Value Creation and Value Capture in Indian Garment Sector Bargaining

Russell Morton*

April 6, 2025

Click Here for Latest Version

Abstract

This paper examines creation and distribution of surplus from global value chains (GVCs) in low- and middle-income country (LMIC) domestic supply chains. While GVC participation can enhance growth and productivity, low prices paid to small input suppliers raise concerns about the concentration of gains among large exporters. But prices alone are insufficient for welfare evaluation as many transactions occur in bargained agreements (rather than spot markets) with non-price terms that affect surplus, such as quantity stability and other insurance-like terms. I enrich a Nash bargaining model to study how both *i)* value creation through insurance-like agreement terms that mitigate spot market frictions and *ii)* value capture via buyers threatening to replace external suppliers with in-house production decrease prices paid to small, risk-averse suppliers. Using novel transaction data from an Indian garment manufacturer and its nearly 500 fabric suppliers, I estimate a structural model to decompose discounts into value creation and capture. Results indicate that discounts reflect value creation from profit-smoothing agreement terms rather than buyer capture; difference-in-differences estimates yield consistent findings. Counterfactual analyses highlight that increasing buyer competition has limited effects on supplier prices, whereas introducing profit insurance substantially increases prices for small, risk-averse suppliers.
JEL Classification: D22, L14, O12

*Yale University, russell.morton@yale.edu. I am extremely grateful to Lauren Falcao Bergquist, Zach Brown, Ying Fan, and Dean Yang for their guidance and support. I also thank Achyuta Adhvaryu, Jie Bai, Hoyt Bleakley, Mayara Felix, Bob Gibbons, Max Huppertz, Chuqing Jin, Amit Khandelwal, Asim Ijaz Khwaja, Francine Lafontaine, David A. Miller, Charlie Murry, Ameet Morjaria, Emir Murathanoglu, Thi Mai Anh Nguyen, Anant Nyshadham, Michael Raith, and Yulu Tang for their excellent feedback and comments, as well as seminar participants at CSAE 2024, EEA 2024, MWIEDC 2024, NEUDC 2024, OECD, PacDev 2025, Paris School of Economics, SIOE 2024, University of Michigan, University of Rochester-Simon Business School, University of Southern California, and Yale University. 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

While global value chains (GVC) can contribute to development in low- and middle-income countries (LMICs), reducing poverty and raising productivity (World Bank, 2020), their ultimate welfare effects in LMIC depend on surplus creation and distribution within LMIC domestic supply chains. Surplus distribution within LMIC domestic supply chains addresses concerns that the benefits of GVCs accrue primarily to the few large exporters, rather than their many small input suppliers, who may receive low prices (Boudreau et al., 2023; Dominguez-Iino, 2023; Zavala, 2022). However, input prices alone provide an incomplete picture of surplus creation and distribution because transactions in export-oriented domestic supply chains frequently occur inside bargained agreements rather than spot markets.[1] Bargained agreements incorporate both price and non-price features, and non-price features can affect surplus. For example, small suppliers may accept lower prices in exchange for quantity stability ("*quantity assurance*"). Despite the empirical relevance of LMIC bargained agreements including non-price features,[2] we have little understanding of how much surplus these agreements create and how buyers and suppliers share it.

This paper analyzes how two common features of LMIC interfirm bargaining affect surplus creation ("*value creation*") and distribution ("*value capture*") using novel transaction data from the universe of bargained agreements between a large exporter—an Indian garment manufacturer–and 471 fabric mill suppliers. First, agreements create value by mitigating frictions in the spot market. Indian fabric mills (suppliers) selling in spot markets face quantity (and therefore profit) risk because orders are indivisible (*i.e.,* cannot be split across suppliers). Given limited access to credit, saving, and insurance markets, suppliers may benefit from agreements stabilizing quantities and smoothing profits even when bargained prices are low. Therefore, the magnitude of value creation depends upon both the extent of profit smoothing inside the agreement and supplier preferences for smooth profits (*i.e.*, risk aversion).[3] Second, buyers might try to capture value through strategic behavior to

[1]The high prevalence of LMIC transactions outside spot markets reflects the increased relevance of spot market frictions (Macchiavello, 2022). Although data limitations prevent explicitly showing that transactions for the same product are less likely to occur on the spot market in LMICs, examples of bargained agreements, especially in export-oriented domestic supply chains, abound.

[2]Quantity assurance features in agreements between Kenyan rose exporters and their buyers (Macchiavello and Morjaria, 2015) and between fishing firms and ice suppliers in Sierra Leone (Ghani and Reed, 2022). Agreements can also address other frictions, such as financial market imperfections. For example, Rwandan coffee mills provide inputs on credit to their farmer suppliers (Macchiavello and Morjaria, 2020) and Ecuadorian manufacturers extend trade credit to buyers (Brugues, 2020).

[3]Examples of risk aversion shaping firm behavior in LMIC include coffee mills (Blouin and Macchiavello, 2019) and Ghanaian agriculture producers (Karlan et al., 2014).

improve their bargaining position. Because the buyer I study sources some inputs through vertically integrated supply, like many other large LMIC exporters,[4] it could capture value by threatening to replace agreement suppliers with in-house production.

My empirical setting exhibits three key features well-suited for studying value creation and value capture. First, the buyer builds a new in-house supplier, generating variation in the buyer's ability to capture value through replacement threats with in-house production. Second, the buyer's large size and simultaneous use of agreement, spot, and integrated suppliers for the same product (fabric) at the same time facilitates directly comparing transactions by supplier types. Third, garment manufacturing is a central LMIC export sector, employing 60 million workers across South Asia (International Labour Organization, 2022).

I build a model that incorporates the richness of the vertical contracts observed, including both *i)* value creation through insurance-like agreement terms (*e.g.,* stable quantities) and *ii)* value capture through strategic threats to replace suppliers in bargained agreements with in-house production, inside a Nash bargaining framework. The model formalizes how *value creation* reduces prices. As the buyer benefits from the agreement through discounts, lower prices represent the buyer's share of agreement surplus. Therefore, discounts from value creation increase in supplier risk aversion because the value of insurance-like agreement terms also increases in supplier risk aversion.

The model produces new implications regarding *value capture* from strategic threats to replace an agreement supplier with in-house production. Standard bargaining models of replacement threats (Ho and Lee, 2019; Ghili, 2022) show that external suppliers offer better terms to buyers when buyers threaten replacement with other external suppliers. My model both extends this result, showing that buyers can gain from replacement threats using *internal* rather than *external* suppliers, and introduces a novel result: contracting frictions can eliminate the benefits of replacement threats when suppliers exhibit low risk aversion. Replacement threats can break agreements if no surplus remains relative to the replacement. As agreement surplus increases in supplier risk aversion, my model yields a cutoff level of supplier risk aversion below which replacement threats break agreements. When contracting frictions prevent easy formation of new agreements, breakdowns require the buyer to increase

[4]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 that also relies on suppliers in agreements. As integration suggests the presence of frictions, many external suppliers for partially integrated firms are likely in bargained agreements. Morton (2023) even provides empirical evidence that vertical integration and these bargained agreements, which are relational contracts, are treated as substitutes.

spot sourcing.[5] However, the buyer prefers to preserve pre-integration agreements (with spot as the outside option) over spot sourcing to conserve the pre-integration agreement discount. The model rationalizes the buyer as if strategically choosing its outside option: the buyer captures value through replacement threats from sufficiently risk-averse suppliers, but maintains pre-integration agreements with suppliers with risk aversion below the cutoff.[6] Therefore, like value creation, value capture reduces prices especially for risk-averse suppliers.

I build and estimate a structural bargaining model to decompose observed discounts into value creation and capture. The key parameter shaping both value creation and capture is supplier risk aversion, estimated separately for each agreement supplier using supplier-by-quarter panel data. Risk aversion determines both how much value the agreement creates and whether the buyer captures value through replacement threats with in-house production. I develop a novel identification approach to disentangle risk aversion from bargaining parameters. To do so, I leverage plausibly exogenous demand shocks. Intuitively, price adjustments to shocks help suppliers smooth profits, and profit smoothing is especially valuable to risk-averse suppliers. Therefore, larger price responses to similarly sized demand shocks imply higher supplier risk aversion. Conversely, as bargaining parameters shift prices regardless of demand, their identification relies on average prices.

Model estimates illustrate that value creation alone, rather than value capture, explains the entirety of discounts from agreement suppliers. This result reflects low, but positive, estimated supplier risk aversion parameters; the median estimated supplier risk aversion is similar to that of LMIC coffee mills (Blouin and Macchiavello, 2019).[7] On average, agreement suppliers offer discounts that are 1.5% of spot market prices. This discount is meaningful

[5]Agreement breakdowns increase spot sourcing rather than in-house production because of capacity constraints at the integrated supplier. The buyer prefers to use in-house production to displace spot supply because spot suppliers are the suppliers least preferred by the buyer.

[6]I establish that the buyer endogenously choosing its outside option is a reduced-form representation of two distinct extensions of a Nash bargaining game with a contractual equilibrium that includes both cooperative and non-cooperative stages (Miller and Watson, 2013). First, the buyer and supplier engage in a signaling game about the buyer's residual in-house capacity, which determines the buyer's actual outside option, before bargaining. Private information about residual capacity is consistent with the setting given that contract terms with other suppliers have limited, if any, observability during bargaining with any individual supplier. Second, the buyer could commit to a specific level of residual in-house capacity before arriving at the bargaining table. Such commitment matches the buyer's commitment elsewhere in the agreement, such as paying agreement prices over spot market prices in some periods.

[7]Equivalently, all agreement suppliers have risk aversion below the cutoff where the buyer makes strategic replacement threats. The estimated risk aversion at the median of agreement suppliers is similar to estimates from Blouin and Macchiavello (2019) for LMIC coffee producers with revenues averaging 3.5 million USD per year. Although these are large firms by LMIC standards, they are much smaller than the buyer, which purchases nearly 230 million USD of fabric annually.

compared to spot markups averaging around 7% of costs and amounts to nearly .87 million USD annually aggregating across agreement suppliers.

I validate that discounts reflect value creation, and not value capture from replacement threats with in-house production, using a difference-in-differences (DiD) design. I leverage both a change in vertical integration–the buyer builds an integrated mill– and heterogeneity in supplier exposure to it.[8] This validation exercise complements the structural estimates because the two approaches rely on different variation for identification; exogenous demand shocks identify risk aversion parameters in the structural model. Consistent with the model—given estimated supplier risk aversion–DiD results show that integration does not affect agreement supplier prices or quantities, but reduces spot supplier quantities. Furthermore, DiD estimates are inconsistent with prices predicted by a standard model of replacement threats (Ho and Lee, 2019). This finding emphasizes how contracting environments, including the presence of co-existing spot markets, shape the effects of replacement threats.

The model also quantifies both *i)* the minimum risk aversion threshold such that the buyer captures value from strategic replacement threats with in-hose production and *ii)* how value capture affects prices. The minimum risk aversion cutoff is not only higher than all observed agreement suppliers but even high relative to risk aversion estimates in the literature, although not orders of magnitude larger. This result emphasizes how the contracting environment limits value capture from replacement threats. However, in other settings where the agreement provides more quantity stability or costs are more convex, the risk aversion cutoff would be lower. In (hypothetical) agreements with suppliers sufficiently risk averse that the buyer makes replacement threats, discounts reach up to 6.7% of spot prices. This magnitude is economically meaningful; it is larger than Indian Value Added Tax for apparel at the time (5%).[9] And, suppliers with the largest risk aversion are likely the smallest firms. As this analysis studies prices with suppliers more risk averse than all observed agreement suppliers, I validate the model's out-of-sample performance. Specifically, I estimate the model using only pre-integration periods and show the model fits out-of-sample post-integration quarters.

Last, I use the structural model to analyze two counterfactual policies to increase prices, emphasizing how supplier risk aversion shapes policy effects. First, policies that increase downstream buyer competition (*e.g.*, entry subsidies) only raise prices for suppliers with low risk aversion. I model the buyer responding to competition by increasing suppliers' surplus share to prevent long-run switches to another buyer. (As establishing agreements

[8]Heterogeneity in supplier exposure to integration derives from suppliers' product mix, with suppliers producing fabric that can be made at integrated supplier classified as exposed to integration.

[9]https://cleartax.in/s/gst-rates-clothing

requires time and effort, suppliers cannot switch in the short run.) For a supplier as risk averse as the median agreement supplier, this policy reduces the bargained discount by 75%. However, effects are negligible for suppliers sufficiently risk averse that the buyer captures value through replacement threats with in-house production. Agreement surplus relative to in-house production is so small that even allocating the entire surplus to the supplier barely shifts prices. Therefore, omitting value capture from replacement threats leads to large overestimates of the benefits of increased competition for the most risk-averse suppliers.

Second, creating the missing product for profit insurance against spot market quantity risk increases prices broadly, and especially for risk-averse small suppliers. Although moral hazard may render profit insurance infeasible, policies could help firms smooth profit through increasing financial market access (facilitating self-insurance), supporting cooperatives to protect prices, and developing products to hedge risk (*e.g.*, derivatives). Insurance enables suppliers to smooth profits without the agreement. Therefore, insurance both reduces value created by the agreement and prevents value capture by improving suppliers' outside options. This result highlights that strengthening small, risk-averse firms' ability to smooth profits can generate benefits beyond equalizing utility across states by improving prices they receive.

This paper primarily contributes to three strands of literature. While a broad empirical bargaining literature evaluates the efficiency of bargaining versus spot markets (*e.g.*, Grennan (2013); Keniston (2017); Freyberger and Larsen (2023)) and the role of behavioral norms (*e.g.*, Backus et al. (2020); Keniston et al. (2022)), among other topics, this paper specifically complements analysis of the determinants of bargained terms. Backus et al. (2019) and Larsen and Zhang (2021) study how asymmetric information about the inside option affects outcomes. I highlight the importance of inside option value created by agreement components that mitigate spot market frictions. Other studies focus on outside option values, such as how endowments shape outside option values (Hardy et al., 2022). Similarly, Ho and Lee (2019) and Ghili (2022) illustrate that firms threaten replacement to improve contract terms. This paper demonstrates how contracting frictions that hinder the easy formation of new agreements, in an environment with an active spot market and capacity constraints, can eliminate incentives for making replacement threats when such threats destroy agreements.

This paper adds to the literature on relational contracts in LMIC–agreements enforced by the value of future surplus rather than formal contracts due to contracting frictions–as the agreements studied are relationally enforced. Prior work explores both *i)* the efficiency of these agreements (Brugues, 2020; Harris and Nguyen, 2023) and *ii)* reactions, such as adapation or breakdown, to shocks to market structure or economic conditions, such as sup-

ply or demand (Macchiavello and Morjaria (2015, 2020, 2023); Ghani and Reed (2022)). I emphasize how relational contracts mitigate spot frictions to create value and how relational value is shared, rather than studying how contracting frictions result in relational instead of formal contract enforcement. My model highlights how supplier risk aversion, reflecting limited access to smoothing institutions (*e.g.*, insurance and financial markets), enables value creation through insurance-like agreement terms and how large integrated buyers might capture value through replacement threats. Furthermore, I illustrate how agreements respond to changes in the *vertical* organization of production rather than shocks to market structure–the *horizontal* organization of production. Broadly, as all relational contracts—including those outside GVCs—generate surplus (or else they would breakdown), value creation and capture apply generally to all relational contracts.

Because the trade-off between incentives (generally strongest in spot markets) and insurance is central to vertical relationships (Lafontaine and Slade, 2007), my empirical model incorporates risk aversion and insurance-like agreement terms. As risk aversion drives value creation, I develop a new approach to estimate risk aversion and separately identify it from bargaining parameters that leverages how demand shocks affect both the level and convexity of the agreement pricing function. Previous empirical work on vertical relationships (*e.g.*, Cuesta et al. (2019); Lee et al. (2021)) abstracts from risk aversion, which can result in underestimating the bargaining parameter for the more risk-averse party. Even with correct parameters, abstracting from insurance-like agreement terms can generate incorrect prices. In my empirical setting, omitting quantity assurance in the agreement—a key feature of the vertical relationship–yields underestimated supplier surplus and overestimated prices.[10]

The remainder of the paper proceeds as follows. Section 2 describes the data and context. Section 3 presents the theoretical bargaining model that incorporates key features of LMIC negotiations. In section 4, I estimate and validate a structural bargaining model to quantify value creation and capture. Section 5 analyzes counterfactual policies to reallocate surplus to small firms. Section 6 concludes and considers broader implications.

2 Data and Context

My empirical setting centers around an Indian garment manufacturer sourcing fabric from agreement suppliers, the spot market, and in-house production. First, I first describe the

[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 even be different orders of magnitude.

data. Second, I provide context on the garment manufacturing process and fabric procurement. Last, I describe the terms of the bargained agreement, highlighting empirical evidence consistent with my characterization of the agreement.

2.1 Data

I use two data sets: transaction data and cost data. The transaction data include the universe of the buyer's fabric purchases between September 1, 2016 and December 31, 2019. The data record the transaction price, date, supplier (including internal suppliers), 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). The buyer sources fabric from two integrated suppliers. The buyer built one integrated supplier during the middle of the data timeframe (June 2018), which generates variation in the buyer's ability to threaten to replace external suppliers with in-house production. The second internal supplier has produced fabric for the buyer for many years prior to the start of the transaction data (exact date unknown). Additionally, the transaction data include estimated production costs for fabrics produced internally. These cost estimates in the transaction data are for a mill that is not capacity constrained and represent average variable costs.

The cost data from the internal mills provide information about 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

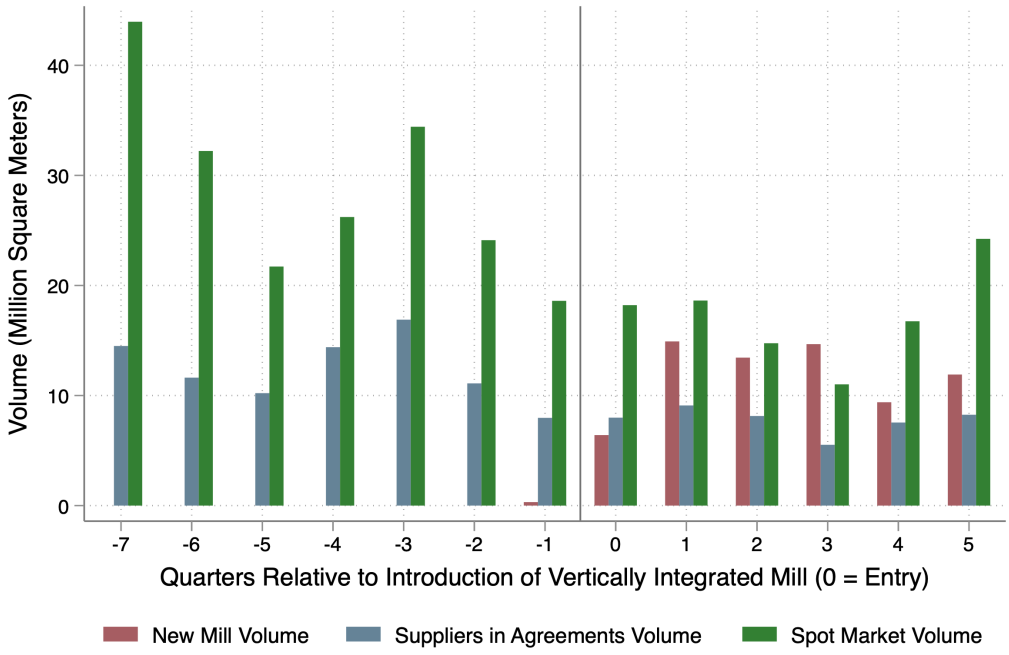
Prior to building the new integrated mill, the buyer sources 250 million square meters of fabric per year (worth 228 million USD over 10,000 distinct orders) both at a pre-existing internal mill (21% on average) and from external suppliers. External suppliers include both agreement suppliers, transacting under terms of bargained agreements, and spot market suppliers.

Figure 1: Integrated Mill Constructed in June 2018



Figure 2 illustrates three key facts about this empirical setting that make it well-suited for analyzing value creation and value capture. First, suppliers in agreements are empirically relevant, providing roughly 25% of fabric soured by the buyer pre-integration. Therefore, the buyer has strong incentives to care about prices it receives from these suppliers. Second, the increase in vertical integration is sufficiently large for the buyer to credibly threaten to replace suppliers with in-house production to capture value. The new mill fulfills roughly 20% of the buyer's fabric needs, which is more fabric than the largest external supplier provides ($\sim 7\%$). Third, as the buyer's demand from end clients (*e.g.,* garment retailers like Old Navy) does not change much (decreasing slightly to 215 million square meters of fabric per year), the buyer must make strategic decisions to reduce the number of external suppliers, the average volume per external supplier, or some combination of both. However, this integration is not so large that the buyer must reduce quantities from agreement suppliers. The buyer could (and, as I later show in the DiD analysis, does) reduce volumes exclusively from spot market suppliers.

Figure 2: Fabric Production Over Time by External Suppliers and New Fabric Mill



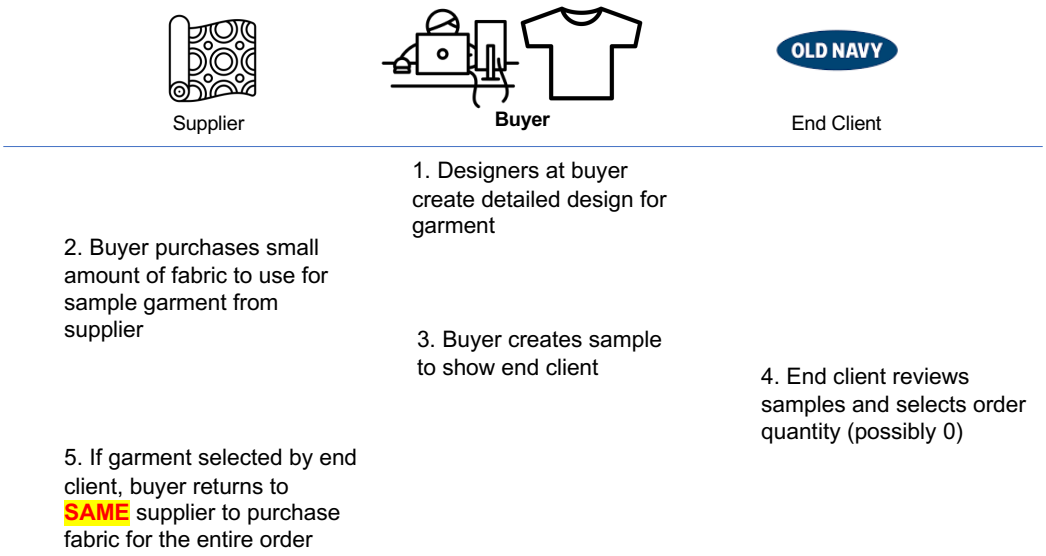
Note: Data from universe of fabric transactions by buyer. Bars represent the quantity of fabric purchased during the quarter from each supplier type.

2.3 Context: Fabric Procurement and Garment Manufacturing

The garment manufacturing process in my empirical setting (shown in Figure 3) begins with designers who work for the garment manufacturer (the buyer) conceptualizing and planning the garment. The designers sketch the shape and color and choose the broad type of fabric. 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, the cost is effectively negligible given the small quantity of fabric ordered. (In fact, fabric procurement costs for sample garments are so trivial that the buyer does not even record them in a centralized database.) Notably, the price for sample does not commit the buyer or supplier to this price for future orders. The buyer then shows the samples to an end client (*e.g.*, Old Navy), who subsequently decides which samples to order and determines quantities for any samples selected.

After the order is placed by the end client, the buyer returns to the identical supplier who provided the sample to purchase the fabric necessary for the entire order of the garment. Importantly, when the buyer picks a supplier for the sample, the buyer knows neither which

Figure 3: Garment Manufacturing Process



samples will be selected nor the order quantity. 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. Typically, there are only about 30 days between the initial fabric purchase for the sample and sourcing fabric for approved orders, in between which the sample is shown to the end client and the quantity is determined.

The key feature of this process is that the buyer commits to order fabric for the entire garment from the supplier that provided fabric for the sample. This commitment reflects the buyer's desire to ensure that the fabric purchased is identical to the sample ordered by the end client.[11] This supplier rigidity reflects that many subtle features of fabric are difficult to specify and reproduce. Consequently, switching suppliers might result in the end client rejecting garments for having (even a slightly) different fabric. Furthermore, short timelines prevent trying a new, cheaper supplier and engaging in trial-and-error to match the fabric. The buyer's emphasis on ensuring that the garment produced for the end client

[11]Based on discussions with fabric procurement managers at the buyer, the probability that a supplier which provided the sample is chosen for the entire order is at least 90%.

perfectly matches the sample aligns with other evidence of the importance of the end client to the buyer. For example, the buyer's production process prioritizes on-time deliveries for end clients over productive efficiency (Adhvaryu et al., 2019). In the few cases where the supplier changes after sampling, production is typically brought in-house, facilitating stringent quality control, and for orders with a long time before delivery, enabling experimentation to match the fabric.

The rigidity in supplier selection has two primary consequences: suppliers *i)* face quantity risk and *ii)* can hold-up the buyer. Suppliers face quantity risk because they will provide fabric for the entire order, but do not know at the time of sampling which orders will be purchased nor what quantities would be for confirmed garments. Additionally, because suppliers know they are the only option for the buyer after sampling, they can charge a high hold-up price to extract surplus (*i.e.,* profit) the buyer receives from the order.[12]

2.4 Context: Agreements in Indian Fabric Procurement

Bargained agreements create value through two insurance-like agreement terms. First, the agreement specifies a pricing function that maps capacities (quantities) to unit prices to help smooth profits. To build intuition about insurance provision from the pricing function, consider a supplier that sells at a constant unit price and has constant marginal costs. Then, the supplier charges a constant markup and profit increases linearly with quantities. A bargained agreement where unit prices decrease linearly with respect to quantity would smooth profits for such a supplier. As fabric suppliers have increasing marginal costs in this setting (*e.g.,* due to capacity constraints, and verified empirically using data from in-

[12]Theoretically, 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. In this setting, the buyer is unlikely to hold up the supplier given the timing of the garment production process. Namely, transaction prices are determined 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, pillow makers who need fabric for stuffing, or others. The buyer's main opportunity to hold up the supplier would occur after the supplier has paid the sunk cost to actually produce the fabric for the entire order. Then, 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. Furthermore, the buyer may have concerns about developing a reputation for not upholding prior agreements, as the buyer is a large firm and suppliers could communicate with each other. Then, holding up a supplier could adversely impact future orders with many suppliers.

house production in Figure 10),[13] the pricing function that best smooths profits is U-shaped with high prices for extreme capacities, small or large. When capacity is low, profit with constant unit prices would be low because there are few units; conversely, when capacity is high, profit with constant unit prices is low because costs per unit are high. Therefore, the U-shaped pricing function has the lowest prices when capacity is at target utilization (*i.e.,* 100%).[14] This profit insurance is valuable to risk-averse suppliers. Supplier risk aversion can reflect credit and financial market imperfections that generate challenges for firms to save and borrow funds. As a result, small firms might prefer stable profits to ensure that they can make required payments for the quarter, such as payroll, from flow profits rather than saving or borrowing. Additionally, as owners are the residual claimants of profits, small firms may have preferences like individuals.

Second, the agreement reduces quantity variability (*i.e.,* provides *quantity assurance*) compared to what suppliers would experience in the spot market. This quantity assurance both stabilizes profits and reduces average costs. Reduced quantity variability lowers profit variance because profits depend directly on quantities; additional insurance beyond the U-shaped pricing function creates value because the pricing function diminishes, but does not eliminate, profit variability.[15] Additionally, cost convexity means that quantity variability reductions also decrease average costs. Formally, this cost reduction follows from Jensen's inequality: a convex function–here, the cost function–evaluated at its expected value is lower

[13]Supplier cost convexity likely reflects that suppliers cannot adjust inputs frictionlessly to demand shocks over time given regulations and financial constraints. For example, the Factories Act of 1948 specifies that employers must pay double the standard wage to overtime workers. The buyer even emphasizes the importance of managing capacity because it affects costs in discussions. For example, a fabric procurement manager at the buyer who works with H&M states that past volume "is the most important factor whenever we address and negotiate the prices and lead time." 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.

[14]A perfect insurance contract is not feasible in this setting for both incentive and practical reasons. Regarding incentives, this agreement–like many other LMIC agreements–is enforced relationally. Therefore, temptations to deviate constrain the set of feasible contracts. A perfect insurance contract would imply effectively infinite prices for some quantities, in which case the buyer would deviate from the contract. Even expanding the set of contract options to allow for a fixed fee per quarter and a price that varies per order would not resolve this issue, as the buyer would have a strong incentive to deviate whenever required to pay a fee large enough to meaningfully provide any insurance. Therefore, to remain self-enforcing, the contract naturally does not include perfect insurance. Furthermore, in practice, as the agreement is over a quarter rather than for an individual order, it is generally not known what the last order in the quarter will be. Therefore, the buyer does not know what the ultimate revenue and costs will be for the supplier over the quarter, preventing the buyer from spreading out a perfect insurance payment in an incentive-compatible structure over the quarter.

[15]See footnote 14.

than the expected value of the convex function. Quantity variability reductions move average costs closer to costs at the expected value. Therefore, quantity assurance creates value in two ways: *i)* it provides insurance, helping risk-averse suppliers; and *ii)* reduces average costs, benefiting all suppliers regardless of risk aversion.

Because this agreement is not a formal and explicit court-enforceable contract, but is instead enforced by the value of future surplus from future trade, it is a relational contract. While the economic implications of value creation and value capture explored in this paper are generalizable (*i.e.,* they do not depend upon relational rather than formal contract enforcement), various appendices verify that results remain unchanged by incorporating dynamic incentive compatibility constraints to account for relational enforcement. In fact, as discussed in the appendices, results tend to become slightly *stronger* once incorporating relational contracts. Morton (2023) provides additional discussion of the relational enforcement of agreements in this setting. Importantly, this relational contract is *not* a relational *incentive* contract (Levin, 2003), as the relational contract does not (at least primarily) address information problems (*i.e.,* there is no subjective evaluation).

2.4.1 Empirical Patterns Consistent with Hypothesized Agreement

The central features of the bargained agreements, namely the agreement pricing function and quantity assurance, have testable empirical implications. First, prices in agreements should be U-shaped with respect to capacity. However, there is no reason to expect that spot suppliers would also have a U-shaped pricing function. Second, the quantity variability reduction in the agreement implies that volumes for agreement suppliers should be more stable than spot supplier volumes.

As the econometrician does not observe a supplier's classification as an agreement or spot supplier, the terms of the agreement pricing function, nor the level of quantity assurance, it is necessary to define these objects (or proxies for them) to confirm that empirical patterns are consistent with the hypothesized agreement. Therefore, I classify suppliers as transacting under agreement rather than spot market terms if they transact every month. This definition reflects that agreements are enforced relationally, and relational enforcement requires

frequent transactions.[16,17] I classify 19 of 471 external suppliers as agreement suppliers (*i.e.,* relational contract suppliers). These 19 suppliers account for 25% of fabric volume prior to integration. Importantly, while the empirical tests that follow highlight that transactions with agreement suppliers differ from those with spot market suppliers, they do not result mechanically from this supplier classification approach. I also later validate that transaction patterns with agreement suppliers differ from other large spot market suppliers with similar volumes, supporting this supplier classification approach.

To test empirically for a U-shaped pricing function, I estimate the association between prices and capacities for agreement suppliers. 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, much of which is unpredictable.[18] This measurement approach aims to capture the extent to which suppliers hold up the buyer. For example, consider purchases from different suppliers that occur simultaneously 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. Therefore, using unadjusted prices as a measure of supplier pricing would not account for cost heterogeneity across fabrics and over time.[19]

[16]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.

[17]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 agreement suppliers do not allow the increase in vertical integration to affect supplier classification as an agreement supplier. Appendix Figure A.2 shows the distribution of the percentage of months with transactions across suppliers.

[18]While some seasons generally have higher capacity utilization, such as the early fall in anticipation of increased demand in export destinations from December holiday shopping, individual fabrics experience large unpredictable heterogeneity, as shown in Appendix Figure A.3.

[19]Alternatively, I could measure markups directly using the buyer's cost data for fabrics they produce. However, this approach would both limit the sample, as the buyer's cost data are only available for fabrics they produce, and fail to adjust for industry-wide seasonality in capacity.

Figure 4: Pricing Function: Binned

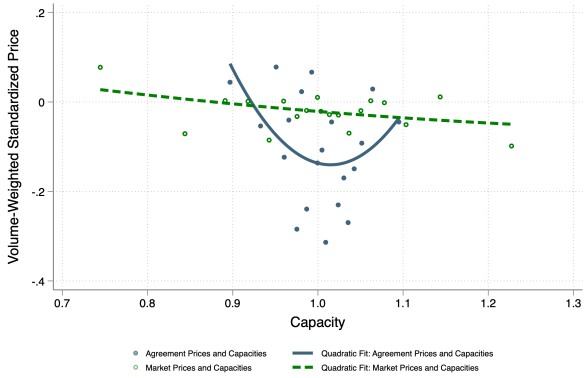
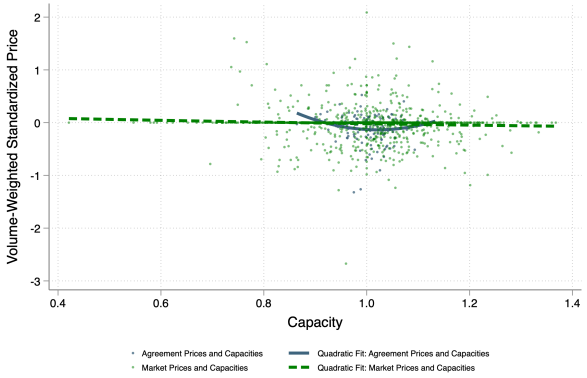


Figure 5: Pricing Function: Raw



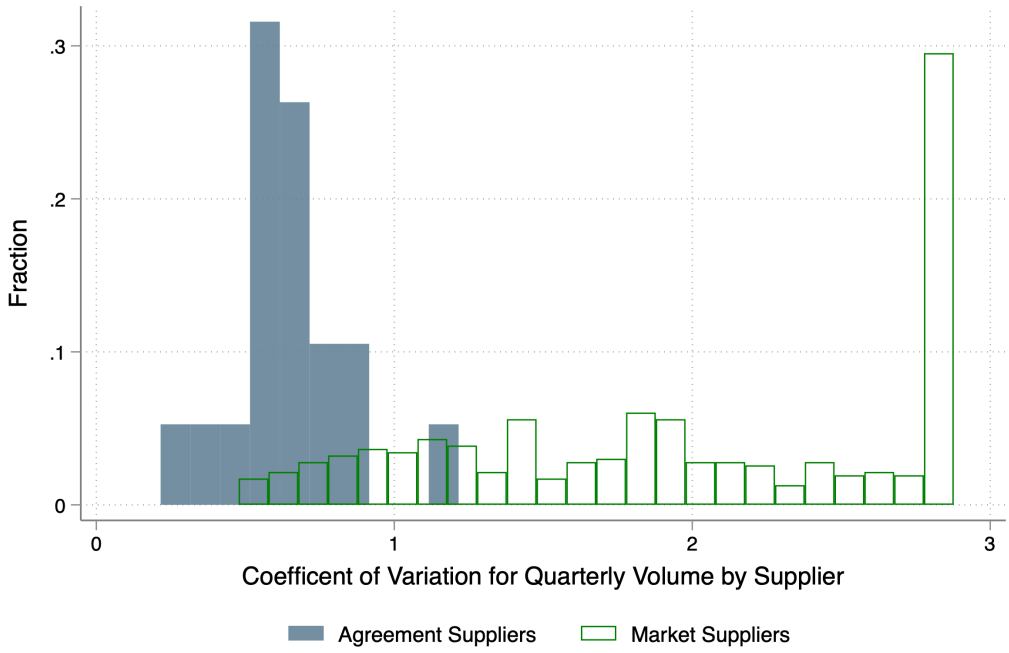
Note: Data from universe of fabric transactions by the buyer, transformed into supplier-by-quarter panel data. Quadratic fit is from a regression of volume-weighted standardized price on capacity and capacity squared. Figure 15 bins prices and capacities using 20 quantiles. The mapping from quantities to capacities is as discussed in 4.2.2.

Figures 15 and 16 illustrate the U-shaped pricing functions with respect to capacity for agreement suppliers, especially compared to the relatively flat spot market prices (although possibly with some quantity discounts). A key difference between the figures is the scale of the y-axis; although the U-shaped pattern is clearer using the binned data, the pattern is still observable even compared to the overall price variation in the data. However, 16 illustrates that spot prices are quite flat, especially when compared to overall variability in prices. These figures also highlight that agreement prices are, on average, lower than spot prices. Appendix Section A.2.1 shows that agreement prices pass statistical tests for being U-shaped, while spot prices do not pass such tests. Note that a pricing function in a competitive environment (*i.e.,* without any market power) that respects order indivisibility would trace out the marginal cost curve. Therefore, as marginal costs increase throughout the domain of the function, the competitive pricing function would be neither U-shaped nor flat. Instead, the first-order determinants of pricing in this setting are consistent with the importance of insurance-like terms with agreement suppliers and hold-up with spot market suppliers.

As quantity assurance (*i.e.,* reduced quantity variability) results in stable agreement quantities over time, quantity variability for agreement suppliers should be lower than for spot suppliers. To bring this hypothesis to the data, I compare the coefficient of variation by supplier type: agreement versus spot. Figure 6 illustrates that agreement suppliers have less volatile demand than spot suppliers. The difference in average coefficient of variation by supplier type is statistically significant ($p < .001$). This result is not entirely mechanical,

as supplier classification as an agreement supplier depends only upon transaction frequency and not volume. It follows that suppliers classified as agreement suppliers, in principle, could have extremely volatile demand, provided that they receive at least one order per month. In Appendix A.3, I show that quantity patterns for agreement suppliers suggest the buyer directs sampling opportunities during the quarter to help stabilize agreement supplier quarterly volumes, but not for spot suppliers.

Figure 6: Coefficient of Variation and Supplier Type



Note: Data from universe of fabric transactions by the 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 (perhaps a price schedule depending on quantity to accommodate cost convexity) should the end client order the garment seems like an attractive alternative to resolve the hold-up problem. In Appendix A.4, I discuss why the buyer and suppliers rely on relational rather than such a rich explicit contract. Separately, as the buyer prefers sourcing fabrics from agreement suppliers to spot suppliers, the buyer would prefer to rely exclusively on agreement suppliers. Appendix A.5 highlights forces that prevent the buyer from sourcing exclusively from agreement suppliers.

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

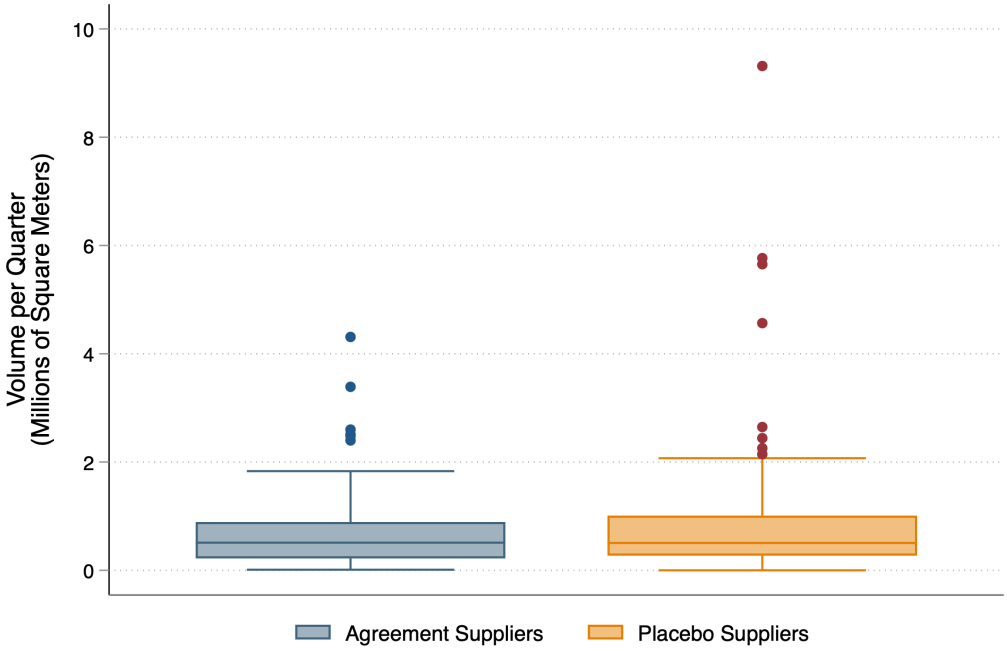
Although the empirical evidence aligns with the hypothesized bargained agreement, I conduct additional analysis to mitigate concerns that unobserved supplier heterogeneity might instead explain observed patterns. For example, some suppliers could be more efficient, potentially due to better management given that technology is largely standardized (at least among the suppliers that produce similar fabrics and meet the buyer's screening criteria). More efficient suppliers might pass on some cost savings to the buyer, resulting in lower prices. Efficient suppliers might also receive more (although not necessarily more stable) volumes. However, this scenario seems unlikely. There are no incentives for efficient spot suppliers to pass on cost savings outside a bargained agreement–they should optimally charge the high hold-up price.

To address possible concerns that unobserved supplier heterogeneity, rather than the contracting under a bargained agreement, explains empirical patterns, I reproduce comparisons of transactions between agreement suppliers and spot suppliers but only consider a restricted subset of spot suppliers that receive large orders ("placebo suppliers"). As relational enforcement of the agreement requires frequent transactions (furthermore, quantity assurance over time in the agreement relies on many transactions in practice), placebo suppliers are plausibly not in an agreement. Importantly, most plausible explanations of unobserved heterogeneity suggest that supplier volume, rather than transaction frequency, should differentiate terms with suppliers. For example, if unobserved supplier efficiency explains lower prices, then suppliers with similar volumes–regardless of transaction frequency–should also charge lower prices. In other words, plausible unobserved supplier heterogeneity is likely highly correlated with supplier volume. Therefore, finding different patterns in the data for placebo and agreement suppliers undermines the idea that unobserved heterogeneity explains differences between agreement suppliers and spot suppliers more broadly. Note that this analysis alone does not suggest that agreement suppliers are more or less efficient than other suppliers. Instead, it highlights that differences in transactions by supplier type are due to the agreement versus spot terms rather than any efficiency advantages they may (or may not) have.

I construct a set of placebo suppliers by finding the spot suppliers with the largest volumes pre-integration that are not classified as agreement suppliers. To ensure the set of placebo suppliers is as directly comparable to the set of agreement suppliers as possible, I select the same number of placebo suppliers as agreement suppliers. I then verify that volumes are indeed similar for agreement and placebo suppliers, noting that this approach to defining

placebo suppliers does not mechanically ensure that placebo suppliers receive similar volumes to agreement suppliers.

Figure 7: Volume Comparison: Agreement and Placebo Suppliers



Note: Data from universe of fabric transactions by the buyer. Boxplots show volume per supplier per quarter using all quarters 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 agreement suppliers.

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

Appendix A.6 shows that placebo suppliers do not have U-shaped pricing functions and have less stable volumes than agreement suppliers.

3 Theoretical Bargaining Model

This section introduces a theoretical bargaining model to analyze how bargained outcomes are shaped by both *i)* value creation through non-price agreement terms that mitigate spot

market frictions and *ii)* value capture from large partially-integrated buyers strategically threatening to replace suppliers with in-house production. To do so, I enrich a standard Nash bargaining framework along two key dimensions. First, the model incorporates value creation through both insurance-like agreement terms: a U-shaped pricing function and quantity assurance. Second, the model evaluates both the conditions under which the buyer would benefit from making replacement threats and the magnitude of the gains from such threats. I demonstrate that the buyer capturing value from replacement threats exclusively from risk averse suppliers is a reduced-form representation of two novel model extensions that effectively endogenize the buyer's outside option.[20]

The U-shaped pricing function and quantity assurance in the agreement create value through two separate channels. First, both agreement terms reduce (but do not eliminate) profit variability, which is valuable to risk-averse suppliers. Second, quantity assurance creates additional value by reducing average production costs due to cost convexity.[21] Cost reductions benefit all suppliers regardless of risk aversion. Therefore, the buyer's discount with a risk-neutral supplier reflects the buyer's share of benefits from cost reduction, as profit smoothing does not create value for risk-neutral suppliers. Additional price reductions as supplier risk aversion increases reflect value created by profit smoothing.

The model also includes value capture through threats to replace a supplier with in-house production.[22] To measure value capture, I calculate model prices under both of the partially-integrated buyer's outside options–the spot market and in-house production. The price response to the change in outside option represents value capture from replacement

[20]Both model extensions consider a contractual equilibrium that includes both a cooperative (Nash bargaining) and non-cooperative stage (Miller and Watson, 2013). The first extension is based on a signaling game about residual in-house capacity (Appendix A.9), while the second features the buyer committing to a level of capacity prior to bargaining (Appendix A.10).

[21]By Jensen's inequality, the expected value of a convex function decreases as variance decreases.

[22]Notably, my empirical setting isolates the change in bargained prices due to bargaining position effects of the replacement threat. This emphasis on bargaining position effects reflects that the vertical integration episode studied, while large enough to credibly threaten displacement, is sufficiently small that agreements need not change (average) quantities after integration. Therefore, assuming that average quantities with agreement suppliers were optimally determined pre-integration (given constraints), I need not model how average quantities respond to integration. I confirm empirically, using the difference-in-differences model, that agreement suppliers do not experience quantity decreases). Isolating the bargaining position effect is useful because average quantity changes might also affect bargained prices, at least by altering the value of quantity assurance included in the agreement. Equivalently, I study value creation and capture in agreements with fixed target quantities (*i.e.,* at steady state) rather than focusing on contracting in earlier stages of the agreement that shapes the level of target quantities in the agreement (see Brugues (2020) for empirical analysis of how relational contracts begin and evolve). That said, the effects of replacement threats due to partial vertical integration broadly apply even in settings when quantities must change. In such settings, prices can adjust due to multiple mechanisms including replacement threats.

threats. Additionally, the model yields a cutoff value of supplier risk aversion below which the buyer refrains from making strategic replacement threats, as replacement threats then break agreements and necessitate increased high-priced spot sourcing.

3.1 Model Description

I model agreement prices as solving a Nash bargaining problem over a pricing function that maps quantities to prices, as expressed in 1. To focus on value creation and value capture, the model abstracts from relational enforcement. Extending the model in Appendix A.7.1 to incorporate dynamics essential for relational enforcement (specifically, the dynamic incentive compatibility constraint as described in Macchiavello and Morjaria (2023)) does not alter any key results for reasonable discount factors[23] As I study the effects of value creation and capture in stable agreements (*i.e.,* at steady state) with quantity assurance, I model quantities as realizations of shocks drawn from a distribution with a pre-existing quantity target set during prior (unmodeled) bargaining stages.[24]

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

The buyer's surplus is the difference in profit when buying fabric under the bargained terms of the agreement–purchasing quantity q^A at price $p^A(q^A)$–and sourcing fabric from the outside option–with value O_B. I model the buyer as risk neutral given its large size.[25] Therefore, buyer surplus from the agreement can be expressed as the difference in costs between purchasing the fabric in the agreement and the costs of purchasing fabric through the outside option. I omit the buyer's revenue from sales to the end client because buyer revenues are independent of fabric prices given the sequence of the bargaining process. Specifically, negotiations over order quantities with the end client precede bargaining for fabric purchases. Therefore, the buyer's revenue is the same regardless of fabric supplier, meaning that revenue

[23] Some exact numbers change as the contract adjusts (and can even break) due to dynamic incentive compatibility constraints. Because contract adjustment to satisfy relational enforcement can generate additional breakdowns of agreements, relational enforcement generally amplifies constraints on the use of replacement threats to capture value.

[24] See Brugues (2020) for a discussion of how quantities evolve to steady state in relational contracts.

[25] Additionally, the buyer requires suppliers to purchase insurance against the main risk of production delays. Moreover, each supplier's quantity is small relative to the buyer's overall volumes. Therefore, the buyer is likely risk neutral, especially for any individual agreement supplier.

does not generate any profit difference between the inside and outside options.

The supplier's surplus is the utility of profit in the agreement–with CARA utility $U(\cdot)$ and production costs a function of 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 may arise from various factors, including owners' preferences over profits mirroring individuals' preferences over consumption because small supplier profits are likely the primary income source for owners. Alternatively, firms may need to make some payments every period (*e.g.,* payroll). Financial market frictions could make it challenging or expensive to borrow or save to pay such expenses. Firms may then prefer to use contemporaneous cash flows to cover these high-frequency expenses. I model the supplier cost function with increasing marginal costs (*i.e.,* convex costs). Marginal costs can increase because machine inputs are fixed in the short run; therefore, scaling up production can generate costly congestion. Additionally, under Indian labor laws, labor costs could increase dramatically with larger quantities–the Factories Act of 1948 mandates double wages for overtime work.[26] Using data from the integrated supplier, I empirically validate cost function convexity (see Figure 10).

The outside option for both the buyer prior to vertical integration and the supplier is to transact in the spot market. Even though the buyer is a large firm, the buyer (and suppliers) are small relative to the Indian fabric market; the buyer's domestic market share is less than 0.5%.[27] Therefore, I model the spot market as perfectly elastic with a constant price p^M, which is higher than production costs due to hold-up. This assumption is consistent with the largely flat spot prices in Figure 16.[28] Equivalently, I explicitly model market power within agreements through bargaining over prices, while capturing the first-order pricing determinant in the spot market, specifically that suppliers hold up the buyer, in a reduced-form way consistent with the observed flat spot market prices (as in Figure 16).

After integration, the buyer's outside option could change to use the same production

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

[27]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).

[28]Incorporating buyer market power when purchasing inputs in the spot market leads to underestimates of value creation but has no effect on value capture. If the buyer has market power in the spot market, then the agreement creates additional value for the buyer by reducing spot prices for fabric purchased in the spot market because the buyer would face an upward-sloping supply curve. For value capture, there is no change in spot quantities when the buyer outside option changes, as sourcing remains from agreement suppliers. Therefore, buyer power has no effect on value capture through replacement threats.

technology as the external supplier. This assumption reflects conversations with staff at the buyer emphasizing that the buyer and agreement suppliers use similar technology. Broadly, these firms all rely on (near-)frontier technology but do not engage in meaningful research and development efforts 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, differences in production technology shift costs for both the inside and outside options fairly similarly. Therefore, payoffs do not change much (for reasonable levels of heterogeneity).[29] This change in the outside option assumes that the buyer treats its integrated supplier and agreement suppliers as alternatives (*i.e.,* substitutes). Morton (2023) shows strong evidence of substitutability in this setting, leveraging exogenous variation at the fabric level. It follows that the integrated supplier serves a credible threat to displace agreement suppliers, as building a mill constitutes a meaningful and visible sunk cost.

Quantities in the model capture the stochastic nature of demand, as end clients determine quantities before the buyer bargains over prices with fabric suppliers. As the bargained agreement provides quantity assurance rather than large volumes, I model spot market quantities as a mean-preserving spread of agreement quantities. I set mean quantities (for both the agreement and spot market) to 100% capacity (and ensure that, given market prices and productive technology, the supplier's optimal capacity utilization is 100%). This assumption need not imply that the supplier exclusively works with the buyer. Instead, quantities can be interpreted relative to capacity reserved by the supplier for the buyer, such as floor space or time saved. The importance of the agreement to the supplier, in comparison to their contracts with other buyers, can then affect the risk aversion parameter (*i.e.,* suppliers with diverse income streams might be less risk averse over production for any individual buyer.)[30] Formally, the outside options are:

[29]Large heterogeneity in costs functions seems implausible–if the integrated supplier were much more efficient, the buyer should completely integrate. Conversely, if external suppliers were much more efficient, then the buyer should never integrate.

[30]In principle, working with other buyers could also affect cost convexity. For example, if a supplier has many buyers, then the supplier might be able to balance demand shocks across different buyers. However, seasonality in industry demand limits the potential to smooth demand by contracting with many buyers (*i.e.*, demand shocks across buyers are correlated). Furthermore, recall that reasonable changes in costs do not affect the magnitude of value creation or capture, as shown in Appendix A.7.2.

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

To close the model, bargaining parameters determine surplus shares in the optimization problem that yields the Nash bargaining solution. The buyer bargaining parameter is α_B. Higher values of α_B indicate that the buyer receives a larger share of value created by the agreement. As the buyer and supplier split the surplus, their bargaining parameters must sum to one. Accordingly, the supplier bargaining parameter is $1 - \alpha_B$.

Appendix A.8 provides details on the exact parameterization of the model and computation of bargained prices. Parameters and functional forms are chosen to illustrate the economics of value creation and capture rather than calibrated to match data or parameter estimates from other studies. To quantify how value creation and value capture affect prices in Indian garment manufacturing, I build and estimate a structural model closely tied to this theoretical model (Section 4).

3.2 Value Creation and Capture: Supplier Risk Aversion

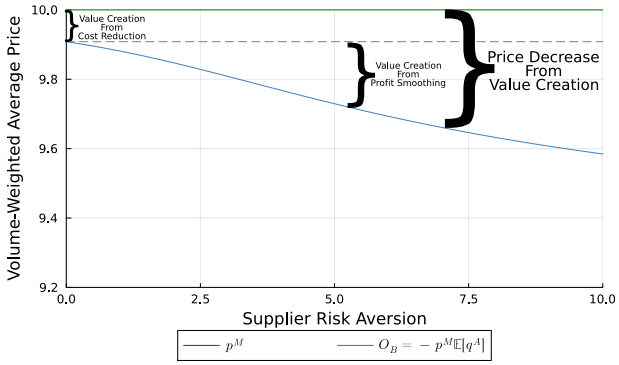


Figure 8: Value Creation

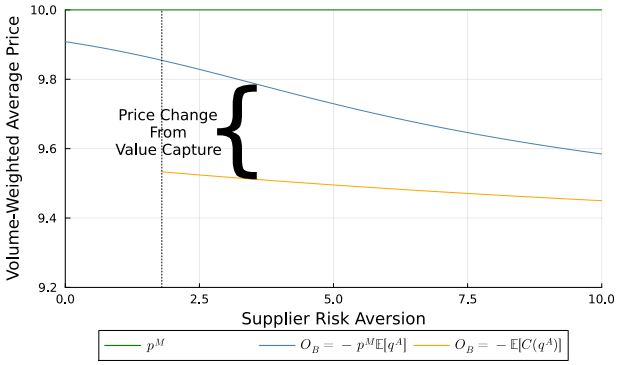


Figure 9: Value Capture

Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.8.

Figures 8 and 9 illustrate how the expected cost to the buyer (*i.e.,* the volume-weighted average fabric price) varies with both supplier risk aversion and the buyer's outside option, holding the buyer bargaining parameter constant. When the buyer purchases from the spot market, there is no bargaining, and the expected cost is unrelated to supplier risk aversion. However, when the buyer's outside option is the spot market, value creation in the agreement results in price discounts for the buyer relative to spot prices. As long as the buyer bargaining parameter is positive, some of the value created in the agreement is shared with the buyer as a price discount. The discount in agreements with risk-neutral suppliers represents the buyer's share of surplus generated by cost reduction from quantity assurance; cost convexity means that quantity variability reductions reduce costs. As supplier risk aversion increases, the agreement creates additional value by smoothing profits through both insurance-like agreement terms: *i)* the quantity assurance and *ii)* the U-shaped pricing function. Consequently, the discount in the agreement increases.

When the buyer's outside option switches to the integrated supplier, value capture from threatening to replace suppliers with in-house production can increase discounts. This additional discount is visible in Figure 9 as the decrease in prices when the buyer's outside option is in-house supply. However, the model also highlights that the contracting environment constrains the benefits of replacement threats. Specifically, there is a cutoff level of risk-aversion below which the agreement breaks down. This breakdown occurs because suppliers with risk aversion below the cutoff derive surplus from the outside option of selling in the spot market, where prices are high due to hold-up. Therefore, when the buyer bargains with such suppliers, no pricing function generates surplus for both the buyer and supplier.

What happens when supplier risk aversion falls below the cutoff, meaning the agreement would breakdowns with integrated supply as the outside option? The buyer would prefer to retain its pre-integration agreement because the agreement features a discount relative to spot prices. But in Nash bargaining models, outside options are not choices but rather primitives that represent the value of the outcome that would occur under disagreement. However, in practice, parties may have incentives to strategically present specific outside options during bargaining, especially when the credibility of that outside option is not visible to the other party. To capture this intuition, I show that the buyer endogenously choosing its outside option can be interpreted as the reduced-form representation of two distinct extensions of the Nash bargaining game. Both extensions consider a *contractual equilibrium* which includes both a cooperative phase (the Nash bargain) and a non-cooperative phase (Miller and Watson, 2013). As Nash bargaining is axiomatic, there is no clear notion of

player's beliefs in Nash bargaining games (although Lee et al. (2021) highlight that the Nash approach has a similar "flavor" to passive beliefs). Therefore, the non-cooperative phase facilitates incorporating beliefs in a principled way.

The first extension incorporates a non-cooperative signaling game about residual capacity at the in-house supplier before Nash bargaining occurs. This approach leverages that incomplete information bargaining is often modeled as a signaling game (*e.g.,* Backus et al. (2019)). Residual capacity determines whether the in-house supplier is the buyer's true outside option. Appendix A.9 establishes the existence of a weak Perfect Bayesian Equilibrium where the buyer claims that its in-house supplier is not capacity constrained to capture value when bargaining with suppliers that highly value the non-price agreement terms. However, this equilibrium only exists when supplier non-price valuation of the agreement is above a cutoff. This cutoff value corresponds to the cutoff level of supplier risk aversion for value capture in the baseline model. Information asymmetry about residual capacity–the key "state of the world" in this environment–is central to the model. This private information about capacity reflects the buyer's own knowledge of both their orders from other end clients and other commitments in-house. Appendix A.10 highlights an alternative model with a simpler non-cooperative game. First, the buyer initially commits to a residual capacity in the non-cooperative stage of the game, the cooperative game occurs, and then the buyer fulfills any additional orders. The key to this model is the buyer's ability to commit to behavior that is statically not optimal. As the agreement is a relational contract, it already includes commitment to behavior that is not statically optimal–for example, the U-shaped pricing function embeds the buyer's commitment to pay agreement prices over spot prices for some realized quantities. As with the signaling game, the commitment game also results in an equilibrium where the buyer commits to saving capacity in-house to capture value, but exclusively when bargaining with suppliers that highly benefit from non-price agreement terms.

This model considers only the bilateral bargaining game, but the buyer bargains multilaterally with many suppliers. In my setting, agreement quantities are inframarginal, mitigating concerns about contracting externalities across agreement suppliers. Through the lens of the signaling game in Appendix A.9, the buyer's fabric purchases from other suppliers (at their equilibrium values) relative to their total demand from end customers shapes residual supply available at the in-house supplier. Alternatively, the game with commitment in Appendix A.10 directly incorporates the spot market for additional sourcing at the end of the game. Extending either model from a bilateral game to a sequential multilateral game that repeats the bilateral game would yield the same equilibrium as the bilateral case in both models

regardless of the bargaining order with agreement suppliers.[31]

This equivalence reflects that, on path, the buyer never actually displaces an agreement supplier with in-house production. Therefore, prior agreements with in-house suppliers do not change the credibility of in-house production as a replacement threat because residual in-house capacity remains unchanged.[32] Similarly, a simultaneous (rather than sequential) multilateral game, such as Nash-in-Nash, would also produce the same multilateral equilibrium. The equivalence of the simultaneous game reflects that neither the signaling game nor the commitment game need create any contracting externalities across negotiations, as none of the bilateral equilibria actually use any in-house capacity.[33]

Although only supplier risk aversion, and not the bargaining parameter, directly affects value creation from insurance-like agreement terms and capture from strategic replacement threats, for completeness, a similar analysis holding supplier risk aversion constant but varying the bargaining parameter is in Appendix A.11.

4 Structural Bargaining Model: Estimation, Identification, and Validation

4.1 Why a Structural Model

I build and estimate a structural model closely resembling the theoretical model to quantify how value creation and value capture each separately affect bargained prices. A structural model is necessary because the surplus generated by non-price agreement features (the U-shaped pricing function and quantity assurance) is not directly observed in the data.

Additionally, the structural model enables analyzing prices, and the roles of value creation and value capture, for any level of supplier risk aversion (and buyer bargaining param-

[31] In the commitment game, it would need to be the case that the buyer bargains with all agreement suppliers before realizing any final capacity shocks. If the buyer could select the bargaining order, it would select this order of bargaining first with agreement suppliers

[32] Only private equilibria can be maintained in the multilateral sequential signaling game. For example, if suppliers have homogenous costs and agreement prices are observable, then a supplier who sees agreement prices consistent with the pooling equilibrium where the buyer is capacity constrained should not believe that the buyer is unconstrained. This condition is stronger than reality, as capacity constraints might vary by fabric, which would make public equilibria more attainable. It follows that this equivalence holds provided at least one of: *i)* agreement prices are private information, *ii)* there is unknown cost heterogeneity, or *iii)* the bargaining order specifies that the buyer first bargains with risk-averse suppliers (which would be the buyer's preference) to ensure that it never signals that it is capacity constrained before it desires to do so.

[33] As with the sequential multilateral game, the commitment game would require bargaining with all agreement suppliers before the ultimate capacity shock is drawn.

eter). Out-of-sample analysis is valuable because I only observe a limited, selected sample of agreement suppliers. However, even very small firms in LMIC GVCs rely on bargained agreements, such as smallholder Rwandan farmers (Macchiavello and Morjaria, 2020). But, given the large size of the buyer I study, it seems unlikely that such small suppliers are observed as the buyer's agreement suppliers. Furthermore, the structural model provides a framework for evaluating policies to increase prices paid to suppliers and how policy effectiveness relates to supplier risk aversion. I consider two policies to increase prices: *i)* increasing downstream buyer competition; and *ii)* creating the missing insurance product so suppliers do not need the agreement to smooth profits.

First, I discuss estimation and identification of the model. Then, I present parameter estimates and model validation against both untargeted moments and out-of-sample fit. Next, I decompose observed agreement discounts into value creation and capture. Taking the model seriously, I also examine how value creation and capture affect prices for heterogeneous buyer and supplier (*i.e.*, relationship) types, including types different from observed agreement supplier types. I recover both *i)* the cutoff level of supplier risk aversion such that the buyer captures value through strategic replacement threats and *ii)* the price response to value capture. Subsequently, I validate the model using a difference-in-differences (DiD) approach that compares the model-implied effect of value capture for agreement suppliers, which depends upon estimated supplier risk aversion parameters, with a direct estimate from the DiD model. Additionally, I show that event study coefficients using model-implied prices match untargeted event study coefficients using actual prices. Notably, the DiD estimation leverages variation not used to estimate the structural model.

Last, I use the DiD estimates to highlight how the bargaining environment restrains the buyer's ability to capture value through strategic replacement threats. I estimate a Nash-in-Nash with Threat of Replacement (NNTR) model in the spirit of Ho and Lee (2019). The NNTR model considers an environment where a buyer can easily form new agreements and abstracts from the simultaneous presence of spot and in-house production alongside bargained agreements. I demonstrate that the prices implied by the NNTR model are inconsistent with the DiD event study estimates.

4.2 Estimation Approach

4.2.1 Buyer and Supplier Primitives

Estimating the structural model recovers three groups of parameters: supplier-specific risk aversion parameters; agreement-specific buyer bargaining parameters; and a tuning parameter common across all agreements. The tuning parameter penalizes deviations from risk neutrality, motivated both by economic intuition, namely that firms should not have large deviations from risk neutrality, and econometric concerns of overfitting, given the relatively few observations per supplier.[34]

I estimate model parameters using pre-integration data and test out-of-sample fit using post-integration data. This approach ensures a hold-out data sample for model validation. I first describe the "inner" optimization that estimates supplier risk aversion and buyer bargaining parameters given the tuning parameter. Subsequently, I discuss the "outer" optimization to estimate the tuning parameter.

The "inner" optimization has three steps. First, I derive the model-implied agreement pricing function $\hat{p}^A(q^A; \mathbf{\Omega}_s)$ for each buyer-supplier pair by finding the function that optimizes the Nash product as in (1). In the expression $\hat{p}^A(q^A; \mathbf{\Omega}_s)$, $\boldsymbol{\Omega_s}$ represents the parameter vector for each relationship: the supplier risk aversion parameter and the buyer bargaining parameter. The agreement pricing function depends on these parameters through the Nash product, as shifting these parameters changes the value of the Nash product and, consequently, the optimal agreement pricing function.[35] Second, I estimate prices for observed agreement quantities using the model-implied pricing function. Last, I compute the objec-

[34] As I study discounts in stable agreements, miscategorization of spot market suppliers as agreement suppliers would lead me to underestimate risk aversion for such suppliers and incorrectly conclude that value creation is not important. This concern motivates my conservative classification of agreement suppliers to ensure that all estimates are truly for such suppliers; as parameters are supplier-specific, possibly excluding some agreement suppliers hs minimal implications for parameter esttimates for suppliers categorized as agreement suppliers. Regardless, the difference in observed contracting patterns between agreement and placebo suppliers–the suppliers most likely to be incorrectly categorized as spot market rather than agreement suppliers–suggests that the classification approach performs well.

[35] I use a second-order Taylor approximation, meaning $\hat{p}^A(q^A; \mathbf{\Omega}_s) = \beta_{0,s} + \beta_{1,s} q^A + \beta_{2,s}(q^A)^2$ to recover the $\beta_{0,s}, \beta_{1,s}, \beta_{2,s}$ that maximize the Nash product. As suppliers typically operate near full capacity, I estimate this second-order Taylor approximation around a capacity of one. Identification of $\beta_{0,s}, \beta_{1,s}, \beta_{2,s}$ follows the argument broadly for identification of buyer bargaining parameter and supplier risk aversion in 4.3. For example, $\beta_{0,s}$ shifts the level of average agreement prices, and $\beta_{2,s}$ largely determines the convexity of the agreement pricing function. Including $\beta_{1,s}$ permits additional flexibility for the agreement 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 agreement price would be the same for both capacities without including $\beta_{1,s}$. It follows that supplier risk aversion, given cost convexity, identify $\beta_{1,s}, \beta_{2,s}$ as these parameters govern how U-shaped the pricing function is ($\beta_{2,s}$) and if it is asymmetric ($\beta_{1,s}$).

tive function for each supplier, which is the root mean squared error of price estimates from the pricing function using only the pre-integration data.

The estimation routine penalizes extreme estimates of supplier risk aversion parameters by adding the absolute value of the risk aversion parameter, weighted by the tuning parameter, to the objective function. In the CARA utility function used, a risk aversion parameter of zero represents risk neutrality. Therefore, this approach penalizes deviations from risk neutrality in either direction equally. This regularization method is similar to a LASSO regression by constraining extreme parameter values. Incorporating the tuning parameter reflects both an economic motivation, as firms are unlikely to have large deviations from risk neutrality, as well as econometric concerns around overfitting, given relatively few observations per supplier.

Formally, the "inner" optimization finds the vector of supplier risk aversion and buyer bargaining parameters that, given the tuning parameter (λ), minimize the sum of the per supplier objective functions:

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

In this optimization problem, S represents the set of suppliers, $\{q^A_{s,t}\}$ represents the set of observed quantities for agreement supplier s across pre-integration periods t,[36] T is the number of periods, $\hat{p}^A_s(q^A_{s,t}; \mathbf{\Omega}_s)$ is the price estimated by the agreement pricing function for supplier s for quantity $q^A_{s,t}$, $p^A_s(q^A_{s,t})$ is that actual realized agreement price for supplier s for quantity $q^A_{s,t}$, λ is the tuning parameter, and θ_s is the supplier risk aversion parameter.

The "outer" optimization recovers the tuning parameter than maximizes fit in post-integration quarters. Specifically, the optimization routine finds the tuning parameter that leads to the bargained pricing function that best fits the post-integration data, noting that the agreement pricing function is estimated using only pre-integration data. The tuning

[36] While it would be appealing to use monthly rather than quarterly data for this analysis to increase the sample size, the agreements appear to operate at the quarterly level rather than the monthly level. Not only are the reduced-form patterns consistent with the agreement clearest at the quarterly level, but anecdotally staff procuring fabric at the buyer emphasize that meetings with key suppliers occur quarterly rather than monthly.

parameter impacts post-integration fit by influencing the estimated agreement pricing function: a larger tuning parameter pushes estimated supplier risk aversion parameters towards risk neutrality. Formally, the "outer" optimization solves:

$$\min_{\lambda} \underbrace{\sum_{s \in S}}_{\text{Sum over suppliers}} \Bigg[\overbrace{\sum_{\{q^A_{s,t}\}}}^{\substack{\text{Sum over} \\ \text{post-integration} \\ \text{quarters}}} \underbrace{\sqrt{\frac{1}{T}\left(\hat{p}^A_s(q^A_{s,t}; \mathbf{\Omega}_s) - p^A_{s,t}\right)^2}}_{\text{RMSE: Model \& Actual Prices}} \Bigg]$$

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:

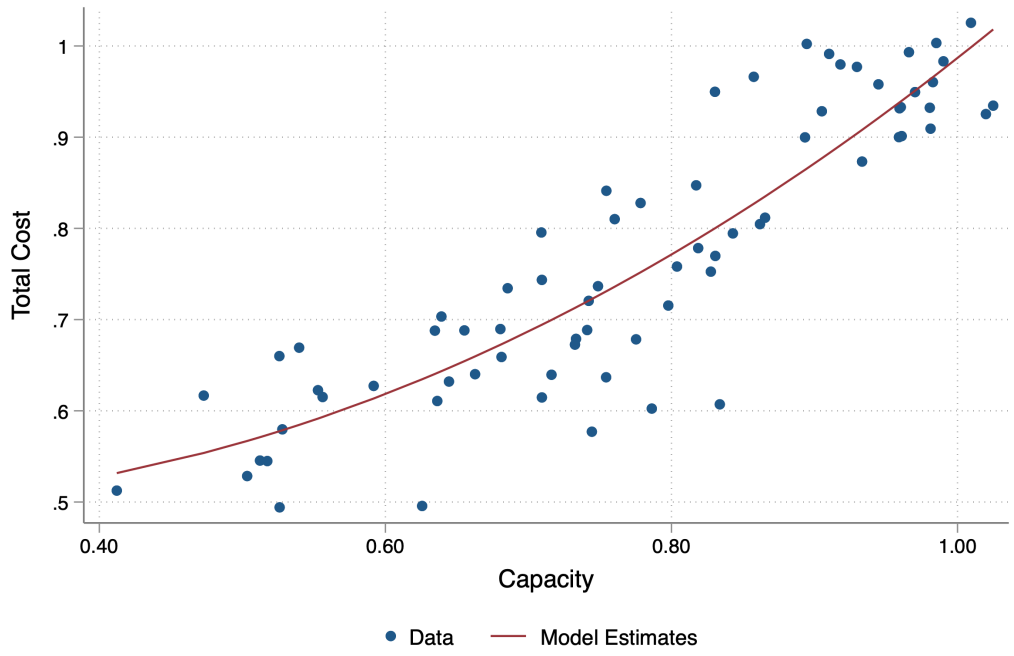
Table 1: Additional Moments for Estimating the Structural Model

Parameter(s) in Model	Moment(s) to Match	Data
Cost Function Convexity	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 Spot Suppliers Compared to Variance in Quantities for Agreement Suppliers	Transaction Data
p^M	Ratio of Prices from Spot Suppliers to Estimates from Integrated Supplier	Transaction Data

I estimate the cost function using panel data on capacity and costs from both integrated suppliers. As procurement staff at the buyer suggest that agreement suppliers use similar technology as the integrated suppliers, their cost functions should be comparable. 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 to capture convexity, but do not impose any structure to ensure convexity (*i.e.,* $\phi_2 > 0$). The estimated cost function is equivalent to a second-order Taylor approximation of the true cost function. The panel data on cost and capacity are at the division-by-month level from April 2019 to March 2020, where divison maps broadly to fabric type (*e.g.,* knits and striped

fabrics).[37] Identification of cost function parameters relies on exogenous capacity shocks. Capacity shock exogeneity derives from the institutional features of garment manufacturing. Specifically, the fabric supplier is determined when the sample garment is constructed, which occurs *before* the final order quantity is established. Therefore, the exact quantity used at the integrated supplier reflects stochastic downstream demand rather than endogenous buyer decisions.[38] Therefore, as quantities are (at least locally) exogenous, OLS regressions provide causal estimates of the effect of capacity on costs.

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.

[37] Covexity estimates leverage the most temporally disaggregated data rather than a collapsed panel. To illustrate why, consider a supplier that operates at 100% capacity for two months and another supplier than operates at 80% capacity one month and at 120% the next. Both suppliers have an average capacity of 100%. Then, aggregated bi-monthly level would suggest different costs for the same capacity. Regressing costs on capacity would not recover the true relationship between them that is clear in the disaggregated data.

[38] The buyer presumably has some sense of how common certain volumes are from various end clients based on previous transactions. Although the buyer likely uses this knowledge when sampling to try to ensure that capacity utilization at the integrated supplier is not too extreme, it is impossible to perfectly achieve targeted capacity usage over short periods. Indeed, empirically, I observe meaningful heterogeneity in capacity utilization.

Estimates demonstrate that the model fits the data well. The R^2 is .795 and the RMSE is .07, which is small relative to overall variance in costs observed. Additionally, the estimated cost function is meaningfully convex, with the coefficient on q^2 at .7815 and a p-value of .02. I transform costs into percentages relative to mean costs for each fabric group, as costs are reported by fabric group, facilitating analysis that combines data across all fabrics. The strong fit of the model suggests that heterogeneity by fabric type is unlikely to be a central feature of the cost function.

As the cost function is measured in capacity units, transaction quantities must be converted into capacity units to estimate the structural model. Two assumptions underlie my conversion. First, I assume that log quantities determine capacity utilization. This assumption reflects that orders with larger quantities typically also have longer time between order and delivery. Therefore, capacity utilization does not increase linearly with quantity (see Appendix Table A.3). Expressing quantities in logs captures this pattern in a reduced-form manner without requiring any additional parameters that govern how the distribution of quantities over time mediates the mapping from quantity to capacity.

The second assumption is that the supplier's capacity, at least for agreement suppliers, is the mean capacity in the pre-period. This assumption reflects that the agreement provides consistent capacities around an agreed-upon capacity target. I model quantity realizations in the agreement as drawn from the gamma distribution that corresponds with the method of moments estimator for the observed agreement quantities. I chose a gamma distribution because it is a simple two parameter distribution with positive support and fits the data well.[39] 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 an identical deviation in actual quantity as a smaller change in capacity.

Additionally, I must estimate the counterfactual capacity variance that agreement suppliers would experience with their outside option–the spot market. As the agreement includes quantity assurance, counterfactual spot market quantity variance should be larger than agreement quantity variance. Therefore, I assume that agreement suppliers' coun-

[39]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.

terfactual distribution of spot capacity realizations would have the same mean as in the agreement but with higher variance, as in the theoretical model. As a supplier's quantities from other buyers are not observed, I leverage differences in quantity variability by supplier type. Specifically, I assume that counterfactual spot market variance for agreement suppliers matches the empirical ratio of average spot market quantity variance to average agreement quantity variance, averaging over suppliers. Including all spot suppliers yields overestimates of the variance that agreement suppliers would face outside the agreement. This overestimate reflects that any individual spot market supplier presumably sells to many buyers. As long as demand across buyers is imperfectly correlated, supplier capacity variance from the spot market overall is smaller than variance in orders from any one individual buyer. Therefore, I only consider placebo suppliers, rather than all spot suppliers, for this calculation.[40] Placebo suppliers are the closest counterfactual for agreement suppliers overall given their similar volumes.[41] Encouragingly, the estimate of counterfactual variance using placebo suppliers is reasonable: spot quantity variance is 8.8 times larger than agreement quantity variance. By contrast, using all spot suppliers to estimate counterfactual variance yields an unreasonable estimate that spot variance is 971 times larger than agreement variance.

Last, I estimate the spot price relative to costs using transaction data. For each transaction, I compute the margin charged by the supplier as the ratio of the transaction price over the estimated production costs for a mill at 100% capacity utilization. Cost estimates derive from production cost data from the integrated suppliers.[42] The spot market price, with incorporates hold up, has an estimated margin of 7 percentage over costs. This estimate also seems quite plausible: it is a large enough markup to create incentives for strategic sourcing through both vertical integration and bargained agreements but not so large as to be implausible.

4.3 Separate Identification of Supplier Risk Aversion from Bargaining Parameters

It is not immediately apparent that separate identification of supplier risk aversion from the buyer bargaining parameter is possible. Increasing either supplier risk aversion or the

[40] As $log(0)$ is undefined, periods with no orders are considered as having zero capacity utilization.

[41] 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.

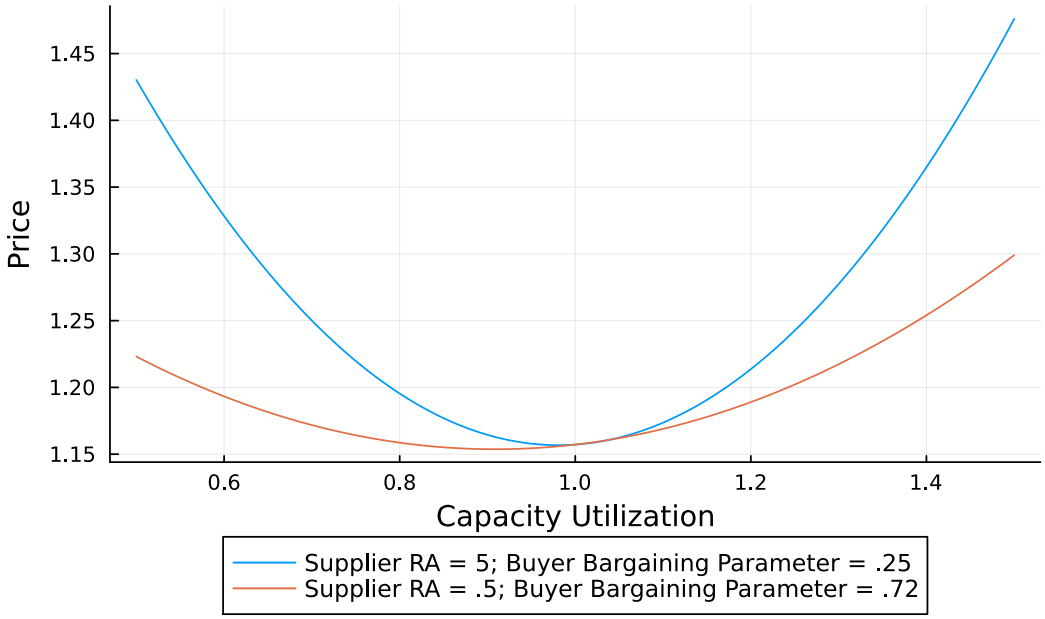
[42] Using internal cost estimates means that margins are constructed from fabrics produced internally because cost estimates are only available for these fabrics.

buyer bargaining parameter alone reduces average bargained prices. As in the theoretical model, when supplier risk aversion increases, holding the buyer bargaining parameter fixed, the value created by the agreement increases. As some of this value created is shared with the buyer (as long as the buyer bargaining parameter is not zero), higher risk aversion results in lower prices. Similarly, when the buyer bargaining parameter is high, the buyer receives a larger share of the value created by the agreement through lower prices.

My novel identification strategy leverages two key sets of variation in the data. First, there is variation in how prices adjust to capacity changes. The covariance of prices and capacity helps identify supplier risk aversion, as changing supplier risk aversion alters the *convexity* of the pricing function (*i.e.,* how U-shaped it is). More convex contracts are associated with more risk-averse suppliers because convex prices reduce profit variance across fluctuations in capacity. This logic reflects that profits are highest when capacity is near one. When capacity is low, even though cost per unit is low, small quantities require high prices to equalize profits. Conversely, when capacity is high, cost per unit is high due to increasing marginal costs, necessitating higher prices to smooth profits. Second, there is variation in the level of average bargained prices relative to spot prices. This variation in the level of average bargained prices is essential for identifying the buyer bargaining parameter, as increasing this parameter reduces prices regardless of capacity.

When only one price-capacity tuple is observed, it is impossible to separately identify supplier risk aversion from the buyer bargaining parameter. Specifically, for any given buyer bargaining parameter, adjusting supplier risk aversion can rationalize any observed price-capacity tuple (with higher risk aversion leading to a lower price and vice-versa). Analogously, given any supplier risk aversion, changing the buyer bargaining parameter can rationalize any observed price-capacity tuple. As illustrated in Figure 11, both a supplier risk aversion of 5.0 with a buyer bargaining parameter of 0.25 and a supplier risk aversion of 0.5 and a buyer bargaining parameter of .72 result in the same price when capacity is one.

Figure 11: Convex Pricing Function for Different Parameters



However, a *second* price-capacity tuple for a different capacity facilitates separately identifying supplier risk aversion from the buyer bargaining parameter. 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. A third price-capacity tuple identifies asymmetry in the U-shaped function that enables prices to smooth profit more effectively, matching observed patterns in agreement prices.

To further visualize this identification strategy, Figures 12 and 13 show the effects of individually shifting supplier risk aversion and the buyer bargaining parameter, respectively. They highlight that shifting the risk aversion parameter mostly shifts the shape of the agreement pricing function while shifting the buyer bargaining parameter mostly shifts price levels.

Figure 12: Shift Supplier Risk Aversion

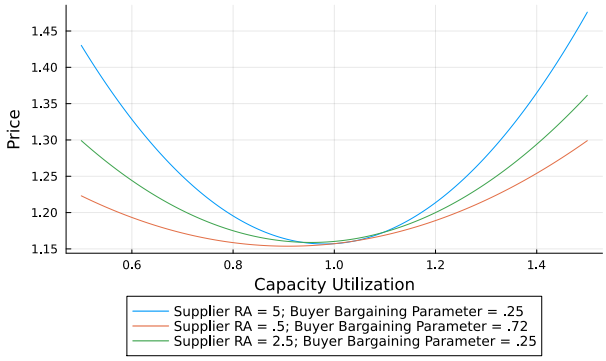
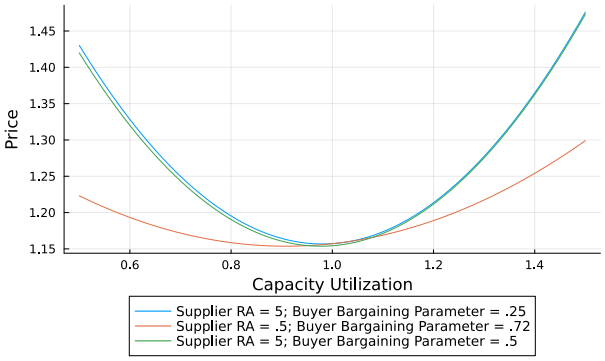


Figure 13: Shift Buyer Bargaining Parameter



To bring this identification argument to the data, I leverage supplier-specific capacity shocks that originate from stochastic end client demand. Although the bargained agreement has a target capacity, institutional features of the garment production process–namely, that neither which samples will be ordered nor what quantities would be are known at the time of sampling–prevent the bargained agreement from eliminating capacity variance across periods. (Instead, the bargained agreement reduces capacity variance). In Appendix A.12.1, I discuss these demand shocks in more detail and provide supporting evidence that they generate exogenous variation. Specifically, I show results inconsistent with the most plausible threat to exogeneity–that the agreement includes dynamic quantity assurance across quarters (*i.e.,* quarters following an extreme capacity are closer to the capacity target).

4.4 Supplier Risk Aversion and Buyer Bargaining Parameter Estimates

Table 2: Estimated Supplier Risk Aversion and Buyer Bargaining Parameters

	Supplier Risk Aversion	Buyer Bargaining Parameter
Minimum	.0000	.0094
25th Percentile	.0026	.0301
Median	.0085	.6214
75th Percentile	.0223	.9155
Max	14.6811	.9985
Mean	.8282	.5113
Standard Deviation	3.4573	.4286
Correlation Between Supplier Risk Aversion & Buyer Bargaining Parameter	.2847	

Note: Estimates from structural model.

Parameter estimates in Table 2 (confidence intervals in Appendix Table A.14) show that risk aversion estimates for agreement suppliers resembles coffee mills studied in Blouin and Macchiavello (2019) (.0085 for the median agreement suppliers and .0068 for coffee mills).[43] These coffee mills, described as "large firms by developing country standards," average over $3.5 million a year in sales and about $2 million in total assets. Given similarities in size, sophistication, and setting–with respect to both market institutions (as firms in LMIC) and stakes of riskiness (contracts with buyers)–to the fabric suppliers studied in this paper, it is promising that risk aversion estimates are similar. Importantly, even though these results suggest that agreement suppliers are large firms (especially by LMIC standards), the buyer studied is much larger. For example, the buyer's fabric purchases, which are smaller than revenue, surpass $200 million per year. The estimated tuning parameter (λ) is small at .002, meaning that low estimated supplier risk aversion parameters are not primarily due to a large tuning parameter.

Estimated buyer bargaining parameters imply that the buyer typically receives slightly more than half of surplus created by the agreement,[44] This result aligns with the risk aversion estimates, which also suggest that agreement suppliers are large firms (with the exception of one outlier supplier with high risk aversion and a high buyer bargaining parameter). Together, these results underscore that suppliers in agreements are large, strategically sophisticated firms, consistent with positive selection into the agreement. This positive selection matches the intuition that creating and maintaining agreements, especially complex agreements such as the one studied here, requires advanced organizational capabilities within the firm (Macchiavello and Morjaria, 2023).

4.5 Model Validation and Fit

As the estimation routine includes outside estimates, I prioritize model validation to ensure that the model reasonably captures agreement prices. First, I demonstrate that the model matches an untargeted moment. Second, I show that the model outperforms two alternatives in post-integration quarters, which are largely out-of-sample as the estimation routine

[43] Blouin and Macchiavello (2019) and other studies estimate risk aversion using different utility functions. For comparability, I convert reported risk aversion parameters to a CARA coefficient given their estimate and utility functions. Appendix Table A.6 shows (converted) CARA risk aversion parameters.

[44] I explore one possible microfoundation for heterogeneity in buyer bargaining parameters in A.23.

primarily uses pre-integration quarters.[45] Model validation using a difference-in-differences (DiD) approach follows the discussion of value creation and capture. As the DiD provides a direct estimate of value capture from replacement threats with in-house production, I postpone this validation exercise until after discussing the model estimate of value capture.

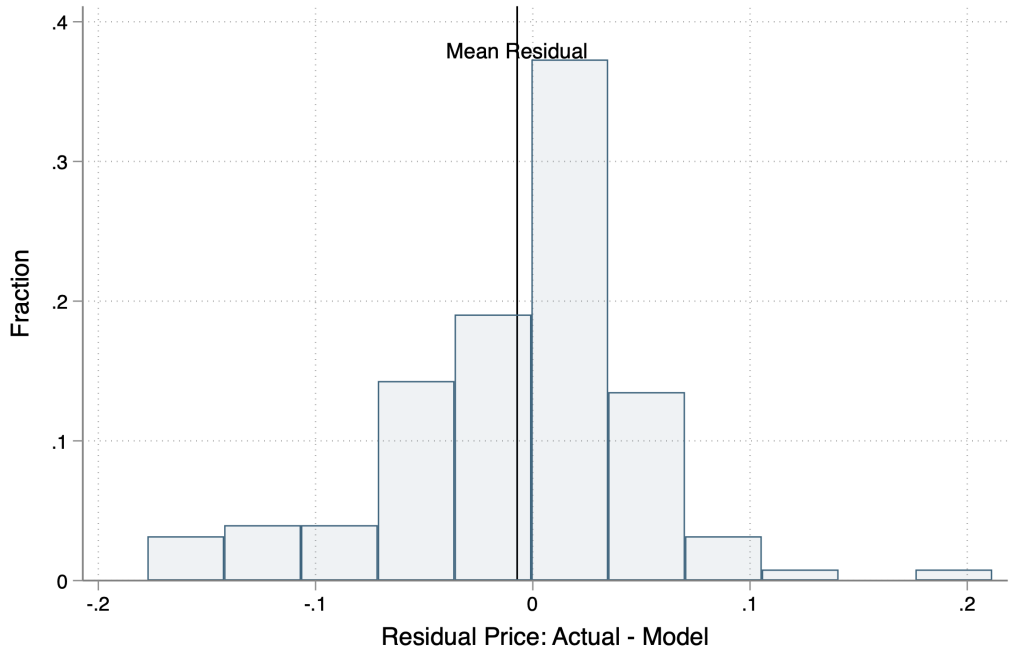
4.5.1 Model Validation: Untargeted Moment

As the estimation routine minimizes root-mean-squared-error (RMSE) between model-implied and actual prices, estimated prices could be biased. The possibility of bias emerges because minimizing RMSE targets a combination of both variance and bias. For example, an outlier high price can lead the model to increase all price estimates because RMSE penalizes large deviations heavily. Then, the mean (and median) residual price would be biased, as the model overestimates other prices to reduce the residual for the outlier high price. Increasing other prices results in residuals that are not symmetrically distributed around zero.[46] Figure 14 demonstrates that the distribution of residuals is symmetric with the mean residual near 0 at $-.007$. I conclude that the model performs well for an untargeted moment: the mean residual. Moreover, the median residual, which is also untargeted, is even closer to 0 at .002.

[45] Supplier risk aversion and buyer bargaining parameter estimates use only pre-integration data, conditional on the tuning parameter. However, the tuning parameter estimation does depend upon post-integration quarters.

[46] Simulations of this exact scenario verify that the model can yield biased average prices. The direction of the bias (positive or negative) depends on which quantities have outlier prices. For example, when outlier high prices are in the middle of the quantity distribution, then the model reduces supplier risk aversion to reduce convexity. Accordingly, the model also reduces the buyer bargaining parameter to keep average prices constant. As price levels increase on average due to the buyer bargaining parameter change, bias is positive. Conversely, when high outlier prices occur for extreme capacities, then the model overestimates supplier risk aversion to increase convexity. Correspondingly, overestimated risk aversion yields lower buyer bargaining parameters. As price levels decrease due to the parameter shifts, bias is negative.

Figure 14: Untargeted Moments: Bias



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

4.5.2 Model Validation: Out-of-Sample Fit

I also analyze out-of-sample fit, comparing the distribution of supplier RMSE in the post-integration period from the model with the analogous distribution from two alternative estimators, as shown in Table 3.

First, I compare model estimates to the mean price for the supplier in pre-integration quarters ("pre-integration mean"). Note that the pre-integration mean approximates price estimates from a constrained version of my model that abstracts from risk neutrality. This approximation reflects that the bargaining parameter is largely determined by the mean price. Therefore, a model with just a bargaining parameter would largely target the mean price. My structural bargaining model performs better than the mean pre-integration price broadly throughout the distribution of supplier RMSE, including at the minimum, the max, and all quartiles. The mean and median supplier RMSE are seven percent higher using the mean pre-integration price compared to my bargaining model.

Second, I estimate out-of-sample prices using an OLS regression of prices on capacity (including the intercept) in the pre-integration period. The OLS model has two parameters

per supplier–the coefficient on capacity and the constant term–matching the number of parameters per supplier in the bargaining model (buyer bargaining parameter and supplier risk aversion). Consequently, performance gains relative to the OLS model highlight the benefits of the structural bargaining model for fitting the data. Median and mean RMSE using the OLS estimates are, respectively, 23 and 12 percent higher than using my bargaining model. Appendix Figure A.18 shows the distribution of RMSE across suppliers for both estimators.

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 lacks an economic motivation for the tuning parameter, whereas economic intuition suggests extreme values of firm risk aversion are unlikely, my preferred OLS comparisons exclude the tuning parameter. However, there is an econometric rationale to include it–to reduce overfitting. Therefore, Appendix A.16 shows results for an OLS model that includes a tuning parameter that penalizes deviations from zero for the coefficient on capacity. The distribution of supplier RMSE for the OLS estimates with a tuning parameter is almost always between the pre-integration supplier mean and standard OLS estimates without the tuning parameter. Overall, fit improves by less than 1% relative to the pre-integration supplier mean.

4.6 Decomposing Discounts into Value Creation and Value Capture

Table 4 presents the decomposition of observed discounts from agreement suppliers into value creation and value capture. First, I determine whether a supplier is sufficiently risk averse for the buyer to threaten replacement without breaking the agreement. To do so, I calculate whether any pricing function yields weakly positive surplus for both the buyer and supplier when the buyer's outside option is the integrated supplier given estimated supplier

risk aversion. This approach reflects that the buyer does not capture value through strategic replacement threats when doing so breaks the agreement–as occurs with suppliers below the risk-aversion cutoff. With such suppliers, all observed discounts are due to value creation. Conversely, for sufficiently risk-averse suppliers, the observed discount reflects both value creation and value capture. Then, the discount when the buyer's outside option is the spot market derives from value creation, while value capture generates the additional discount when the buyer's outside option is in-house production.

Estimates illustrate that observed discounts from agreement suppliers are entirely attributable to value creation, with no evidence of value capture. This outcome reflects the low, although still present, risk aversion of agreement suppliers. As suppliers have low risk aversion, the buyer does not make replacement threats to capture value. The average discount across agreement suppliers is 1.5% of spot prices, a meaningful reduction given the nearly 7% markup in the spot market. These discounts collectively lower the buyer's annual fabric procurement costs by approximately $0.87 million. I further decompose value creation in the agreement into two components: cost reductions and profit smoothing (*i.e.,* insurance). To isolate the role of cost reductions alone, I compute discounts from a risk-neutral version of the supplier, holding all else constant. Profit smoothing then explains the remainder of the discount. Estimates show that cost reductions and profit smoothing each comprise roughly half of the discount.

To assess robustness to the one outlier supplier with high risk aversion, I compute discounts for hypothetical suppliers with *i)* median and *ii)* mean values of the supplier risk aversion and buyer bargaining parameters.[47] The primary finding–that value creation alone accounts for all observed discounts–remains unchanged. As all suppliers fall below the risk-aversion cutoff for the buyer to capture value through replacement threats, both the median and mean risk aversion also fall below the cutoff. However, both the level of the discount and the decomposition into cost reduction and profit smoothing change. The level of discount decreases because parameters have non-linear effects on discounts. For example, doubling both the risk aversion and bargaining parameter more than doubles the discount.[48] Additionally, as profit smoothing is especially valuable for risk-averse suppliers, reducing the influence of one outlier supplier with high risk aversion reduces the relative contribution of

[47] I model hypothetical suppliers as having the average agreement capacity variance of agreement suppliers. Average discounts for each agreement supplier but using the median and mean supplier risk aversion and buyer bargaining parameters are almost identical (less than .0001 p.p. different).

[48] This non-linearity has multiple sources. Two central reasons are that *i)* the Nash product is multiplicative, not additive, and *i)* doubling both parameters increases both the surplus size and the share accruing to the buyer.

insurance to the discount.[49]

Table 4: Discounts from Agreement Supplier

Supplier Type	Total Discount	Discount Decomposition:		
		Value Creation		*Value Capture*
		From Cost Reduction	*From Insurance*	
Mean Across Suppliers	1.51%	0.77%	0.74%	0.00%
Median Supplier Risk Aversion & Bargaining Parameter	0.95%	0.96%	0.00%	0.00%
Mean Supplier Risk Aversion & Bargaining Parameter	0.82%	0.79%	0.04%	0.00%

Note: Total discount is calculated as difference between spot market price and model-implied volume-weighted average agreement price. Then, I decompose the discount into value creation and value capture. The discount from value creation calculation holds the buyer's outside option fixed as the spot market, while the discount from value capture is the additional price change, if any, from the buyer as if choosing strategically whether its outside option is the spot market or in-house production. The decomposition of the discount from value creation into cost reduction finds the discount for a risk-neutral version of the supplier, holding all else constant. Any remaining discount from value creation reflects profit smoothing from insurance-like agreement terms.

4.7 Model Validation using Difference-in-Differences Estimates

To complement the validation exercises in Section 4.5, I also validate the structural model using a difference-in-differences (DiD) design. The DiD directly estimates the effects of value capture from threats to replace suppliers with in-house production. The first difference is time: before and after the buyer adds in-house capacity by building a new fabric mill. The second difference leverages heterogeneity in supplier exposure to vertical integration based on fabric production. Suppliers are classified as exposed to vertical integration (*i.e.,* "treatment" suppliers) if their pre-integration fabrics overlap with production at the integrated supplier. Therefore, the buyer could credibly threaten to switch to in-house production to displace such suppliers. Alternatively, suppliers whose pre-integration fabrics are not made at the integrated supplier are classified as unexposed (*i.e.,* "control" suppliers). Importantly, all comparisons are within supplier type, ensuring that agreement suppliers are only compared to other agreement suppliers. Recall that neither of these sources of variation help identify parameters estimated in the structural model.

[49]Given the identification argument about the shape of the agreement pricing function, I confirm that removing the most risk-averse supplier flattens, but does not eliminate, the U-shape of the agreement pricing function (Appendix Figure A.20).

Second, I leverage the DiD for validation by using event study coefficients as untargeted moments. Specifically, I show that event study coefficients using model-implied prices are consistent with those using actual prices. These event study coefficients estimated using actual prices are untargeted moments because the structural model never incorporates heterogeneity in supplier exposure to integration based on fabrics.

Last, I compare the DiD estimates with the price effects of replacement threats implied by a Nash-in-Nash with threats of replacement bargaining model in the spirit of Ho and Lee (2019)). Price decreases predicted by such a model are inconsistent with the DiD estimates. This result emphasizes the importance of the contracting environment for determining the effects of replacement threats.

4.7.1 Difference-in-Differences Estimation Approach

In a standard DiD design applied to this setting, the identifying assumption is that counterfactual potential outcomes for treatment suppliers evolve in a parallel fashion as compared to control suppliers. Identification does not require random assignment of suppliers to treatment. Indeed, as shown in Morton (2023), fabrics were not randomly selected to be integrated. Fabrics with more agreement suppliers are less likely to be brought in-house because vertical integration and agreements (relational contracts) serve as substitutes. However, suppliers are not selected into agreements based on individual fabrics (as discussed in Morton (2023)).[50] Therefore, suppliers are unlikely to be intentionally selected to be exposed to integration (even if fabrics are). Furthermore, I compare trends in pre-treatment outcomes to support the validity of the parallel trends assumption.

The data highlight that fabrics were selected to be brought in-house based on trends in fabric demand over time. The newly constructed mill is more likely to produce fabrics with increasing aggregate demand prior to the mill's construction. This endogenous fabric choice creates non-random assignment in exposure to vertical integration at the supplier level based on each supplier's fabric mix. To address this, I adopt a DiD design, detailed in Appendix A.18.2, that accommodates non-random exposure to treatment through a propensity score adjustment (Abadie, 2005). This approach weakens the parallel assumptions

[50]There are two primary reasons why economic logic makes it unlikely that a supplier would be selected for any individual fabric. First, the garment industry experiences short-run trends but relational contracts are long-run agreements; therefore, selecting a supplier for an individual fabric is a poor long-term strategy. Second, it is easier to fulfill quantity guarantees when suppliers sell a wide array of fabrics, suggesting the buyer selects suppliers based on their portfolio of fabrics. Consistent with the logic of agreement suppliers selling many distinct fabrics, the median number of distinct fabrics sold by agreement suppliers in the 6 months prior to integration is 42.

trend to conditional parallel trends; in my empirical setting, comparisons are conditional on supplier exposure to trends in fabric demand. Aggregate changes in fabric demand across all suppliers reflect evolving fashion tastes that shift demand for specific fabrics. Falsification tests show that the propensity score adjustment based on these aggregate fabric demand trends achieves balance pre-treatment. Notably, propensity scores are not calculated using supplier-level fabric demand. Therefore, these falsification tests do not simply guarantee balance by construction. Furthermore, because outcome variables are *not* explicitly matched pre-integration, pre-trends provide diagnostic information about the validity of the parallel trends assumption.[51]

4.7.2 Validation: Effects of Replacement Threats

The model generates sharp predictions for the effects of vertical integration on both prices and quantities, given supplier risk aversion estimates. As even the most risk-averse agreement supplier has risk aversion below the cutoff for the buyer to make replacement threats (see Appendix Figure A.21),the model does not predict any meaningful change in prices for agreement suppliers. Agreement quantities also should not change, with volumes reduced instead from spot suppliers (recalling that the new in-house supplier is not associated with an increase in demand from end clients).

Figure 15: DiD Estimates: Price Effects

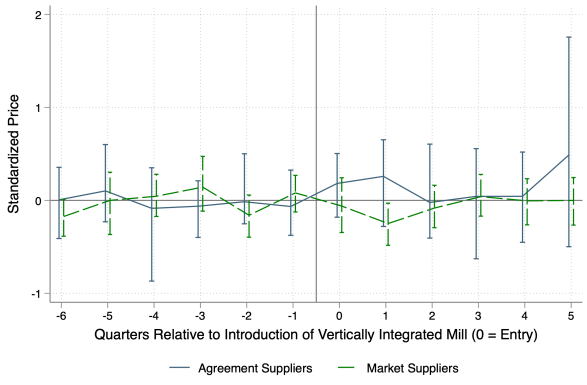
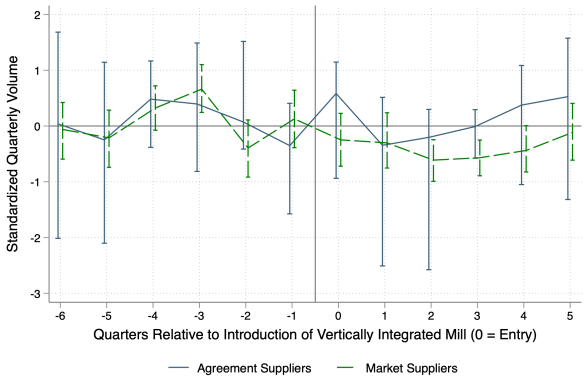


Figure 16: DiD Estimates: Quarterly Volume Effects



[51] A doubly-robust approach (as in Callaway and Sant'Anna (2021)) produces similar point estimates but is less precise as there is uncertainty about both the propensity score and the regression outcome adjustment. As controlling for propensity scores alone achieves balance, I report results for propensity score adjustment alone. Results using the doubly-robust approach are available from the author upon request.

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

The DiD results align with model predictions. Figure 15 shows that vertical integration does not affect prices for either supplier type. Figure 16 highlights that agreement supplier quantities (standardized by supplier) do not decrease but spot supplier quantities decline. The significant results for spot suppliers suggest that the DiD design is sufficiently powered to detect meaningful effects. There are no readily apparent pre-trends for either outcome variable across both supplier types.[52]

As the small number of agreement suppliers limits statistical power, I also examine the average treatment effect pooled across all pre- and post-integration periods in Table 5.[53] The pooled results show a statistically significant and economically meaningful decrease in quantities for spot suppliers. The decline averages .4 standard deviations, amounting to approximately 120,000 square meters of fabric per quarter. Conversely, agreement suppliers do not have any statistically significant treatment effect for quantities or prices. Point estimates are economically close to zero, although coefficients are not always precisely estimated. Additionally, I confirm robustness to an alternative measure of volume with desirable properties, including scale-invariant t-statistics: the quartic root (Thakral and Toh, 2023).[54]

I also verify that vertical integration does not have unanticipated effects on other outcome variables. First, I study reliability, measured as the percent of on-time deliveries. Second, I analyze transaction counts, as the decrease in spot supplier volumes could reflect fewer transactions, reallocation to orders with smaller volumes, or both. I do not find statistically significant results for any pooled coefficient for these outcome variables.[55]

[52]The most visible pre-trend is that volumes for agreement suppliers were *decreasing* pre-integration. As this pre-trend is fairly weak, and there is no clear reason why it would occur, the results provide strong evidence that agreement supplier volumes did not decrease due to vertical integration.

[53]Bootstrapped confidence intervals account for both clustering at the supplier level, as highlighted in Bertrand et al. (2004), and uncertainty in the propensity score.

[54]The positive and marginally significant effect for spot 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 at 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.

[55]Full event study plots for these outcome variables are available in Appendix A.19.1.

Table 5: Pooled Pre- and Post-Treatment Effects

	Standardized Volume	Volume (Quartic Root)	Volume-Weighted Standardized Price	Percent of Volume Delivered On-Time	Transaction Count
Agreement Suppliers Pre	.07	.85	−.02	−.02	−.28
	[−.28, .23]	[−.83, 1.57]	[−.12, .12]	[−.09, .13]	[−3.54, 7.45]
Agreement 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 semiparametric 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.

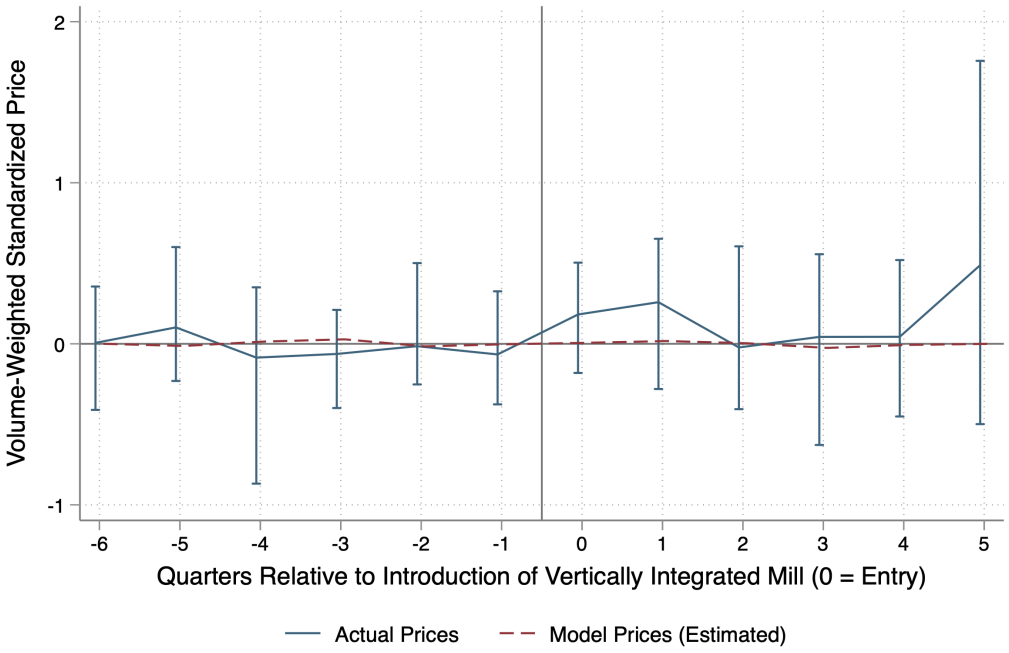
In Appendix A.19.2, I document robustness to the synthetic DiD approach proposed in Arkhangelsky et al. (2021). This methodology reproduces the results that quantities decrease only for spot suppliers and prices do not change for either agreement or spot suppliers.

4.7.3 Validation: Difference-in-Difference as Untargeted Moments

As structural model estimation does not use the variation leveraged in the DiD design,[56] I conduct a validation exercise that compares DiD event study coefficients using model-implied prices with untargeted analogous estimates using actual prices. Event study estimates using model-implied prices in Figure 17 are all within the 95% confidence intervals of event study coefficients using actual prices.

[56]Recall that DiD estimates are based on *i)* heterogeneity in exposure to vertical integration based on fabrics an *ii)* time (before vs. after the buyer builds the new mill).

Figure 17: Untargeted Moments: Difference-in-Differences

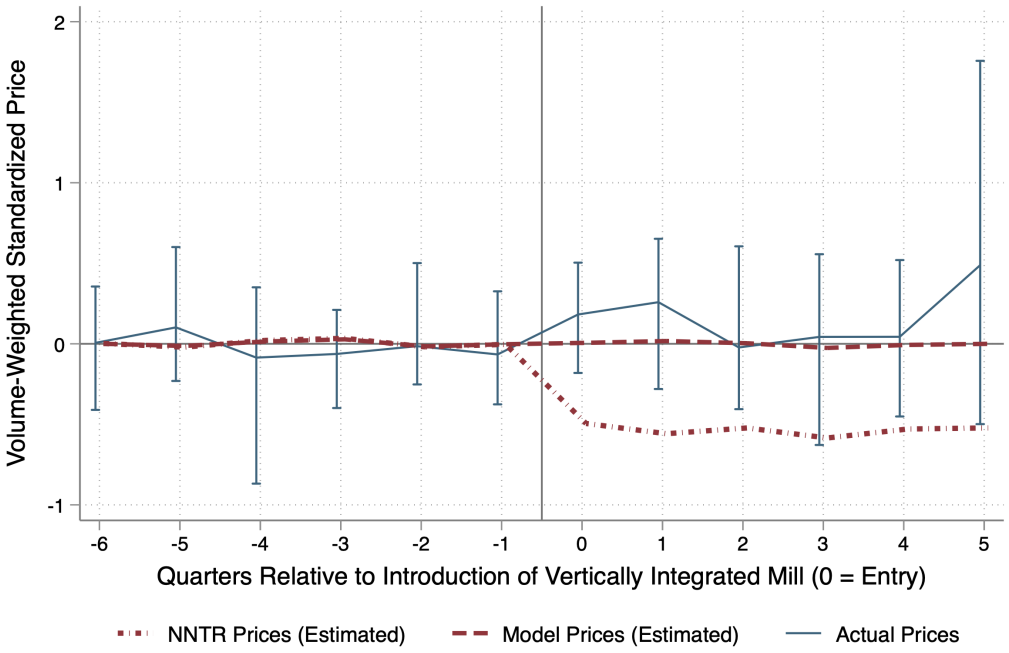


Note: Data from universe of fabric transactions by buyer. Semiparametric difference-in-differences event study estimates with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

4.8 Difference-in-Differences Estimates Inconsistent with Nash-in-Nash with Threats of Replacement Model

I compare the DiD event study estimates with prices estimated using a standard model of replacement threats: the Nash-in-Nash with Threat of Replacement (NNTR) bargaining model as proposed in Ho and Lee (2019). Ho and Lee (2019) study a bargaining environment without a spot market and without integrated supply–the market for employer-sponsored health insurance in California. Therefore, NNTR does not incorporate constraints to forming new agreements (reflecting relational enforcement) nor replacement threats from in-house production. Differences between the predictions of my model and NNTR emphasize the importance of the contracting environment for the effects of replacement threats (rather than a critique of NNTR models generally). As illustrated in Figure 18, the price decreases predicted by the NNTR model are inconsistent with the observed price changes due to value capture from replacement threats.

Figure 18: Comparison with NNTR Model



Note: Data from universe of fabric transactions by buyer. Semiparametric dfference-in-differences event study estimates with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design. Nash-in-Nash with Threats of Replacement (NNTR) prices are computed as described briefly below and in Appendix A.20.

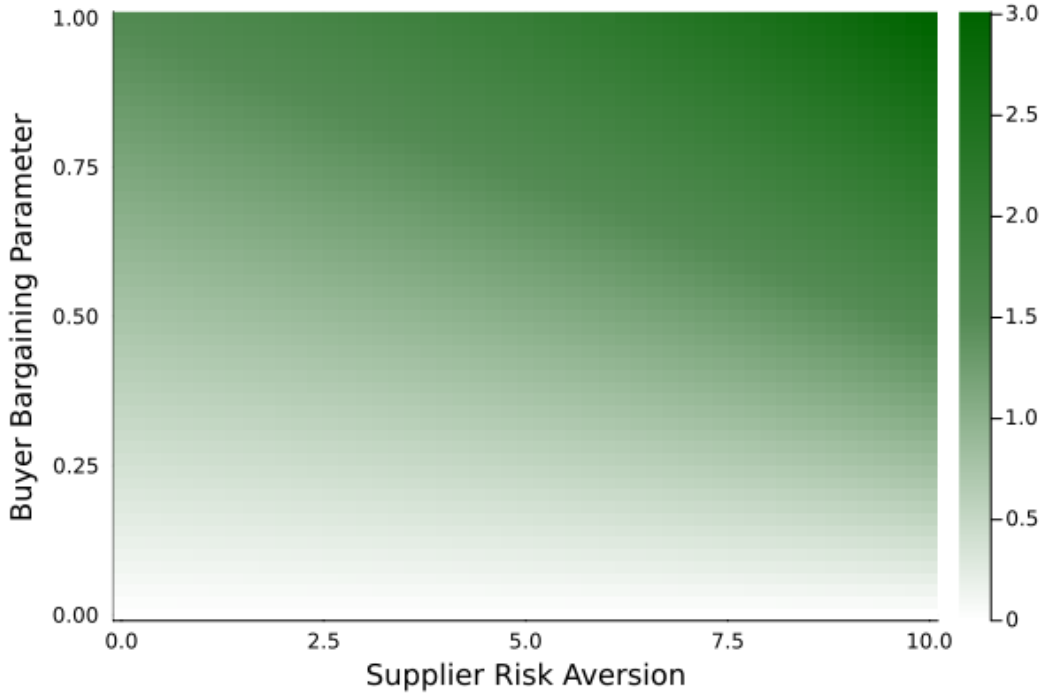
My NNTR estimation approach retains as much as possible from my model, to ensure comparability, but adopts the economic environment in Ho and Lee (2019). First, because there is no spot market in the NNTR bargaining environment, I model the buyer's outside option for not trading with any party as the loss in profit from no trade. Profit loss uses a calibrated profit margin based on Cajal-Grossi et al. (2023), who estimate margins for Bangladeshi garment manufacturers.[57] I model the integrated supplier as one alternative supplier available to the buyer, where the integrated supplier has a risk aversion parameter of zero and the buyer captures all surplus (*i.e.,* the buyer bargaining parameter is one). Additional details on the NNTR model and estimation are in Appendix A.20.

[57]Specifically, the markups as reported in Online Appendix Table E.7.

4.9 Model Estimates: Bargained Prices by Supplier Risk Aversion and Buyer Bargaining Parameter

Taking the model seriously, as validation exercises suggest the model reasonably captures bargained outcomes, I compute expected agreement prices for a wide array of supplier risk aversion and buyer bargaining parameters. This exercise facilitates studying value creation and value capture for suppliers of many different types, as observed agreement suppliers are likely larger than many policy-relevant small suppliers. Figure 19 reports the discount the buyer receives in agreements, as a percent of spot prices, for different supplier risk aversion and buyer bargaining parameters.

Figure 19: Discount in the Bargained Agreement (% of market price)



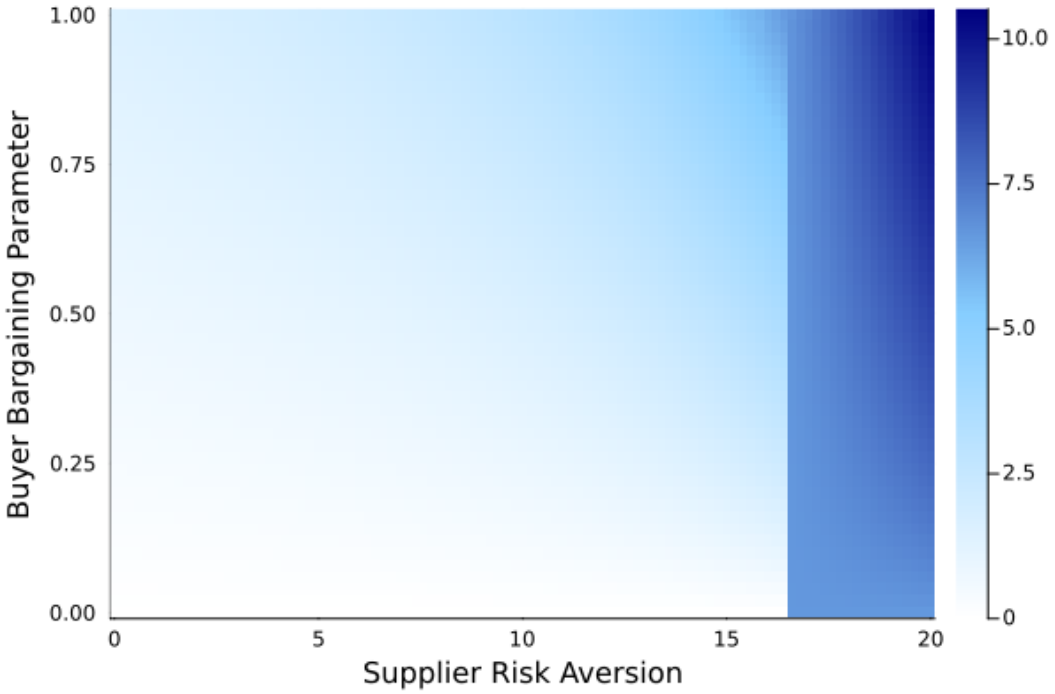
Note: Model estimates for the discount in the agreement relative to the spot market as a percent of the spot price. The heatmap legend is to the right of the heatmap and is in percentage units. Model estimates use the average capacity and quantity assurance across in-sample agreement suppliers.

Figure 19 illustrates that discounts increase in both supplier risk aversion and the buyer bargaining parameter. However, the additional discount from higher risk aversion is modest. For example, the discount for a risk-neutral supplier is less than one percentage point smaller

than the discount from a supplier at the median of agreement supplier risk aversion. Similarly, increasing risk aversion from the median to the 75th percentile increases the discount by less than one percentage point. For all parameters included in this heatmap, discounts reflect value creation, rather than value capture. This result reflects that the buyer always strategically uses the spot market as its outside option with suppliers with risk aversion parameters shown. Furthermore, as discounts do not change much with respect to risk aversion, discounts largely reflect cost reductions rather than profit smoothing.

Figure 20 replicates the analysis above, but for a larger set of risk supplier aversion parameters to show results for suppliers sufficiently risk averse for the buyer to capture value through strategic replacement threats.

Figure 20: Discount in the Bargained Agreement (% of market price)



Note: Model estimates for the discount in the agreement relative to the spot market as a percent of the spot price. The heatmap legend is to the right of the heatmap and is in percentage units. Model estimates use the average capacity and quantity assurance across in-sample agreement suppliers.

The minimum estimated level of risk aversion for the buyer to make replacement threats is 16.6, slightly exceeding the upper bound of estimates in the literature (12.9 in Gandelman and Hernandez-Murillo (2014)). This finding underscores the strong constraints imposed by the contracting environment on the effectiveness of replacement threats. However, the

relevance of replacement threats may vary across settings, as the risk-aversion threshold is not a structural parameter but a function of them. Specifically, the cutoff level of risk aversion decreases when cost functions are more convex or agreements provide more risk reduction relative to the spot market (*i.e.,* more quantity assurance). Quantity assurance in agreements can create more value by mitigating risk in both demand and supply, as occurs for agricultural workers facing productivity risk in Jayachandran (2006).[58] LMIC agricultural settings with smallholder farmers may satisfy these conditions. Furthermore, smallholder farmers may rely on bargained agreements (Macchiavello and Morjaria, 2020). Such farmers presumably have highly convex production costs given that a key input is largely fixed (land) and are likely more risk-averse than agreement suppliers in this study.

Second, the discount increases meaningfully once supplier risk aversion surpasses the threshold where the buyer makes replacement threats. For example, when the buyer bargaining parameter implies agreement value is shared equally between the parties, replacement threats increase the discount by 76% from 3.8% to 6.7% of spot prices. This difference is economically meaningful, as bargained prices for agreement suppliers average only 1.5% of spot prices. At 6.7%, the discount exceeds the 5% Indian VAT for apparel during the study period.[59] Additionally, the discount is large relative to markups in the spot market where suppliers hold up the buyer, with spot prices estimated at 107.36% of costs.

Appendix A.21 incorporates relational enforcement for a conservative calibration of the discount factor, documenting negligible differences in estimated prices. This result suggests that suppliers are at least as patient as the calibrated discount factor, which is based on microfinance interest rates. A high discount factor for agreement suppliers is consistent with low estimates of supplier risk aversion, as both suggest that agreement suppliers are large firms. [60]

[58] Adverse productivity shocks both reduce demand for worker labor and increase supply from competing laborers, magnifying the variance in income caused by productivity shocks.

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

[60] The model features a cutoff level of the discount factor above which dynamic incentive compatibility constraints are slack in this model. Higher discount factors above this cutoff do not affect prices. When dynamic incentive compatibility constraints are not slack, they "flatten" the agreement pricing function to ensure that relational enforcement holds. Flattening the U-shaped pricing function, which reduces insurance provision in the agreement, diminishes the difference between agreement and spot prices. Since differences between agreement and spot prices create static incentives for one firm to deviate (for the supplier when agreement prices are below spot prices and for the buyer when agreement prices exceed spot prices), flatter pricing functions can satisfy relational enforcement constraints even with lower discount factors.

5 Policy Evaluation: Increasing Prices Suppliers Receive

Both value creation from non-price agreement terms and value capture by strategic threats to replace suppliers with in-house production can result in low average agreement prices paid to small, risk-averse suppliers. These low prices raise distributional concerns about surplus sharing, not only within global value chains (GVCs) but broadly across the firm size distribution as the most risk-averse firms are also likely the smallest. Distributional implications motivate policy remedies, consistent with omnipresent policy initiatives targeting small and medium enterprises (SMEs), such as the World Bank's policy advising on finance tools for SMEs[61] and the IMF's 2019 conference on SME financing.[62] Policymaker focus on small firms may reflect their significant representation in the firm size distribution and role as primary employers in low-income countries (Hsieh and Olken, 2014). Prioritizing small firms may also align with preferences to reduce inequality, as redistributing profit to small firms could increase aggregate welfare when low profit and income imply high marginal utility. Furthermore, reallocating to small firms could improve efficiency dynamically when credit and savings market frictions lead small firms to rely on current profits to finance growth.

I evaluate two policies to redistribute surplus to small, risk-averse suppliers using the structural model. First, motivated by concerns about market power of the few, large buyers (exporters), I increase competition between them for output from suppliers. Second, given the importance on missing markets in development economics, I create the missing insurance product so risk-averse firms can smooth profit outside the agreement.

5.1 Increase Downstream Buyer Competition

I first study the effect of increasing downstream buyer competition, estimating agreement prices when additional buyers compete to form agreements with suppliers. Policymakers can support the development of large exporting firms through many initiatives, including reducing red tape and other administrative costs associated with exporting, facilitating connections with international end clients ("export promotion"), reducing export tariffs, and even subsidizing firms to export.

I model increased downstream buyer competition as reducing the buyer bargaining pa-

[61] World Bank SME Finance
[62] IMF SME Financing

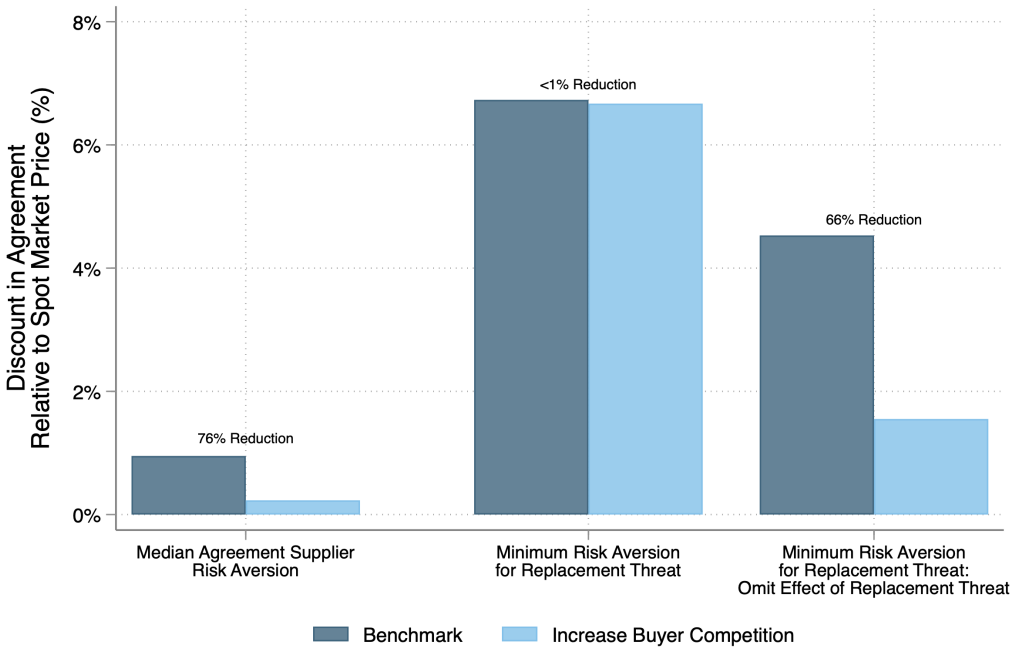
rameter.[63] While competition from additional buyers should improve suppliers' outside options in the long-run, competition does not change suppliers' short-run outside option of selling on the spot market. The short-run outside option does not change because suppliers cannot quickly establish a new agreement with relational enforcement. Relational enforcement requires time and effort to develop and maintain Cajal-Grossi et al. (2023); Brugues (2020). Therefore, the buyer responds to competition by offering suppliers a larger share of surplus to ensure their loyalty. This mechanism mirrors efficiency wages, as both incentivize loyalty through a higher share of surplus (Katz, 1986).

Suggestive empirical patterns support modeling buyer competition as reducing the buyer bargaining parameter; specifically, more valuable suppliers receive larger surplus shares (*i.e.,* have lower estimated buyer bargaining parameters).[64] Agreement suppliers producing fabrics with higher concentration, compared to other agreement suppliers, are likely more valuable to the buyer. For example, when the buyer seeks to procure a specific fabric at a discount, it prefers sourcing that fabric from an agreement supplier. Additionally, supplier fabric concentration relates to how easily the buyer can meet quantity targets in agreements. To illustrate, consider the extreme case where all suppliers sell the same fabric. Then, when the aggregate quantity for that fabric is low, the buyer cannot provide quantity assurance to all suppliers because assigning an order to one supplier precludes assigning that order to another supplier. It follows that the buyer more easily fulfills the agreement's quantity assurance terms with suppliers in relatively unconcentrated fabric markets (relative to other agreement suppliers).

[63] Specifically, I halve the buyer bargaining parameter.

[64] Appendix Section A.23 describes the approach to establish that supplier concentration in fabric markets, relative to other agreement suppliers, is significantly negatively correlated with estimated buyer bargaining parameters. Additionally, this proxy for agreement supplier value is not significantly correlated with estimated risk aversion parameters, mitigating concerns that the correlation between supplier value and buyer bargaining parameters is spurious.

Figure 21: Downstream Buyer Competition and Agreement Prices



Note: Model estimates. Figure shows change in agreement discount from halving the buyer bargaining parameter relative to the estimated median for agreement suppliers. Risk aversion for the median agreement supplier is the median of estimated risk aversion parameters. Risk aversion for the risk-averse supplier is the minimum risk aversion such that the buyer strategically decides to make replacement threats. Model estimates use the capacity and quantity assurance for the median agreement supplier.

Figure 21 illustrates that increased downstream buyer competition meaningfully reduces the discount from agreement suppliers with median risk aversion but has negligible effects for more risk-averse suppliers. Increased competition decreases agreement discounts for a supplier with median risk aversion by 76%, from .95% to 0.23%, where the discount is defined as the difference between the average volume-weighted agreement price and the spot price, as a percent of the spot price. However, for suppliers with the minimum risk aversion for the buyer to use in-house supply as a replacement threat, the discount falls by less than 1%, from 6.73% to 6.67%. Omitting value capture from replacement threats *overestimates* the benefits of increased competition. Without value capture, increased competition has much larger effects, reducing the discount by 66%–from 4.5% to 1.5%.

Intuitively, increasing downstream buyer competition proves ineffective as a policy because, once the buyer makes a replacement threat, the value created is small relative to the buyer's participation constraint. When the apparent surplus is small, changing the sharing

rule has minimal effects on prices. Not only does this intuition clarify why abstracting from value capture generates incorrect policy conclusions for small risk-averse firms, but it highlights why this policy is effective at increasing prices paid to large suppliers with low risk aversion. When the buyer's threat point is the spot market, there is enough value created by the agreement that changing surplus sharing meaningfully increases agreement prices.[65]

5.2 Create Missing Insurance Product

The importance of missing markets in LMIC settings, and the limited impact of increasing downstream buyer competition, motivate analyzing the effects of adding the missing profit insurance product. I evaluate how access to actuarially fair insurance for profits from the spot market changes agreement discounts. This insurance product mitigates how supplier risk aversion increases discounts by eliminating profit risk in the spot market.[66] As all uncertainty in this economic environment arises from capacity and costs, insurance over prices (or, equivalently, revenue) would achieve the same effect.[67]

Although governments are unlikely to directly offer such insurance given moral hazard concerns, policymakers have many tools to provide limited insurance.[68] Policies that enhance access to financial markets can enable small firms to hedge risk without relying on the agreement. Specifically, increasing credit access can assist suppliers with borrowing to cover expenses during low profit periods rather than accepting low agreement prices. Policies that enable small firms to receive credit at favorable rates are common, even in high-income countries. Beyond credit, improving access to savings can support suppliers to accumulate reserves over time, effectively allowing them to self-insure. Furthermore, policies can support markets for derivatives or other financial products to enable direct risk hedging, though

[65] Appendix A.22 considers the effects of horizontal mergers of suppliers. Although horizontal consolidation does both reduce the value created by the agreement, as the merged entity is less risk averse, and restrain the buyer's ability to capture value through threats of replacement, this counterfactual effectively invokes a suite of unspecified policies that remove the myriad frictions that lead to the existence of many small constrained firms that motivates many distributional concerns.

[66] Blouin and Macchiavello (2019) also explore how missing insurance markets interact with contracting behaviors, highlighting how suppliers' incentives to strategically default leads trading partners to use index-priced contracts to remove such temptations. However, index-price contracts do not provide insurance, meaning that strategic default incentives reduce insurance provision. While Blouin and Macchiavello (2019) focus on the efficiency implications of the missing insurance markets, I focus on effects on surplus creation and distribution.

[67] However, the insurance contract over prices or revenue would need to be non-linear, as profits are non-linear in capacity due to cost convexity.

[68] Many observable policies effectively provide some limited forms of insurance. For example, common price floors for agricultural goods insure producers against adverse prices.

financial literacy and transaction costs may constrain take-up. Additionally, policies can support trade associations and cooperatives to provide price or revenue insurance. These organizations may be able to mitigate moral hazard, such as by indexing the quantity of insurance offered to a running average of historic capacity or revenue.

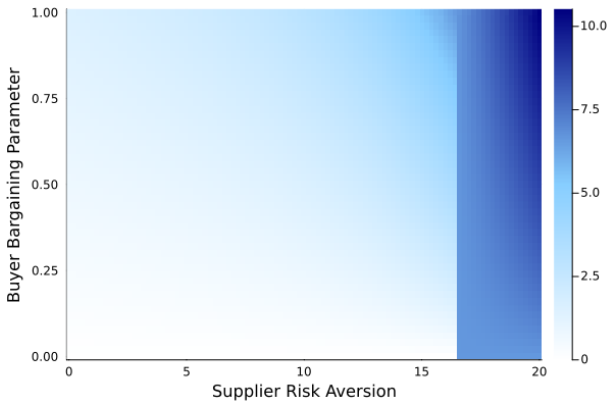


Figure 22: Without Insurance

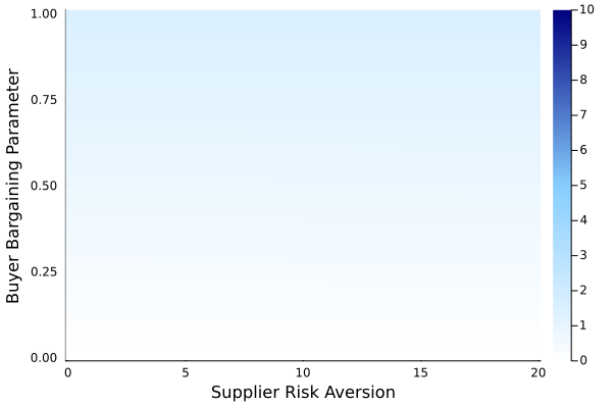


Figure 23: With Insurance

Note: Model estimates for the bargained 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 insurance product reduces discounts and eliminates the buyer's incentives to capture value through strategic replacement threats. As the buyer cannot capture value through replacement threats, risk-averse suppliers no longer receive much lower prices. Intuitively, this result reflects that an effective approach to counteract the buyer's improved bargaining position from the change in its outside option is symmetric: strengthen the supplier's bargaining position by improving its outside option.

6 Conclusion

This paper studies surplus creation and distribution in an empirically relevant LMIC domestic supply chain. I analyze how bargained outcomes depend on both value created by non-price agreement terms that mitigate spot market frictions and value capture from threats to replace trading partners with in-house production. I write down a theoretical bargaining model to formalize how value creation and capture affect agreement prices. The model incorporates the agreement's empirically observable insurance-like terms that create value: a U-shaped pricing function and quantity assurance. In the model, value creation can lower

average agreement prices because the buyer receives a share of the value created as a discount. Additionally, the model shows that the buyer's bargaining position can improve by using integrated capacity as a threat to replace external suppliers, consistent with standard bargaining models. However, such threats can break agreements with suppliers with low risk aversion, as the agreement may not create value relative to in-house production with such suppliers. Then, the buyer must increase expensive spot sourcing given contracting frictions that prevent forming new agreements and cost convexity. Therefore, the buyer as if strategically chooses to capture value through replacement threats, but only from sufficiently risk-averse suppliers that agreements do not break. Overall, this analysis underscores how the contracting environment constrains the strategic use of replacement threats.

I build and estimate a structural model closely tied to the theoretical model to decompose observed discounts into value creation and capture using supplier-by-quarter panel data. The key identification challenge is to separate supplier risk aversion from buyer bargaining parameters. To do so, I leverage exogenous demand shocks, as the price response to demand shocks increases in supplier risk aversion. Model estimates illustrate that value creation, rather than value capture, explains the entirety of observed discounts. This result derives from low estimated risk aversion for agreement suppliers. The model also illustrates that contracting frictions restrain the buyer's ability to capture value through replacement threats. Namely, the supplier risk aversion cutoff for the buyer to make such threats is not only higher than all estimated risk aversion coefficients, but also high relative to estimates from other settings. In agreements with suppliers with risk aversion above the cutoff, such that the buyer makes replacement threats, value capture results in discounts that reach 6.7% of spot prices[69]–larger than Indian VAT for apparel at the time. As small firms are both economically and policy relevant as the majority of firms and employers in LMIC and receive the lowest prices due to both value creation and capture, I evaluate two policies to reallocate surplus to them. First, motivated by concerns about limited competition among the few big exporters in GVCs, I analyze the effects of increasing buyer competition. Although this policy helps suppliers with risk aversion similar to agreement suppliers, it has a trivial effect for the most risk-averse firms. Importantly, not accounting for the effects of value capture by replacement threats 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 insurance product that allows small risk-averse firms to smooth profit without relying on the bargained

[69]Estimates use a buyer bargaining parameter of .5

agreement. This policy increases prices for all suppliers and eliminates value capture through strategic replacement threats.

These results contribute to the growing evidence that challenges to smoothing, especially relevant in LMIC settings, not only decrease welfare by generating different marginal utility levels across states, but also have important price effects. For example, LMIC agricultural producers buy and sell at the same time as other nearby producers given challenges storing agricultural output. Simultaneous sales generate negative price effects–producers "sell low and buy high" (Burke et al., 2018). Similarly, challenges smoothing over productivity shocks leads wage workers in agricultural settings to increase labor supply simultaneously, magnifying the decrease in their wages (Jayachandran, 2006). In my model, small risk-averse firms accept lower average prices in order to smooth profits becuase of missing risk-smoothing institutions. Together, these results highlight that policies that address exposure to risk may have both positive *efficiency* effects–by equating marginal utility (and revenue, etc.) across states–but also benefit *equity* by improving the outside option for those most exposed to risk.

References

Abadie, Alberto, "Semiparametric Difference-in-Differences Estimators," *The Review of Economic Studies*, 01 2005, *72* (1), 1–19.

Adhvaryu, Achyuta, 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.

Backus, Matthew, Thomas Blake, and Steven Tadelis, "On the Empirical Content of Cheap-Talk Signaling: An Application to Bargaining," *Journal of Political Economy*, 2019, *127* (4), 1599–1628.

_ , _ , Brad Larsen, and Steven Tadelis, "Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions," *The Quarterly Journal of Economics*, 02 2020, *135* (3), 1319–1361.

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.

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.

Dominguez-Iino, Tomas, *Efficiency and Redistribution in Environmental Policy: An Equilibrium Analysis of Agricultural Supply Chains*, Working Paper, 2023.

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.

Freyberger, Joachim and Bradley Larsen, "How Well Does Bargaining Work in Consumer Markets? A Robust Bounds Approach," *Working Paper*, 2023.

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.

Ghili, Soheil, "Network Formation and Bargaining in Vertical Markets: The Case of Narrow Networks in Health Insurance," *Marketing Science*, 2022, *41* (3), 501–527.

Grennan, Matthew, "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices," *American Economic Review*, February 2013, *103* (1), 145–77.

Hardy, Morgan, Gisella Kagy, and Lena Song, "Gotta Have Money to Make Money? Bargaining Behavior and Financial Need of Microentrepreneurs," *American Economic Review: Insights*, March 2022, *4* (1), 1–17.

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.

Ho, Kate and Robin S. Lee, "Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets," *American Economic Review*, February 2019, *109* (2), 473–522.

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

International Labour Organization, *Employment, wages and productivity trends in the Asian garment sector: data and policy insights for the future of work* 2022.

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.

Keniston, Daniel, "Bargaining and Welfare: A Dynamic Structural Analysis of the Autorickshaw Market," *Working Paper*, 2017.

_ , Bradley Larsen, Shengwu Li, J.J. Prescott, Bernardo Silveira, and Chuan Yu, "Fairness in Incomplete Information Bargaining: Fairness in Incomplete Information Bargaining: Theory and Widespread Evidence from the Field," *Working Paper*, 2022.

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

Larsen, Bradley and Anthony Lee Zhang, "Quantifying Bargaining Power Under Incomplete Information: A Supply-Side Analysis of the Used-Car Industry," *Working Paper*, 2021.

Lee, Robin S., Michael D. Whinston, and Ali Yurukoglu, *Handbook of Industrial Organization (Structural Empirical Analysis of Contracting in Vertical Markets)*, Vol. 4, Elsevier, 2021.

Levin, Jonathan, "Relational Incentive Contracts," *American Economic Review*, June 2003, *93* (3).

Lind, Jo Thori and Halvor Mehlum, "With or Without U? The Appropriate Test for a U-Shaped Relationship," *Oxford Bulletin of Economics and Statistics*, 2010, *72* (1), 109–118.

Macchiavello, Rocco, "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.

Miller, David A. and Joel Watson, "A Theory of Disagreement in Repeated Games With Bargaining," *Econometrica*, 2013, *81* (6), 2303–2350.

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.

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.

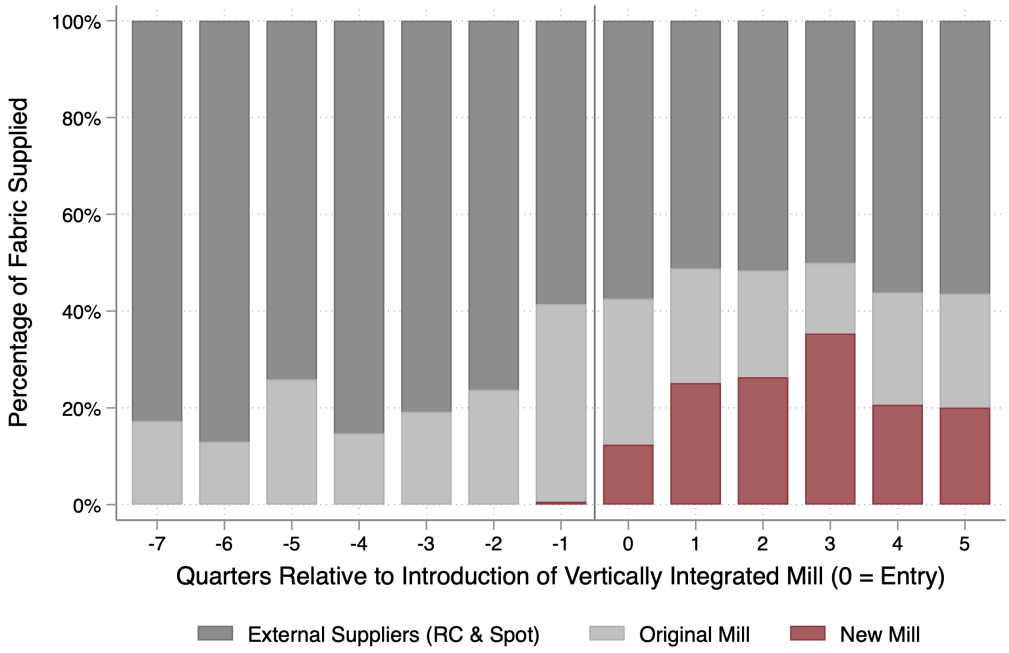
World Bank, *Trading For Development In the Age of Global Value Chains* 2020.

Zavala, Lucas, "Unfair Trade? Monopsony Power in Agricultural Value Chains," *Working Paper*, 2022.

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 the buyer. Bars represent the percentages of fabric volumes purchased during the quarter from each supplier type.

A.2 Fabric Procurement Agreements

A.2.1 Price Function Convexity

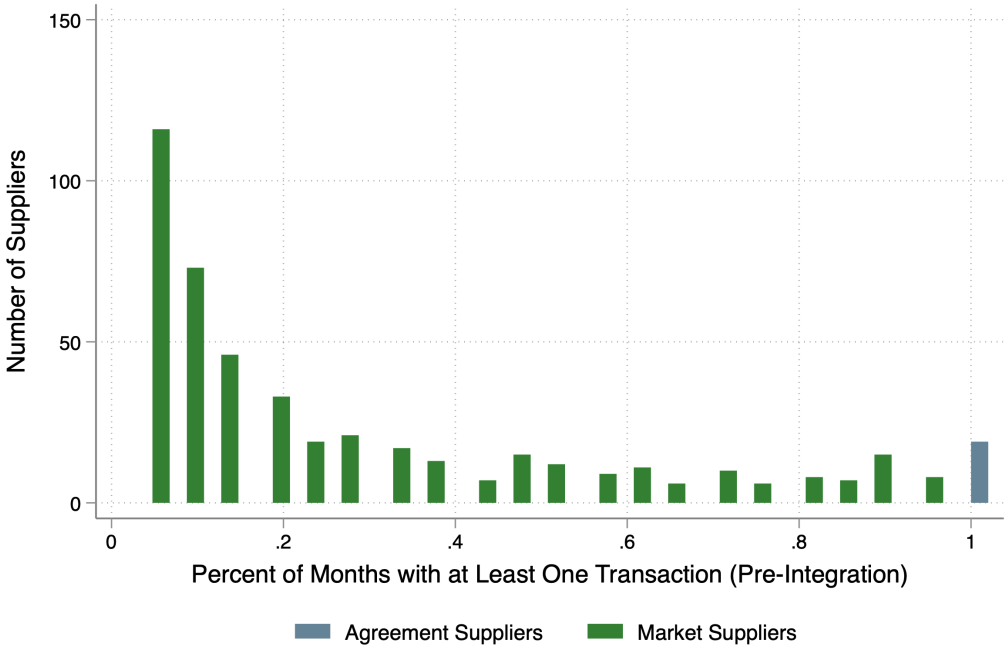
I formally test the U-shape of the pricing function following the methodological approach in Lind and Mehlum (2010). The approach includes four tests. First, the coefficient on the quadratic term must be positive and significant. Second, the slope at the minimum must be negative and significant. Third, the slope at the maximum must be positive and significant. Fourth, the implied minimum must be in the interior of the domain. Note that, in my setting, as the minimum price should be at target capacity, it is natural to expect that the minimum price would occur at a capacity of one. Table A.1 shows that the prices for agreement suppliers meet the criteria for being U-shaped, while prices for spot suppliers do not.

Table A.1: Pricing Function Convexity by Supplier Type

	Volume-Weighted Standardized Price					
Capacity	−26.438**	−.229	−36.723**	−.755	−32.61**	−.79
	(10.986)	(.546)	(14.455)	(.717)	(14.646)	(.721)
Capacity^2	12.95**	.043	18.113**	.311	16.169**	.331
	(5.562)	(.27)	(7.16)	(.359)	(7.252)	(.362)
Supplier Type	Agreement	Spot	Agreement	Spot	Agreement	Spot
Supplier FE	N	N	Y	Y	Y	Y
Quarter FE	N	N	N	N	Y	Y
Implied Quantity for Minimum Price	1.021	2.655	1.014	1.212	1.008	1.193
95% CI	[.981, 1.06]	[−17.697, 23.008]	[.994, 1.033]	[.676, 1.747]	[.984, 1.033]	[.715, 1.67]
Slope at Minimum Capacity	−4.055	−.192	−5.415	−.492	−4.662	−.511
95% CI	[−6.828, −1.282]	[−.823, .438]	[−9.541, −1.29]	[−1.31, .326]	[−8.865, −.459]	[−1.331, .309]
Slope at Maximum Capacity	2.914	−.111	4.333	.097	4.04	.116
95% CI	[−.349, 6.178]	[−.52, .298]	[.778, 7.887]	[−.445, .639]	[.431, 7.648]	[−.434, .666]
R^2	.034	.003	.462	.326	.497	.329
N	133	1470	133	1333	133	1333

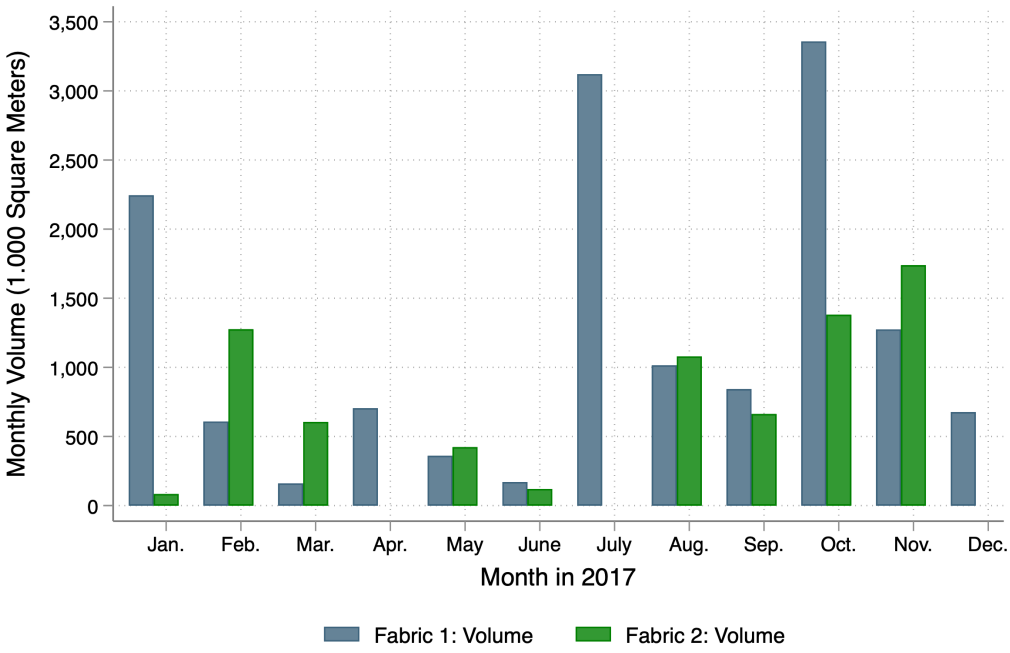
Note: Data from universe of fabric transactions by the buyer, transformed into supplier-by-quarter panel data. Prices are the volume-weighted standardized prices.

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



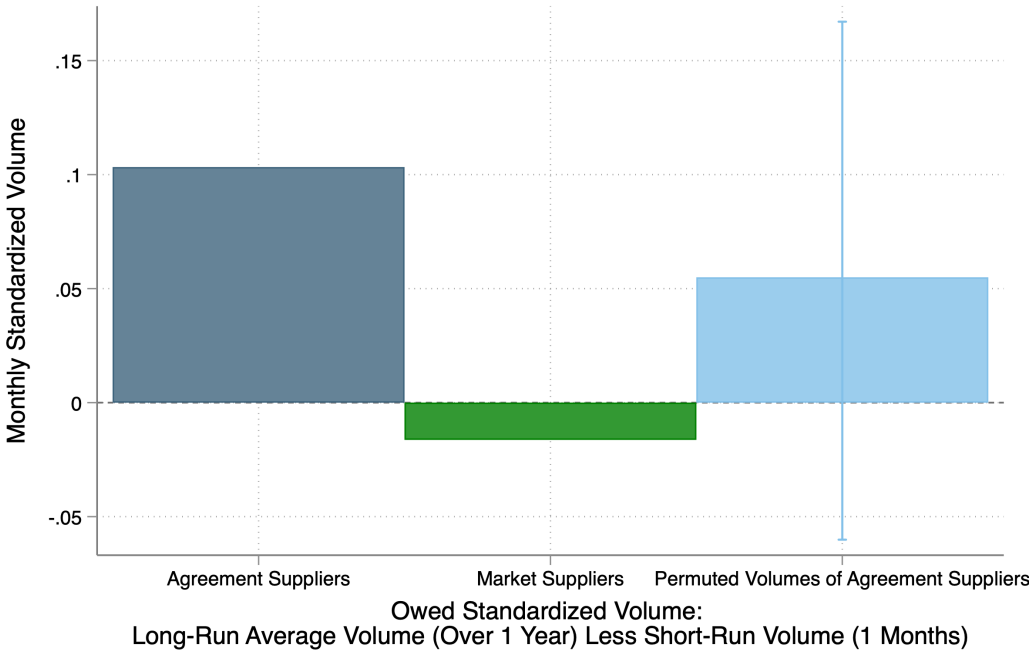
Note: Data from universe of fabric transactions by the buyer. Bars represent the number of suppliers with the relevant percent of months with transactions.

Figure A.3: Variation in Monthly Quantity by Fabric



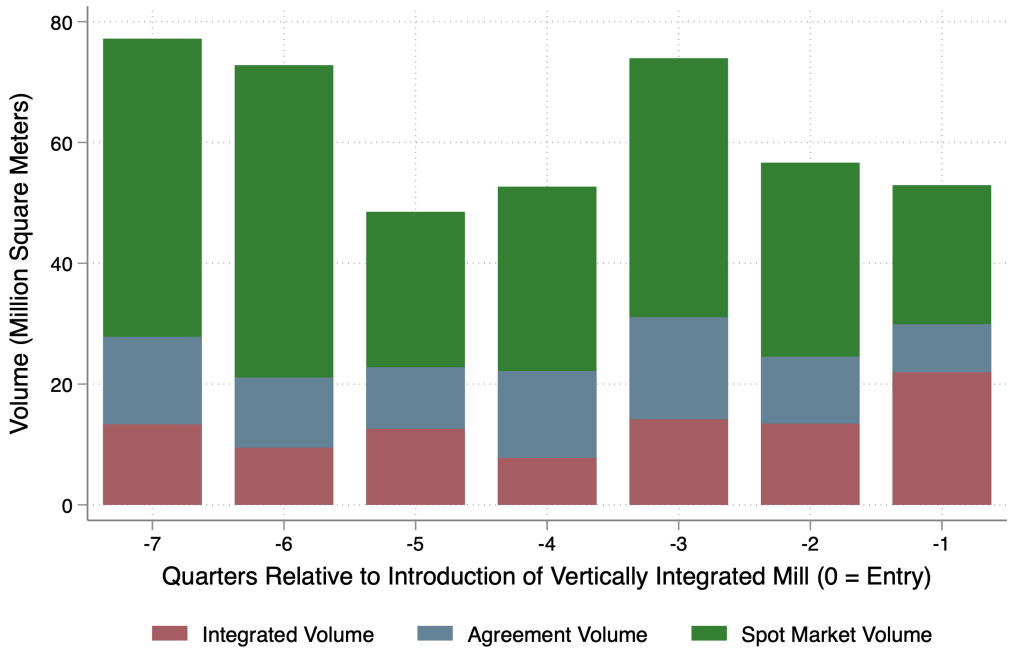
Note: Data from universe of fabric transactions by the buyer. Bars represent the quantity of fabric purchased during the month in 2017 for two selected fabrics. Fabrics were selected based on data availability to illustrate variation.

Figure A.4: Coefficient on Owed Quantity by Supplier Type



Note: Data from universe of fabric transactions by the buyer. Bars show the coefficient for owed standardized volume from a regression of current volume on owed volume 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 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 agreement suppliers represents the mean across permutations; the 90% empirical CI is also shown.

Figure A.5: Volume by Supplier Type (Pre-Integration)



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

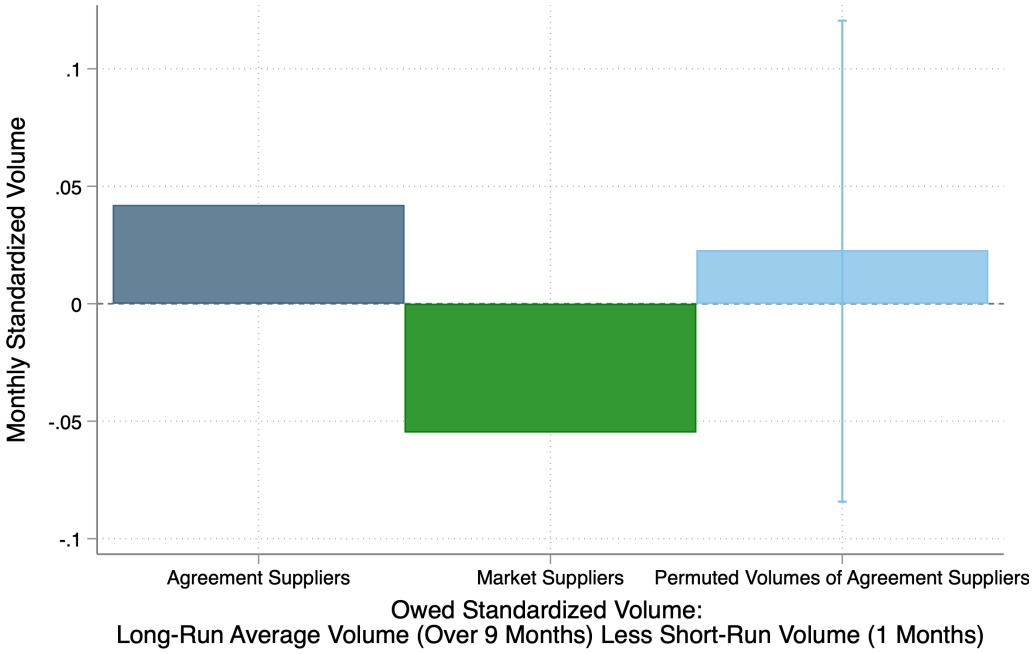
A.3 Owed Quantity Guarantees

In addition to showing that agreement suppliers have more stable quantities in general, I explore the association between a proxy for the target quantity owed in the agreement and the volume received in the current period. For agreement suppliers, owed quantity should be positively correlated with current volume: when the agreement owes more volume to the supplier in order to stabilize volumes (*i.e.,* quantities so far in the quarter are below the long-run average), 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 agreement with month fixed effects. A positive association between owed volume and current volume would be consistent with the hypothesized agreement. Long-run monthly average volume serves as a proxy for the quantity guarantee. Owed quantity is then measured as the difference between the proxy for the long-run quantity guarantee and the quantity fulfilled recently (*i.e.,* in the prior month). Alternatively, as there is no agreement with spot market suppliers, there is no reason to expect a relationship between owed volume and current volume for spot market suppliers.

Figure A.6 illustrates that a one standard deviation increase in owed volume for an

agreement supplier is associated with a .04 standard deviation increase in current volume.[70] 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.6: Coefficient on Owed Quantity and Supplier Type



Note: Data from universe of fabric transactions by the buyer. Bars show the coefficient for owed standardized volume from a regression of current volume on owed volume 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 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 agreement 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 agreement suppliers over time, which should also be mean reverting by construction. Additionally, the most reasonable alternative mean-reverting contracting strategy distinct from the agreement 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

[70]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 shifters that are incorporated into the agreement but not the analysis.

could potentially pass on some cost savings to the buyer. However, this strategy would likely 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.4).

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 agreement suppliers and the restriction to only using pre-integration data to avoid incorporating any potential effects of integration on the agreement.

A.4 Why Not an Explicit Contract?

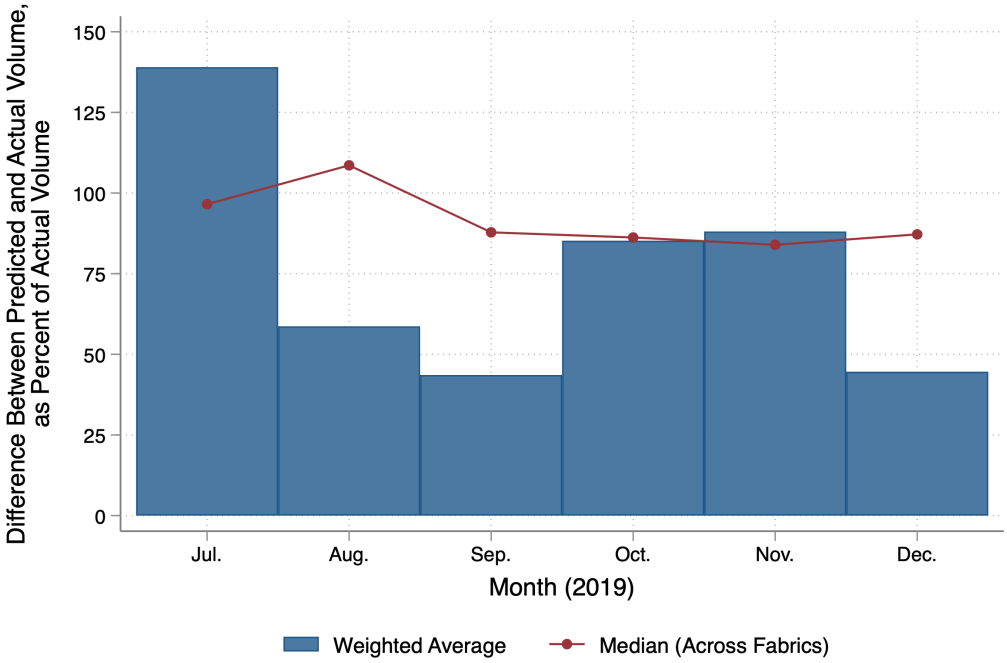
There are (at least) four factors that can prevent explicit contracts from resolving the hold-up problem in this setting. First, if the explicit contract is not enforceable because of weak contracting institutions, then the explicit contract is effectively irrelevant. Research in India suggests that contracting parties may be concerned that courts are 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 both supplier and industry capacity utilization. A contract with constant prices that does 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). 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 that fabric suppliers receive from other buyers, such a contract would likely be difficult (or at least expensive) to write, even abstracting from enforcement concerns. The contract would be challenging to specify because it would require agreeing to a 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 Agreement Suppliers?

Given the evidence that agreement suppliers offer lower prices than market suppliers, the buyer would minimize costs by sourcing exclusively through agreements. However, although agreement suppliers receive 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 agreements is the challenge of fulfilling owed quantity guarantees in this setting (*i.e.,* challenges to the buyer in meeting its obligations under the agreement). 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 the buyer can direct to any individual supplier.

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 agreements. The data highlight the volatility and unpredictability of volume, especially at the fabric level. Figure A.7 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 average monthly volume. 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 larger 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 agrement and integrated suppliers appears to be quite stable pre-integration, with market suppliers absorbing stochastic variation in quantities. This "buffer-stock" sourcing approach is visible in A.5.

Figure A.7: Forecast Error in Volume Predictions



Note: Data from universe of fabric transactions by the 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.8: Pricing Function (Placebo Comparison): Binned

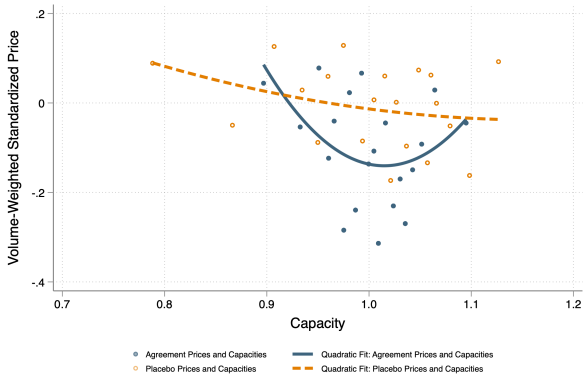
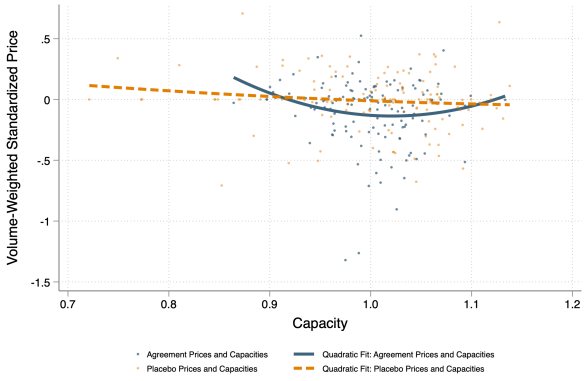


Figure A.9: Pricing Function (Placebo Comparison): Raw



Note: Data from universe of fabric transactions by the buyer, transformed into supplier-by-quarter panel data. Quadratic fit is from a regression of volume-weighted standardized price on capacity and capacity squared. Figure A.8 bins prices and capacities using 20 quantiles. The mapping from quantities to capacities is as discussed in 4.2.2.

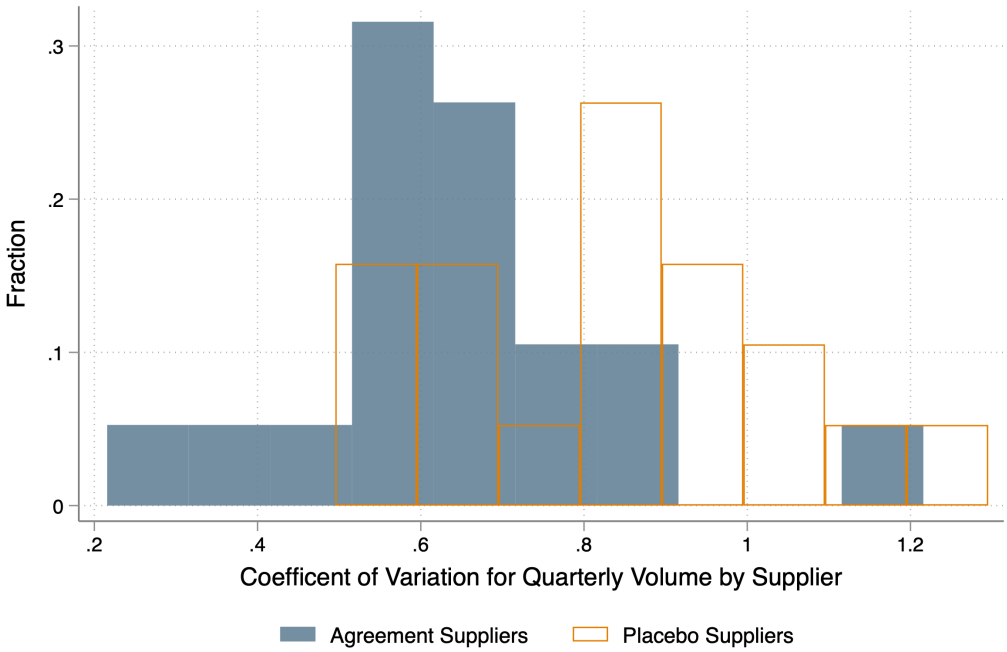
Table A.2: Pricing Function Convexity by Supplier Type (Placebo Comparison)

	Volume-Weighted Standardized Price					
Capacity	−26.438**	−1.345	−36.723**	−6.454	−32.61**	−7.361
	(10.986)	(3.682)	(14.455)	(4.495)	(14.646)	(4.643)
Capacity^2	12.95**	.519	18.113**	3.182	16.169**	3.663
	(5.562)	(1.952)	(7.16)	(2.391)	(7.252)	(2.471)
Supplier Type	Agreement	Placebo	Agreement	Placebo	Agreement	Placebo
Supplier FE	N	N	Y	Y	Y	Y
Quarter FE	N	N	N	N	Y	Y
Implied Quantity for Minimum Price	1.021	1.295	1.014	1.014	1.008	1.005
95% CI	[.981, 1.06]	[−1.365, 3.956]	[.994, 1.033]	[.88, 1.148]	[.984, 1.033]	[.894, 1.116]
Slope at Minimum Capacity	−4.055	−.596	−5.415	−1.864	−4.662	−2.079
95% CI	[−6.828, −1.282]	[−2.354, 1.162]	[−9.541, −1.29]	[−3.958, .23]	[−8.865, −.459]	[−4.236, .079]
Slope at Maximum Capacity	2.914	−.164	4.333	.788	4.04	.973
95% CI	[−.349, 6.178]	[−1.756, 1.429]	[.778, 7.887]	[−1.14, 2.715]	[.431, 7.648]	[−1.024, 2.97]
R^2	.034	.015	.462	.25	.497	.281
N	133	130	133	130	133	130

Note: Data from universe of fabric transactions by the buyer, transformed into supplier-by-quarter panel data, restricting spot suppliers to the subset of placebo suppliers. Prices are the volume-weighted standardized prices.

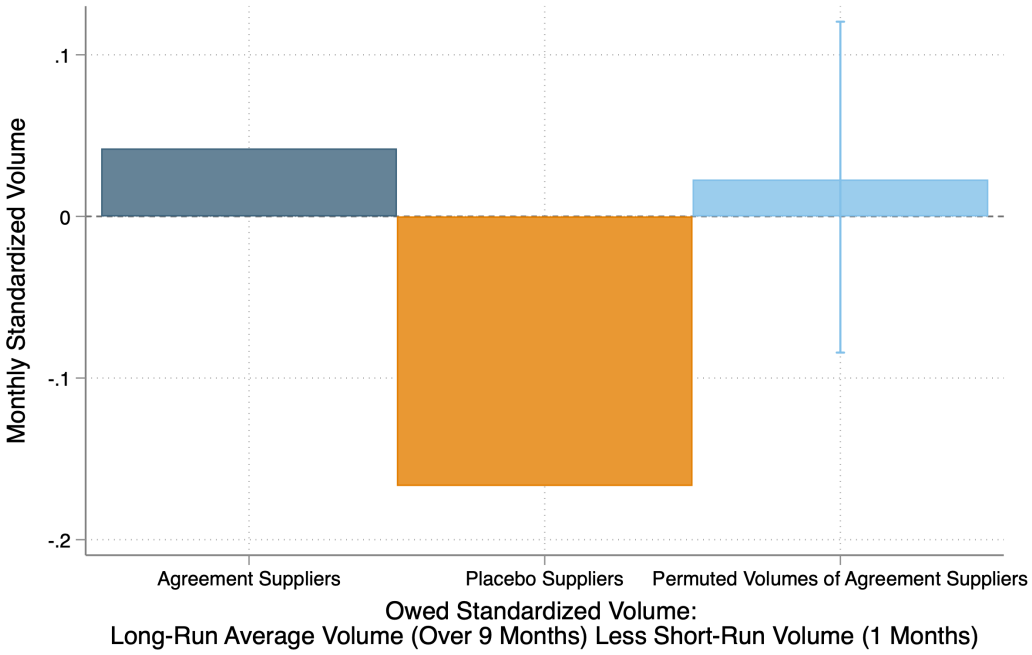
Table A.2 illustrates prices with placebo suppliers never meet more than one of the four conditions for being U-shaped.

Figure A.10: Coefficient of Variation by Supplier Type (Placebo Comparison)



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

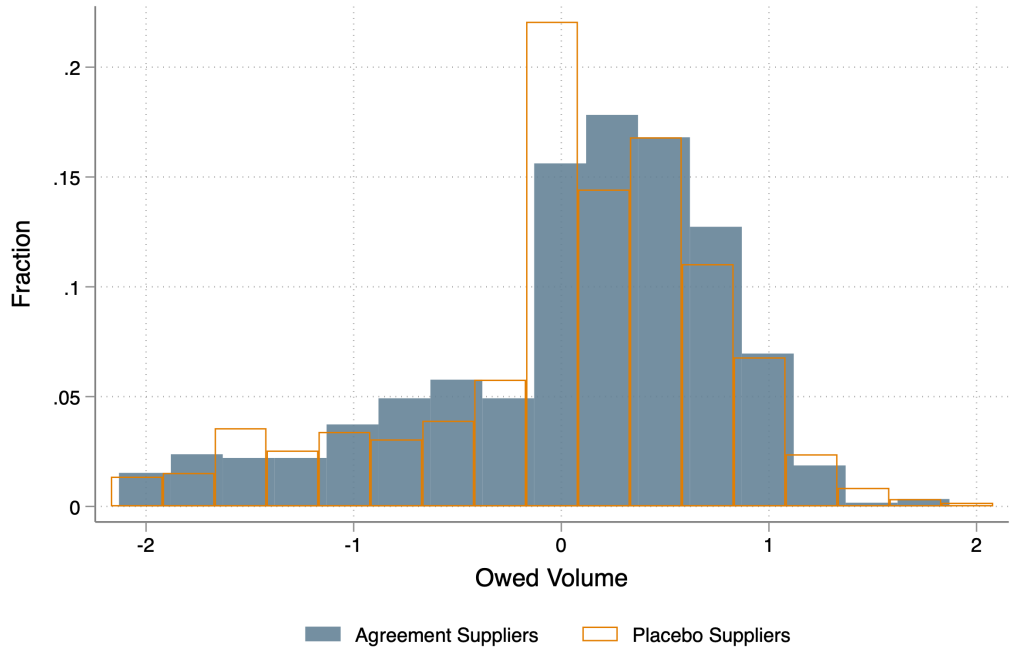
Figure A.11: Coefficient on Owed Quantity by Supplier Type (Placebo Comparison)



Note: Data from universe of fabric transactions by the buyer. Bars show the coefficient for owed standardized volume from a regression of current volume on owed volume 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 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 or group-based identity (as in Sanchez de la Sierra (2021)). However, agreement pricing 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. Such a contract seems unlikely given that social 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 agreements or influence which suppliers are selected as agreement suppliers. Importantly, in these scenarios, the agreement is still governed by its terms, under relational enforcement, with their associated impacts on transactions.

Figure A.12: Distribution of Owed Quantity by Supplier Type (Placebo Comparison)



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

A.7 Robustness Checks for Theoretical Bargaining 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 changes to deviation payoffs. 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 spot 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 lowest:

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

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 spot market price. 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 occurring when

$p^{RC}q^{RC} - p^M q^{RC}$ is the largest. In this case, the buyer's outside option is:

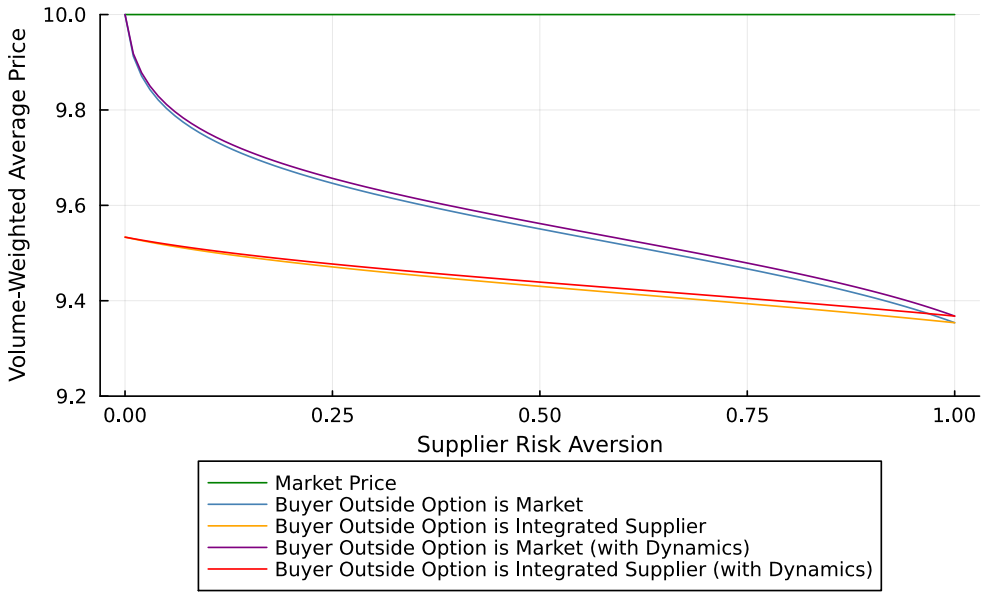
$$O_B = (\delta^B)\Big(- p^M \mathbb{E}_{q^{RC}}[q^{RC}]\Big) + (1-\delta^B)\Big(- p^M q^M\Big) \text{ if not integrated}$$

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

Note that for δ^S, δ^B sufficiently large, dynamic incentives are irrelevant as only long-run payoffs shape behavior.

The analysis below shows that results do not change meaningfully when incorporating dynamics.[71]

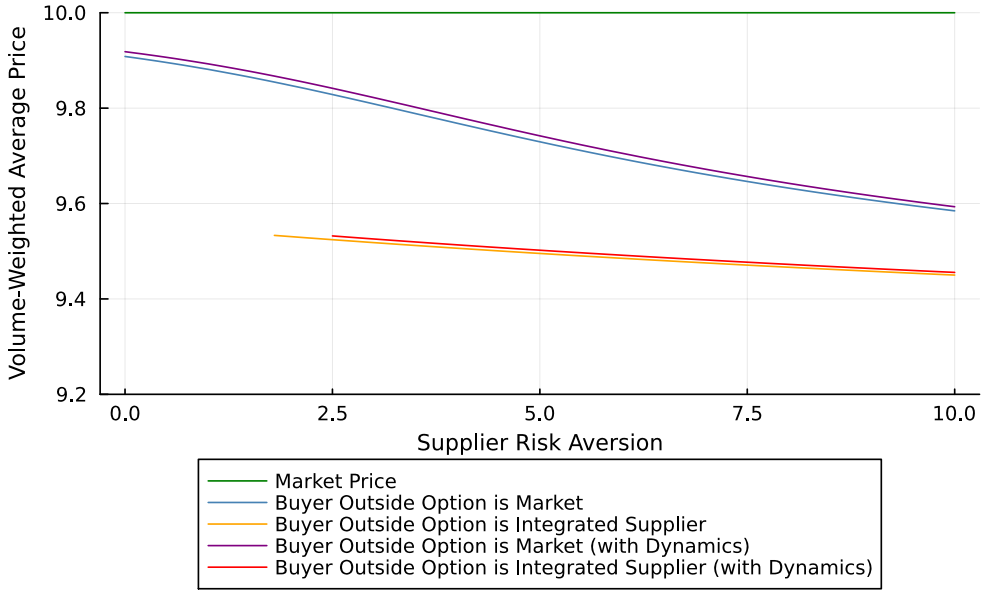
Figure A.13: Relational Contract Prices and the Buyer Bargaining Parameter with Dynamics



Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.8.

[71] For this analysis, I set $\delta^S = \delta^B = .9$. As empirical analysis of this model is at the quarter level, this quarterly discount factor 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 is conservative.

Figure A.14: Relational Contract Prices and Supplier Risk Aversion with Dynamics

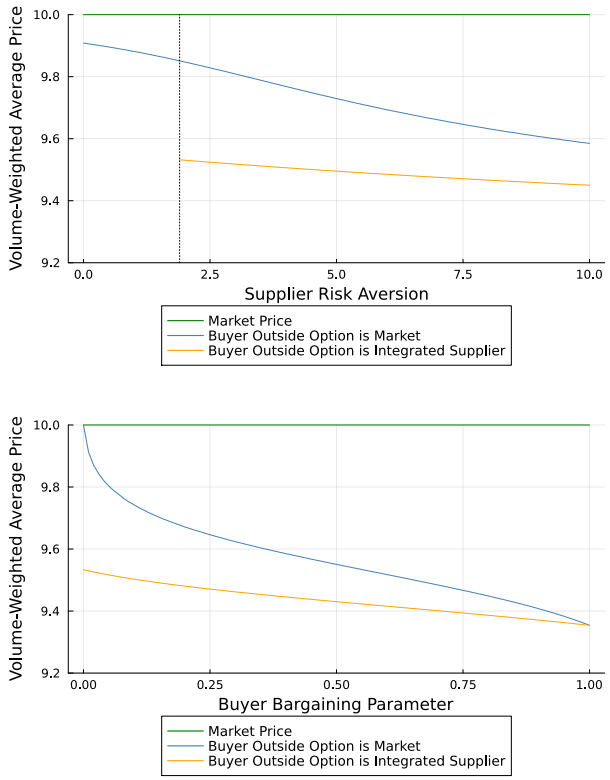


Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.8.

A.7.2 Cost Heterogeneity Between External and Internal Supplier

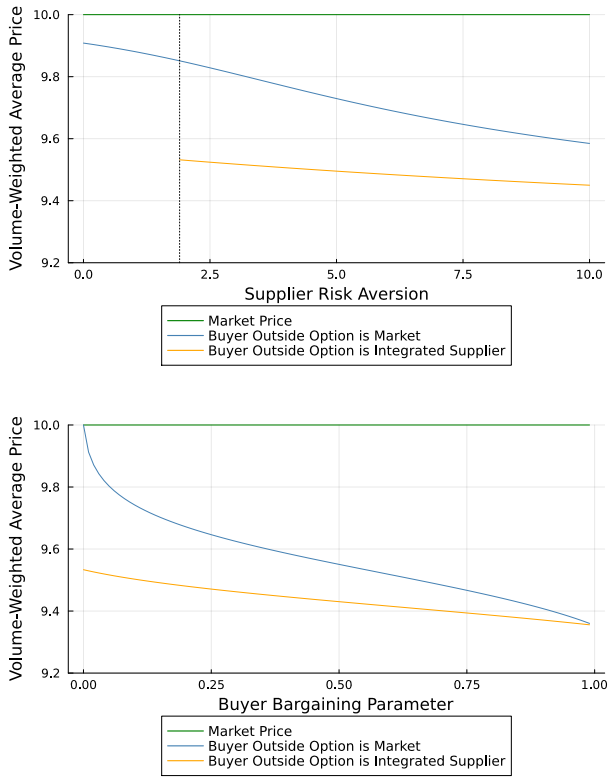
Results are robust to the external supplier having lower or higher costs than the integrated supplier, shifting costs by 10%. To ensure that optimal quantity is still $q = 1$ without changing market prices, I only shift the fixed part of costs by 10%, keeping convexity the same. Results (available upon request) shifting all costs are similar. Large 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 outside option (assuming the supplier can see the integrated mill and make a reasonable inference about is productivity).

Figure A.15: Agreement Prices, Supplier Risk Aversion, and Buyer Bargaining Parameter 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.8.

Figure A.16: Agreement Prices, Supplier Risk Aversion, and Buyer Bargaining Parameter 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.8.

A.8 Parameterization and Solution of the Theoretical Model

The model is parameterized as:

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

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

A.9 Replacement Threats in a Signaling Game

Consider a signaling game that proceeds as follows:

1. Nature draws a state of the world, which represents the residual capacity remaining at the buyer's in-house supplier. This approach captures exogenous demand shocks that the buyer faces for fabrics at the integrated suppliers from end client demand. I binarize the state of residual capacity to highlight the key cutoff: whether it would be lower cost for the buyer to produce the relevant volume internally or to source it through the spot market. Therefore, $\tilde{q} = 0$ means that the in-house supplier is relatively capacity unconstrained, such that production is lower cost in-house, while $\tilde{q} = 1$ means that the in-house supplier is relatively capacity constrained, such that production is more costly in-house than sourcing from the spot market. I express the buyer's cost function as $C_B(\tilde{q}, 1)$ to denote the cost of producing one unit given residual capacity is $\tilde{q}$. It follows that $C_B(1, 1) > p^M > C_B(0, 1)$.

2. Then, the buyer announces whether it is capacity unconstrained ($\hat{q} = 0$) or capacity constrained ($\hat{q} = 1$).

3. The supplier then chooses whether to accept the buyer's claim about its capacity constraint or to reject it (as if the supplier then makes a final offer).

4. The cooperative stage of bargaining occurs, where the buyer and supplier engage in Nash bargaining relative to whatever outside option the supplier considers in its final offers.

5. The buyer's payoff depends upon the state of the world $\tilde{q}$. If the agreement breaks down, as could occur if there is no surplus, then the buyer pays fabric at the cost of the outside option ($C_B(0, 1)$ if $\tilde{q} = 0$ and the buyer sources internally, and p^M if $\tilde{q} = 1$ and the buyer sources from the spot market). If the agreement remains, then the buyer's payoff is $-p^A(\hat{q})$ if the supplier accepts $\hat{q}$ and $-p^A(1 - \hat{q})$ if the supplier rejects $\hat{q}$.

6. The supplier's payoff inside an agreement is $p^A(\hat{q}) + \theta - C_S(1)$ if the supplier accepts the buyer's $\hat{q}$. In the supplier's payoff, θ represents the non-price benefits of the agreement, such as both profit smoothing and cost reductions from the U-shaped pricing function and quantity assurance. Therefore, in the empirical setting of the agreements studied, θ should be higher for suppliers with higher risk aversion. I also normalize θ by subtracting off the expected payoff from selling to the spot market, such that θ represents the net benefits of the agreement compared to spot market *without* accounting for the revenue or production cost in the agreement (such that the supplier's outside option value is 0). $C_S(1)$ is the production cost for the order. If the supplier rejects $\hat{q}$,

it receives a payoff of $p^A(1-\hat{q}) + \theta - C(1)$ as long as the agreement remains. Given the normalization in θ, the supplier receives 0 if the agreement breaks down.

7. Additionally, if the supplier rejects the buyer's claim ($\hat{q}$), then costly verification (*e.g.,* a factory tour) reveals the true state of the world $\tilde{q}$. In this case, a transfer γ is made within the relationship that punishes either the buyer for claiming an inaccurate $\hat{q}$ or the supplier for rejecting an accurate $\hat{q}$. In the model, γ broadly captures (in a reduced-form way) any incentives that prevent parties from trying to misrepresent the state of the world. Note that no model results depend upon the γ being a transfer; although not shown, results are qualitatively identical if γ is exclusively a penalty.[72]

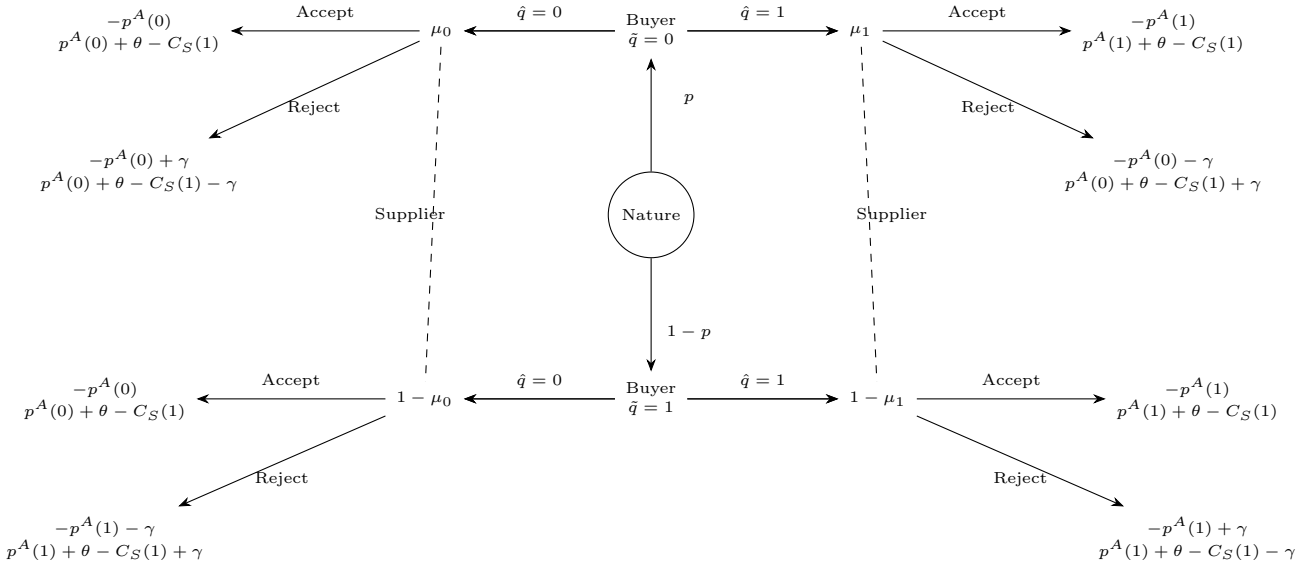
8. Nash bargaining prices are computed assuming that the buyer and supplier equally split the surplus (in other words, bargaining parameters are .5). I assume that the transfer γ does not occur if the agreement dissolves.

Below, I show that there are two equilibria. First, there is an equilibrium where, with a sufficiently risk-averse supplier (*i.e.,* high enough θ), the buyer pools and communicates $\hat{q} = 0$. However, this equilibrium no longer holds for low θ. Second, there is an equilibrium where, with low θ, the buyer pools and communicates instead $\hat{q} = 1$.

As I use weak Perfect Bayesian Equilibrium as the solution concept, denote the supplier's belief that they are at the node where $\tilde{q} = 0$ as μ_0 after the buyer plays $\hat{q} = 0$ and as μ_1 after the buyer plays $\hat{q} = 1$.

A.9.1 Replacement Threat: Pooling Equilibrium at $\hat{q} = 0$

First, consider the case where θ is sufficiently large that the agreement persists for both states of nature. Then, the game can be expressed as:



[72] While the exact restrictions on parameters for certain equilibrium to hold do shift slightly, the qualitative conclusions are unchanged.

In this game, the supplier's best response is to accept as long as:

$$\begin{aligned}
p^A(0) + \theta - C_S(1) &\geq \mu_0(p^A(0) + \theta - 1 - \gamma) + (1-\mu_0)(p^A(1) + \theta - C_S(1) + \gamma) \\
(1-\mu_0)(p^A(0) + \theta - C_S(1)) + \mu_0\gamma &\geq (1-\mu_0)(p^A(1) + \theta - C_S(1) + \gamma) \\
p^A(0) + \theta - C_S(1) + \frac{\mu_0}{1-\mu_0}\gamma &\geq p^A(1) + \theta - C_S(1) + \gamma \\
\frac{2\mu_0 - 1}{1-\mu_0}\gamma &\geq p^A(1) - p^A(0)
\end{aligned}$$

Now, solve the Nash bargaining problem:

$$p^A(0) = max_{\bar{p}_0}(-\bar{p}_0 - (-C_B(0,1)))^{.5}(\bar{p}_0 + \theta - C_S(1))^{.5}$$

The F.O.C. then implies:

$$\begin{aligned}
0 &= (.5)(-\bar{p}_0 + C_B(0,1))^{-.5}(-1)(\bar{p}_0 + \theta - C_S(1))^{.5} \\
&\quad + (-\bar{p}_0 + C_B(0,1))^{.5}(.5)(\bar{p}_0 + \theta - C_S(1))^{-.5} \\
\bar{p}_0 + \theta - C_S(1) &= -\bar{p}_0 + C_B(0,1) \\
\bar{p}_0 &= \frac{C_S(1) - \theta + C_B(1)}{2}
\end{aligned}$$

Similarly,

$$p^A(1) = max_{\bar{p}_1}(-\bar{p}_1 - (-p^M))^{.5}(\bar{p} + \theta - C_S(1))^{.5}$$

The F.O.C. then implies:

$$\begin{aligned}
0 &= (.5)(-\bar{p}_1 + p^M)^{-.5}(-1)(\bar{p}_1 + \theta - C_S(1))^{.5} \\
&\quad + (-\bar{p}_1 + p^M)^{.5}(.5)(\bar{p}_1 + \theta - C_S(1))^{-.5} \\
\bar{p}_1 + \theta - C_S(1) &= -\bar{p}_1 + p^M \\
\bar{p}_1 &= \frac{C_S(1) - \theta + p^M}{2} \\
\therefore p^A(1) - p^A(0) &= \frac{p^M - C_B(0,1)}{2}
\end{aligned}$$

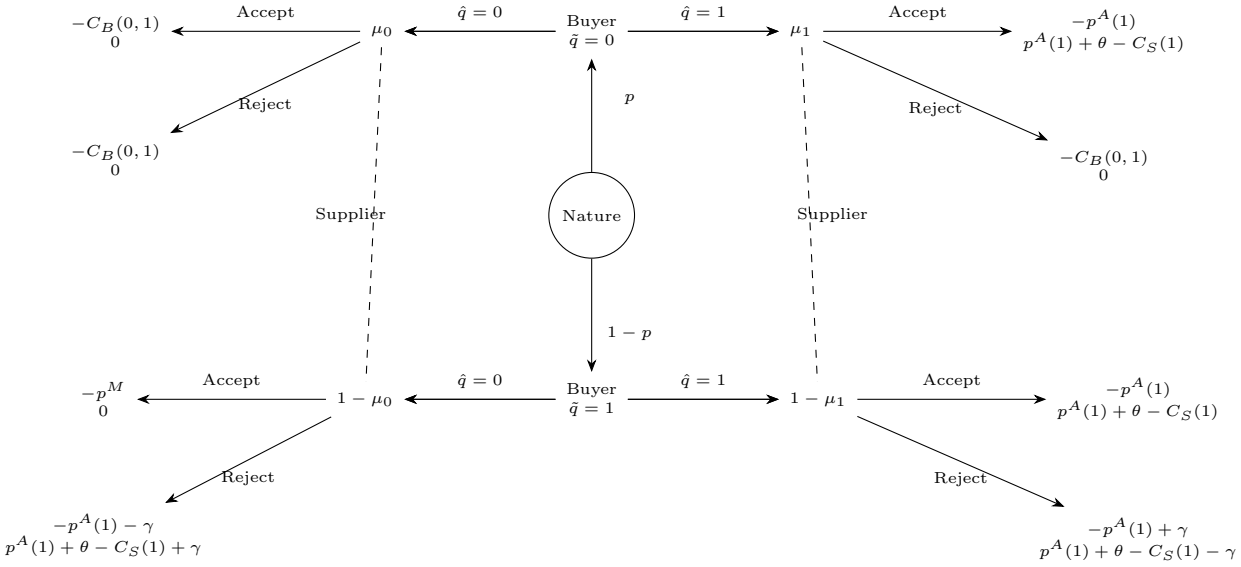
Intuitively, as long as either *i)* it is likely that the buyer is indeed unconstrained or *ii)* the transfer from costly verification is sufficiently large, then the supplier accepts. Note that this condition does have bite: as $p^A(1) > p^A(0)$, the existence of this equilibrium is not trivially satisfied. Because the buyer pays lower agreement prices when bargaining relative to its unconstrained outside option, the buyer also prefers this option.

Formally, the equilibrium is that the buyer plays $\hat{q} = 0$ for both types, the supplier accepts after $\hat{q} = 0$, and $\mu_0 = p$. The supplier can accept or reject after $\hat{q} = 1$ without changing the equilibrium, and has the beliefs consistent with whichever strategy is specified for the supplier after $\hat{q} = 1$. (There are multiple possible equilibria as the information set after the supplier plays $\hat{q} = 1$ is not reached given that the buyer prefers to play $\hat{q} = 1$).

The key requirement is that $p, \gamma, C_B(0,1), p^M : \frac{2p-1}{1-p}\gamma \geq \frac{p^M - C_B(0,1)}{2}$. Note that γ can be arbitrarily small and this will still be satisfied as long as at least one of: *i)* the buyer is not capacity constrained with sufficiently high probability (as shown in figure 10) or *ii)* the buyer's cost advantage relative to the spot market is small. While the exact level of p required depends upon the γ, it is clear that for $2p-1$ to be positive it must be the case that $p > .5$. Does this match with the data? Indeed, integrated suppliers operate at under 100% of capacity in 96% of months. That said, it seems natural to think that the relevant question is whether the marginal order would push the integrated supplier to a capacity level over 100%. I find that that the median capacity is 78%; therefore, as long as the marginal order does not use 22% of capacity internally (which would be inconsistent with volume data), it follows that the data are plausibly consistent with a sufficiently high p for this equilibrium to hold.

A.9.2 No Replacement Threat: Pooling Equilibrium at $\hat{q} = 1$

Now, consider the case where θ is small enough, representing a supplier with low risk aversion, such that there is no surplus relative to $\tilde{q} = 0$ but large enough that there are gains relative to $\tilde{q} = 1$. The payoffs then change, with the agreement breaking down and the buyer sourcing internally after declaring $\hat{q} = 0$. This result reflects that the supplier's participation constraint is $p^A(0) : p^A(0) + \theta - C(1) \geq 0$. As θ decreases, no surplus remains. Then, the game can be expressed as:



Now, I show that the pooling equilibrium at $\hat{q} = 0$ may break. The supplier's best

response is to accept after the buyer plays $\hat{q} = 0$ as long as:

$$0 \geq \mu_0(0) + (1 - \mu_0)(p^A(1) + \theta - C_S(1) + \gamma)$$
$$0 \geq p^A(1) + \theta - C_S(1) + \gamma$$

As $\gamma > 0$,

$$0 > p^A(1) + \theta - C_S(1)$$

But, this contradicts the definition of θ as large enough that there are gains relative to $\tilde{q} = 1$.

Instead, there is a pooling equilibrium at $\hat{q} = 1$ where the supplier accepts. Accepting is a strictly dominant strategy for the supplier as its best outcome is the Nash bargain when the buyer's outside option is the spot market because that bargain has the highest price. And, when θ is low, then the incentive to reject after $\hat{q} = 1$ to earn γ, but at the cost of a lower prices, goes away because the agreement dissolves when bargaining occurs with the buyer's unconstrained in-house supply as the outside option.

Formally, the equilibrium is that the buyer plays $\hat{q} = 1$, the supplier accepts after $\hat{q} = 1$ and rejects after $\hat{q} = 0$, and has the beliefs that $\mu_1 = p$ and $\mu_0 \in [0, 1]$ (there are no constraints on μ_0 because reject dominates accept for the supplier after $\hat{q} = 0$ and the information set is not reached in on-path, so neither belief consistency nor sequential rationality discipline μ_0.)

A.9.3 A Numerical Example

I now provide a numerical example to provide additional intuition. To do so, I select specific values for the cost functions and the market price. First, let $C_S(q) = q^2$, $c_B(\tilde{q}, q) = (\tilde{q} + q)^2$, and $p^M = 2$, where q is the order size (and I analyze one order).

Consider when the buyer bargains with a supplier with $\theta = 1$, as occurs for a highly risk-averse supplier. Then, there is surplus regardless of the buyer's residual capacity state: when $\tilde{q} = 0$, as the gains from collaboration are θ. Then, the Nash bargaining prices are:

$$p^A(0) = .5$$
$$p^A(1) = 1$$

Then, there is a pooling equilibrium at $\hat{q} = 0$ as long as:

$$\frac{2p - 1}{1 - p}\gamma \geq .5$$

Although I do not observe p directly from the data, note that $\frac{2p-1}{1-p}$ is increasing in $p \forall p \in [0, 1)$, it follows that for p, γ sufficiently high, that this inequality holds. Additionally, for p^M close to $C_S(1)$, then the right-hand side decreases to something quite small.

Formally, the equilibrium features the buyer playing $\hat{q} = 0$, while the supplier accepts after $\hat{q} = 0$

Alternatively, consider when the buyer bargains with a supplier with $\theta = -.5$, as occurs for a supplier with low risk aversion. (Recall that θ is normalizes such that it represents the value of the non-price agreement components less profits in the spot market, so negative values are reasonable). Then, there is surplus only when the buyer's residual capacity is $\tilde{q} = 1$. Then, the surplus generated by the agreement is $\theta - C_S(1) + p^M = -.5 - 1 + 2 = .5$. However, when $\tilde{q} = 0$, the surplus from cooperation is $\theta - C_S(1) + C_B(1) = -.5 - 1 + 1 < 0$. It follows that the agreement breaks down in such a case

$$p^A(1) = 1.75$$

Now, I show that the supplier does not accept when the buyer plays $\hat{q} = 0$. The supplier's payoff from accepting $\hat{q} = 0$ is 0 given that the agreement breaks down. However, the payoff from rejecting is $p(0) + (1-p)(p^A(1) + \theta - C_S(1) + \gamma) = (1-p)(.25) > 0 \forall p \in (0,1)$. Therefore, rejecting dominates accepting. It follows that the supplier rejects after $\hat{q} = 0$. And, as accepting after $\hat{q} = 1$ is a dominant strategy, the supplier will do so.

A.9.4 Discussion

This analysis shows a game where the buyer is able to threaten replacement with in-house production–by signaling that residual capacity in-house is low–in order to improve its bargaining position, but this equilibrium only exists when the buyer bargains with risk-averse suppliers. Although not shown, the pooling equilibrium at $\hat{q} = 1$ also exists when θ is high. However, if the buyer–as the first mover (or larger firm)–can select which equilibrium it prefers, it would always prefer to pool at $\hat{q} = 0$ because $p^A(0) < p^A(1)$ as long as $C_B(1,1) > C_B(0,1)$.

A.10 Replacement Threats in a Game with Commitment

Consider a game with commitment that proceeds as follows, using notation as in A.9.

1. The buyer commits to a capacity level at its integrated supplier of $q^C \in \{0,1\}$ where $q^C = 0$ means that the in-house supplier is relatively capacity unconstrained, such that production is lower cost in-house, while $q^C = 1$ means that the in-house supplier is relatively capacity constrained, such that production is more costly in-house than on the spot market. Mapping to the empirical setting, $q^C = 1$ represents the buyer committing to use the integrated supplier to displace some other spot supply. This commitment, in the context of this empirical setting, means that the buyer saves capacity at the in-house supplier. Maintaining this reserve of in-house capacity can have costs, such as the opportunity cost of unused in-house capacity. This opportunity cost is included inside this model.

2. Then, the cooperative stage of bargaining occurs, where the buyer and supplier engage in Nash bargaining relative to the buyer's outside option. The agreement can breakdown if no surplus remains.

3. The buyer then receives an exogenous capacity shock, which represents a new order being selected by an end client. With probability p, it does not receive an additional order and with probability $1-p$ it does receive an additional order.

4. The buyer's payoffs are the production costs for both the order volume associated with the initial commitment as well as the stochastic additional order.

5. The supplier's payoff inside an agreement is $p^A(q^C) + \theta - C_S(1)$ if there is surplus in the agreement. In this function, θ represents any net non-price benefits of the agreement, such as profit smoothing through a U-shaped pricing function. Therefore, in the context of the agreement I study, θ could represent (in a reduced-form way) higher risk aversion. I also normalize θ by subtracting off the expected payoff from selling to the spot market, such that θ represents the net benefits of the relationship compared to spot market *without* the revenue or production cost in the agreement (such that the supplier's outside option value is 0). $C_S(1)$ is the production cost for the order.

6. Nash bargaining prices are computed assuming that the buyer and supplier equally split the surplus. In the Nash bargain, the buyer's outside option is whichever of the spot and in-house supply can produce the order at lowest cost. In-house supply is lower cost if residual capacity is $q^C = 0$ and spot supply if $q^C = 1$.

A.10.1 Replacement Threat: Committing to $q^C = 0$

When θ is sufficiently high that Nash bargaining never breaks down, the buyer commits to $q^C = 0$. Specifically, the buyer's payoff from selecting $q^C = 0$ is $-p^A(0) - (1-p)C_B(0,1)$, where $C_B(0,1)$ is the cost of producing one unit at the in-house supplier when residual capacity is 1 (*i.e.,* is capacity constrained). Solving the Nash bargaining problem, $p^A(0) = \frac{C_S(1)-\theta+C_B(0,1)}{2}$. Therefore, the buyer's payoff accounting for both orders is $-\frac{C_S(1)-\theta-C_B(0,1)}{2} - (1-p)C_B(0,1)$. Alternatively, if the buyer selects $q^C = 1$, its payoff is $-p^A(1)-(1-p)C_B(1,1)$. It follows that as long as $p^A(0) < p^A(1)$, the buyer prefers to play $q^C = 0$. This result reflects that the agreement with the supplier always occur, meaning that $C_B(0,1) = C_B(1,1)$ in this case. And, as $p^A(1) = \frac{C_S(1)-\theta+p^M}{2}$ and $p^M > C_B(0,1)$, it follows that $p^A(0) < p^A(1)$ and the buyer's optimal choice is $q^C = 0$.

A.10.2 No Replacement Threat: Committing to $q^C = 1$

Now, consider the case when θ is lower, as occurs for a less risk-averse supplier, such that surplus does not remain relative to $q^C = 0$, although does relative to $q^C = 1$. For this scenario to be true, it must be the case that $\theta : p^M + C_s(1) > \theta > C_B(0,1) + C_s(1)$.

Then, if the buyer selects $q^C = 0$, its payoff is $-C_B(0,1) - (1-p)p^M$, as the buyer has some chance of receiving an additional order that it will have to source through the spot market (it only wants to use in-house supply once given cost capacity). However, if the buyer selects $q^C = 1$, its payoff is $-p^A(1) - (1-p)C_B(0,1) = -\frac{C_S(1)-\theta+p^M}{2} - (1-p)C_B(0,1)$, as it has saved the internal capacity to use in the event that it receives an additional order. Therefore, the buyer prefers $q^C = 1$ as long as:

$$\begin{aligned}
-\frac{C_S(1)-\theta+p^M}{2} - (1-p)C_B(0,1) &\geq -C_B(0,1) - (1-p)p^M \\
\frac{\theta - C_S(1) - p^M + 2C_B(0,1)}{2} &\geq (1-p)(C_B(0,1) - p^M) \\
\frac{(C_B(0,1)+\theta-C_S(1)) + (C_B(0,1)-p^M)}{2} &\geq (1-p)(C_B(0,1) - p^M) \\
\frac{C_B(0,1)+\theta-C_S(1)}{2} &\geq (.5-p)(C_B(0,1) - p^M)
\end{aligned}$$

Recall that $\theta > C_B(0,1) + C_s(1) \implies C_B(0,1) + \theta - C_S(1) > C_B(0,1) + C_B(0,1) + C_s(1) - C_s(1) = 2C_B(0,1)$. Therefore, the left-hand side is at least $C_B(0,1)$, which is strictly positive. However, the right-hand side is negative by definition of $C_B(0,1)$. It follows that this inequality holds. As the buyer prefers $q^C = 0$, the equilibrium of the commitment game is for the buyer to select $q^C = 1$.

A.10.3 Discussion

A natural question in this model is whether the commitment is credible and believed by the suppliers. Intuitively, the buyer might not want to commit to $q^C = 0$ because doing so might leave some in-house capacity slack. Through the lens of the model, this intuition suggests that the buyer might prefer to pay $C_B(0,1)$ for the guaranteed volume and then, in expectation, pay $(1-p)p^A(1)$ for any additional order that arrives. Note that this sourcing strategy violates commitment because the supplier commits to having slack in-house capacity, but does not actually have it (although still Nash Bargain relative to the fact that it no longer has capacity). It follows that the buyer will not face commitment problems when $C_B(0,1) + (1-p)p^A(1) \geq p^A(0) + (1-p)C_B(0,1)$. Formally, this condition can be expressed as:

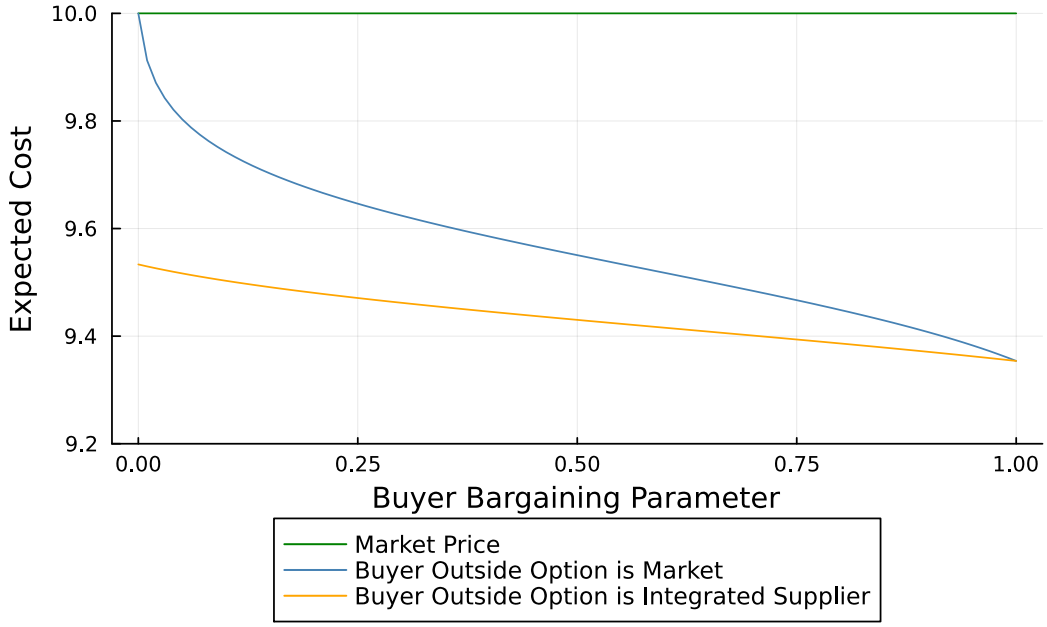
$$
\begin{aligned}
C_B(0,1) + (1-p)p^A(1) &\geq p^A(0) + (1-p)C_B(0,1) \\
1-p &\geq \frac{p^A(0) - C_B(0,1)}{p^A(1) - C_B(0,1)} \\
p &\leq \frac{p^A(1) - p^A(0)}{p^A(1) - C_B(0,1)} \\
p &\leq \frac{.5(p^M - C_B(0,1))}{p^A(1) - C_B(0,1)}
\end{aligned}
$$

As p^M increases relative to $p^A(1)$, the temptation to deviate goes away. This object increases in θ, meaning that there is little incentive to deviate when the agreement has low prices because suppliers highly value the non-price component, leading to lower prices. However, when θ is small, then this condition may not be satisfied and inded there is a temptation to deviate. But, if the buyer is considering deviating, it must be bargaining with a supplier where $p^A(0)$ is not just the outside option. Therefore, for the buyer to consider deviating, it must be bargaing with a high θ (formally, $\theta : \theta - C_S(1) - C_B(0,1) > 0$).

Although commitment is not likely to bind, the model does not rule the existence of θ where the buyer would face commitment problems. In these cases, commitment may be credible in the context of the agreement because the agreement already embeds commitment. Inside the agreement, statically, the buyer has incentives to deviate when paying agreement prices over spot market prices. The relational contract embeds commitment in this setting, enforcing the buyer's commitment in the agreement to pay high prices. As the relational contract already embeds some commitment about prices, adding additional comittment about outside options is likely plausible in this setting with relational enforcement.

A.11 Value Creation and Capture: Buyer Bargaining Parameter

Figure A.17: Bargained 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.8.

Figure A.17 shows how the buyer's expected fabric procurement cost in the agreement 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. This result reflects that suppliers capture all value when the buyer bargaining parameter is zero; therefore, the bargained price is identical to the spot market prices and the supplier extracts all benefit from the profit insurance. However, as the buyer bargaining parameter increases, the buyer captures value by receiving a share of the value created by the non-price features of the agreement. Because the buyer benefits from the agreement only through price reductions, the buyer captures value through lower prices. When the buyer's outside option is the integrated supplier, as the buyer's threat point improves (*i.e.,* is lower cost), prices also decrease. As the buyer bargaining parameter increases, the supplier's outside option plays an increasingly important role in determining prices. Therefore, when the buyer bargaining parameter is one, the supplier's outside option determines prices and the buyer's outside option is effectively irrelevant, resulting in the same bargained prices (and expected cost) regardless of buyer outside option.

The price change due to value capture is visible in the graph as the difference between the bargained price when the outside option is the market (the blue line) and the price when the

outside option is the integrated supplier (the orange line).Value capture can be quantitatively important (over 5% of market prices) when the buyer's outside option influences bargained prices. However, when the buyer has a high bargaining parameter, changing the buyer's outside option has a negligible effect on bargained 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, value capture 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 bargained outcomes.

A.11.1 Log Quantities

Table A.3: 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 the buyer. Coefficients reported are from a regression of log days between fabric order date and delivery date on log quantity.

A.12 Structural Bargaining Model: Evidence Supporting Model Assumptions

A.12.1 Quantity Exogeneity

Although capacity utilization for any supplier is certainly not random, as the buyer provides quantity assurance in the agreement, supplier-period deviations from the targeted capacity in the agreement 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 agreement suppliers as compared to other suppliers.

However, buyer behavior that reduces *dynamic* capacity variation would bias estimates, as I would underestimate how much quantity assurance the agreement 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).[73]

To mitigate these concerns, I directly test the most likely examples of such behavior in Table A.4. 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 agreement 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.4: 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

Note: Data from universe of fabric transactions by the 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 agreement target quantities are adjusted after vertical integration; however, this possibility seems unlikely as it is not in the buyer or supplier's interest to decrease the amount sourced from the agreement 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 systematically overestimate capacity shocks in the post-period. I test whether average capacity in agreements decreases in the post period in Table A.5 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.

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

Table A.5: 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 the buyer. Coefficients reported are from a regression using quarter-by-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.13 Structural Bargaining Model: Comparison of Risk Aversion Estimates

Table A.6: 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.14 Confidence Interval for Buyer Bargaining Parameter and Supplier Risk Aversion

Table A.7 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.

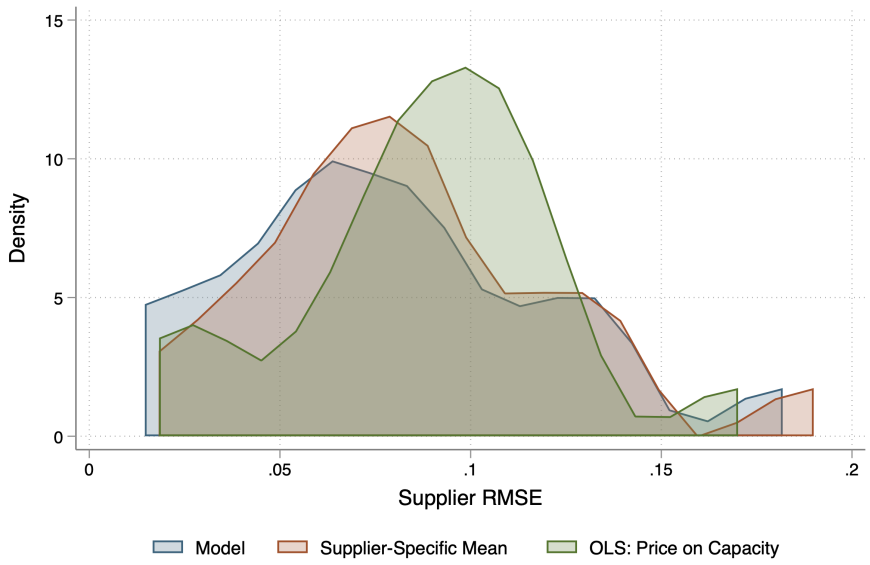
Table A.7: Buyer Bargaining Parameter and Supplier Risk Aversion: 95% CI

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.15 Model Validation: Out-Of-Sample Fit RMSE Density

Figure A.18: 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.

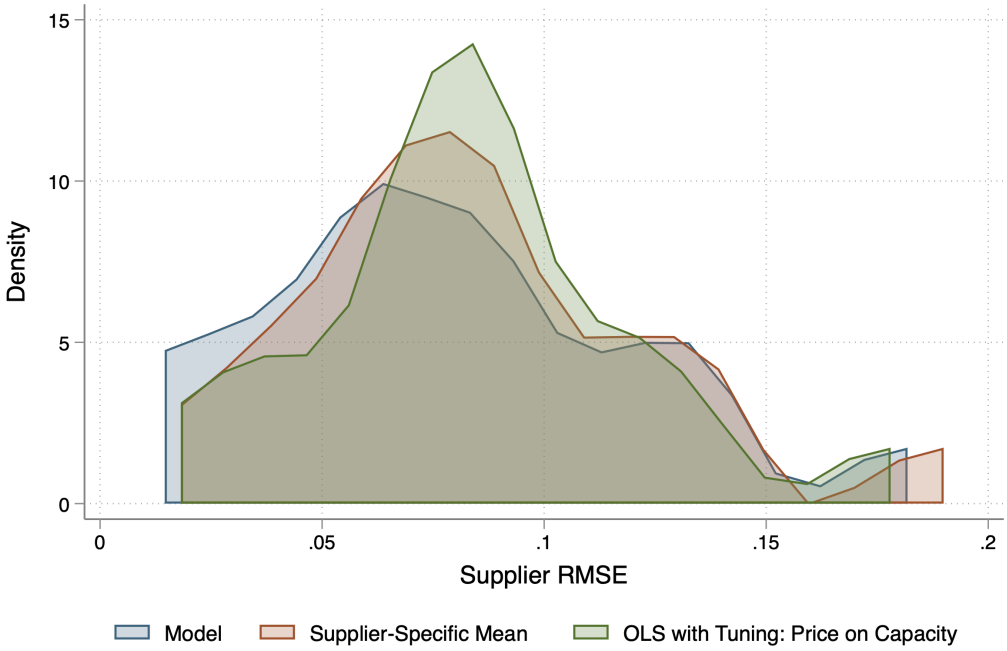
A.16 Model Validation: Out-Of-Sample Fit Comparison with OLS with Tuning Parameter

Table A.8: 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.

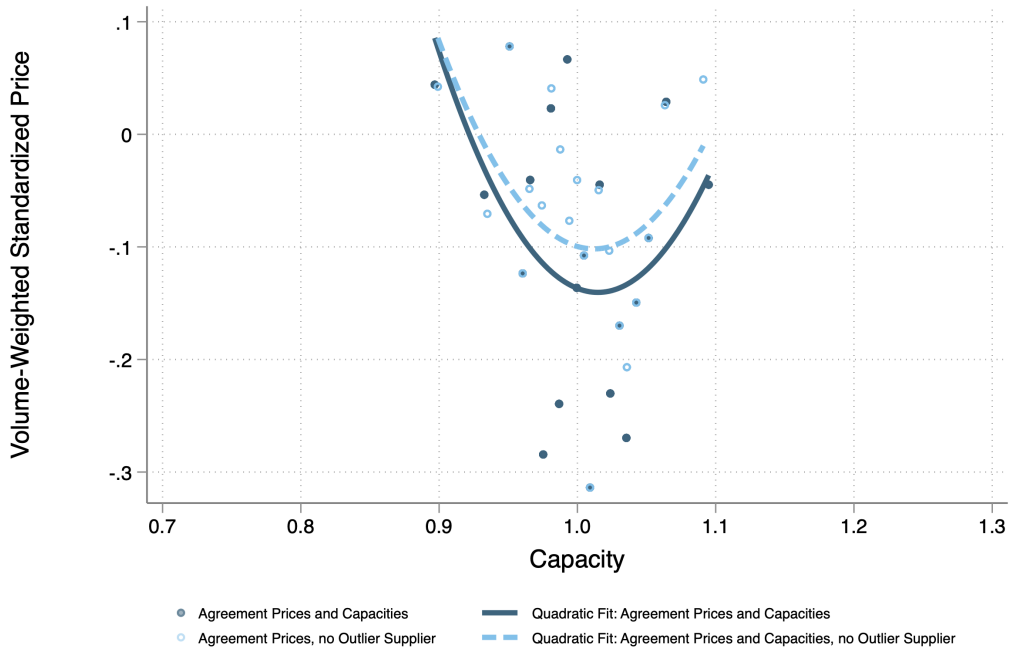
Figure A.19: Distribution of Supplier RMSE by Estimation Method



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

A.17 Pricing Function Robustness to Outlier Supplier

Figure A.20: Pricing Function & Outlier Supplier: Binned

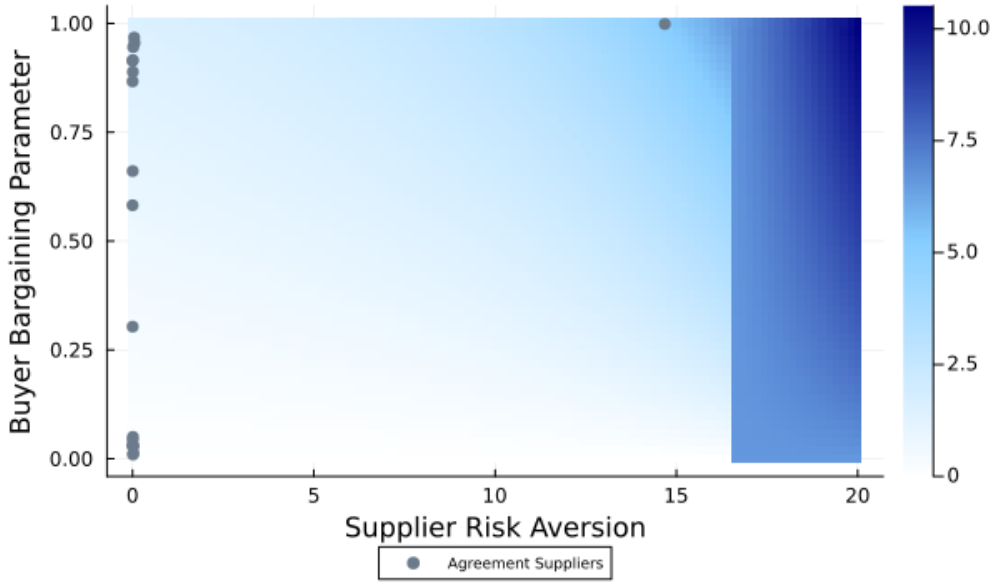


Note: Data from universe of fabric transactions by the buyer with agreement suppliers, transformed into supplier-by-quarter panel data. Quadratic fit is from a regression of volume-weighted standardized price on capacity and capacity squared. The figure bins prices and capacities using 20 quantiles. The mapping from quantities to capacities is as discussed in 4.2.2.

A.18 Validation using Difference-in-Differences Estimates

A.18.1 Model-Implied Value Capture from Agreement Suppliers

Figure A.21: Value Capture from Agreement Suppliers

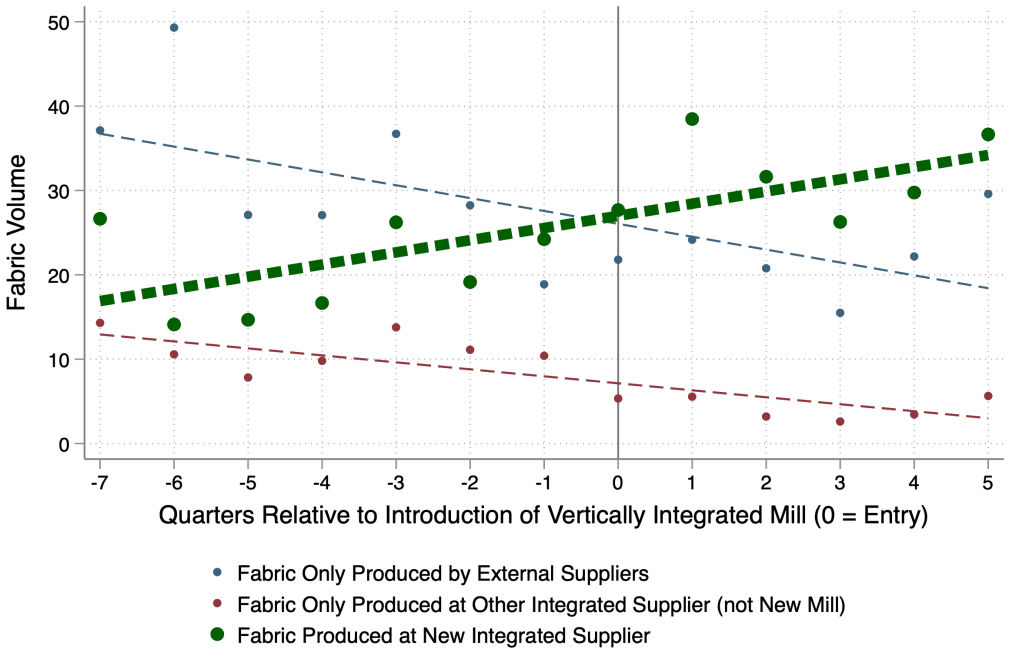


Note: Model estimates.

A.18.2 Semiparametric Difference-in-Differences Design and Identification

In addition to bringing less production in-house for fabrics with more agreement suppliers, Figure A.22 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.22: Trends in Fabric Demand and Vertical Integration



Note: Data from universe of fabric transactions by the 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 semiparametric difference-in-differences approach as in Abadie (2005), which weakens the parallel trends assumption to *conditional* parallel trends.[74] 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.[75] Formally, this assumption can be expressed as:

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

[74] 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.

[75] Specifically, the 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 $\boldsymbol{X_i}$,[76] 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 semiparametric approach, the estimated average treatment effect on the treated is unbiased if the propensity score adjustment accounts for selection into treatment appropriately. 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. 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.9 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,[76] it is not mechanical that this propensity score approach would achieve balance. I also document in Figure A.23 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.[77]

Table A.9: Balance on Untargeted Outcomes

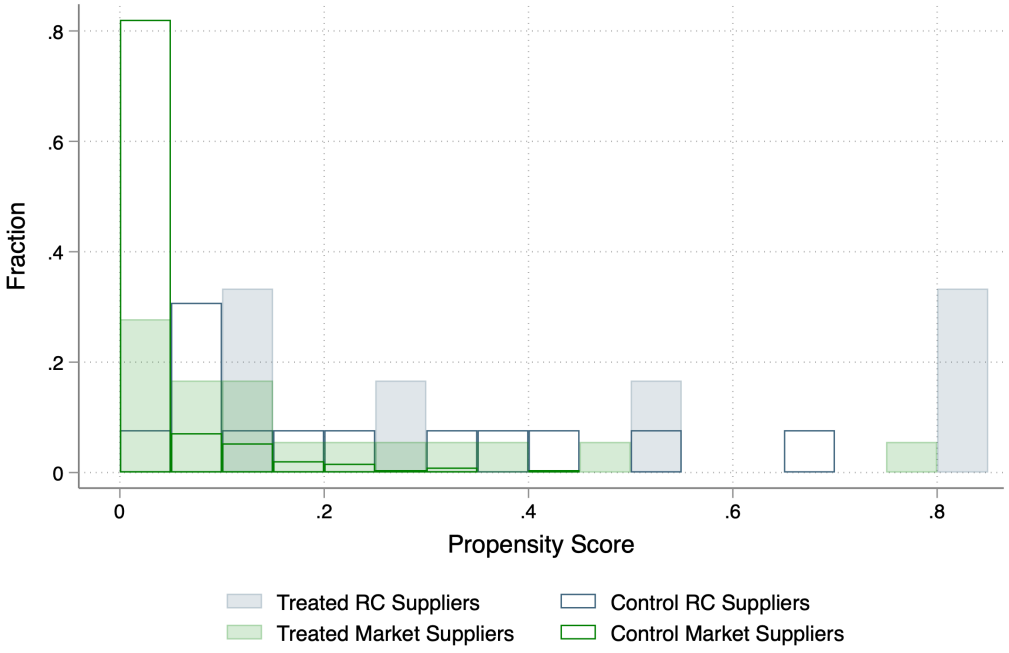
	Coefficient on Treatment Indicator:	
Outcome Variable	Without IPW Adjustment	With IPW Adjustment
Pre-Integration Volume	3033768	1085459
	[1189072, 5489585]	[−1004042, 3383491]
Pre-Integration Transaction Count	72.93	−3.35
	[44.49, 106.45]	[−46.3, 15.1]
Pre-Integration Fabric Count	26.71	−.98
	[17.08,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

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

[76] 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 produced by the supplier in the pre-integration quarters, an indicator for agreement suppliers, and the interaction of agreement supplier indicator and the linear term for fabric growth and the count of distinct fabrics.

[77] 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.23: Distribution of Propensity Scores by Supplier Type

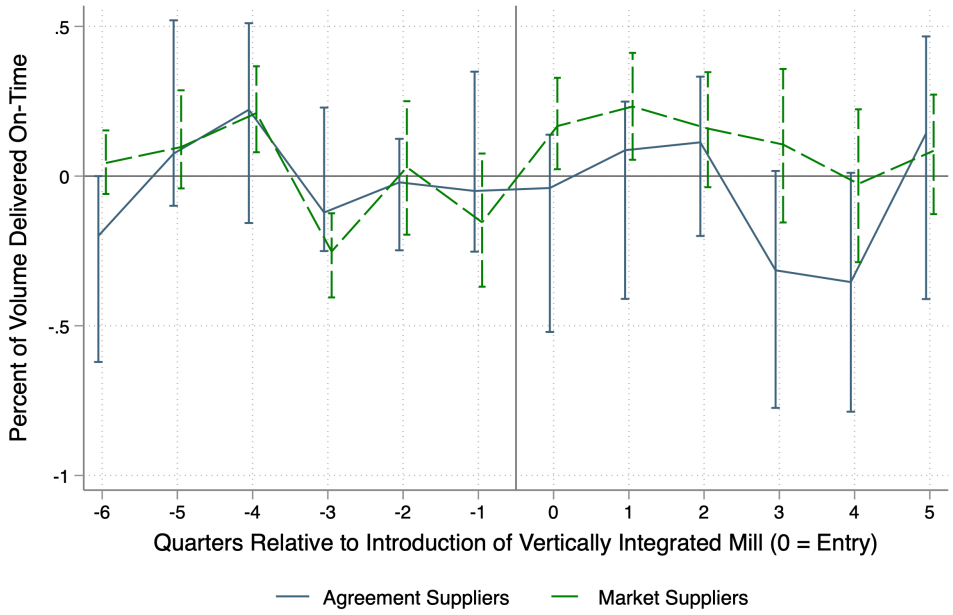


Note: Data from universe of fabric transactions by the buyer, restricting to use only pre-integration data.

A.19 Difference-in-Differences Supplemental Analysis and Robustness Checks

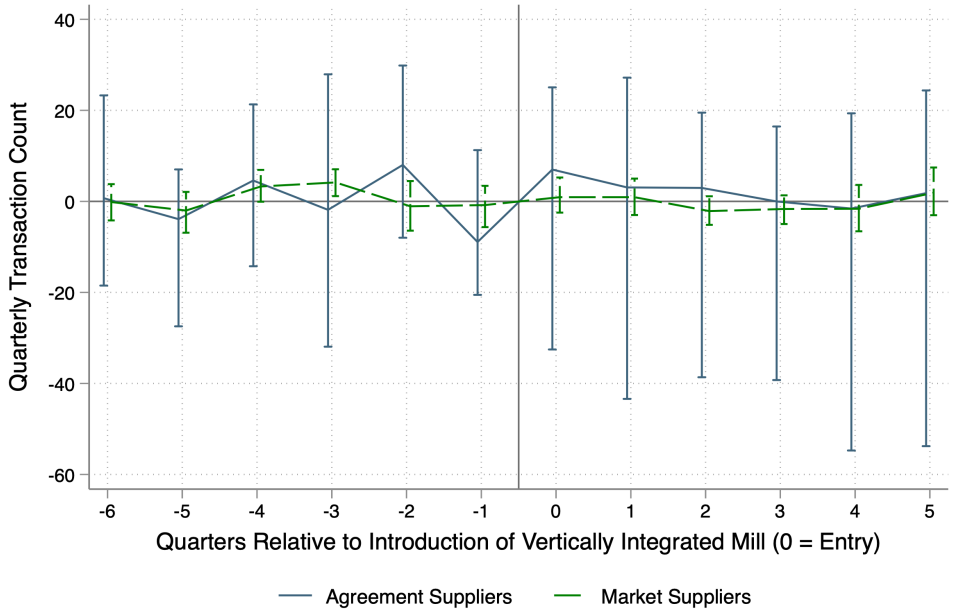
A.19.1 Difference-in-Differences Additional Event Study Plots

Figure A.24: Difference-in-Difference Estimates for Effect of Vertical Integration on Reliability



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

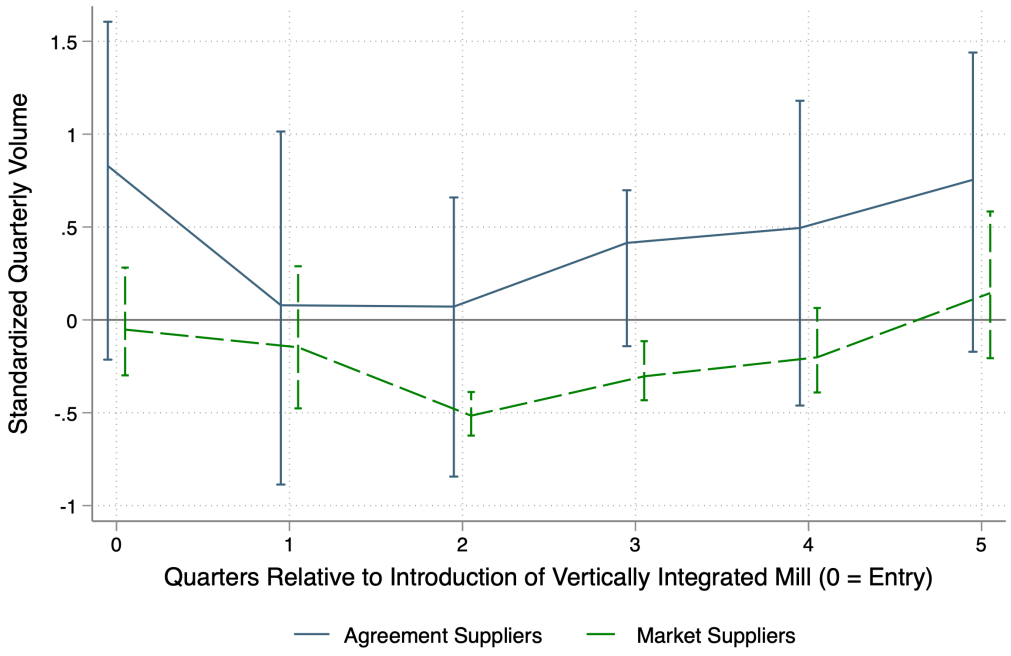
Figure A.25: Difference-in-Difference Estimates for Effect of Vertical Integration on Transaction Count



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

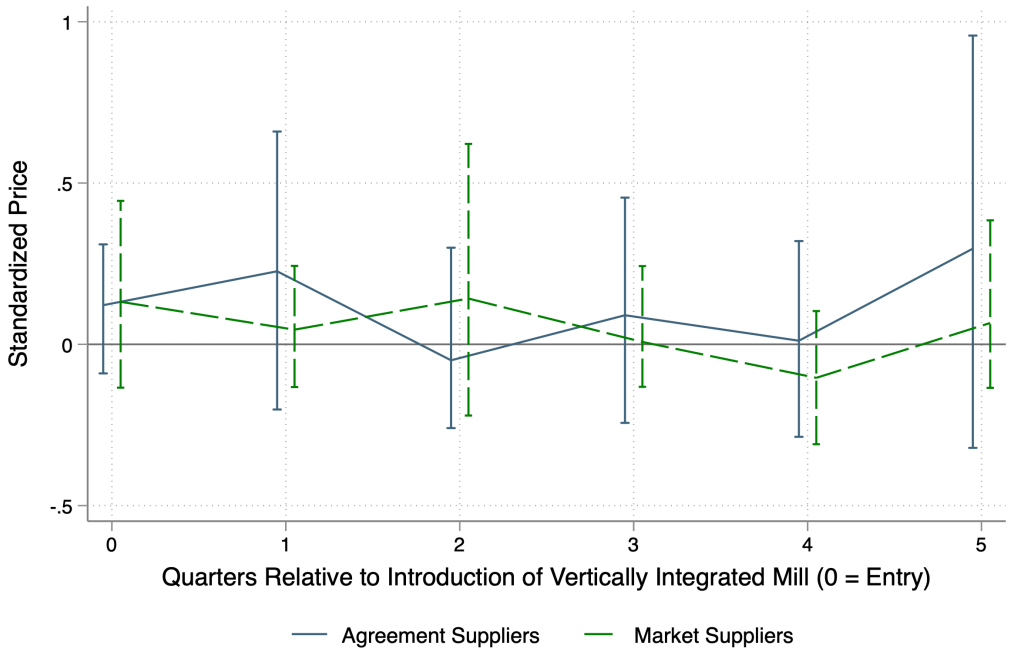
A.19.2 Synthetic Difference-in-Differences

Figure A.26: Synthetic Difference-in-Differences Estimates for Effect of Vertical Integration on Quarterly Volume



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

Figure A.27: Synthetic Difference-in-Differences Estimates for Effect of Vertical Integration on Standardized Prices



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

A.20 NNTR Implementation

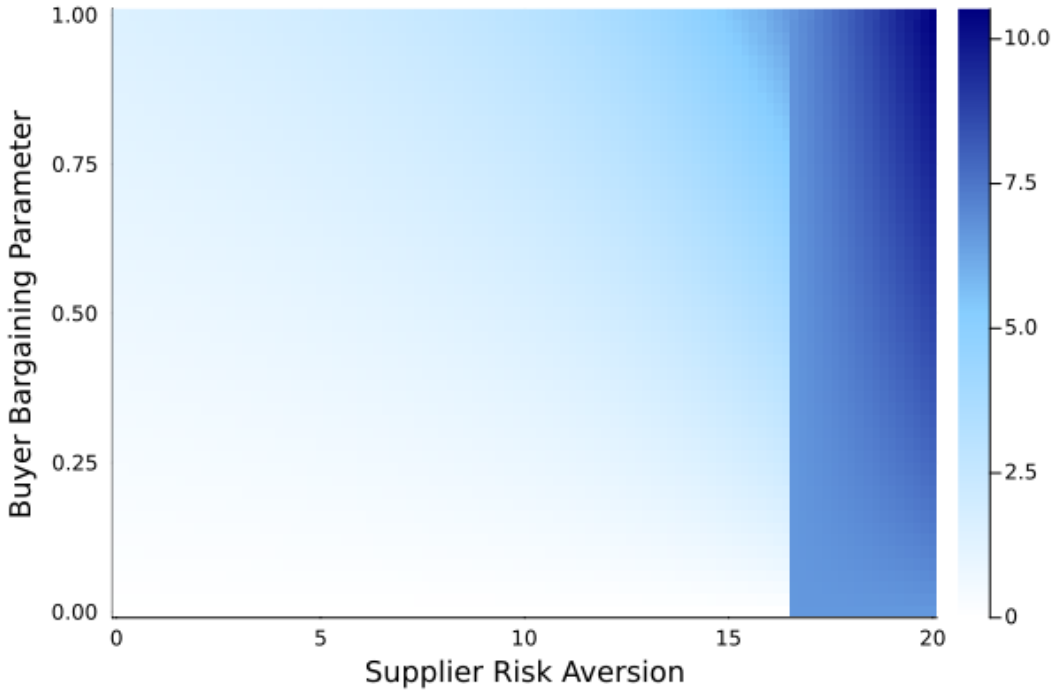
My NNTR estimation approach aims to capture the strategic incentives for threats of replacement in Ho and Lee (2019) but in a bargaining environment similar to their setting. Namely, buyers and suppliers can freely form new agreements and there is neither spot market nor integrated production. To focus on the buyer's strategic incentives, I hold the supply side constant. However, I update the buyer model in three key ways. First, I model the buyer's outside option of no agreement as losing the profit from the order entirely. This assumption reflects that, absent the spot market, the buyer cannot source fabric to fulfill the order. To do so, I calibrate profit margins for a garment manufacturing using estimates from Cajal-Grossi et al. (2023).[78] Second, I model the integrated supplier as an additional external supplier, as the bargaining environment in Ho and Lee (2019) does not include any integrated production. I assume that the integrated supplier has zero risk aversion, as it is part of a large firm, and the buyer can extract all surplus from it. Then, I implement the bargaining model as in NNTR, finding the average price that would make the buyer indifferent between keeping a supplier and switching to the excluded supplier for each excluded

[78]Specifically, the markups as reported in Online Appendix Table E.7, adjusting for fabric costs.

supplier. Inside agreements, suppliers charge the minimum of the bargained price and the agreement price. I find that the buyer actually uses the integrated supplier as the excluded supplier, as doing so brings down prices from included suppliers.

A.21 Buyer Gains from Agreement with DICC

Figure A.28: Buyer Gains from Agreement (% of market price)



Note: Model estimates for the discount in the agreement relative to the market as a percent of the market price, incorporating dynamic incentive compatibility constraints. Heatmap legend is to the right of the heatmap. Model estimates use the average capacity and quantity assurance across agreement suppliers. The discount factor is calibrated based on microfinance institution insurance rates in India; to be conversative, I use a low discount factor associated with 150% of standard market interest rates in Indian microfinance such that the quarterly discount factor is .88 (implying a yearly discount factor of .60).

A.22 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%.[79]

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 quantity 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.23 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.10: Buyer Bargaining Parameter and Supplier Concentration

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

[79]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.

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

I find in Appendix Table A.10 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.11: 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 the buyer and model estimates.