.14.1, which also document no treatment effect on these outcomes.

Table 4: Pooled Pre- and Post-Treatment Effects

	Standardized Volume	Volume (Quartic Root)	Volume-Weighted Standardized Price	Percent of Volume Delivered On-Time	Transaction Count
RC Suppliers Pre	.07	.85	$-.02$	$-.02$	$-.28$
	$[-.28, .23]$	$[-.83, 1.57]$	$[-.12, .12]$	$[-.09, .13]$	$[-3.54, 7.45]$
RC Suppliers Post	.15	.91	.16	$-.06$	2.21
	$[-1.38, .79]$	$[-9.98, 6.59]$	$[-.42, .59]$	$[-.41, .17]$	$[-49.59, 20.66]$
Market Suppliers Pre	.09	1.16	$-.02$	$-.01$	.69
	$[-.05, .22]$	$[.01, 2.18]$	$[-.16, .07]$	$[-.24, .03]$	$[-.55, 1.83]$
Market Suppliers Post	$-.38$	-4.88	$-.05$	.14	$-.28$
	$[-.75, -.01]$	$[-7.77, -1.78]$	$[-.21, .21]$	$[-.17, .28]$	$[-3.48, 3.04]$

Note: Data from universe of fabric transactions by buyer. Estimates represent the sum of doubly-robust difference-in-differences pre- and post-integration event study coefficients with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

In Appendix A.14.2, I document robustness to using the synthetic difference-in-differences

[43] The positive and significant effect for market suppliers pre-integration for the quartic root is consistent with a small anticipation effect. This anticipation effect corresponds with the integrated supplier starting to produce fabrics a the end of the last pre-integration quarter before it started running fully at the beginning of the quarter. However, this result is only marginally significant, so I do not interpret it as suggestive of a meaningful decrease between the pre periods and the quarter right before the integrated supplier is fully operational.

approach proposed in Arkhangelsky et al. (2021). This methodology reproduces the results that quantities decrease only for market suppliers and prices do not change for either relational contract or market suppliers.

4.7.3 Validation: Difference-in-Difference as Untargeted Moments

As the model does not use the difference-in-differences moments,[44] specifically the estimation does not leverage any variation in supplier exposure to integration based on fabrics produced, I compare difference-in-differences event study coefficients using model prices with untargeted analogous estimates using actual prices. The model predicted prices in Figure 19 are all within the 95% confidence interval.

Figure 19: Untargeted Moments: Difference-in-Differences

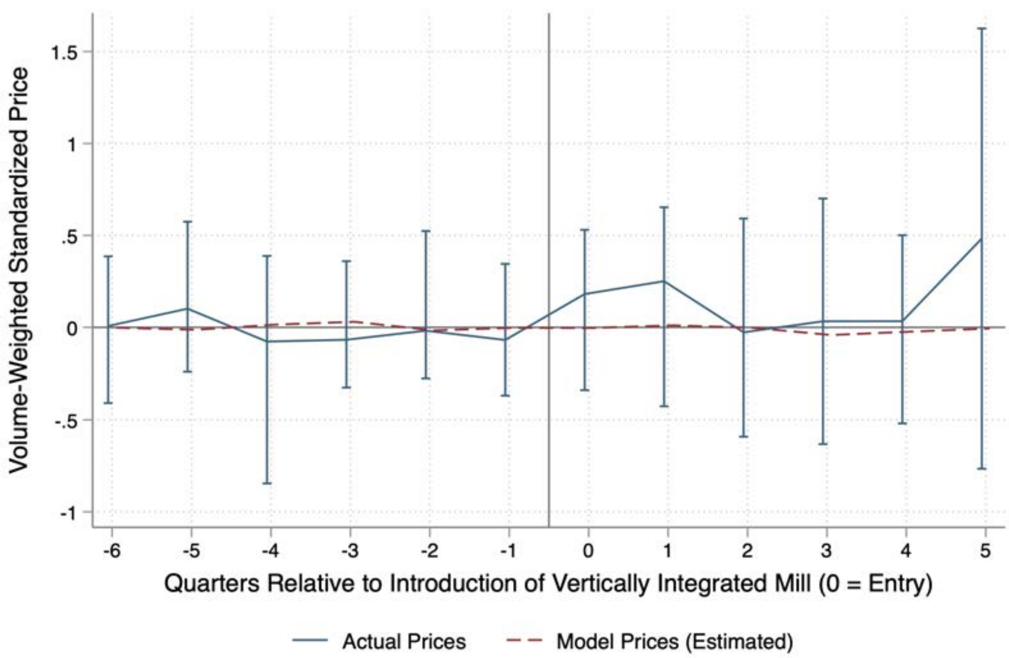

Note: Data from universe of fabric transactions by buyer. Doubly-robust difference-in-differences event study estimates with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

[44]Recall that these estimates are based on exposure to vertical integration based on fabrics before and after integration.

5 Policy Implications: Reallocating Surplus to Small Suppliers

The threat point effect decreases relational contract prices most for small constrained (*i.e.*, risk-averse) suppliers, motivating policy remedies given that policy makers likely prioritize small firms due to their large presence in the firm size distribution and role as the majority of employers in low- and middle-income countries (Hsieh and Olken, 2014). Prioritization of small firms could also reflect preferences to redistribute to reduce inequality, as reallocating profit to small firms could increase aggregate welfare if low levels of profit (and employee income) result in high marginal utility from additional profit. Last, preferences favoring small firms could reflect dynamic efficiency concerns in environments where productive small firms need to finance growth from contemporaneous profits given credit market frictions. Policy focus on small firms is evident empirically as initiatives targeting small and medium enterprises (SMEs) are omnipresent, such as the World Bank's policy advising on finance tools for SMEs[45] and the IMF's 2019 conference on SME financing.[46]

I consider two avenues for policy makers interested in redistributing surplus to small, risk averse suppliers. First, motivated by concerns common in industrial organization about market power, I consider increasing downstream buyer competition. Second, given the focus on missing markets in development economics, I create the missing market for insurance as insurance allows small firms to smooth profit without relying on the relational contract.

5.1 Increase Downstream Buyer Competition

I first study the effect of increasing downstream buyer competition, estimating relational contract prices in a different market structure where there are additional potential relational contract buyers that the supplier could trade with instead of the buyer studied. Policy makers have many tools to support the development of large exporting firms, including reducing red tape to export, supporting network formation between large firms and international end clients, changing tariffs, and even possibly subsidizing large firms to begin exporting. Moreover, policy makers could target barriers that constrain small firm growth, such as helping them access capital and other inputs necessary for growth.

To interpret the policy of increased downstream buyer competition through the lens of the model, I conceptualize increased downstream buyer competition as decreasing the buyer

[45] World Bank SME Finance

[46] IMF SME Financing

bargaining power parameter.[47] More relational buyers should improve the supplier's long-run outside option, but not the short-run outside option of selling on the market because it is not a credible threat for the supplier to add a relational contract buyer in the short-run. As long as the increase in buyer competition is known by the buyer, I propose that the buyer responds to competition by increasing the share of surplus suppliers receive to ensure supplier loyalty. This argument reflects similar economics as efficiency wages because both incentivize loyalty by offering a higher share of the surplus (Katz, 1986).

Suggestive empirical patterns based on the correlation between supplier attributes and estimated buyer bargaining parameters support this interpretation. Suppliers who provide fabrics with higher concentration (compared to other relational contract suppliers) are likely most valuable to the buyer. For example, if the buyer needs to procure a fabric for a specific garment at a discount relative to the market, then it needs to have a relational contract for that fabric. Additionally, suppliers who provide fabrics with higher concentration presumably make it easier for the buyer to fulfill the quantity guarantees in the relational contracts. To illustrate this idea, consider an extreme case where all suppliers sell the same fabric such that concentration is low. In this setting, if the aggregate quantity for that fabric is low, then the buyer will not be able to fulfill its quantity guarantees to all suppliers because assigning an order to one supplier to meet its quantity guarantee inherently means not assigning that order to another supplier. Therefore, suppliers who produce more concentrated fabrics, calculating concentration only among relational contract suppliers, need not compete with other suppliers for the same volumes.

I find that suppliers who sell products with low concentration are in relationships with lower buyer bargaining parameters, consistent with the buyer ensuring that these relational suppliers remain loyal as they are especially valuable. Appendix Section A.16 provides more information on the empirical approach used to show that that supplier concentration in fabric markets, relative to other relational contract suppliers, is negatively correlated with the estimated buyer bargaining parameters. Additionally, this proxy for relational contract supplier value is not significantly correlated with estimated risk aversion parameters, suggesting that the correlation between it and buyer bargaining parameters is not simply spurious.

[47] I halve the buyer bargaining parameter.

Figure 20: Downstream Buyer Competition and Relational Contract Prices

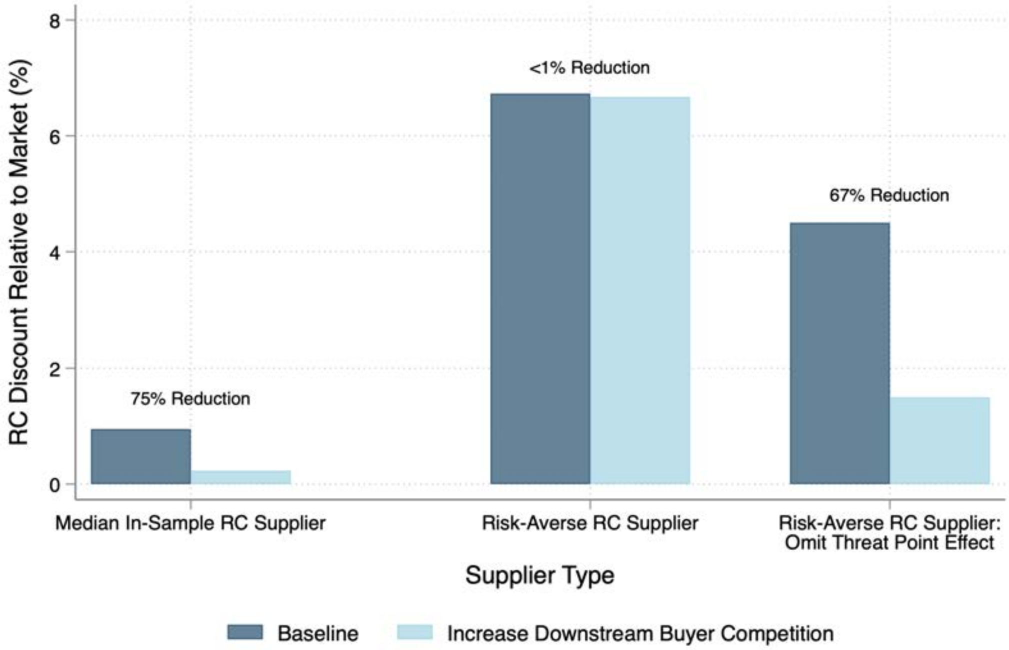

Note: Model estimates. Figure shows change in relational contract discount from halving the buyer bargaining parameter from the estimated median for relational contract suppliers. Risk aversion for the median in-sample relational contract supplier is the median of estimated risk aversion parameters. Risk aversion for the risk-averse supplier is the minimum risk aversion such that there is a threat point effect. Model estimates use the capacity and demand assurance for a median in-sample relational contract supplier.

Figure 20 illustrates the change in discount in the relational contract relative to the market, expressed as a percent of the market price, induced by increased downstream buyer competition. The effect of this policy is small, especially for risk-averse suppliers. Specifically, while halving the median estimated buyer bargaining parameter results in a large 75% reduction in discount for a supplier at the median risk aversion of in-sample suppliers–from .95% to 0.23%–the same reduction in buyer bargaining parameter has a negligible effect for suppliers with the minimum risk aversion for the threat point effect to matter, reducing the discount by less than 1% from 6.73% to 6.67%. Importantly, ignoring the threat point effect *overestimates the benefits* of reducing the buyer bargaining parameter, as without the threat point effect the discount would decrease by 66% from 4.5% to 1.5%.

Intuitively, increasing downstream buyer competition is not an effective policy because, once the buyer uses the integrated supplier as its threat point, prices are so close to the buyer's participation constraint that there is little room to shift relational contract prices. It follows that changing the surplus sharing rule has minimal effects on prices. This intuition clarifies why omitting the threat point effect results in incorrect policy conclusions for small risk-averse firms and why this policy is effective for large suppliers–when the buyer's threat point is the market, there is enough surplus that changing surplus sharing effectively reduces relational contract prices.

Appendix A.15 considers the effects of horizontal mergers of suppliers. Although horizontal consolidation does reduce the relational contract discount, consolidation effectively assumes away the problem of the existence of small constrained firms in the first place.

5.2 Create Missing Market for Insurance

The importance of missing markets in low- and middle-income country settings, as well as the limited effects of increasing downstream buyer competition, motivate analyzing the effects of adding the missing market for insurance. Specifically, I evaluate the change in the relational contract discount that would occur if suppliers had access to actuarially fair insurance over profits when trading with the market. This insurance would effectively blunt the effect of risk aversion on discounts the buyer receives because suppliers can smooth profit now without relying on the relational contract.[48] As all uncertainty is on the capacity and cost side in this setting, insurance over prices (or, equivalently, revenue) would function equally well, although such a contract would need to be non-linear as profits are not linear in capacity due to cost convexity.

While the government is presumably unlikely to offer such insurance (although price floor policies for agricultural goods are common and effectively insure agricultural producers against some adverse prices and could also improve producers' threat points), policies could support trade associations and cooperatives that provide price or revenue insurance (or at least support), and potentially even implement policies to mitigate moral hazard, such as pegging the volume of insurance offered to a running average of historic capacity or revenue.[49] Alternatively, the government could support policies to improve access to financial markets for suppliers to reduce their exposure to profit fluctuations. For example, improving access to credit and savings markets might help suppliers self-insure by building savings over time,

[48] Blouin and Macchiavello (2019) also shows how missing insurance markets interact with contracting behaviors, focusing on how supplier incentives to strategic default lead to lower levels of insurance provided by using index-priced contracts. While Blouin and Macchiavello (2019) focus on efficiency effects of the missing insurance markets, I focus on its implications for the distribution of surplus.

[49] As all uncertainty in this model for suppliers is from revenue, and not costs, price insurance would be isomorphic to profit insurance, although the insurance contract would need to non-linear in capacity to smooth profit given convex costs.

increasing credit access for suppliers could allow them to borrow rather than need relational contracts, and supporting the development of markets for derivatives or other financial products could allow suppliers to hedge their risk more effectively (although, admittedly, financial literacy and transaction costs could limit take-up). And, policies supporting small firms to access financing at more favorable rates are common in many settings, including high-income countries.

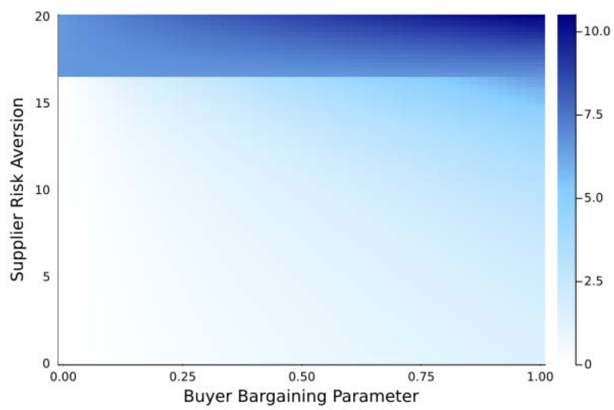

Figure 21: Without Insurance

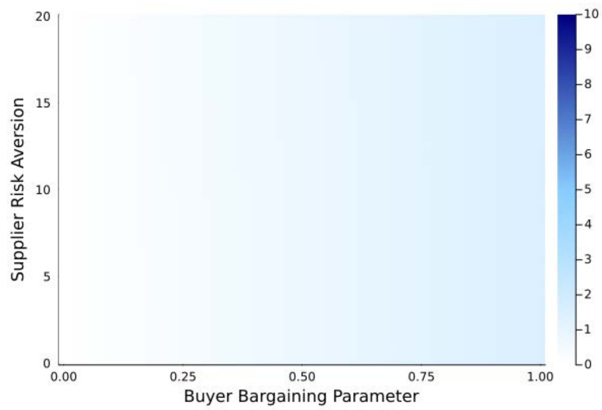

Figure 22: With Insurance

Note: Model estimates for the discount in the relational contract relative to the market as a percent of the market price. Heatmap legend is to the right of the heatmap. Model estimates use the average capacity and demand assurance across in-sample relational contract suppliers.

Creating the missing market for insurance reduces the discount offered by risk-averse suppliers in the relational contract. Furthermore, it effectively eliminates the threat point effect, with risk-averse suppliers no longer receiving significantly lower prices. Adding insurance is effective because it prevents the buyer from using the integrated supplier as its threat point to prevent the relational contract from breaking, as the supplier's threat point improves. In other words, the most effective approach to counteract the buyer's improved bargaining position due to a change in the buyer's threat point is symmetric: improve the supplier's bargaining position by improving its threat point.

6 Conclusion

This paper highlights that firms can vertically integrate to improve their bargaining position with external firms. The intuition is that vertical integration changes a firm's threat point, which improves its bargaining position vis-à-vis external firms. This improvement in bargaining position can improve contracting with external firms, allowing the integrated firm to extract more surplus–I call this change in contracts caused by the change in threat point due to vertical integration *the threat point effect*. I build a conceptual model studying how vertical integration affects relational contracts. The details of the model are motivated by the empirical setting of a large Indian garment manufacturer which added an integrated fabric supplier bargaining with external relational contract suppliers. The conceptual model highlights that the magnitude of the threat point effect depends upon the sharing rule for and level of surplus in the relational contract. The threat point effect can reach over 5% of market prices when suppliers are risk averse, as supplier risk aversion means supplier surplus is large even when prices are low. Therefore, the buyer can leverage its threat point change to further increase the relational contract discount. However, if the integrating party extracts all surplus pre-integration, then there is no threat point effect as there is no additional surplus available for the buyer to extract. If the level of surplus in the relational contract is so small pre-integration such the relational contract would break with the integrated firm as the threat point, then there also is no threat point effect. In this case, the integrating firm would prefer to keep the pre-integration relational contract, so the relational contract equilibrium would not change.

I estimate the model, finding that the threat point effect magnitude is quantitatively important for small risk-averse firms, reaching up to 6.7% of pre-integration relational contract prices. Given that small firms are both economically and policy relevant as the majority of firms and employers in low- and middle-income countries and experience the largest threat point effect, I evaluate two policies to reallocate surplus to them. First, motivated by concerns about competition and market power common in industrial organization, I analyze the effects of increasing downstream buyer competition. Although this policy helps large firms, it has a trivial effect for small firms. Importantly, not accounting for the threat point effect results in both *underestimates* of the level of the discount offered by small risk-averse firms and *overestimates* of the benefits of this policy. Second, motivated by the focus in development economics on missing markets, I create the missing market for insurance to allow small risk-averse firms to smooth profit outside the relational contract. This policy improves small firm surplus and eliminates the threat point effect.

More broadly, my results contribute to the growing body of evidence that challenges to smoothing profits in low-income country settings not only decrease welfare through preventing equalizing marginal utility across states, but also have important price effects. For example, agricultural producers in low-income country settings tend to buy and sell at the same time as other nearby producers due to limitations in their ability to store agricultural output, resulting in negative price effects as they "sell low and buy high" (Burke et al., 2018). And, inability to smooth over productivity shocks leads wage workers in low-income agricultural settings to increase their hours, magnifying the decrease in their wages as productivity shocks are correlated over space and time (Jayachandran, 2006). In my setting, small risk-averse firms accept lower average prices in order to smooth profits. Therefore, policies reducing exposure to profit (or income) fluctuations can have important welfare benefits not only through equalizing marginal utility but also by improving prices. Future research should seek to identify which policies are most effective at assisting small firms and producers in low-income settings to cope with risk, measuring policy effectiveness with respect to both smoothing across states and prices.

Last, my results also emphasize that increasing competition in low- and middle-income country markets with relational contracts may not always achieve desired policy goals of improving market performance and enhancing welfare. For example, increased competition reduces output and small producer (*i.e.,* farmer) welfare in Rwandan coffee markets (Macchiavello and Morjaria, 2020). Similarly, Brugues (2020) argues that mitigating market power alone would lead to welfare losses given contracting frictions in the context of relational contracts in the Ecuadorian manufacturing supply chain. My counterfactual analysis suggests that policies directly targeting the underlying missing markets that create demand for the relational contract may be more effective at meeting policy goals than increasing competition. For example, in the same empirical setting in Macchiavello and Morjaria (2020), improving farmers' access to financial markets to procure inputs could reduce their need for trade credit and allow them to benefit from downstream price competition for their coffee. The standard intuition of the benefits of competition would likely apply once firms rely less on relational contracts. Keeping in mind concerns that changes in second-best institutional settings can have unintended adverse effects and that demand for relational contracts in some settings emerges due to transaction rather than market features, future research should identify and evaluate policies creating missing markets (and, more broadly, mitigating market failures) that generate demand for relational contracts from small firms.

References

Acemoglu, Daron, Simon Johnson, and Todd Mitton, "Determinants of Vertical Integration: Financial Development and Contracting Costs," *The Journal of Finance*, 2009, *64* (3), 1251–1290.

Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham, "Returns to On-The-Job Soft Skills Training," *Journal of Political Economy*, 2023.

_ , Vittorio Bassi, Anant Nyshadham, and Jorge Tamayo, "No Line Left Behind: Assortative Matching Inside the Firm," *Working Paper*, 2019.

Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager, "Synthetic Difference-in-Differences," *American Economic Review*, December 2021, *111* (12), 4088–4118.

Asker, John, "Diagnosing Foreclosure due to Exclusive Dealing," *The Journal of Industrial Economics*, 2016, *64* (3), 375–410.

Atalay, Enghin, Ali Hortaçsu, and Chad Syverson, "Vertical Integration and Input Flows," *American Economic Review*, April 2014, *104* (4), 1120–48.

Banerjee, Abhijit V. and Esther Duflo, "Reputation Effects and the Limits of Contracting: A Study of the Indian Software Industry," *The Quarterly Journal of Economics*, August 2000, *115* (3), 989–1017.

_ , Greg Fischer, Dean Karlan, Matt Lowe, and Benjamin N. Roth, "Does the Invisble Hand Efficiently Guide Entry and Exit? Evidence from a Vegetable Market Experiment in India," *Working Paper*, 2022.

Bergquist, Lauren Falcao and Michael Dinerstein, "Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya," *American Economic Review*, December 2020, *110* (12), 3705–47.

_ , Craig McIntosh, and Meredith Startz, "Search Costs, Intermediation, and Trade: Experimental Evidence from Ugandan Agricultural Markets," *Working Paper*, 2023.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *The Quarterly Journal of Economics*, February 2004, *119* (1), 249–275.

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

Boudreau, Laura, Julia Cajal-Grossi, and Rocco Macchiavello, "Global Value Chains in Developing Countries: A Relational Perspective from Coffee and Garments," *Journal of Economic Perspectives*, August 2023, *37* (3), 59–86.

Breza, Emily and Andres Liberman, "Financial Contracting and Organizational Form: Evidence from the Regulation of Trade Credit," *The Journal of Finance*, 2017, *72* (1), 291–324.

Brugues, Felipe, "Take the Goods and Run: Contracting Frictions and Market Power in Supply Chains," *Working Paper*, 2020.

Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel, "Sell Low and Buy High: Arbitrage and Local Price Effects in Kenyan Markets," *The Quarterly Journal of Economics*, December 2018, *134* (2), 785–842.

Cajal-Grossi, Julia, Rocco Macchiavello, and Guillermo Noguera, "Buyers' Sourcing Strategies and Suppliers' Markups in Bangladeshi Garments," *The Quarterly Journal of Economics*, June 2023, pp. 1–60.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230. Themed Issue: Treatment Effect 1.

Cuesta, Jose Ignacio, Carlos Noton, and Benjamin Vatter, "Vertical Integration between Hospitals and Insurers," *Working Paper*, 2019.

de la Sierra, Raul Sanchez, "Whither Formal Contracts?," *Econometrica*, 2021, *89* (5), 2341–2373.

Forbes, Silke J. and Mara Lederman, "Adaptation and Vertical Integration in the Airline Industry," *American Economic Review*, December 2009, *99* (5), 1831–49.

Freyaldenhoven, Simon, Christian Hansen, and Jesse M. Shapiro, "Pre-event Trends in the Panel Event-Study Design," *American Economic Review*, September 2019, *109* (9), 3307–38.

Galenianos, Manolis and Alessandro Gavazza, "A Structural Model of the Retail Market for Illicit Drugs," *American Economic Review*, March 2017, *107* (3), 858–96.

Gandelman, Nestor and Ruben Hernandez-Murillo, "Risk Aversion at the Country Level," *Federal Reserve Bank of St. Louis Working Paper Series*, 2014.

Garg, Shresth, Pulak Ghosh, and Brandon Joel Tan, "Within firm supply chains: Evidence from India," *Journal of International Economics*, 2023, *144*, 103793.

Ghani, Tarek and Tristan Reed, "Relationships on the Rocks: Contract Evolution in a Market for Ice," *American Economic Journal: Microeconomics*, February 2022, *14* (1), 330–65.

Gil, Ricard, Myongjin Kim, and Giorgio Zanarone, "Relationships Under Stress: Relational Outsourcing in the U.S. Airline Industry After the 2008 Financial Crisis," *Management Science*, 2022, *68* (2), 1256–1277.

Hansman, Christopher, Jonas Hjort, Gianmarco León-Ciliotta, and Matthieu Teachout, "Vertical Integration, Supplier Behavior, and Quality Upgrading among Exporters," *Journal of Political Economy*, 2020, *128* (9), 3570–3625.

Harris, Adam and Thi Mai Anh Nguyen, "Long-Term Relationships and the Spot Market: Evidence from US Trucking," *Working Paper*, 2023.

Hawkins-Pierot, J. T. and Katherine Wagner, "Technology Lock-In and Optimal Carbon Pricing," *Working Paper*, 2022.

Hortaçsu, Ali and Chad Syverson, "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices," *Journal of Political Economy*, 2007, *115* (2), 250–301.

Hsieh, Chang-Tai and Benjamin A. Olken, "The Missing "Missing Middle"," *Journal of Economic Perspectives*, September 2014, *28* (3), 89–108.

Jayachandran, Seema, "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries," *Journal of Political Economy*, 2006, *114* (3), 538–575.

Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Christopher Udry, " Agricultural Decisions after Relaxing Credit and Risk Constraints," *The Quarterly Journal of Economics*, February 2014, *129* (2), 597–652.

Katz, Lawrence F., "Efficiency Wage Theories: A Partial Evaluation," *NBER Macroeconomics Annual*, 1986, *1*, 235–276.

Lafontaine, Francine and Margaret Slade, "Vertical Integration and Firm Boundaries: The Evidence," *Journal of Economic Literature*, September 2007, *45* (3), 629–685.

Macchiavello, Rocco, "Financial Development and Vertical Integration: Theory and Evidence," *Journal of the European Economic Association*, April 2012, *10* (2), 255–289.

_ , "Relational Contracts and Development," *Annual Review of Economics*, 2022, *14* (1), 337–362.

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

_ **and** _ , "Competition and Relational Contracts in the Rwanda Coffee Chain," *The Quarterly Journal of Economics*, December 2020, *136* (2), 1089–1143.

_ **and** _ , "Relational Contracts: Recent Empirical Advances and Open Questions," *Working Paper*, 2023.

_ **and Josepa Miquel-Florensa**, "Vertical Integration and Relational Contracts: Evidence from the Costa Rica Coffee Chain," *Working Paper*, 2017.

Morton, Russell, "Vertical Integration & Relational Contracts: Empirical Evidence on Substitutability from Indian Garment Manufacturing," *Working Paper*, 2023.

Rao, Manaswini, "Front-line Courts As State Capacity: Evidence From India," *Working Paper*, 2022.

Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe, "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *Working Paper*, 2022.

Sant'Anna, Pedro H.C. and Jun Zhao, "Doubly robust difference-in-differences estimators," *Journal of Econometrics*, 2020, *219* (1), 101–122.

Startz, Meredith, "The value of face-to-face: Search and contracting problems in Nigerian Trade," *Working Paper*, 2021.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199. Themed Issue: Treatment Effect 1.

Thakral, Neil and Linh Toh, "Scale Equivariance in Regressions," *Working Paper*, 2023.

Williamson, Oliver E., "Transaction-Cost Economics: The Governance of Contractual Relations," *The Journal of Law & Economics*, 1979, *22* (2), 233–261.

A Appendix

A.1 Input Sourcing by Supplier Type

Figure A.1: Fabric Sourcing by Supplier Type

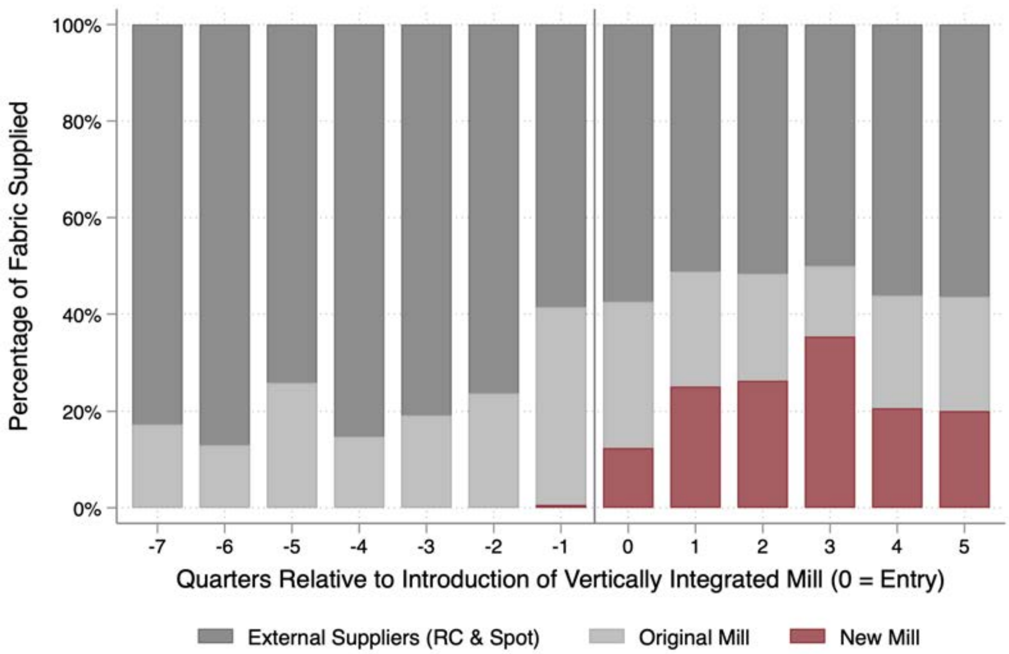

Note: Data from universe of fabric transactions by buyer. Bars represent the percentages of fabric volumes purchased during the quarter from each supplier type.

A.2 Relational Contracts in Fabric Procurement

Figure A.2: Distribution of Percentage of Months with Transactions (Pre-Integration)

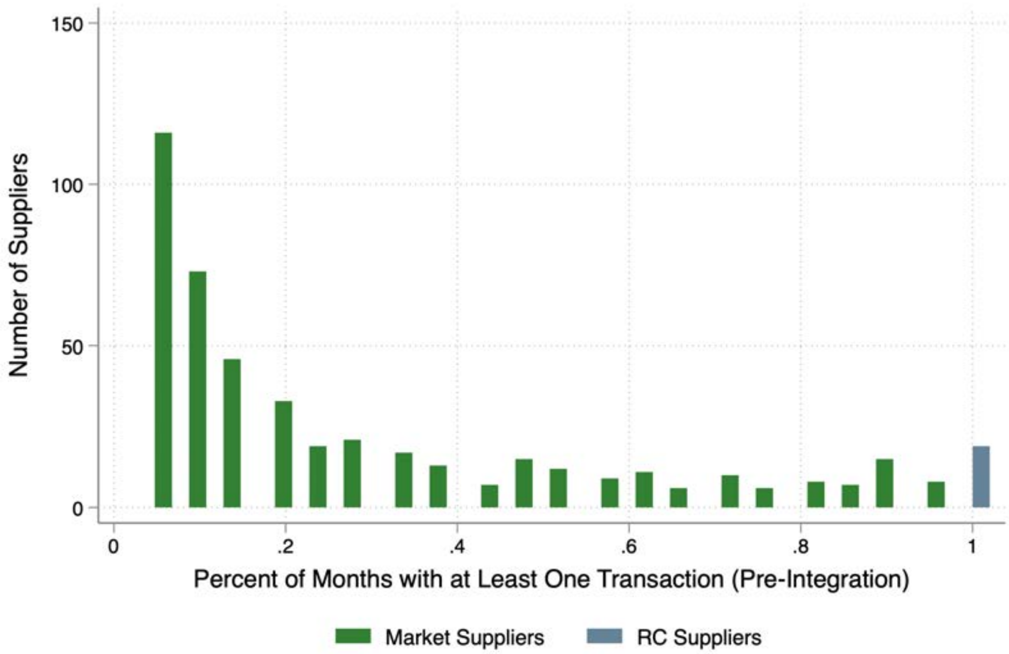

Note: Data from universe of fabric transactions by buyer. Bars represent the number of suppliers with the relevant percent of months with transactions.

Figure A.3: Variation in Monthly Quantity by Fabric

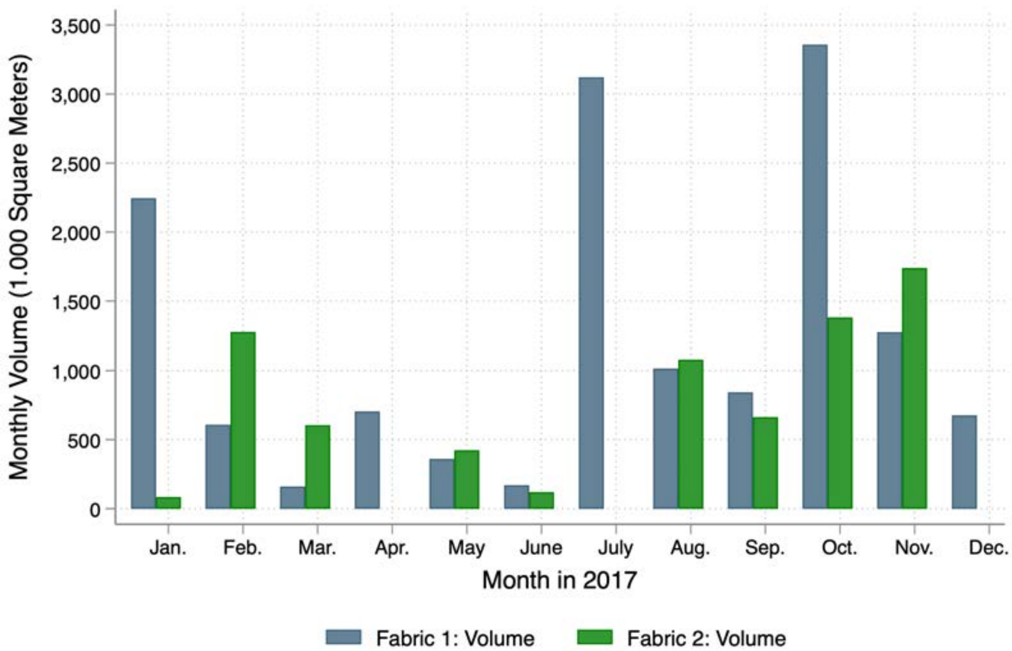

Note: Data from universe of fabric transactions by buyer. Bars represent the quantity of fabric purchased during the month in 2017 for two selected fabrics. Fabrics were selected based on data availability to illustrate variation.

Figure A.4: Discounts and Relational Contract Suppliers

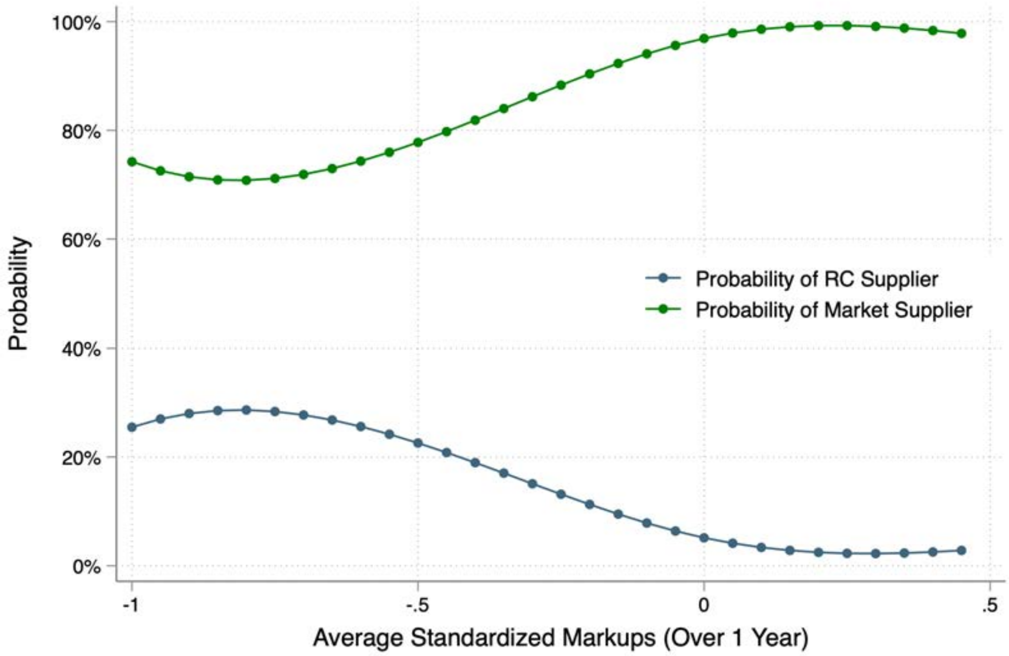

Note: Data from universe of fabric transactions by buyer. Probability estimates represent marginal effects from a non-parameteric series estimator calculated using transactions from the year prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.5: Discounts and Supplier Type

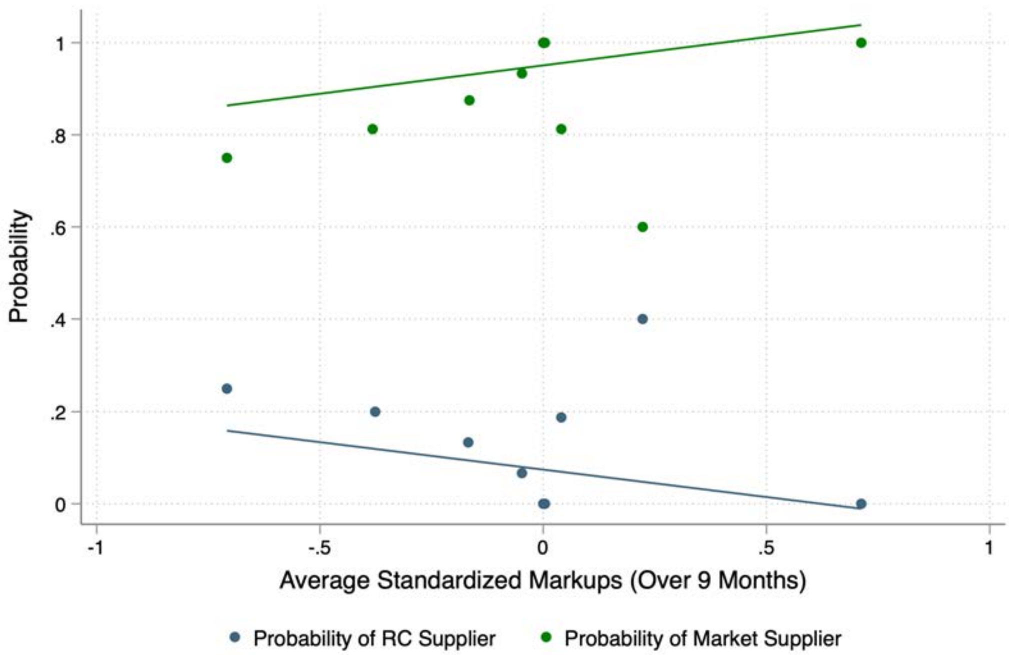

Note: Data from universe of fabric transactions by buyer. Probability estimates from a binscatter calculated using transactions from the three quarters prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.6: Discounts and Supplier Type

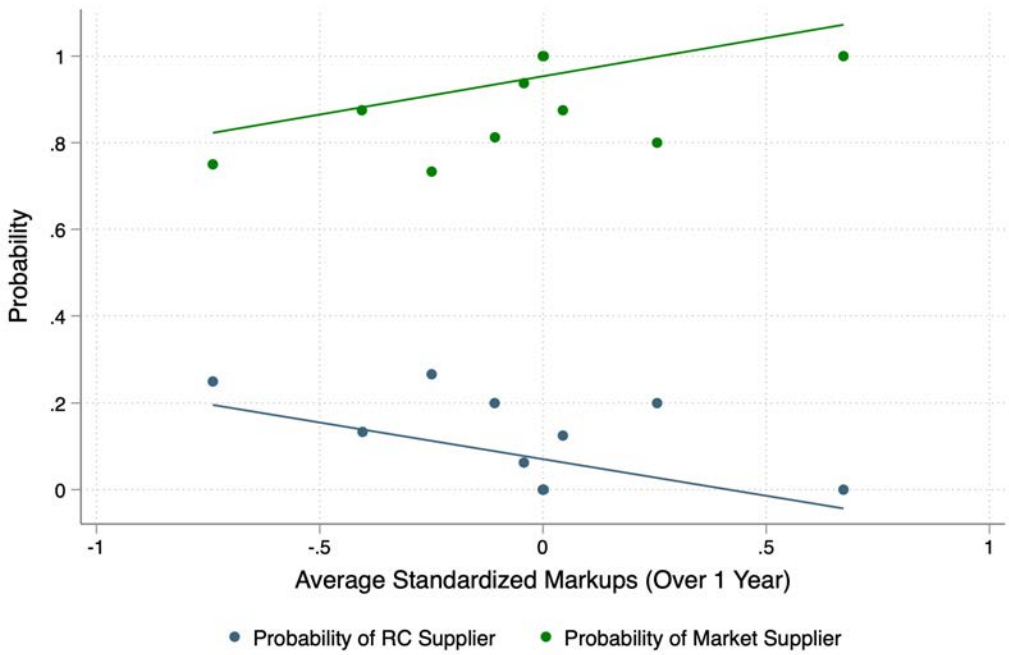

Note: Data from universe of fabric transactions by buyer. Probability estimates from a binscatter calculated using transactions from the year prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.7: Coefficient on Owed Quantity by Supplier Type

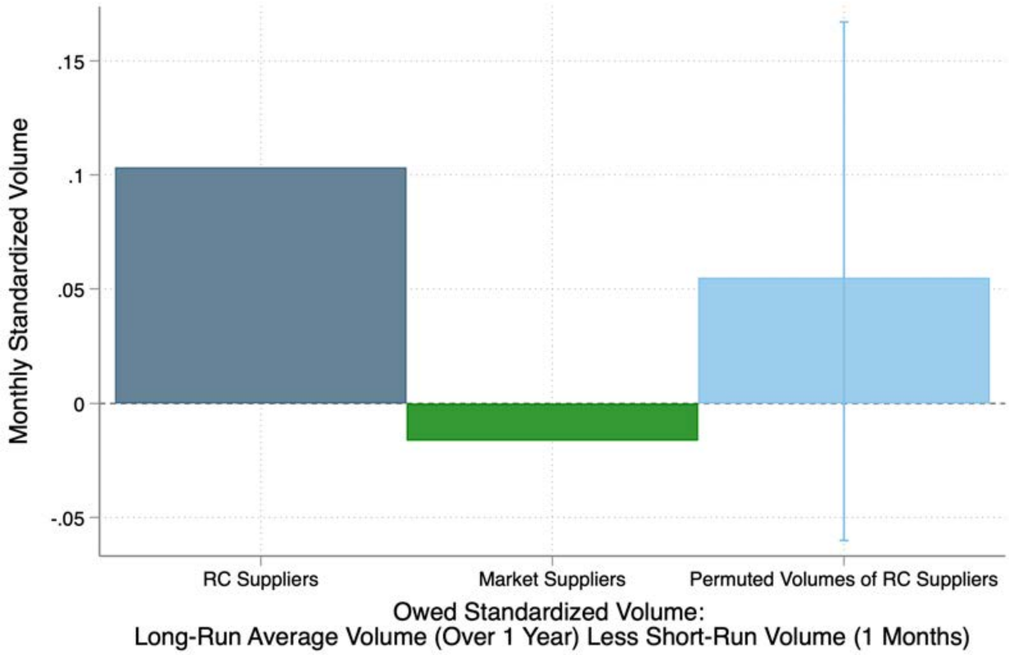

Note: Data from universe of fabric transactions by buyer. Results are the coefficient for owed standardized volume from a regression of current volume on owed volue and month fixed effects. Owed standardized volume is calculated for each supplier quarter as the standardize volume in the prior month less the average of monthly standardize volume in the prior year, not including the prior month. Permuted volumes randomly reshuffle standardized volumes, then calculate owed standardized volume using the permuted data across 500 permutations. The value shown for permuted volumes for relational contract suppliers represents the mean across permutations; the 90% empirical CI is also shown.

Figure A.8: Volume by Supplier Type (Pre-Integration)

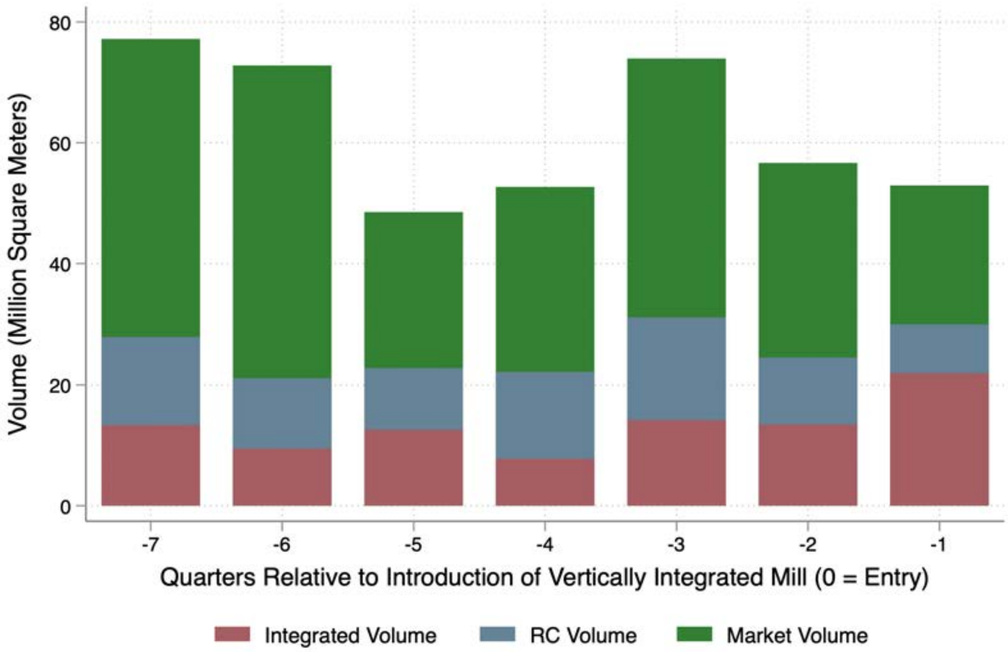

Note: Data from universe of fabric transactions by buyer. Quarterly volume by suppliers considers all transactions in the quarter by supplier type.

A.3 Owed Quantity Guarantees

In addition to showing that relational suppliers have more stable quantities in general, I explore the association between a proxy for the quantity owed in the relational contract and the volume received in the current period. For relational contract suppliers, owed quantity should be positively correlated with current volume: when the relational contract owes more volume to the supplier, the supplier should receive more volume. To test this hypothesis in the data, I regress volume in a month on a proxy for the volume owed in the relational contract with month fixed effects. A positive association between owed volume and current volume would be consistent with the hypothesized relational contract. Long-run monthly average volume serves as a proxy for the quantity guarantee, as the relational contract delivers the guaranteed quantity to suppliers, conditional on the quantity guarantee not changing in an economically meaningful way over the course of the sample. Owed quantity is then measured as the difference between the proxy for the long-run quantity guarantee and the quantity fufilled recently (*i.e.,* in the prior month). Alternatively, as there is no relational contract for market suppliers, there is no reason to expect a relationship between owed volume and current volume for market suppliers.

Figure A.9 illustrates that a one standard deviation increase in owed volume for a relational contract supplier is associated with a .04 standard deviation increase in current volume.[50] Importantly, the same result does not hold for market suppliers, where an increase in owed quantity is associated with a *decrease* in current volume of -.05 standard deviations.

Figure A.9: Coefficient on Owed Quantity and Supplier Type

Note: Data from universe of fabric transactions by buyer. Results are the coefficient for owed standardized volume from a regression of current volume on owed volue and month fixed effects. Owed standardized volume is calculated for each supplier quarter as the standardized volume in the prior month less the average of monthly standardize volume in the three prior quarters, not including the prior month. Permuted volumes randomly reshuffle standardized volumes, then calculate owed standardized volume using the permuted data across 500 permutations. The value shown for permuted volumes for relational contract suppliers represents the mean across permutations; the 90% empirical CI is also shown.

Although this result is consistent with any mean-reverting process, I find that the magnitude of the effect is larger than what would happen if the buyer randomly directed the *same* quantities to the suppliers over time, which should also be mean reverting by construction. Additionally, the most reasonable alternative mean-reverting contracting strategy distinct from the relational contract described would likely direct volumes based on the manufacturer's knowledge from past transactions. Specifically, if the buyer has directed lower volumes than usual to the supplier, then it should have lower capacity utilization and could potentially pass on some cost savings to the buyer. However, this strategy would likely

[50]The pass-through of less than one could reflect that quantity guarantees operate over a longer time frame than one month, measurement error, or market-level demand shifts that are incorporated into the relational contract but not the analysis.

be ineffective given that supplier capacity can change over time without being observed by the buyer due to transactions with other buyers. Furthermore, if this contracting narrative were correct, the point estimate for the shorter long-run version (nine as opposed to twelve months) should likely be higher, as more recent volume is presumably more relevant for current capacity utilization than older orders (see Appendix Figure A.7).

Empirically, the evidence also suggests that the observed association is neither random nor mechanical. Specifically, a placebo exercise which randomly permutes volume across months for the same supplier, such that the set of volumes associated with each supplier is the same across all simulations but the mapping between month and volume is changed, has a smaller mean point estimate across 500 randomizations than the observed association, although confidence intervals are large given the small number of relational contract suppliers and the restriction to only using pre-integration data to avoid incorporating any effects of integration on the relational contract.

A.4 Why Not an Explicit Contract?

There are four factors that presumably lead explicit contracts to fail to resolve the hold-up problem in this setting. First, if the explicit contract is not enforceable because of the institutional environment for contracting, then the explicit contract is effectively irrelevant. Research in India suggests that contracting parties may be concerned that courts will be unable to effectively enforce an explicit contract (Rao (2022); Banerjee and Duflo (2000)). Second, an effective long-term contract would require specifying the price as a function of capacity utilization, as transaction prices shift due to supplier and industry capacity utilization. A contract with constant prices that do not adjust for capacity utilization would result in extreme profit variance, with some orders potentially even unprofitable when supplier capacity usage is sufficiently high; suppliers would be unlikely to agree to any long-term contract that did not condition prices on capacity (at least not without extremely high prices in general). As there are many possible future capacity states due to both the unpredictable stochastic demand from the end client, as well as the unpredictable demand a supplier would face from other buyers, such a contract would likely be difficult (or at least expensive) to write, even abstracting from enforcement concerns, as it would require agreeing on a different price for a large number of capacity states. The just-in-time nature of the procurement process only exacerbates this problem as the complex contingent contract would need to be written in a short time period. Third, even if a contract could specify a price as a function of the supplier's capacity, it may introduce supplier moral hazard that would undermine the effectiveness of the contract. Specifically, as contracted prices increase in capacity, such a contract reduces supplier incentives to reject sampling opportunities (and, thereby orders), from other buyers to save capacity as the supplier would receive similar markups for all capacity states. It follows that suppliers would pass on costs to the buyer. Note that this concern is relevant as staff at the buyer suggest that suppliers do indeed try to target their capacity to match the buyer's needs. Fourth, capacity utilization is likely not contractible, as it is not observable by the buyer.

A.5 Why Not Exclusively Relational Contracts?

Given the evidence that relational contracts offer improved pricing relative to market suppliers, the buyer would minimize costs by sourcing exclusively through relational contract suppliers. However, in practice, although relational contract suppliers provide much more volume than a typical supplier, as they are only 4% of suppliers but provide 25% of volume prior to integration, they do not even provide half of all volume. One plausible explanation for the seemingly low level of overall sourcing through relational contracts reflects that challenges of fulfilling owed quantity guarantees in this setting. Given that the buyer does not know what downstream end client demand will be in any period, it can only credibly commit to the lower bound of estimates of future volume. Additionally, suppliers can typically only produce a subset of fabrics, further decreasing the volume that could credibly be sourced through relational contracts.

Although this hypothesis is not directly testable, a necessary condition is that future volume must be unpredictable; if the buyer could perfectly forecast volume, it should theoretically be able to source exclusively through relational contracts. The data highlight the volatility and unpredictability of volume, especially at the fabric level. Figure A.10 shows the average difference between predicted and actual volume, as a percent of actual volume, weighted by the fabric's share of total volume. These results highlight that errors are, on average, always more than the volume of the typical order. Predictions for monthly volume per fabric are estimated from a model with month fixed effects, fabric fixed effects, 12 lags of volume for all fabrics, and 12 lags of volume for the fabric using data for January 2017 through June 2019. Note that the weighting approach means that large fabrics, which have better prediction accuracy, have heigher weights. Accordingly, estimated median error sizes are even larger. While the buyer may be able to use additional expertise to improve upon the estimates from the model, even decreasing the magnitude of errors by 50% would still result in economically large forecast errors. Consistent with this hypothesis, the volume from relational contract and integrated suppliers appears to be quite stable pre-integration, with market suppliers absorbing stochastic variation in quantities in A.8.

Figure A.10: Forecast Error in Volume Predictions

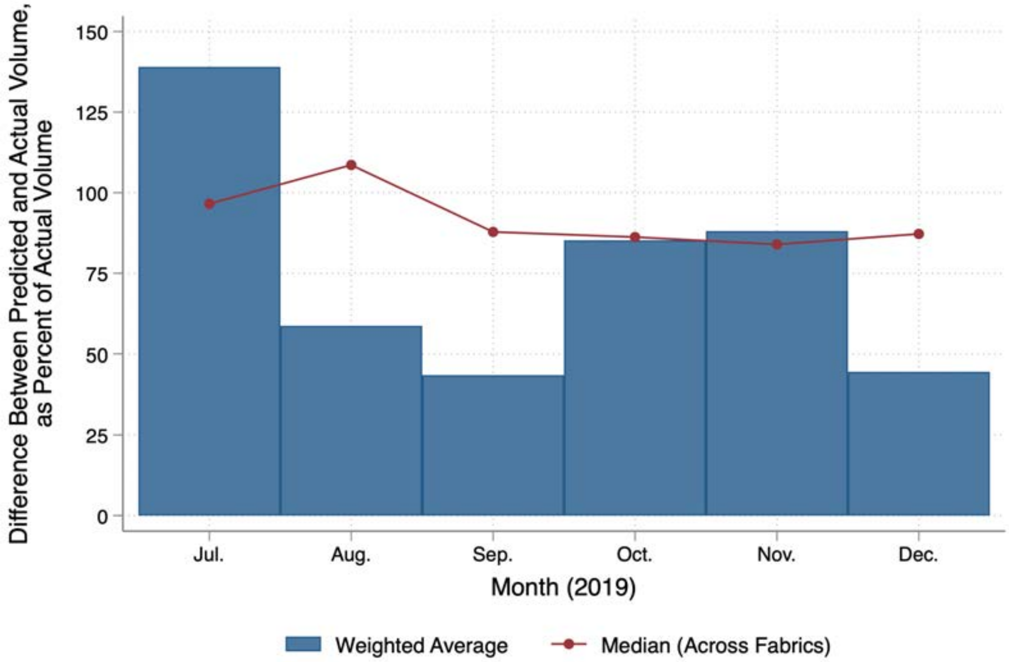

Note: Data from universe of fabric transactions by buyer. Monthly fabric volume forecasts are predictions from a regression of monthly fabric volume on 12 lags of monthly fabric volume, 12 lags of monthly volume across all fabrics, and both fabric and month fixed effects using all but the last six months of data. Errors are the residual fabric volume, expressed as a percent of the monthly fabric volume. Aggregation of errors across fabrics uses the volume-weighted average where volume-weights are the running average percent of volume for the specific fabric over the past year.

A.6 Placebo Analysis: Results

Figure A.11 shows that suppliers that offer discounts are much more likely to be relational contract rather than placebo suppliers, while suppliers that charge high markups are more likely to be placebo suppliers. Additionally, placebo suppliers have more quantity volatitility than relational contract suppliers, and the difference is statistically significant ($p < .01$), as illustrated by Figure A.15. The results for owed volume in Figure A.16 also highlight the difference between relational contract and placebo suppliers, with a one standard deviation increase in owed volume for relational contract suppliers associated with a .04 standard deviation increase in current volume, compared to a -.17 decrease for placebo suppliers. Importantly, given that placebo suppliers and relational contract suppliers have similar monthly average volumes and a similar distribution of owed volumes (see A.17), this difference reflects that current volumes for placebo suppliers are much less associated with owed volume. It follows that it is unlikely that placebo suppliers receive a quantity guarantee.

Figure A.11: Discounts and Supplier Type: Placebo Suppliers

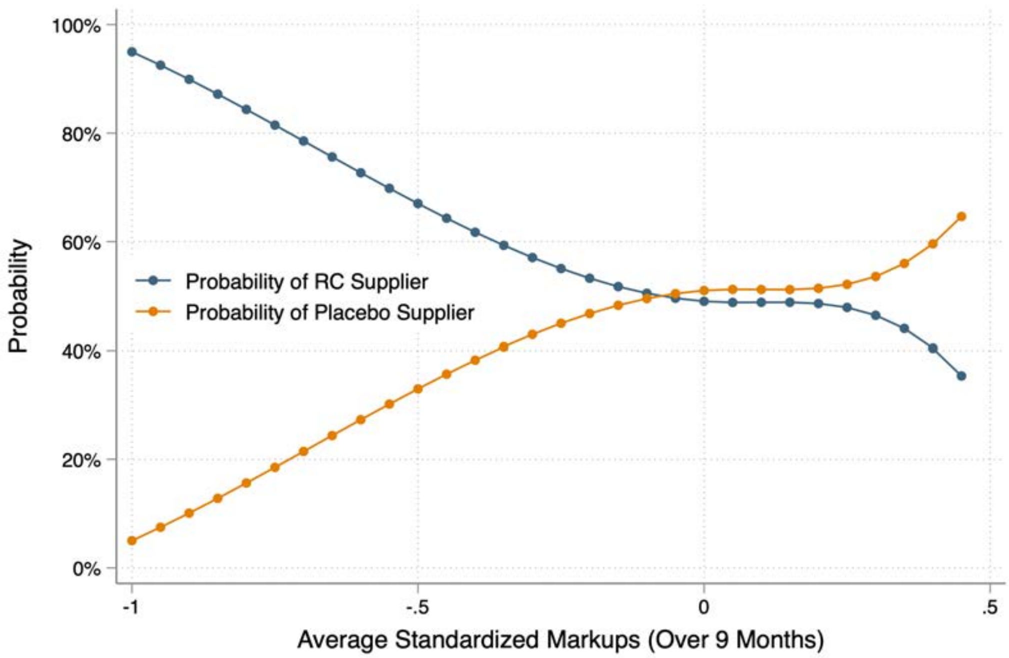

Note: Data from universe of fabric transactions by buyer. Probability estimates represent marginal effects from a non-parameteric series estimator calculated using transactions from the three quarters prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.12: Discounts and Supplier Type: Placebo Suppliers

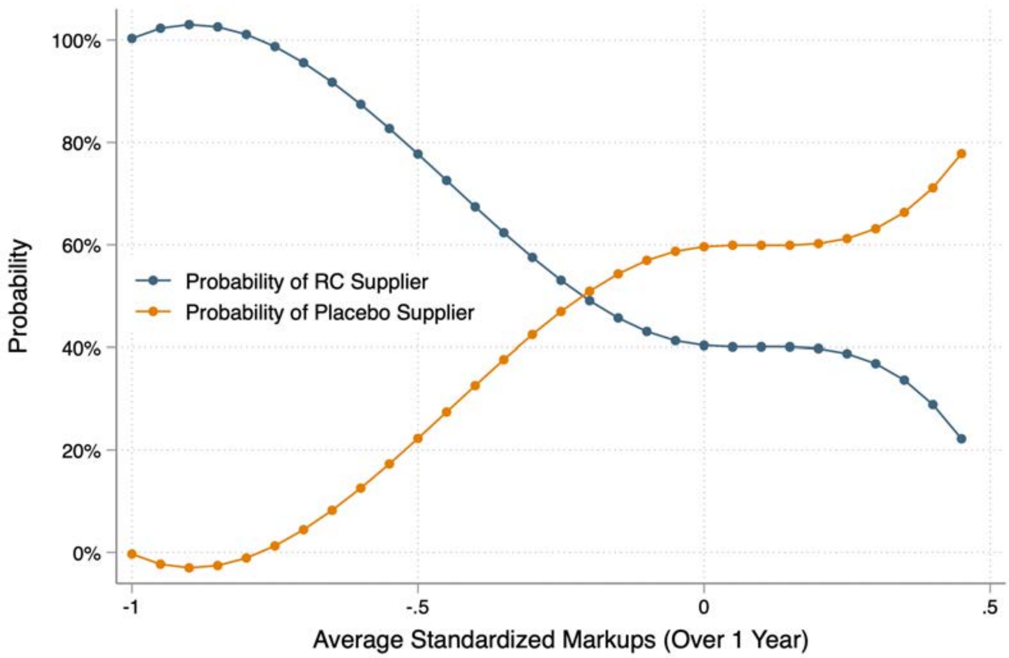

Note: Data from universe of fabric transactions by buyer. Probability estimates represent marginal effects from a non-parameteric series estimator calculated using transactions from the year prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.13: Discounts and Supplier Type: Placebo Suppliers

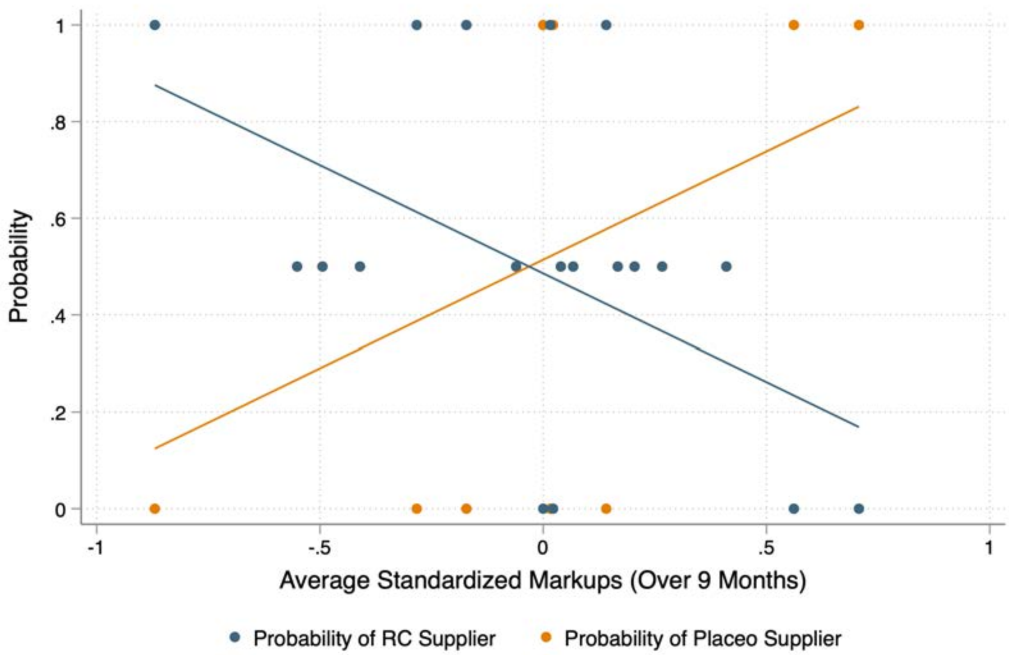

Note: Data from universe of fabric transactions by buyer. Probability estimates from a binscatter calculated using transactions from the three quarters prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.14: Discounts and Supplier Type: Placebo Suppliers

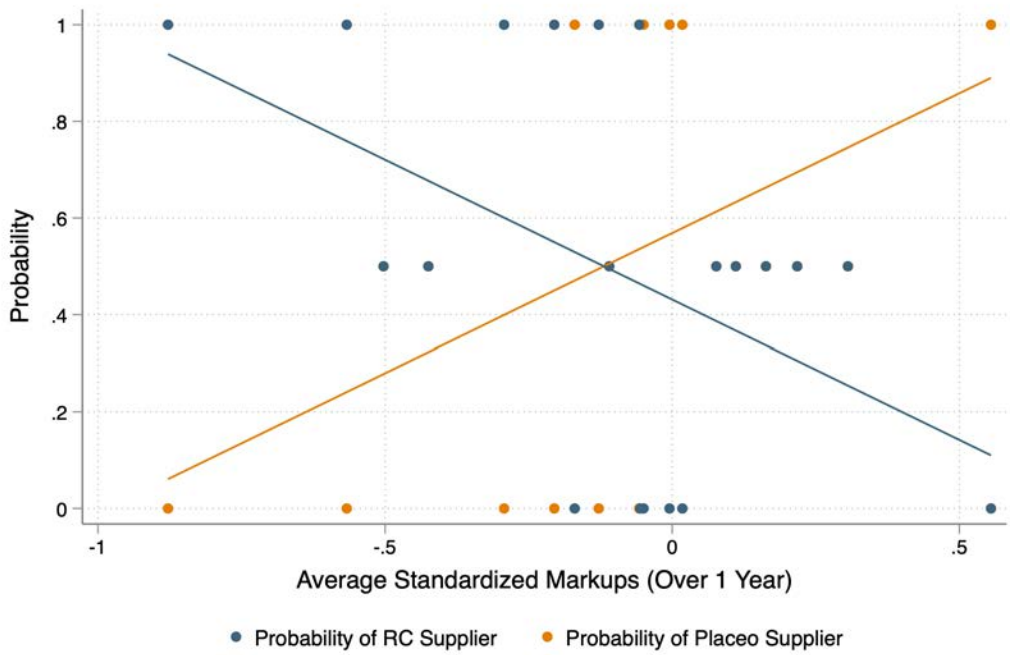

Note: Data from universe of fabric transactions by buyer. Probability estimates from a binscatter calculated using transactions from the year prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.15: Coefficient of Variation by Supplier Type: Placebo Suppliers

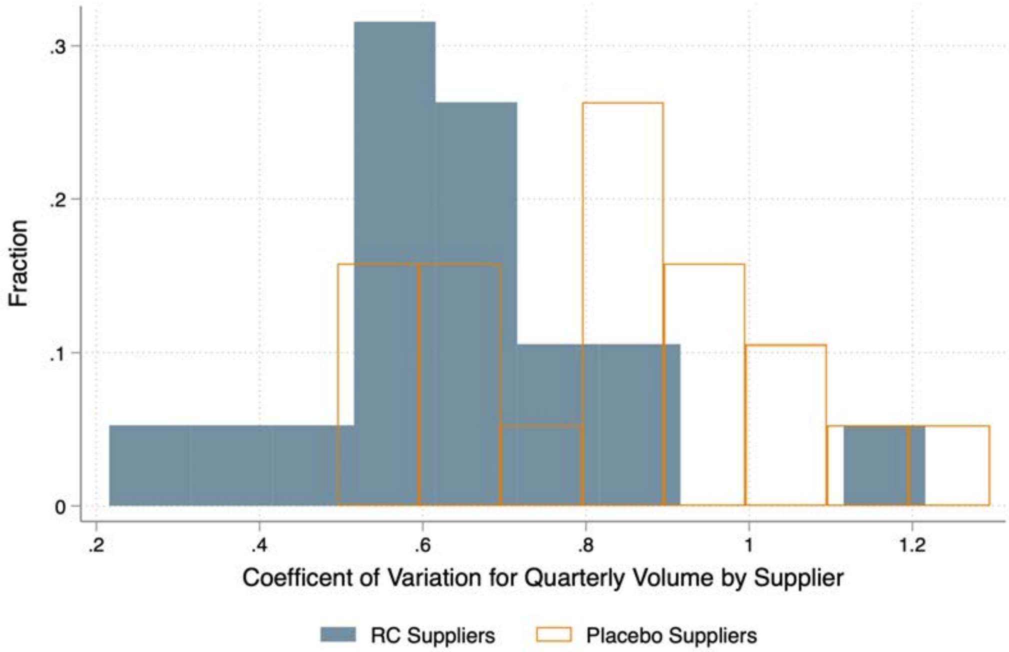

Note: Data from universe of fabric transactions by buyer. Coefficients of variation are calculated using transactions from all quarters prior to integration.

Figure A.16: Coefficient on Owed Quantity by Supplier Type: Placebo Suppliers

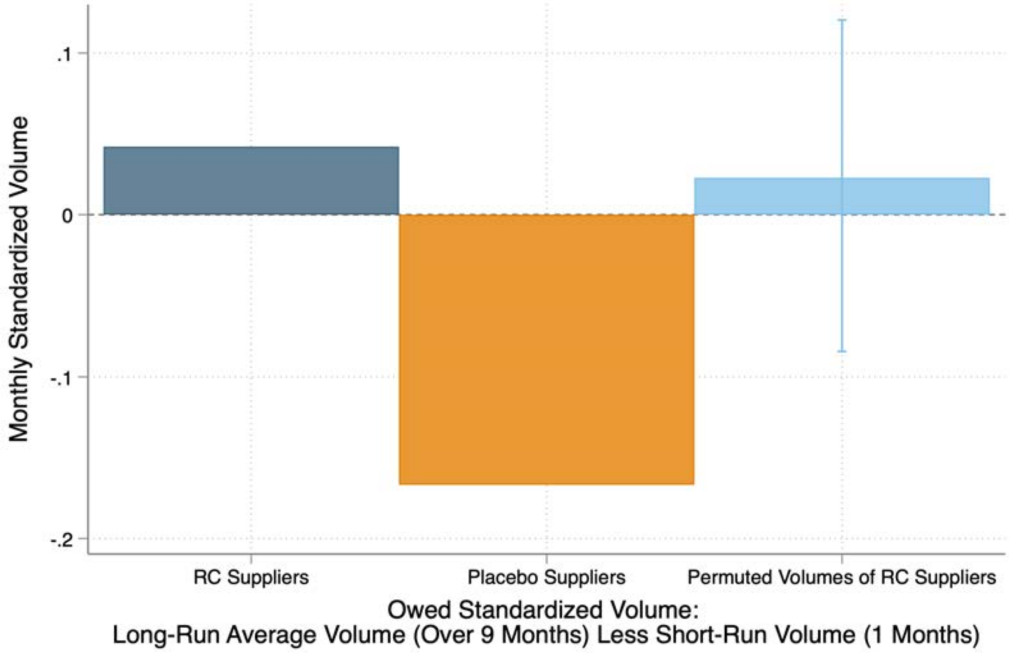

Note: Data from universe of fabric transactions by buyer. Results are the coefficient for owed standardized volume from a regression of current volume on owed volue and month fixed effects. Owed standardized volume is calculated for each supplier quarter as the standardize volume in the prior month less the average of monthly standardized volume in the three prior quarters, not including the prior month. Permuted volumes randomly reshuffle standardized volumes, then calculate owed standardized volume using the permuted data across 500 permutations. The value shown for permuted volumes for relational contract suppliers represents the mean across permutations; the 90% empirical CI is also shown.

Rather than a self-enforcing relational contract, the contracting features observed in the data could be enforced by some other feature in the environment, such as social pressure through kin, ethnic, religious, or other social identity or group based identity (as in Sanchez de la Sierra (2021)). However, pricing in the relational contract seems inconsistent with motives based on such social membership, as the discount received by the larger sophisticated firm (*i.e.,* the buyer) means that the contractual enforcement by social group membership also increases in-group inequality, which seems unlikely as such groups presumably have preferences to redistribute to the less well-off members of the group. To be more concrete, it seems unlikely that kin ties effectively pressure the buyer to work with a supplier, but those same kin ties also require the supplier to offer the buyer a discount. Note that it need not be the case that social group membership plays no role–group membership may facilitate the creation of self-enforcing relational contracts or influence which suppliers are selected as relational contract suppliers. Importantly, in these scenarios, the relational contract still operates as a relational contract with its associated impacts on transactions.

Figure A.17: Distribution of Owed Quantity by Supplier Type

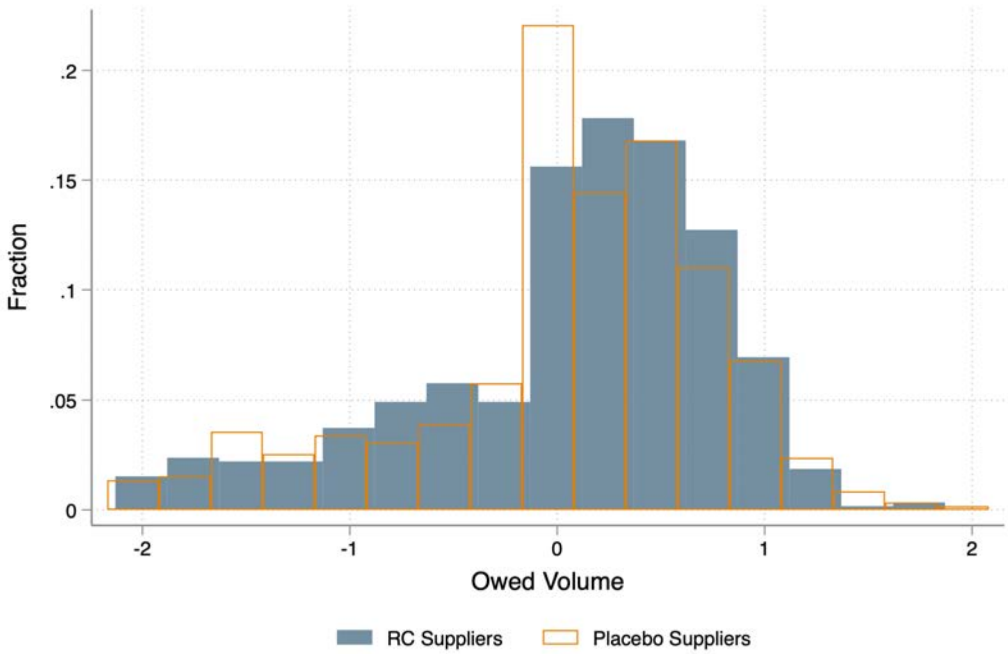

Note: Data from universe of fabric transactions by buyer. Coefficients of variation are calculated using transactions from all quarters prior to integration.

A.7 Robustness Checks for Conceptual Model

A.7.1 Relational Contracting with Dynamics

In a dynamic model that more closely resembles typical models of relational contracts, including a DICC as in Macchiavello and Morjaria (2023), the deviation payoffs change. Note that the on-path payoffs do not change, so I only discuss how the deviation payoffs change in this model. This reflects the interpretation of the relational contract as specifying a pricing rule that is agreed upon before quantities are realized. In the dynamic model, the supplier's outside option is to hold up the buyer making it pay the high market price. Therefore, I search over deviations and find the best static deviation, which occurs when $p^M q^{RC} - p^{RC} q^{RC}$ is the largest:

$$O_S = (\delta^S)\bigg(\mathbb{E}_{q^M}[U(p^M q^M - C(q^M))]\bigg) + (1-\delta^S)\bigg(U(p^M q^{RC} - C(q^{RC}))\bigg)$$

Similarly, the buyer can break the relational contract. The buyer could have an incentive to do so when the relational contract price is higher than the market price, would can occur because the relational contract is designed to smooth profits. For high cost or low quantity states, relational contract prices are high to facilitate supplier profit smoothing. Therefore, I search over static deviations, with the buyer's best deviation occuring when $p^{RC} q^{RC} - p^M q^{RC}$

is the largest. In this case, the buyer's outside option is:

$$O_B = (\delta^B)\Big(- p^M \mathbb{E}_{q^{RC}}[q^{RC}]\Big) + (1-\delta^B)\Big(- p^M q^M\Big) \text{ if not integrated}$$

$$O_B = (\delta^B)\Big(- \mathbb{E}_{q^{RC}}[C(q^{RC})]\Big) + (1-\delta^B)\Big(- C(q^{RC})\Big) \text{ if integrated}$$

Note that for δ^S, δ^B sufficiently large, the dynamic version approaches the long-run version considered in the main text.

The analysis below shows that results do not change meaningfully when incorporating dynamics.[51]

Figure A.18: Relational Contract Prices and the Buyer Bargaining Parameter with Dynamics

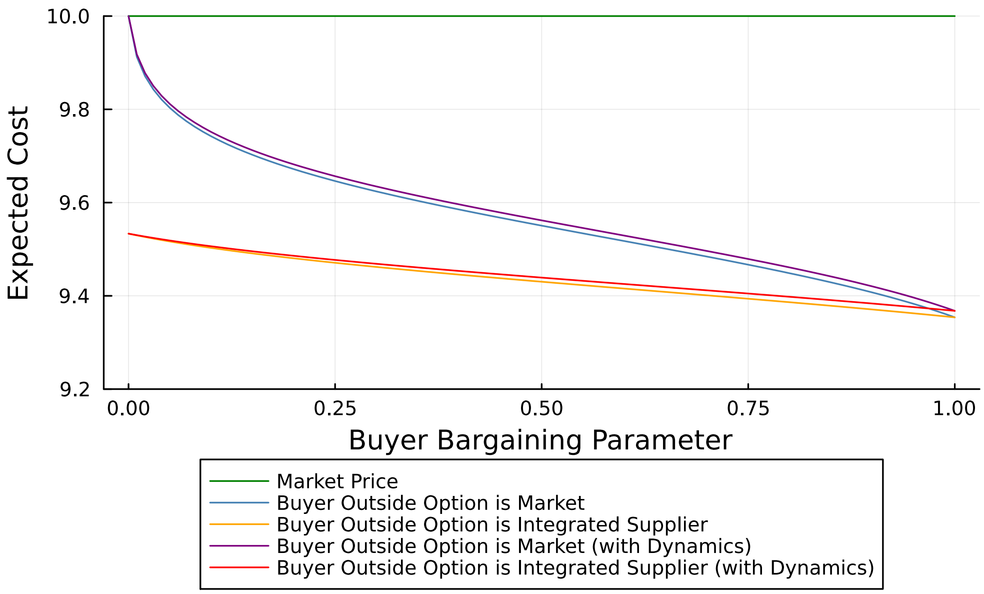

Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.7.3.

[51] For this analysis, I set $\delta^S = \delta^B = .9$. As analysis of this model is at the quarter level in general, this implies a yearly discount factor of .66. Given that as $\delta \to 1$ this analysis collapses to the static case, selecting a low yearly discount factor should be conservative.

Figure A.19: Relational Contract Prices and Supplier Risk Aversion with Dynamics

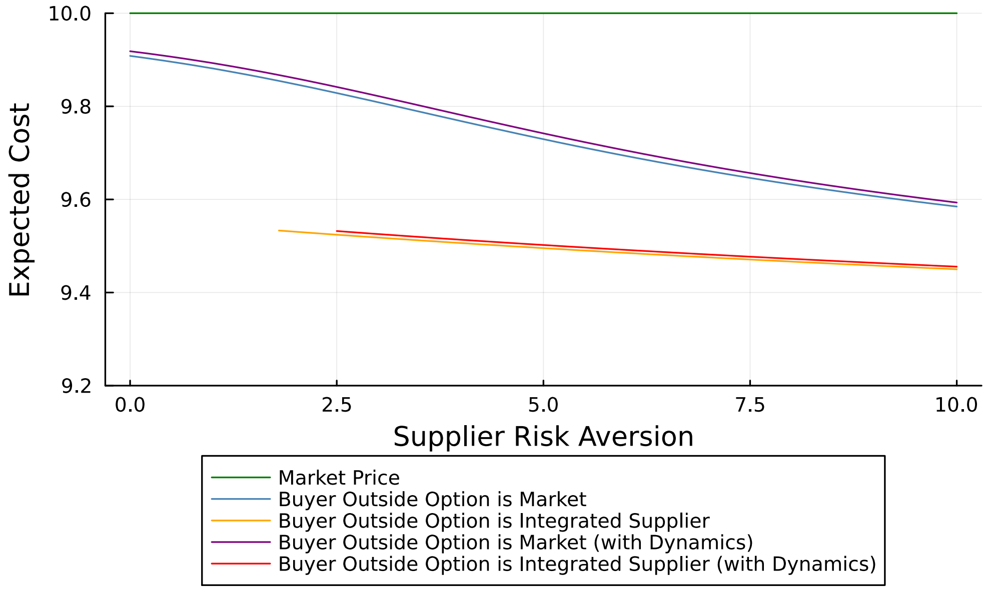

Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.7.3.

A.7.2 Cost Heterogeneity Between External and Internal Supplier

I show that results are robust to the external supplier having lower or higher costs than the integrated supplier, shifting costs by 10%. To ensure that optimal quantity is still $q = 1$ without changing market prices, I only shift the fixed part of costs by 10%, keeping convexity the same. Results (available upon request) shifting all costs are similar. Larger heterogeneity in production costs between the buyer and supplier seem unlikely. If buyer production costs are much lower than supplier production costs, the buyer should fully integrate. On the other hand, if supplier production costs are much lower than buyer production costs, the buyer should never integrate and the integration should not even credibly shift the threat point (assuming the supplier can see the integrated mill and make a reasonable inference about is productivity).

Figure A.20: Relational Contract Prices, the Buyer Bargaining Parameter, and Supplier Risk Aversion when External Supplier has Lower Costs

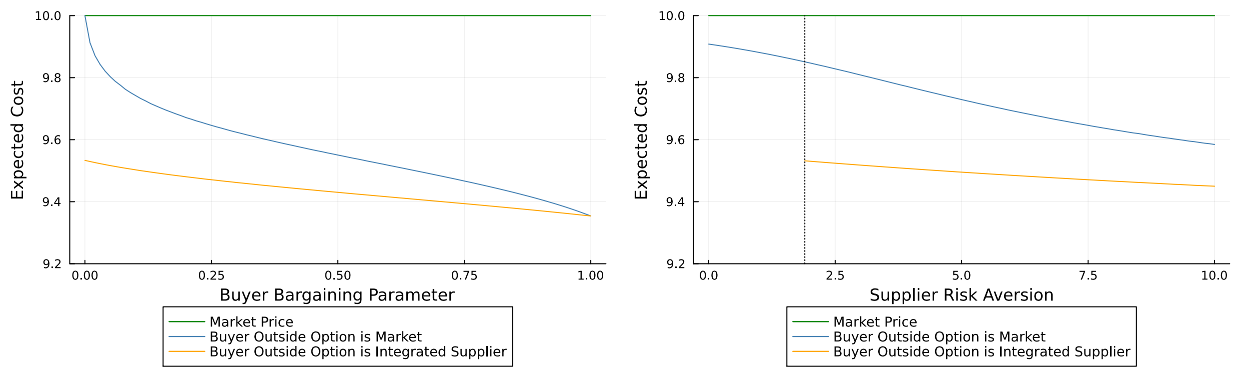

Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.7.3.

Figure A.21: Relational Contract Prices, the Buyer Bargaining Parameter, and Supplier Risk Aversion when External Supplier has Higher Costs

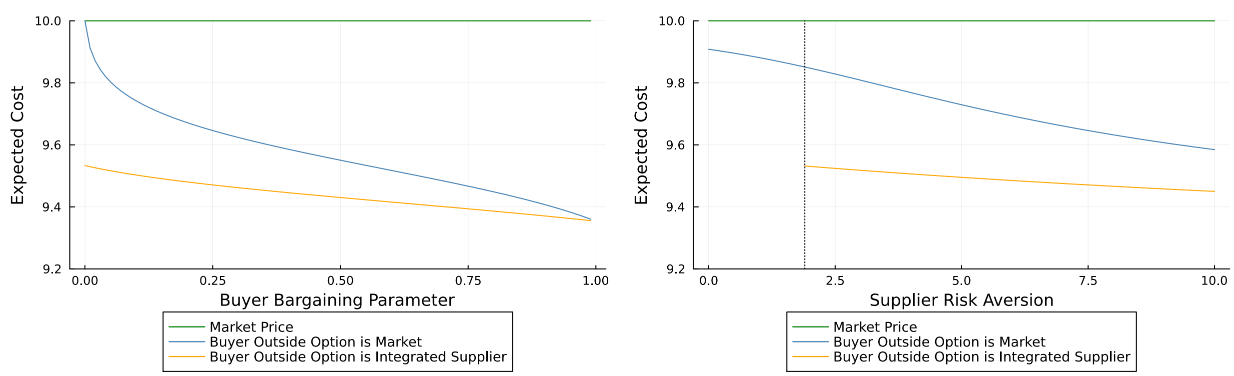

Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.7.3.

A.7.3 Parameterization and Solution of Conceptual Model

The model is parameterized as:

$$
\begin{aligned}
u(c) &= 1 - exp(-\theta * c) \; (CARA \; utility) \\
c(q) &= 1.35 - .5q + .5q^2 \\
\text{RC quantities} : q^{RC} &\in \{.9, 1, 1.1\}, f(q^{RC}) = \frac{1}{3} \\
\text{Market quantities} : q^{M} &\in \{.6, .8, 1.0, 1.2, 1.4\}, f(q^{M}) = \frac{1}{5} \\
p^M &= 10 \\
\theta = 7.5; \alpha_B &= .25
\end{aligned}
$$

I solve the model by finding the relational contract price schedule, specifically the price for each quantity in the relational contract, that solves the Nash bargaining problem.

A.7.4 Log Quantities

Table A.1: Log Quantity and Capacity

	Log Days Between Order and Delivery			
	(1)	(2)	(3)	(4)
Log Quantity	1.8016***	2.4814***	2.9646***	2.9409***
	(.27)	(.2813)	(.6061)	(.5144)
Supplier FE	N	Y	Y	Y
Fabric Group FE	N	N	Y	Y
Month FE	N	N	N	Y
N	16506	16443	14499	14499
R^2	.0008	.0263	.0941	.0979

Note: Data from universe of fabric transactions by buyer. Coefficients reported are from a regression of log days between fabric order date and delivery date on log quantity.

Table A.2: Log Quantity and Prices

	Price (Standardized)			
	(1)	(2)	(3)	(4)
Log Quantity	−.07***	−.1203***	−.2322***	−.2335***
	(.0065)	(.009)	(.0145)	(.0145)
Supplier FE	N	Y	Y	Y
Fabric Group FE	N	N	Y	Y
Month FE	N	N	N	Y
N	5456	5438	5396	5396
R^2	.0206	.1081	.2057	.2253

Note: Data from universe of fabric transactions by buyer. Coefficients reported are from a regression of log price on log quantity.

A.8 Structural Model of Relational Contract: Evidence Supporting Modelling Assumptions

A.8.1 Quantity Exogeneity

Although capacity utilization for any supplier is certainly not random, as the buyer provides demand assurance in the relational contract, supplier-period deviations from the targeted capacity in the relational contract are plausibly exogenous. Specifically, these deviations reflect stochastic end client demand, as the buyer knows neither which sample garments will be selected nor what quantities will be conditional on an end client purchasing a sample. It follows that the buyer strategically behaving to reduce capacity variance within a period does not invalidate capacity deviation exogeneity, but, rather, creates the data patterns that explain why capacity variance is lower for relational contract suppliers as compared to other suppliers.

However, buyer behavior that reduces *dynamic* capacity variation would bias estimates, as I would underestimate how much demand assurance the relational contract provides. It follows that the estimation approach would overestimate at least one of risk aversion and the supplier bargaining parameter (which is one less the buyer bargaining parameter).[52]

To mitigate these concerns, I directly test the most likely examples of such behavior in Table A.3. Specifically, if the buyer provides additional dynamic capacity variance reduction, it would presumably take the form of ensuring that extreme capacity shocks are unlikely to occur repeatedly. I directly test this hypothesis by examining the autocorrelation between extreme capacity realizations, which I define as absolute capacity deviations at 90th percentile of deviations or higher among relational contract suppliers. While I find that prior extreme capacity realizations are associated with lower probability of an extreme capacity realization, the magnitude is small and not statistically significant. Furthermore, the data suggest that the buyer does not even try to ensure that positive capacity shocks are followed by negative capacity shocks (and vice-versa).

Table A.3: Capacity Patterns

	Extreme Capacity (90th Percentile)		Capacity	
	(1)	(2)	(3)	(4)
Lagged Extreme Capacity	$-.043$	$-.056$		
	(.059)	(.048)		
Lagged Capacity			.024	.086
			(.059)	(.109)
Cluster Bootstrap 95% CI	$[-.132, .064]$	$[-.117, .114]$	$[-.163, .173]$	$[-.202, .173]$
Pre-Integration Only	N	Y	N	Y
R^2	.172	.082	.297	.171
N	216	108	216	108

[52] The model would overestimate risk aversion if incorporating additional demand assurance in the relational contract shrinks the variance in the distribution of relational contract capacities. Alternatively, it would overestimate the supplier bargaining parameter if incorporating additional demand assurance in the relational contract increases the variance in the distribution of market capacities.

Note: Data from universe of fabric transactions by buyer. Coefficients reported are from a regression using quarter by relational contract supplier panel data of extreme capacity indicator on lagged extreme capacity in columns 1-2 and capacity on lagged capacity in columns3-4. Robust standard errors are in parentheses and the 95% empirical confidence interval bootstrapped by supplier are shown.

Additionally, if parameter estimates are biased, then model prices should be biased. In this case, the untargeted moment validation exercise in 4.5.1 should not find that model prices are unbiased.

Another possible concern is that relational contract target quantities are adjusted after the vertical integration; however, this possibility seems unlikely as it is not in the buyer or supplier's interest to decrease the amount sourced from relational contracts given the large volume of fabric still procured from market suppliers that the buyer can displace with integrated supply. Regardless, if the target did decrease, then the estimation approach would systemtically overrestimate capacity shocks in the post-period, which would tend to increase the tuning parameter estimate. And, the increased tuning parameter estimate would lead to underestimates of supplier risk aversion. Although this concern seems unlikely given the small tuning parameter estimate (.002), I nevertheless test to see if relational contract average capacity decreases in the post period in Table A.4 by regressing capacity on an indicator for the post-period and total quantity of fabric procured across all other suppliers. I do not find any evidence that capacities are significantly different post-integration.

Table A.4: Capacity Stability

	Capacity
Post-Period Indicator	$-.0099$
	(.0141)
	[.0114]
Leave-Out Volume	.0003***
	(.0001)
	[.0001]
Supplier FE	Y
R^2	.338
N	234

Note: Data from universe of fabric transactions by buyer. Coefficients reported are from a regression using quarter by relational contract supplier panel data of capacity on a post-period indicator and total volume at the buyer from all but the supplier. Robust standard errors are in parentheses and standard errors clustered by supplier are in brackets.

Additionally, if parameter estimates are biased, then model prices should be biased. In this case, the untargeted moment validation exercise in 4.5.1 should not find that model prices are unbiased. Furthermore, out-of-sample fit should be poor if the model is systematically underestimating capacity utilization in the post-period (unless it is perfectly offset by creating bias in risk aversion of bargaining parameters, in which case the untargeted moment validation should fail). However, I find that strong out-of-sample fit in 4.5.2.

A.9 Structural Model of Relational Contract: Comparison of Risk Aversion Estimates

Table A.5: Comparison with Other Risk Aversion Estimates

Source	CARA Coefficient	Risky Activity
Study Fabric Suppliers	*Median:* .0085	*Profit from fabric sales*
Cohen & Einav (2007)	Median: .000034	Small Stakes: Deductible for one year
Handel (2015)	Median : .000422	Small Stakes OOP for one year
Blouin & Macchiavello (2019)	Implied from $u(c) = c^{1-\alpha}$ at mean c: .0068	Medium Stakes: Coffee Contract Revenue
Barsky *et al.* (1997)	Ranges from .3 to 1.04	Large Stakes: Annual Income
Gandelman & Hernández-Murillo (2014)	Implied from CRRA at median c for this study: 12.90	Large Stakes: Investment Choices

A.10 Confidence Interval for Buyer Bargaining Parameter and Supplier Risk Aversion

Table A.6 shows 95% confidence intervals for the structural parameters from 100 bootstrap replications. The results indicate that supplier risk aversion estimates tend to be reasonably precisely estimated–even using the upper bound of every confidence interval would not change the interpretation that in-sample suppliers have similar risk aversion as the coffee mills in Blouin and Macchiavello (2019). However, buyer bargaining parameters are not precisely estimated. This result reflects that when surplus is small, as is the case when risk aversion is low, changing the buyer bargaining parameter has small effects on prices. Note that the buyer bargaining parameter is precisely estimated for the one outlier supplier that is quite risk averse such that surplus is sufficiently large for changes in the buyer bargaining parameter to have large effects on prices.

Table A.6: Buyer Bargaining Parameter and Supplier Risk Aversion: 95% CI

Buyer Bargaining Parameter		*Supplier Risk Aversion*	
Estimate	95% CI	Estimate	95% CI
0.8885	[0.0091, 0.9713]	0.0084	[0.0014, 0.0558]
0.9674	[0.0277, 0.9897]	0.0442	[0.0028, 0.0804]
0.5822	[0.0262, 0.8975]	0.0000	[0.0000, 0.0283]
0.0503	[0.0335, 0.9656]	0.0102	[0.0010, 0.0392]
0.0432	[0.0093, 0.4332]	0.0058	[0.0039, 0.0176]
0.0131	[0.0003, 0.0503]	0.0223	[0.0000, 0.0545]
0.0289	[0.0050, 0.9538]	0.0232	[0.0001, 0.0529]
0.6607	[0.0089, 0.9868]	0.0046	[0.0001, 0.0285]
0.9155	[0.0031, 0.9728]	0.0080	[0.0000, 0.0380]
0.0301	[0.0004, 0.9585]	0.0086	[0.0000, 0.0314]
0.0094	[0.0029, 0.9062]	0.0121	[0.0001, 0.0275]
0.9454	[0.3305, 0.9913]	0.0174	[0.0001, 0.1555]
0.8667	[0.0471, 0.9896]	0.0012	[0.0001, 0.0299]
0.9552	[0.0105, 0.9689]	0.0551	[0.0004, 0.1330]
0.3034	[0.0028, 0.9682]	0.0024	[0.0000, 0.0642]
0.9140	[0.0303, 0.9426]	0.0002	[0.0000, 0.0522]
0.0301	[0.0022, 0.9754]	0.0026	[0.0002, 0.0692]
0.9985	[0.9490, 1.0000]	14.6811	[0.1239, 31.2573]

Note: Model estimates.

A.11 Model Validation: Out-Of-Sample Fit RMSE Density

Figure A.22: Distribution of Supplier Squared Residuals by Estimation Method

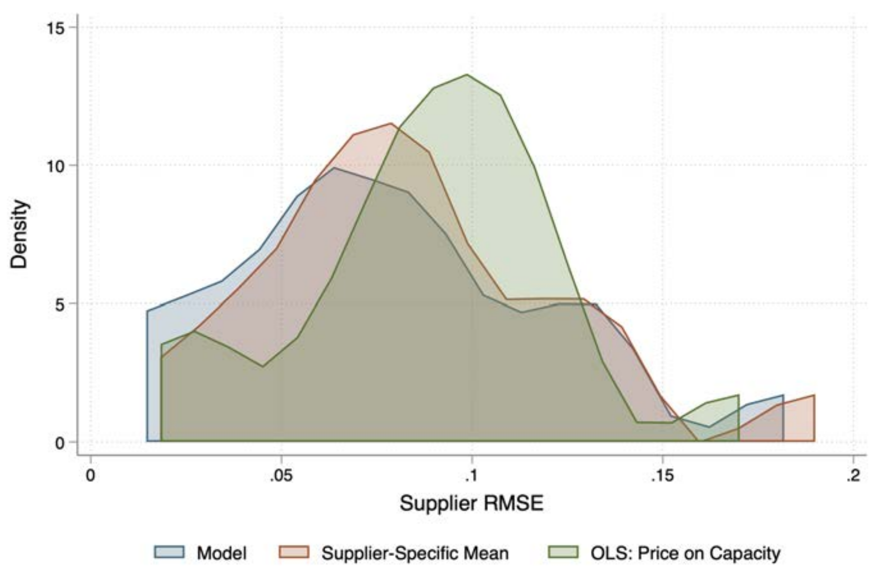

Note: Density of distribution of supplier root-mean squared error in the post- integration error by estimator.

A.12 Model Validation: Out-Of-Sample Fit Comparison with OLS with Tuning Parameter

Table A.7: Out-of-Sample Fit Comparison (OLS with Tuning)

	Structural Model	Supplier Mean	OLS with Tuning
Minimum	.0146	.0182	.0182
25th Percentile	.0482	.0551	.0681
Median	.0741	.0790	.0812
75th Percentile	.1003	.1004	.1014
Max	.1818	.1900	.1780
Mean	.0787	.0844	.0843
Standard Deviation	.0441	.0417	.0379

Note: Comparison of distribution of supplier root-mean squared error in the post- integration error by estimator.

Figure A.23: Distribution of Supplier RMSE by Estimation Method

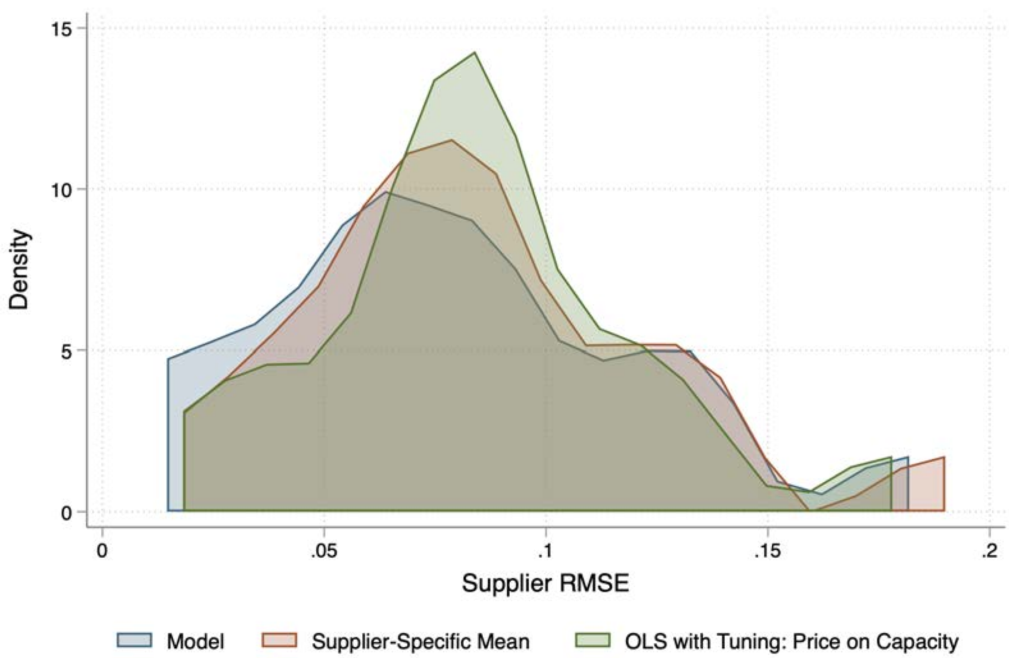

Note: Density of distribution of supplier root-mean squared error in the post- integration error by estimator.

A.13 Validation using Difference-in-Differences Estimates

A.13.1 Model-Implied Threat Point Effect for In-Sample Suppliers

Figure A.24: Threat Point Effect for In-Sample Suppliers

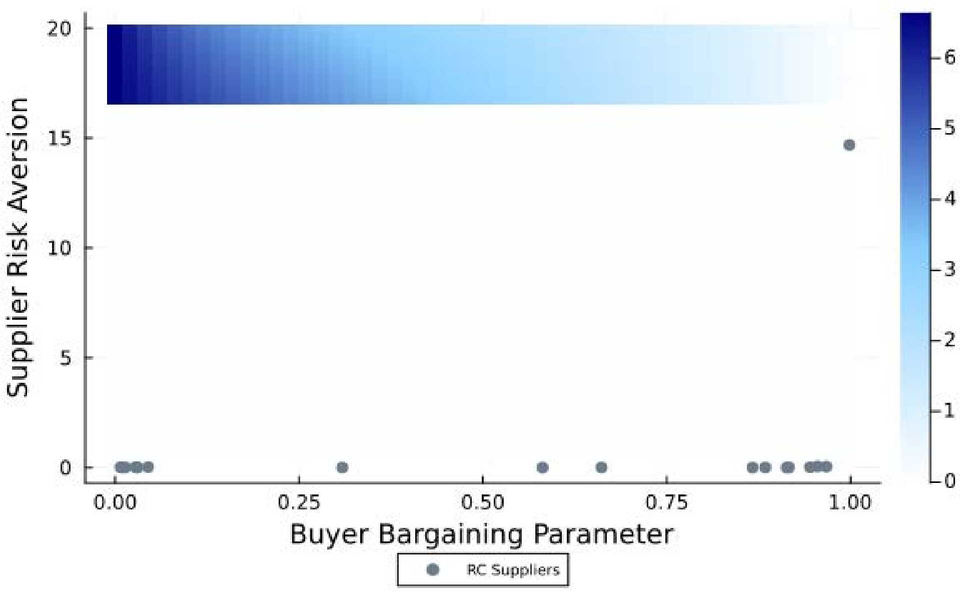

Note: Model estimates.

A.13.2 Doubly Robust Difference-in-Differences Design and Identification

In addition to bringing less production in-house for fabrics with more relational contract suppliers, Figure A.25 illustrates that fabrics with positive trends in demand were also more likely to be brought in-house. Therefore, as suppliers produce different fabrics, some suppliers are more or less likely to face exposure to vertical integration due to industry wide trends in fabric demand rather than supplier performance. To fix ideas using an extreme example, there could be a fantastic supplier which stopped receiving orders after integration due to the fabrics that they produce losing popularity. As trends in fabric demand are correlated with both a supplier's contract quantities and its exposure to vertical integration, counterfactual outcomes of treated and control suppliers may not evolve in a parallel fashion, violating the parallel trends assumption. Additionally, the importance of trends around construction and initiation time for long-run supply decisions, such as building an integrated mill, is consistent with evidence from other settings, such as initial electricity prices having long-run effects on electricity consumption of manufacturing plants (Hawkins-Pierot and Wagner, 2022).

Figure A.25: Trends in Fabric Demand and Vertical Integration

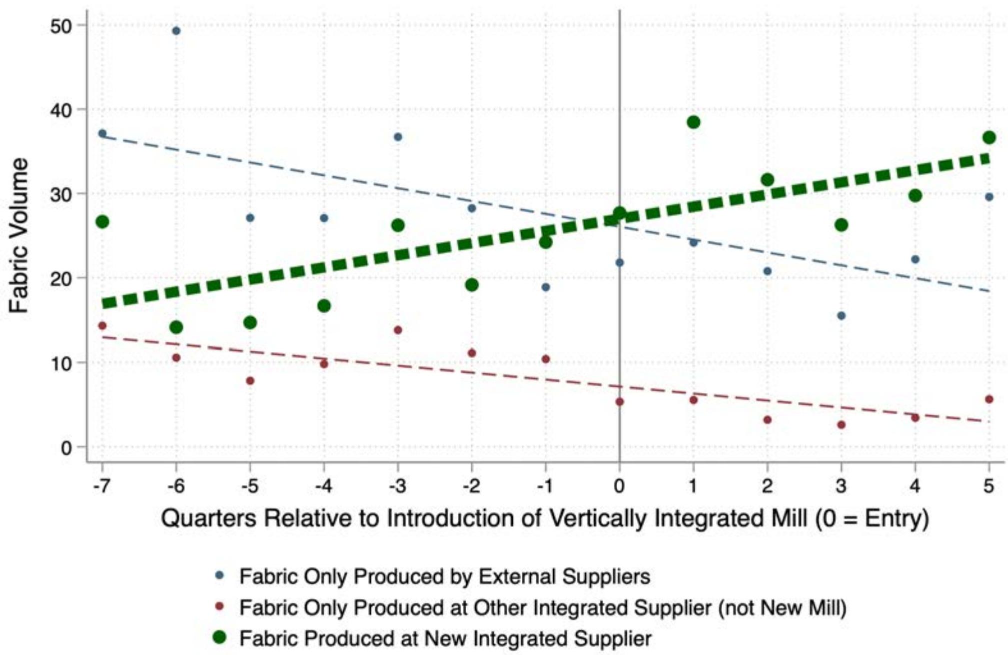

Note: Data from universe of fabric transactions by buyer. Observations represent the volume of fabric purchased (in millions) across fabric groups, where fabrics are grouped based on whether the fabric is brought in-house or not. Trend lines are from linear trends.

The issue of trends in fabric demand motivates a doubly-robust difference-in-differences approach as in Sant'Anna and Zhao (2020), which weakens the parallel trends assumption to *conditional* parallel trends.[53] In this context, the *conditional* parallel trends assumption is that the change in potential outcomes for suppliers exposed to vertical integration is the same as the change in potential outcomes for suppliers unexposed to vertical integration *after* conditioning on trends in fabric demand that influence supplier exposure to vertical integration.[54] Formally, this assumption can be expressed as:

$$\mathbb{E}[Y_i^{post}(0) - Y_i^{pre}(0)|D_i = 1, \boldsymbol{X_i}] = \mathbb{E}[Y_i^{post}(0) - Y_i^{pre}(0)|D_i = 0, \boldsymbol{X_i}] \tag{2}$$

[53]An alternative option would be to use the method in Freyaldenhoven et al. (2019). Unfortunately, this approach would not work well in this setting, as the covariate adjustment is based on long-running trends rather than short-term contemporaneous movements. In other words, the covariate adjustment is mostly at the supplier, rather than supplier by time, level, which is not ideal variation for this method.

[54]Specifically, linear trend in fabric demand for each fabric from a regression of fabric volume across all suppliers on the time period for all quarters in the pre-period. Supplier exposure to trends is calculated as the volume weighted average of the trends for the fabrics the supplier produces. It follows that the supplier's own volume, which is an outcome variable, should not drive results as trends are based on volumes for all suppliers, not just the individual supplier.

In equation 2, trends in fabric demand are included in $\boldsymbol{X_i}$,[55] and $Y_i^{post}(0)$ denotes the untreated potential outcome in the post period with $D_i = 1$ indicating treatment units and $D_i = 0$ indicating control units.

With the doubly-robust approach, the estimated average treatment effect on the treated is unbiased if at least one of the propensity score model or the regression outcome model is correct. As there is only one treatment event, it is not necessary to adjust for heterogeneity in treatment timing, as in Callaway and Sant'Anna (2021); Sun and Abraham (2021). Importantly, unlike a synthetic control methodology to account for trends, this approach does not use any supplier-level outcome variables when constructing the propensity scores or the regression outcome model. It follows that pre-trends retain their diagnostic value, as it is not mechanical that treatment effects are zero in pre-treatment periods. Furthermore, Table A.8 provides evidence that adjusting for propensity scores alone achieves balance between treatment and control suppliers by showing that treatment is no longer correlated with supplier-level outcome variables after adjusting for the propensity scores using inverse probability weighting. Because the propensity scores use fabric, and not supplier, trends,[55] it is not mechanical that this propensity score approach would achieve balance. I also document in Figure A.26 that there is overlap in propensity scores between treatment and control suppliers and that no units have extremely high propensity scores, satisfying the "strong overlap" condition.[56]

Table A.8: Balance on Untargeted Outcomes

	Coefficient on Treatment Indicator:	
Outcome Variable	Without IPW Adjustment	With IPW Adjustment
Pre-Integration Volume	3033768	1085459
	$[1189072, 5489585]$	$[-1004042, 3383491]$
Pre-Integration Transaction Count	72.93	-3.35
	$[44.49, 106.45]$	$[-46.3, 15.1]$
Pre-Integration Fabric Count	26.71	$-.98$
	$[17.08, 37.84]$	$[-9.02, 4.48]$

Note: Data from universe of fabric transactions by buyer. Bootstrapped 90% confidence intervals in brackets.

[55] The full set of variables in $\boldsymbol{X_i}$ are a third order polynomial of the fabric growth per supplier measure described earlier, the count of distinct fabrics seven quarters prior to integration that the supplier supplies, an indicator for being a relational contract supplier given that fabrics with relational contracts are less likely to be integrated, and the interaction of the relational contract indicator and the linear term for fabric growth and the count of distinct fabrics.

[56] Formally, as stated in Roth et al. (2022), "the conditional probability of belonging to the treatment group, given observed characteristics, is uniformly bounded away from one, and the proportion of treated units is bounded away from zero. That is, for some $\epsilon > 0, \mathbb{P}[D_i = 1|X_i] < 1 - \epsilon$, almost surely and $\mathbb{E}[D_i] > 0$.

Figure A.26: Distribution of Propensity Scores by Supplier Type

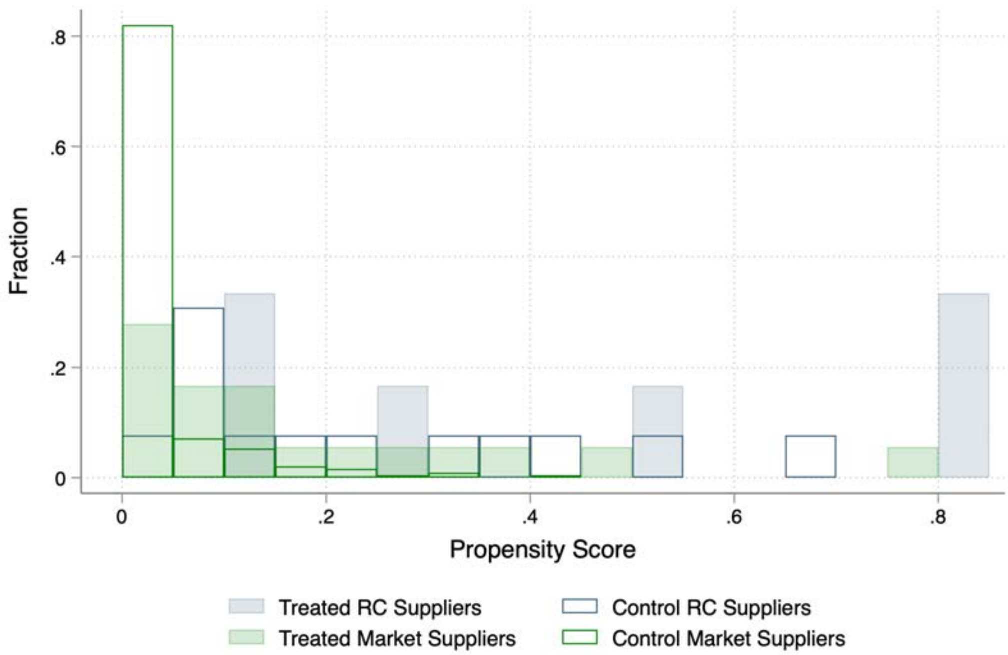

Note: Model estimates using pre-integration data.

A.14 Difference-in-Differences Supplemental Analysis and Robustness Checks

A.14.1 Difference-in-Differences Additional Event Study Plots

Figure A.27: Difference-in-Difference Estimates for Effect of Vertical Integration on Reliability

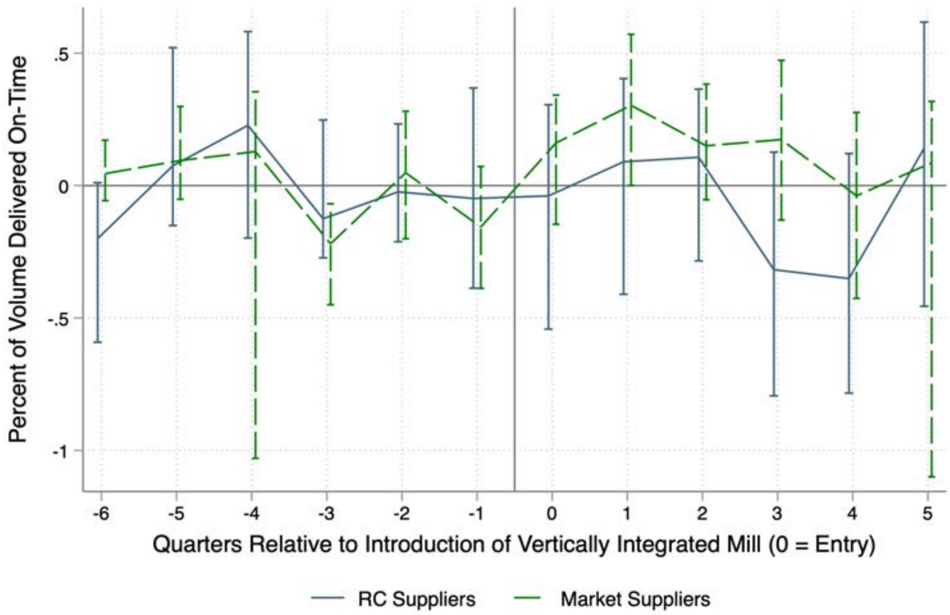

Note: Data from universe of fabric transactions by buyer. Doubly-robust difference-in-differences event study estimates with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

Figure A.28: Difference-in-Difference Estimates for Effect of Vertical Integration on Transaction Count

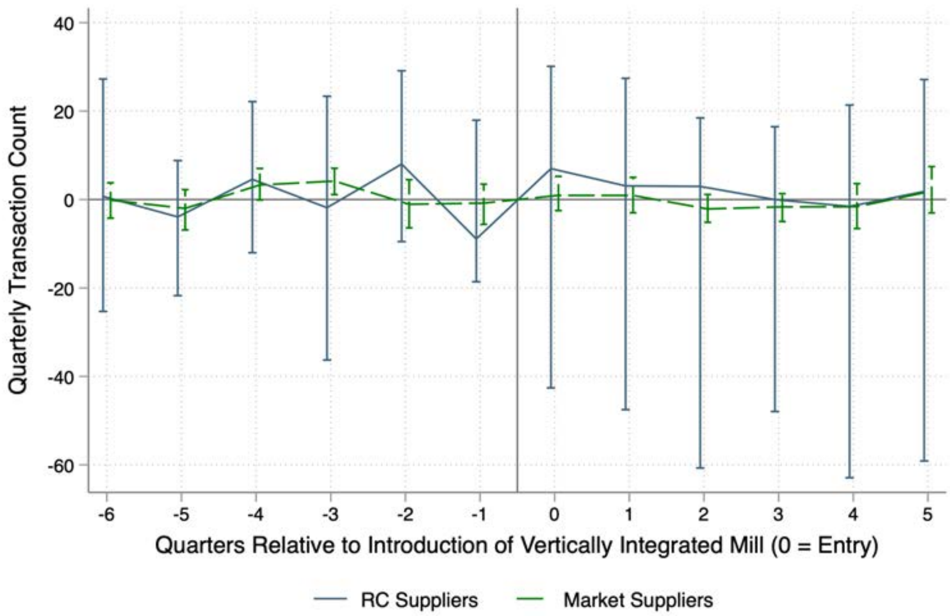

Note: Data from universe of fabric transactions by buyer. Doubly-robust difference-in-differences event study estimates with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

A.14.2 Synthetic Difference-in-Differences

Figure A.29: Synthetic Difference-in-Differences Estimates for Effect of Vertical Integration on Quarterly Volume

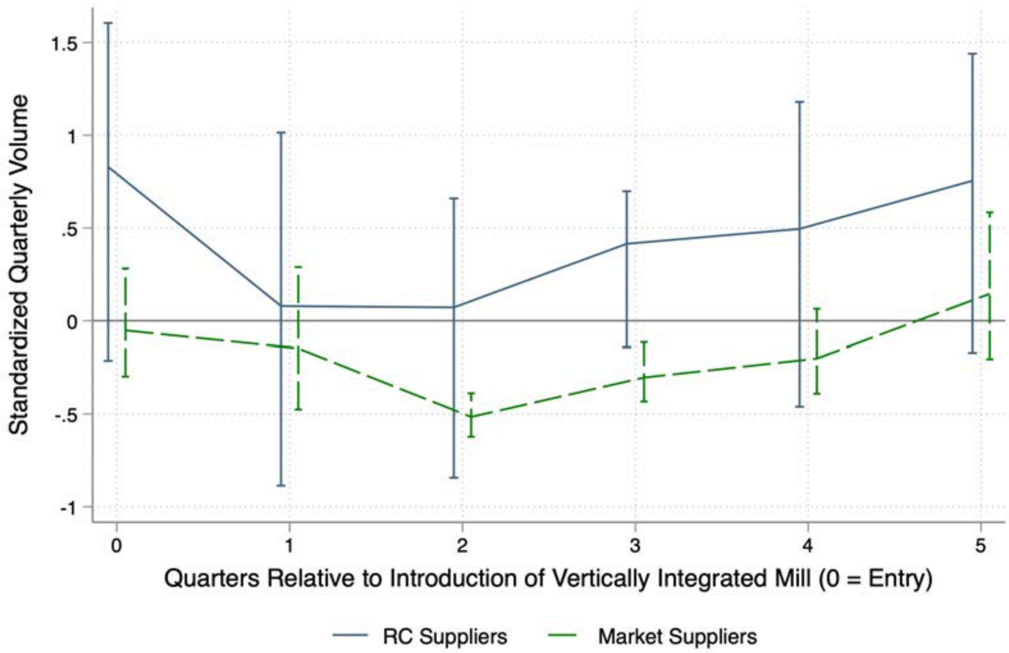

Note: Data from universe of fabric transactions by buyer. Doubly-robust difference-in-differences event study estimates with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

Figure A.30: Synthetic Difference-in-Differences Estimates for Effect of Vertical Integration on Standardized Prices

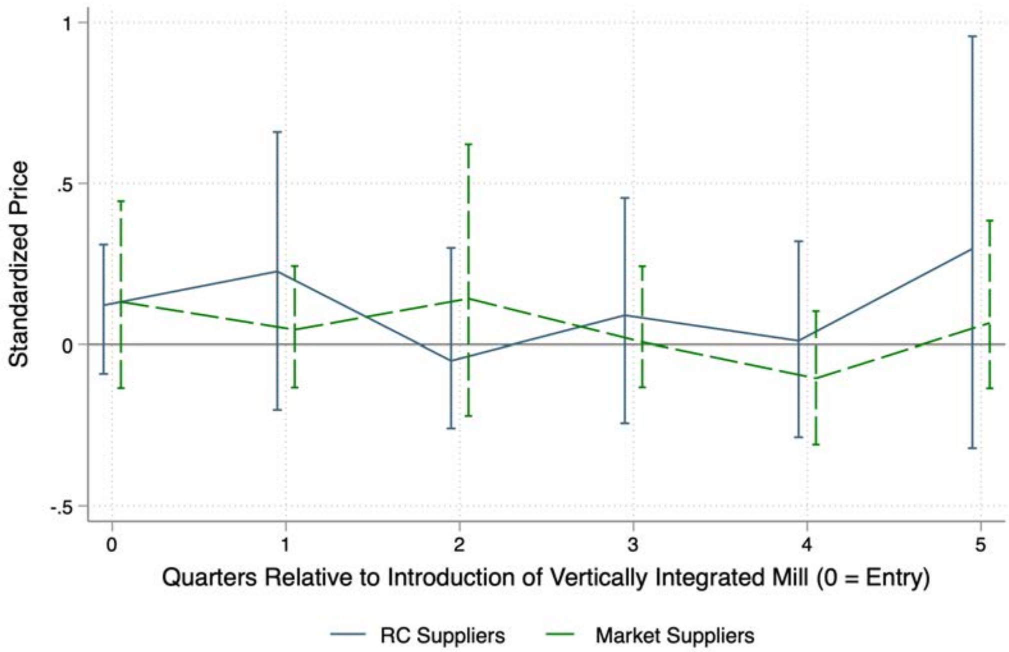

Note: Data from universe of fabric transactions by buyer. Doubly-robust difference-in-differences event study estimates with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

A.15 Horizontal Merger of Upstream Suppliers

I consider the effects of a horizontal acquisition where the least risk-averse supplier purchases the most risk-averse supplier, consistent with the idea of increasing countervailing horizontal market power. This policy can shift prices through three channels: *i)* increasing scale economies, thereby reducing variance within the relational contract (although not the ratio of variance inside versus outside the relational contract), *ii)* reducing supplier risk aversion to the level of the least risk-averse firm, and *iii)* shifting the buyer bargaining parameter to the level of the least risk-averse firm. This policy reduces prices by 17.6% of the pre-integration relational contract prices for the least risk-averse supplier. Decomposing the change, reducing the variance of capacity due to increased scale alone almost achieves the full gain at 16.7%, reducing risk-aversion has a similar effect at 15.6%, and changing only the bargaining parameter increases prices by only 3.7%. Note that the big supplier also benefits from the merger due to capacity variation reduction, but the effect is small–less than 1%.[57]

[57]Even holding prices constant, horizontal integration can increase profits by increasing quantities; that said, the fairly small benefits that accrue to the larger firm help offer an explanation as to the persistence of small firms in low-income countries.

I also consider the effects of a horizontal acquisition where the acquired firm is a small constrained supplier such that the threat point effect does indeed reduce relational contract prices before the horizontal acquisition. I find that prices increase by 24.6%. Decomposing the change as above, I find that reducing the variance increases price by 23.6% and reducing risk aversion similarly increases prices by 22.5%. These results have similar magnitudes because both of these effects of the merger are sufficient to stop the buyer from threatening to use the integrated supplier as the threat point. This result highlights the role of scale economies in effectively reducing the value of demand assurance, as unfavorable capacity states are less likely to be realized after integration due to the increase in quantities.

I find that the minimum risk aversion for the threat point effect to influence prices after horizontal integration increases by a factor of 3.25. For a supplier with such an extreme value of risk aversion, the horizontal merger has extremely large effects, increasing prices by 113.8%. Most of this change is due to the risk aversion reduction, which stops the threat point effect–a 110.1% increase from this part of the merger alone. Changing only variance has a large effect as well, doubling prices (*i.e.*, a 100.0% increase). Switching to the post-merger bargaining parameter has a small effect, increasing prices by only 10.6%.

A.16 Buyer Bargaining Parameter and Supplier Market Power

I measure concentration in supplier's product markets by, first, computing supplier-fabric HHI as the HHI for the fabric across suppliers one quarter prior to integration and, second, aggregating the supplier-fabric HHI to the supplier level as the volume-weighted average of the individual supplier-fabric HHIs. As market power only pertains to relational contract suppliers, I only consider volumes from relational contract suppliers when computing supplier-fabric HHI (results using volumes from all suppliers to compute the supplier-fabric HHI are similar: a point estimate of -1.116 significant at the 5% level).

Table A.9: Buyer Bargaining Parameter and Supplier Concentration

	Buyer Bargaining Parameter
Supplier Concentration	-1.622^{**}
(Volume-Weighted Supplier-Fabric HHI)	(.82)
N	18
R^2	.146

Note: Data from universe of fabric transactions by buyer and model estimates.

I find in Appendix Table A.9 that the proxy for supplier market power is indeed strongly, and statistically significantly (using bootstrapped standard errors), correlated with the buyer bargaining parameter. Specifically, higher concentration in a supplier's product markets is negatively correlated with the buyer bargaining parameter, suggesting that suppliers with more market power receive more of the surplus created by the relational contract.

Note that this pattern does not hold for risk aversion, where the estimate is not statistically significant.

Table A.10: Supplier Risk Aversion and Supplier Concentration

Supplier Risk Aversion		
		No Outlier
Supplier Concentration	−20.417	−.015
(Volume-Weighted Supplier-Fabric HHI)	(13.89)	(.061)
N	18	17
R^2	.355	.006

Note: Data from universe of fabric transactions by buyer and model estimates.