Value Creation & Value Capture in Indian Garment Sector Bargaining

Russell Morton*

September 1, 2025

Click Here for Latest Version

Abstract

This paper examines how prices along export-oriented low- and middle-income country (LMIC) supply chains reflect surplus creation and capture by buyers and suppliers. While exporting can enhance growth and productivity, low prices paid to small suppliers raise concerns that gains accrue primarily to large exporters (buyers). However, when input markets, or complementary markets like risk or credit, are frictional, complex vertical contracts govern prices that reflect more than standard oligopsony markdowns. I quantify how two features of such settings shape prices. First, contracts create surplus by mitigating frictions—for example, specifying stable quantities to reduce profit variability when insurance markets function poorly—so low prices may reflect the buyer's share of gains. Second, buyers reduce dependence on frictional markets by partially sourcing in-house, potentially enabling surplus capture from threats to replace external suppliers with in-house production. I estimate a structural bargaining model using novel data on a major Indian garment manufacturer's universe of fabric purchases from ~ 500 suppliers. I find that surplus creation from insurance-like contract terms explains low prices, with buyers receiving $\sim 60\%$ of gains, rather than capture from replacement threats with in-house production; difference-in-differences estimates yield consistent findings. Counterfactuals illustrate that while greater buyer competition has limited effects on prices paid to small risk-averse suppliers, profit insurance substantially increases prices.

JEL Classification: D23, F63, L14, O12

*University of Rochester, russell.morton@rochester.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, Francesco Amodio, Gianluca Antonecchia, Jie Bai, Hoyt Bleakley, Binta Zahra Diop, Mayara Felix, Rosa Ferrer, Bob Gibbons, Max Huppertz, Chuqing Jin, Amit Khandelwal, Asim Ijaz Khwaja, Francine Lafontaine, David A. Miller, Charlie Murry, Ameet Morjaria, Emir Murathanoglu, Anh Nguyen, Anant Nyshadham, Ellie Prager, Michael Raith, Eleanor Wiseman, and Yulu Tang for their excellent feedback and comments, as well as seminar participants at Chicago DIL, CSAE, EEA, Firms, Labor Markets, & Development, IIOC, LSE IO-Devo Workshop, MWIEDC, NEUDC, Northwestern, OECD, PacDev, Paris School of Economics, SIOE, SITE, University of Michigan, University of Rochester (Simon Business School), USC, Wayne State University, and Yale University. I appreciate the team at Shahi Exports and Good Business Lab for sharing their expertise about fabric procurement and providing data. All errors are my own.

1 Introduction

While exporting can reduce poverty and increase productivity in low- and middle-income countries (LMICs) (World Bank, 2020), low prices often received by small suppliers raise concerns among policymakers that exporters, who are the buyers domestically, capture most of the gains (Boudreau et al., 2023). This concern seems especially compelling in LMIC supply chains, which often feature only a few large exporters purchasing from many small suppliers (Dominguez-Iino, 2023; Zavala, 2022). But, in the presence of frictions in markets for the input itself, or in complementary markets like risk and credit, supply chain transactions are governed by multifaceted vertical contracts. Importantly, prices in such vertical agreements reflect more than standard oligopsony markdowns.

This paper uses proprietary data on the universe of fabric purchases by a large export-oriented Indian garment manufacturer to quantify how two key features of rich vertical contracts shape prices when markets function poorly. First, agreements create surplus by mitigating frictions in the market for the input or auxiliary markets, such as risk and credit. Fabric mills (the suppliers) face variable, lumpy orders that expose them to quantity and profit risk. Agreements can specify stable quantities ("*quantity assurance*"), providing insurance against such risk.[1] In the absence of insurance markets, suppliers may accept lower prices in exchange for quantity assurance; then, price levels may be low because they bundle the transfer for the input with an insurance premium. (In other settings, agreements address different frictions, such as providing trade credit given financial market imperfections;[2] then, prices reflect both the transfer for the input and interest.) Second, poorly functioning markets lead many firms—including the buyer studied—to source some of an input in-house.[3] Partial vertical integration could enable threats to replace external suppliers with in-house production, strengthening the buyer's bargaining position and ability to capture surplus.

Although both mechanisms reflect poor market performance and can rationalize low prices for suppliers, distinguishing between them matters because their distributional implications differ. If low prices embody surplus created by mitigating frictions inside the vertical

[1] Agreements between Kenyan rose buyers and exporters (Macchiavello and Morjaria, 2015) and fishing firms and ice manufacturers in Sierra Leone (Ghani and Reed, 2022) also include quantity assurance.

[2] Rwandan coffee mills (Macchiavello and Morjaria, 2020) and Ecuadorian manufacturers (Brugues, 2025) provide trade credit.

[3] In Karnataka, India, most integrated firms still source inputs externally (Garg et al., 2023). Breza and Liberman (2017) study a Chilean retailer that sources inputs both internally and through agreements. As transaction costs theory predicts integration as a response to poor spot market performance (Williamson, 1979), external suppliers of partially integrated firms are likely agreement, rather than spot, suppliers. Morton (2023) provides empirical evidence that integration and agreements act as substitutes

agreement, they represent the buyer's share of that surplus. Conversely, replacement threats reduce prices exclusively through rent extraction. Buyer market power arises not only from replacement threats, but also shapes agreement surplus sharing. Crucially, buyer power over surplus creation aligns incentives for agreement terms that mitigate frictions, such as insurance-like terms, and generates the discount from surplus creation. Because replacement threats purely enable rent extraction, I reserve "capture" for their effects on prices.

Broadly, this paper develops a framework to interpret vertical contract prices when agreements embed non-price terms (*e.g.,* insurance or credit). By decomposing prices into surplus creation and capture, the framework illustrates how standard oligopsony markdowns present an incomplete picture of price levels.

I must overcome two identification challenges to separately quantify how prices reflect surplus creation from insurance-like terms and capture from replacement threats with in-house production. First, because agreements are informal (*i.e.*, without an official contract), I do not directly observe contract terms. Instead, I infer terms by comparing transactions across agreement and spot suppliers with otherwise identical price-relevant characteristics. Comparisons leverage the richness of the data, which include ∼25,000 transactions and highly granular observables (*e.g.*, ∼6,200 product categories). Two descriptive facts emerge: *i)* agreements reduce, without eliminating, quantity variability; and *ii)* agreements feature non-linear prices that are highest when profits would be lowest with constant prices (and *vice-versa*), while spot suppliers receive flat prices. (I show that unobservables, such as supplier efficiency, are unlikely to explain such differences.) As both quantity assurance and non-linear prices smooth profits, I interpret them as insurance-like agreement terms.

The second identification challenge arises because prices in bargained agreements simultaneously reflect three components: *i)* surplus from mitigating market frictions (*i.e.*, the inside option), which are gains from insurance-like agreement terms in my setting; *ii)* the value of the next best alternative (*i.e.*, the outside option), which incorporates replacement threats from in-house production; and *iii)* the surplus sharing rule, summarized by the buyer bargaining parameter. While the institutional context enables direct estimation of the outside option,[4] I develop a new approach to separate the *level* of surplus generated by insurance-like agreement terms from the *sharing rule* for it. My strategy leverages exogenous demand shocks over time within each supplier.[5] As price adjustments to demand shocks provide insurance, the magnitude of price adjustment identifies supplier preferences for insurance-like

[4]The outside option is the spot market because *i)* agreements with quantity *stability* cannot easily expand and *ii)* new agreements cannot emerge in the short run given the time and effort required to establish them.

[5]Demand shocks reflect institutional features of garment manufacturing (described in Section 2.2).

terms. Conversely, average price levels identify buyer bargaining parameters.

I build a model of the vertical contract that embeds both surplus creation from insurance-like terms and rent extraction from threats to replace external suppliers with in-house production into Nash bargaining. The model formalizes how supplier preferences for smooth profits shape the price effects of insurance-like agreement terms. Such preferences can represent many microfoundations,[6] including borrowing and savings market frictions that push suppliers to fund frequent outflows (*e.g.*, payroll) from current profits; differences in supplier financial access can generate preference heterogeneity. For brevity, I hereafter refer to such preferences as "risk aversion." As more risk-averse suppliers derive greater value from insurance-like terms, agreement surplus increases in risk aversion. Because the buyer receives its surplus share through discounts, greater surplus reduces prices. Therefore, surplus creation from insurance-like terms implies prices decrease in risk aversion.

The model both *i)* extends standard replacement threat results (Ho and Lee, 2019; Ghili, 2022) by showing threats from in-house, rather than external, suppliers improve bargaining positions, and *ii)* introduces a novel cutoff: such threats only affect prices when suppliers are sufficiently risk averse. When supplier risk aversion falls below the cutoff, agreements break down under replacement threats as insurance-like terms generate insufficient surplus. Because agreement and integrated quantities are inframarginal,[7] breakdown increases spot sourcing rather than internal production (given convex costs in-house).[8] But the buyer prefers pre-integration agreements to spot supply since agreements reduce average prices. Therefore, the buyer "as if" endogenously chooses its outside option, using in-house supply as a replacement threat with more risk-averse suppliers while maintaining agreements (with spot supply as the outside option) with less risk-averse suppliers to preserve low prices.[9] Overall, because agreement and integrated sourcing is inframarginal, with marginal supply

[6]Small firms may inherit risk aversion from the owner—the residual claimant of profits.

[7]Similarly, Darmouni et al. (2023) model agreements with quantity assurance as inframarginal, with marginal sourcing from the spot market.

[8]Integrated supplier cost convexity makes the buyer prefer increasing spot sourcing rather than expanding production from capacity-constrained in-house suppliers.

[9]I show formally that the buyer endogenously choosing its outside option is a reduced-form representation of two extensions of a Nash bargaining game that consider a contractual equilibrium with both cooperative and non-cooperative stages (Miller and Watson, 2013). The first extension considers a model with a signaling game about the buyer's residual in-house capacity, which determines the buyer's actual outside option, before cooperative bargaining. Residual capacity as private information is consistent with the setting because transactions with other suppliers and in-house have limited, if any, observability to any individual external supplier; furthermore, residual capacity has stochastic variation because orders from end clients are lumpy. The second extension considers the buyer committing to a level of residual in-house capacity before bargaining with agreement suppliers. Such commitment mirrors commitment elsewhere in the agreement, such as the buyer paying agreement prices over spot market prices in some periods.

from the spot market, the buyer refrains from threatening to replace suppliers with low risk aversion. This restraint prevents agreement breakdown that would force costly spot sourcing. Therefore, this model highlights a novel channel through which replacement threats may not reduce prices arising from the interaction of in-house, agreement, and spot supply.

I then use the model to quantify the relative importance of surplus creation from insurance-like agreement terms and surplus capture from replacement threats with in-house production. I find that surplus creation alone, rather than rent extraction from replacement threats, generates all price reductions relative to spot prices. However, the buyer's market power enables it to capture ∼60% of surplus created, but not through replacement threats. This result reflects that all supplier risk aversion estimates fall below the model-implied cutoff where the buyer makes replacement threats. The median estimated risk aversion aligns with estimates for LMIC coffee mills (Blouin and Macchiavello, 2019).[10] On average, agreements yield a 1.5% discount relative to spot market prices, saving the buyer ∼0.9 million USD annually. This discount is meaningful compared to spot markups of ∼7% over costs.

To validate that replacement threats from in-house production do not affect prices, I use a difference-in-differences (DiD) design. Identification leverages both *i)* a change in vertical integration—the buyer builds an integrated mill—and *ii)* heterogeneity in supplier exposure to it.[11] This validation complements estimates from the structural model because the two approaches rely on distinct sources of variation: exogenous demand shocks identify risk-aversion parameters in the model, while DiD exploits quasi-experimental changes in integration. Consistent with model predictions as estimated supplier risk aversion falls below the cutoff for value capture from replacement threats, DiD estimates show no effects of integration on prices. Furthermore, DiD estimates are inconsistent with prices predicted by a standard replacement threat model in a setting *without* a co-existing spot market where buyers leverage replacement threats to lower prices (Ho and Lee, 2019).

Through the lens of the model, I quantify both *i)* the cutoff level of risk aversion above which the buyer makes replacement threats with in-house production and *ii)* the price effects of such threats. The model-implied risk aversion cutoff exceeds estimates both for agreement suppliers and from other settings, although is not orders of magnitude larger.[12]

[10]The LMIC coffee mills in Blouin and Macchiavello (2019) report revenues that average ∼ 3.5 million USD per year. Although these coffee mills are large firms by LMIC standards, they are much smaller than the buyer, which purchases nearly 180 million USD of fabric annually.

[11]Heterogeneity in supplier exposure to integration derives from suppliers' product mix, with only external suppliers who produce fabric that can be made at the integrated supplier considered exposed to integration.

[12]Settings that feature agreements with more quantity stability or higher cost convexity would feature a lower cutoff level of risk aversion.

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 exceeds Indian Value Added Tax for apparel at the time (5%).[13] And, the most risk-averse suppliers are likely small firms with policy and empirical relevance; 95% of Indian industrial units have 10 employees or fewer (Rajagopalan and Shah, 2024).[14]

Last, I use the estimated model to evaluate two counterfactual policies, focusing on supplier risk aversion. First, I consider increased buyer competition (*e.g.*, from export promotion), modeled as a higher supplier surplus share to prevent long-run switches to another buyer. (Establishing agreements requires time and effort, preventing short-run switches.) For a supplier with median estimated risk aversion, increased competition reduces the agreement discount relative to spot prices by 75%. However, for suppliers with risk aversion above the cutoff where the buyer extracts surplus from replacement threats, agreement surplus relative to in-house production is so small that even allocating all to the supplier barely affects prices. (As surplus relative to the spot market, instead of in-house production, is large, omitting replacement threats overstates the benefits of increased competition.)

Second, introducing the missing product for profit insurance against spot market quantity risk increases prices broadly, and especially for risk-averse suppliers. Insurance enables suppliers to smooth profits outside the agreement, improving their outside option and removing the buyer's ability to extract surplus through replacement threats. While moral hazard may limit the feasibility of profit insurance, policies could promote profit smoothing by increasing financial market access (facilitating self-insurance), supporting cooperatives to stabilize profits, and developing products to hedge risk (*e.g.*, derivatives). The effectiveness of insurance highlights that helping risk-averse firms to smooth profits generates benefits not only through equalizing utility across states but also by raising prices; therefore, insurance improves equity by counteracting market power from replacement threats.

While broadly at the intersection of empirical industrial organization and development (*e.g.,* Jimenez-Hernandez and Siera (2022); Chen and Xu (2023); De Simone and Brugues (2024); Hsiao (2024)), this paper primarily advances three specific literatures. First, it contributes to prior empirical bargaining work (*e.g.*, Grennan (2013); Keniston (2017); Backus et al. (2020); Gaineddenova (2022); Keniston et al. (2022); Freyberger and Larsen (2023)) by studying agreement features that shape inside and outside options when markets function poorly. I highlight inside option surplus created by agreement terms that mitigate spot

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

[14] As a validation exercise, I estimate the model using only pre-integration periods and show that it fits well in post-integration periods, supporting the model's out-of-sample performance.

market frictions. Regarding outside options, while Ho and Lee (2019) and Ghili (2022) illustrate that firms threaten replacement to improve their bargaining position, my model shows that buyers sourcing marginal quantities from spot markets *refrain* from replacement threats when such threats break agreements.

Because the trade-off between incentives and insurance is at the center of vertical relationships (Lafontaine and Slade, 2007), I incorporate insurance-like agreement terms and risk aversion into my empirical model and develop a new approach to separately identify risk aversion from bargaining parameters. My approach leverages how demand shocks affect both the *level* and *curvature* of the agreement pricing function. Empirical models of vertical relationships (*e.g.*, Cuesta et al. (2019); Lee et al. (2021)) that abstract from risk aversion underestimate the more risk-averse party's bargaining parameter. Even with correct parameters, omitting insurance-like terms can generate incorrect prices and surplus estimates. In my empirical setting, omitting quantity assurance—central to these vertical contracts—yields underestimated supplier surplus and overestimated prices.[15]

This paper contributes to the literature on relational contracts—agreements that facilitate cooperation through future surplus rather than formal contracts—as the agreements studied are relationally enforced (Harris and Nguyen, 2023; Macchiavello and Morjaria, 2015, 2020, 2023; Ghani and Reed, 2022; Brugues, 2025). I emphasize the spot market frictions that relational contracts mitigate—and quantify resulting price effects—rather than the contracting frictions that lead to relational enforcement. I also measure relational surplus sharing, which is understudied. To do so, I build and estimate a bargaining model; standard models analyzing relational enforcement consider take-it-or-leave-it offers, attributing all surplus to one party. Last, I illustrate how relational contracts respond to changes in the *vertical* organization of production, building on prior studies of changes to market structure—the *horizontal* organization of production. Not only is the co-existence of vertical integration and relational contracts likely empirically relevant, as both suggest weak spot markets, but it also introduces strategic considerations since vertical integration—unlike market structure—is a choice potentially facilitating rent extraction via replacement threats.

The remainder of the paper is organized as follows. Section 2 introduces the setting and data. Section 3 develops the model. Section 4 presents estimation and results. Section 5 evaluates policy counterfactuals, and Section 6 concludes.

[15]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.

2 Data, Context, & Empirical Facts of Buyer-Supplier Agreements

My empirical setting centers around an Indian garment manufacturer (the buyer) sourcing fabric simultaneously from agreement suppliers, the spot market, and in-house production. First, I describe the data. Second, I provide context on fabric procurement. Third, I empirically describe bargained agreements, documenting two stylized facts about agreement terms: *i)* agreements provide quantity assurance; and *ii)* agreement prices have a U-shaped relationship with quantities, including a discount on average. I interpret these facts as agreements providing insurance to suppliers. Fourth, I highlight an observed increase in vertical integration from the buyer building its own new mill. Last, I describe the buyer's overall fabric sourcing, illustrating that *marginal* quantities are from spot suppliers.

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 include not only the transaction price and currency, quantity, and date but also detailed information about the specific fabric (6,176 codes) and supplier identity. The 34,681 transactions represent over one billion square meters of fabric purchased from 622 distinct suppliers (471 pre-integration). The buyer also uses two integrated suppliers: one built during the middle of the study timeframe (June 2018), generating variation in the buyer's ability to threaten to replace external suppliers with in-house production, and the other producing fabric for the buyer long before the start of the transaction data (exact date unknown). Additionally, the transaction data include estimated production costs for fabrics produced internally. Cost estimates are for a mill that is not capacity constrained and represent average variable costs.

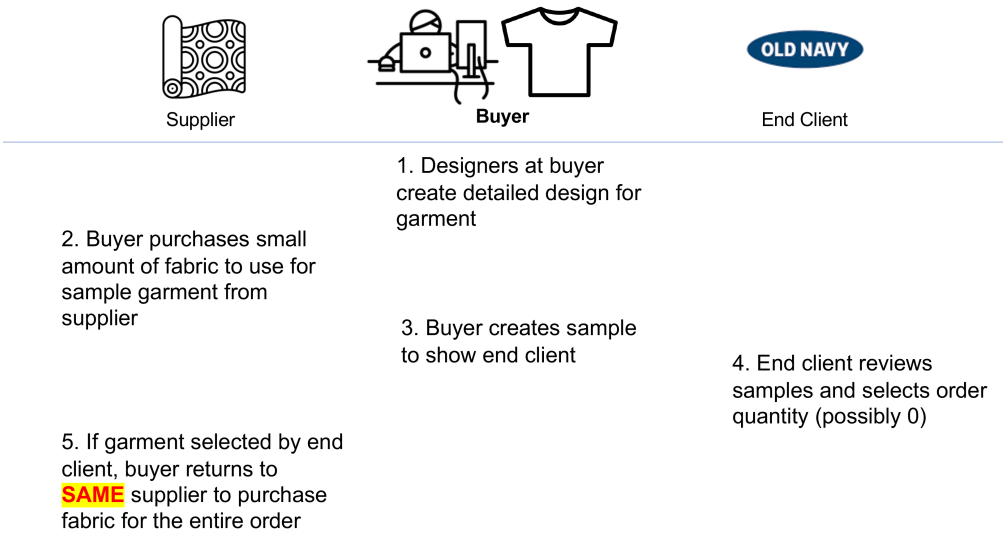
Cost data from the integrated mills provide information about their capacity and costs at the division[16]-by-month level from March 2019 to April 2020. The data include a broad support of capacity levels from 41% to 102%, with capacity over 100% typically indicating that the mill is operating extra shifts on Sunday and/or overnight.

[16]Divisions refer to specific processes for fabric production without substitution across each other, such as weaving, dyeing, and knitting.

2.2 Context: Fabric Procurement and Garment Manufacturing

The garment manufacturing process (in Figure 1) begins with buyer-employed designers sketching the shape and choosing a broad fabric type and color. The buyer then selects a supplier (a mill) to provide fabric for a few prototype garments, called "samples".[17] The end client (*e.g.*, Old Navy) then reviews samples, approves a subset, and specifies quantities. For approved samples, the buyer returns to the *same* supplier from sampling to ensure that fabric exactly matches, as subtle fabric features are hard to recreate and even minor deviations risk rejection by the end client.[18] Furthermore, short timelines preclude trial-and-error with a different supplier—typically, only 28 days separate fabric purchases for samples and the entire order, consistent with "fast fashion."[19]

Figure 1: Garment Manufacturing Process



This process has two primary economic implications: suppliers *i)* face quantity risk and *ii)* can hold-up the buyer. First, suppliers face quantity risk because they provide fabric for the entire order if the sample is chosen, but know neither which designs will be approved nor in what quantities. Second, because suppliers realize that the buyer cannot easily switch

[17] Although the buyer pays for this fabric, the small quantity renders the cost negligible (records are not even maintained).

[18] Discussions with fabric procurement managers at the buyer indicate that the probability that a supplier which provided the sample is chosen for the entire order is at least 90%.

[19] In the unlikely case that the supplier changes after sampling, production is typically brought in-house, facilitating stringent quality control, and is for one of the few orders with a long time before delivery, enabling experimentation to match the fabric perfectly.

suppliers of approved samples, suppliers charge a high hold-up price—equivalently, suppliers act as monopolists for the order.[20] (Online Appendix OA1.4 discusses why the buyer and supplier do not agree to a price for the full order during sampling.)

2.3 Stylized Facts of Buyer-Supplier Agreements

These two features of garment manufacturing—quantity risk and high prices from hold-up—suggest that agreements can generate surplus at lower prices by reducing quantity variability. I document two stylized facts about agreements: (i) they feature quantity stability, and (ii) their prices follow a U-shape in capacity utilization.

Because agreements are informal and unwritten, I recover their terms empirically by comparing quantities and prices across agreement and spot suppliers. To do so, I must first classify supplier type as the data do not specify whether a supplier is agreement or spot. I base classification on transaction frequency because agreements are enforced relationally through future surplus, which requires frequent transactions. This classification strategy mirrors common approaches in the relational contracting literature.[21] It categorizes the 19 external suppliers that transact every month as agreement suppliers.[22] These 19 suppliers account for 25% of pre-integration fabric volume but only ∼4% of the 471 pre-integration external suppliers.

2.3.1 Fact 1: Agreements Provide Quantity Assurance

First, I show empirically that agreements provide quantity assurance (*i.e.,* reduce quantity variability). Figure 2 documents that the coefficients of variation for quarterly volume per supplier are generally lower for agreement than spot suppliers.[23] This result is not

[20]The buyer is unlikely to hold up a supplier given the timing of garment production. Namely, fabric prices are set after the end client has approved the garment and specified the quantity. At this point, the specific fabric is very valuable to the buyer, but the supplier could still sell capacity to other buyers, such as other garment manufacturing firms or furniture producers who use fabric as stuffing. The buyer's primary opportunity to hold up a supplier would presumably occur after a supplier has produced the fabric for the entire order. However, unlike the agreement overall, the specific transaction has a simple explicit contract that is likely court-enforceable. Furthermore, the buyer may have reputational motives to avoid such hold-up.

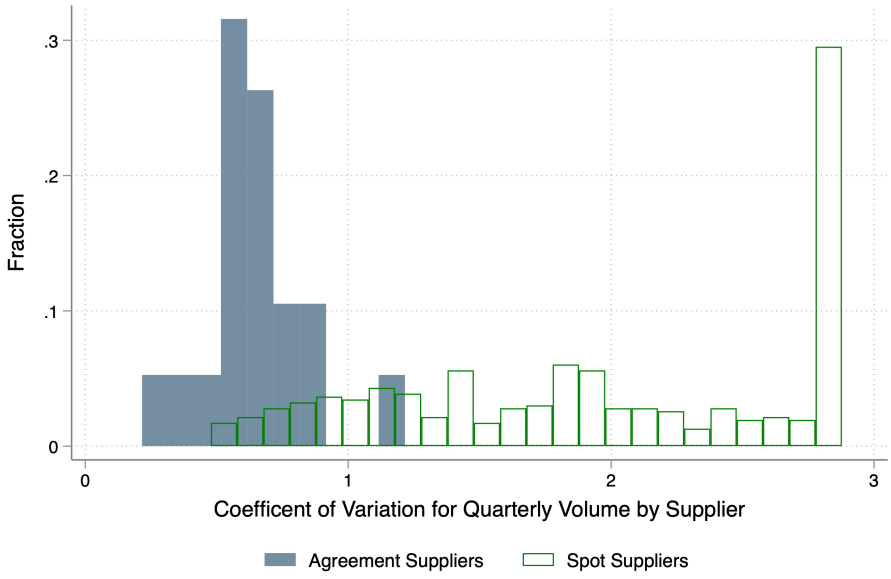
[21]For example, Macchiavello and Miquel-Florensa (2017) similarly define relational contracts based on consistent trading every period ("at least three consecutive seasons") and Macchiavello and Morjaria (2015) also use transaction frequency. This relational contract is *not* a relational *incentive* contract (Levin, 2003), as it does not (at least primarily) address information problems (*e.g.,* subjective evaluation).

[22]I only consider pre-integration for supplier classification to ensure that analyses of the effects of increased vertical integration on agreement suppliers hold supplier classification fixed at baseline. Appendix Figure A.1 shows the distribution of the percentage of months with transactions across suppliers.

[23]The difference in average coefficient of variation by supplier type is statistically significant ($p < .001$).

entirely mechanical: supplier classification depends on transaction frequency, not volume, so agreement suppliers could in principle have volatile quantities. Appendix Section A.1.2 provides additional evidence of quantity assurance by showing that agreement suppliers who are "owed" quantities to meet their typical volumes[24] receive additional volume (and vice-versa). This additional volume exceeds what mean reversion would imply. Conversely, spot suppliers display the *opposite* pattern.[25] The buyer can provide quantity assurance through at least two channels: *i)* increasing the number of samples and *ii*) matching the supplier to end clients with different expected quantities.

Figure 2: 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.

2.3.2 Fact 2: Agreements Specify a U-Shaped Pricing Function

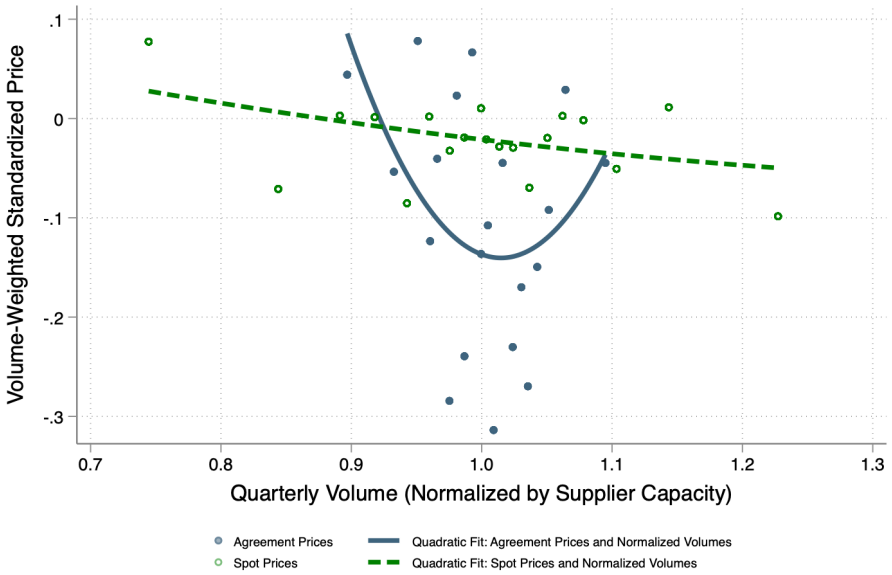
The second fact illustrates empirically that agreements feature a U-shaped pricing function with respect to quantities (normalized by supplier capacity), while spot suppliers receive flat prices. To show this fact, I compute each supplier's volume-weighted average standardized price and examine how it varies with capacity utilization. Standardizing prices by month and fabric isolates the extent to which suppliers hold up the buyer by removing cost differences

[24]Formally, suppliers with lower short-run quantities than typical.

[25]Specifically, spot suppliers experience *less* volume in months with below-typical quantities (and vice-versa).

from seasonal fluctations[26] or varying input prices across fabrics.[27] Quantities are expressed as quarterly volumes relative to supplier productive capacity—effectively, quarterly supplier-specific *capacity utilization*—to account for heterogeneity in supplier scale.

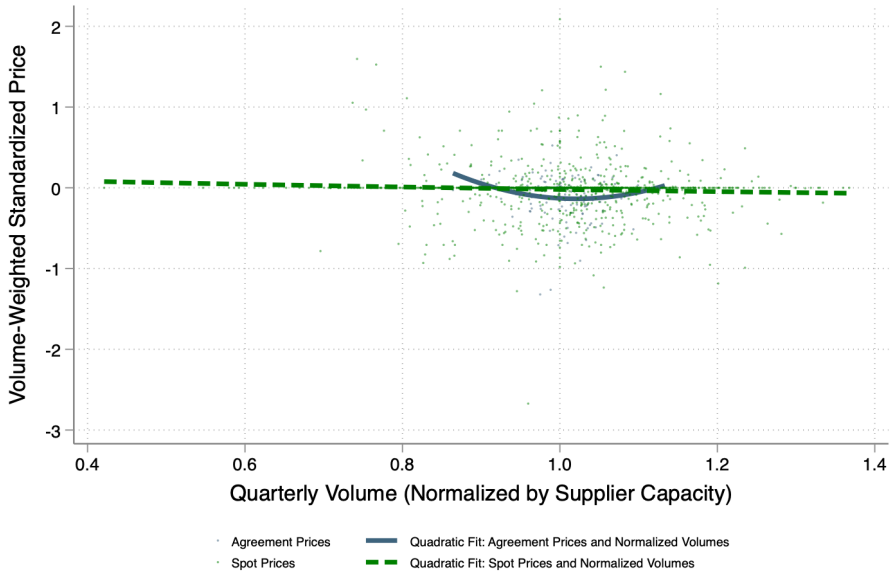
Figure 3: Pricing Function: Binned



[26]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 Online Appendix Figure OA1.1.

[27]To illustrate, consider purchases for different fabrics with different input costs. If the two suppliers charge the same price, the supplier of the high-cost fabric charges a lower markup. Alternatively, I could measure markups directly using the buyer's cost data, but dcost data are only available for the limited set of fabrics they produce.

Figure 4: 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 3 bins prices and capacities using 20 quantiles. The mapping from quantities to capacities is as discussed in 4.1.2.

Figures 3 and 4 illustrate that agreement suppliers receive U-shaped prices with respect to quarterly volumes normalized by supplier capacity (*i.e.,* capacity utilization), while spot prices remain flat. The U-shape is clearest with binned data, the remains visible in the raw scatter despite a much wider y-axis scale. Appendix Section A.1.3 confirms that only agreement, and not spot, prices statistically follow a U-shaped pattern. Furthermore, average agreement prices are lower than spot prices, illustrating that agreements include a discount (shown explicitly in Online Appendix OA1.1).

2.4 Interpretation as Insurance-Like Agreement Terms

I interpret both quantity assurance and the U-shaped pricing function as insurance-like agreement terms that smooth supplier profits. First, quantity variability creates profit variability;[28] therefore, quantity assurance stabilizes profits. However, agreement quantity stabilization is incomplete because orders are lumpy with quantities unknown at sampling and

[28]Quantity variability generates profit variability whenever quantity changes are not perfectly offset by changes in markups.

cannot be shared across suppliers.[29]

Second, the U-shaped pricing function also reduces profit variability compared to flat prices, without necessarily increasing average profits. When quantities are low, profits with constant prices would be low because there are few units; high prices for low quantities help smooth profits. When quantities are high, increasing marginal costs (*e.g.,* due to capacity constraints, and verified empirically using data from in-house suppliers in Figure 10),[30] erode margins under constant prices. High prices at high capacity utilization offset these costs to stabilize profits.[31]

Appendix Section A.1.3 also establishes that agreement prices are lowest when capacity utilization is 100%, supporting the interpretation that U-shaped prices provide insurance. Prices reach their minimum precisely when profits would be highest with constant prices, consistent with agreement prices smoothing profits.[32] (Separately, I interpret spot prices as suppliers holding up the buyer. Therefore, spot prices are determined by the *buyer's* outside option. Their flat relationship with supplier capacity reflects that the buyer's outside option is independent of supplier capacity—any single order is small for the buyer overall.)

Last, I provide evidence that the stylized facts are unlikely to reflect two alternative explanations: *(i)* unobserved supplier heterogeneity drives differences in quantities and prices between agreement and spot suppliers , and *(ii)* misclassification of suppliers as agreement versus spot.

First, I consider unobserved supplier heterogeneity. For example, some suppliers could be more efficient, perhaps due to better management as technology is largely standardized (at least among suppliers that satisfy the buyer's screening). More efficient suppliers might pass on cost savings to the buyer, reducing average prices. However, efficiency alone does not naturally explain why agreements feature more stable, rather than larger, quantities nor a U-shaped pricing function. As a test to mitigate this concern, I replicate the agreement

[29]Limited quantity assurance reflects that the buyer's primary instruments to stabilize quantities—frequency of sampling and the distribution of quantities for a given end client—do not perfectly match to a specific quantity.

[30]Supplier cost convexity likely reflects frictions in adjusting inputs from both regulations and financial constraints. For example, the Factories Act of 1948 requires double wages for overtime hours and buyer discussions emphasize capacity management for cost control.

[31]Complete insurance through prices is infeasible for both incentive and practical reasons. On the incentive side, full insurance would require extreme (even infinite) prices for some quantities, or large fixed quarterly fees, either of which could induce the buyer to deviate. Practically, agreements span over a quarter, so the last order is not known even when it occurs. Therefore, quarterly revenue and costs are uncertain until the very end of the quarter, preventing an incentive-compatible spread of payments over time.

[32]This result assumes that suppliers target capacity such that full capacity utilization balances costs and quantities to maximize profits with constant prices.

versus spot supplier comparisons using a restricted subset of high-volume spot suppliers ("placebo suppliers").[33] Relational enforcement of agreements requires frequent transactions (and quantity assurance likewise relies on many transactions in practice). Therefore, placebo suppliers are plausibly not in agreements despite receiving large volumes. Importantly, plausible explanations of unobserved heterogeneity largely imply that supplier volume, rather than transaction frequency, should shape transactions. For example, if supplier efficiency explains lower prices, then suppliers with similar volumes—regardless of frequency—should charge lower prices. Instead, Online Appendix OA1.2) shows that differences in quantities and prices persist even when comparing with placebo suppliers.[34]

Second, determining whether suppliers are agreement or spot based on transaction frequency may introduce classification error; frequent transactions are necessary but not sufficient for agreements to emerge. In general, such measurement error does *not* imply that the stylized facts about agreements—the quantity assurance and U-shaped pricing function—are incorrect; instead, such error should bias comparisons *against* finding meaningful differences across supplier types. Regardless, the placebo analysis suggests that any misclassification is minor, as many placebo suppliers trade with the buyer almost every month but display different contracting patterns empirically. Moreover, Online Appendix OA1.3 shows that quantity variability and average prices—as the U-shaped pricing functions include a discount—change by a statistically significant amount at the threshold for classification as an agreement supplier.[35] Because subsequent empirical analysis estimates supplier-specific parameters that rationalize observed agreements, misclassifying spot suppliers as agreement suppliers is much more problematic than the reverse. Therefore, I adopt a conservative classification that requires transactions every month.

While interpreting the stylized facts of agreement quantities and prices as insurance-like agreement terms highlights how risk-averse suppliers benefit from the agreement, it raises two natural questions. First, why does the buyer not prefer an explicit contract at the time of sampling that specifies a price (or even a price schedule conditional on quantities to accommodate cost convexity), thereby resolving the hold-up problem by making suppliers compete before sampling? Online Appendix OA1.4 discusses why the buyer and suppliers rely on agreements rather than this explicit contract. Second, if the buyer only receives

[33] Appendix A.2 describes the construction of the set of placebo suppliers.

[34] This analysis 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 rather than any efficiency differences that may, or may not, exist.

[35] Specifically, at the threshold of transacting every month versus nearly every month.

discounts from agreement suppliers, why not source exclusively from them? Online Appendix OA1.5 highlights forces that prevent the buyer from sourcing exclusively from agreement suppliers.

2.5 Context: Vertical Integration at a Large Buyer

Figure 5: Integrated Supplier Constructed in June 2018



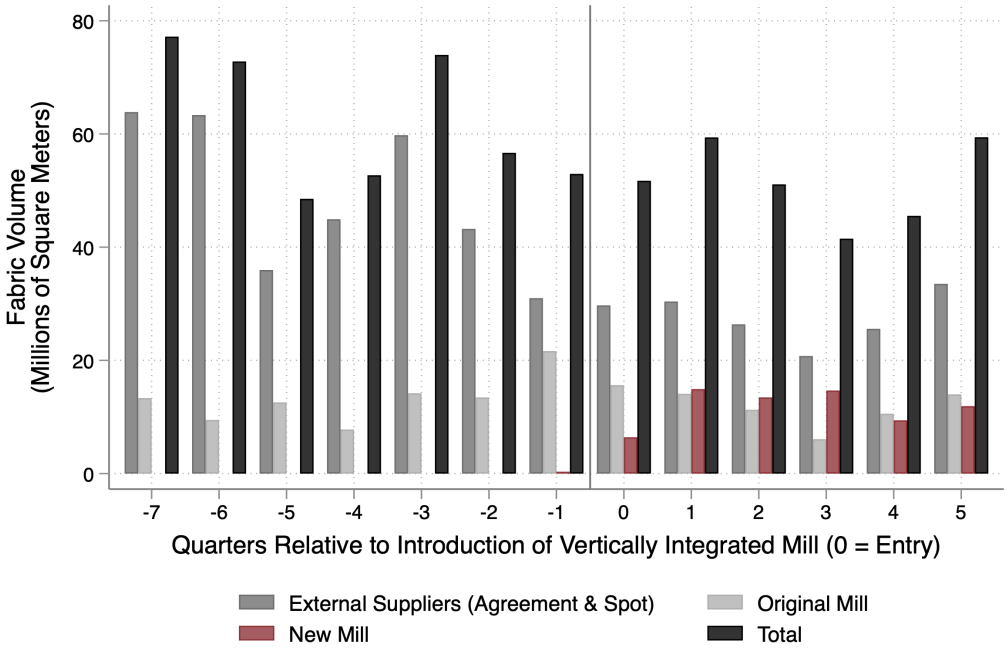
The buyer builds a new mill in the middle of the sample. Anecdotally, the buyer reports that high spot fabric prices motivate integration, consistent with spot suppliers holding up the buyer. (Empirically, cost data from the integrated supplier suggest that spot markups indeed average 7% over average variable costs.) Figure 5 shows the new mill. The organization of the mill emphasizes the buyer's sophistication (and likely scale), in constant to the characterization n Bloom et al. (2012) of large Indian garment manufacturers as messy and disorganized.

2.6 Context: Overall Sourcing

Having described agreement and in-house sourcing, I examine the buyer's overall sourcing across all three channels—agreement, spot, and in-house suppliers. Figure 6 shows volumes by supplier type throughout the sample. The large size of the buyer is immediately apparent: it sources, on average, ∼$180 million of fabric per year (associated with, on average, ∼233 million square meters of fabric per year). For comparison, 90% of Indian garment manufacturers report revenues (which are presumably larger than their fabric procurement costs) under $60,000, emphasizing how large the buyer is compared to the modal Indian garment

manufacturer.[36]

Figure 6: Fabric Sourcing Over Time



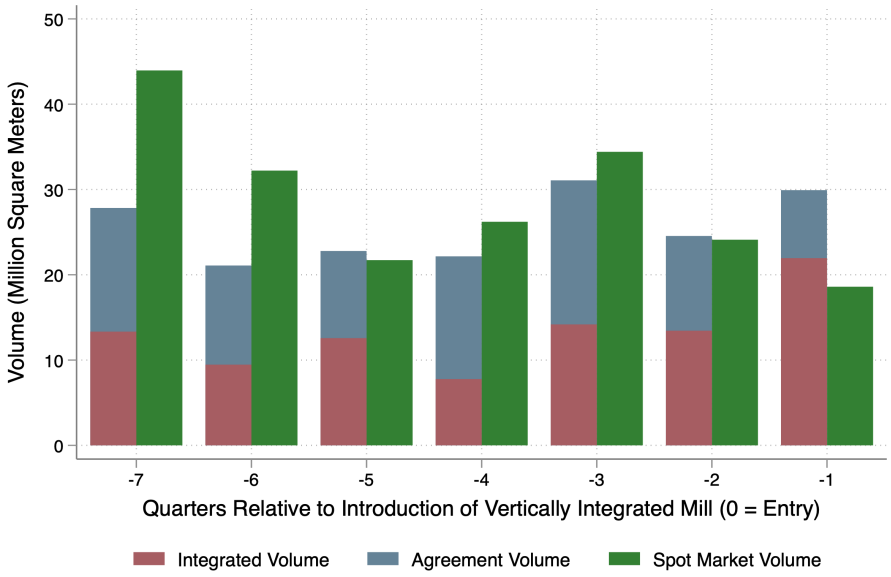
Note: Data from universe of fabric transactions by buyer. Bars represent the quantity of fabric purchased during the quarter from each supplier type.

Three features of this setting make it particularly well-suited for analyzing the price effects of insurance-like agreement terms and replacement threats from in-house supply. First, external suppliers remain empirically central, providing the majority of fabric. Therefore, the buyer has strong incentives to care about prices paid to these suppliers. Second, the increase in vertical integration is sufficiently large to make replacement threats credible: the new mill fulfills roughly 20% of the buyer's fabric needs, compared to only ∼7% from the largest external supplier. Third, demand for end clients (*e.g.,* Old Navy) is relatively stable, decreasing slightly from 250 to 215 million square meters of fabric per year. Therefore, the buyer must make strategic decisions to reduce the number of external suppliers, the average volume per external supplier, or a combination of both.[37]

[36]https://www.kearney.com/industry/consumer-retail/article/creating-a-competitive-advantage-for-india-in-the-global-textile-and-apparel-industry

[37]The integration also does not appear to be in response to a supply-side change in fabric production technology, as a cost function estimated using data from both the old and new integrated suppliers fits well, as shown in section 4.1.2.

Figure 7: Pre-Integration: Fabric Sourcing by Supplier Type



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

Additional analysis suggests the buyer pushes aggregate quantity variability onto spot suppliers, while relying on integrated and agreement supply for stable quantities. Therefore, spot suppliers provide *marginal* quantities, whereas integrated and agreement supply are *inframarginal*. Figure 7 illustrates this pattern: aggregate integrated and agreement volumes are much more stable than spot market supply. Specifically, the coefficient of variation across quarters for agreement and integrated supply is only .152, but the coefficient for spot supply is more than double—.324.[38]

I interpret the use of integrated and agreement supply for *inframarginal* quantities as consistent with two central institutional features. First, increasing marginal costs of internal production generate incentives to keep in-house sourcing volumes stable, rather than

[38]Although the small sample size limits inference, regressing integrated and agreement supply on total volume does not yield a statistically significant coefficient on total volume and produces a small point estimate (0.05) and R^2 (0.03). This result suggests that little of the variation in integrated and agreement supply is explained by aggregate demand, consistent with stable volumes. On the other hand, spot market supply is statistically significantly associated with total volume, with a large point estimate (0.70) and R^2 (0.83), suggesting that aggregate volume shapes spot sourcing. These results are robust to considering the post-integration period instead, noting that pre-integration quarters are analyzed separately to avoid conflating changes in sourcing strategy with changes in capacity by supplier type. Online Appendix OA1.11 shows similar results for the post-integration period, where the coefficent of variation remains almost double for spot (.287) relative to integrated and agreement supply (.154).

increasing with overall demand. Empirically, integrated supply rarely exceeds 100% capacity (see Figure 10), which is consistent with such incentvies. Second, as agreements specify stable quantities, the buyer cannot adjust agreement quantities to offset aggregate demand variability. (Furthermore, because relational contracts require time to establish trust, the buyer cannot easily add new agreement suppliers (Macchiavello and Morjaria, 2015)).

3 Theoretical Bargaining Model

This section introduces a theoretical model to formalize the price effects of both *i)* surplus creation from insurance-like agreement terms and *ii)* surplus capture from threats to replace suppliers with in-house production.

The model features the buyer and supplier in a stationary agreement that represents the steady-state of the relational contract; because the agreement is at steady-state, the buyer and supplier take the amount of quantity assurance provided as given. (Equivalently, I consider the outcome of possible prior negotiations over the amount of quantity assurance as fixed rather than modelling such negotiations explicitly).[39] Therefore, the buyer and supplier bargain over prices to maximize a weighted product of their individual surplus in the agreement relative to their outside options.

I model the buyer as risk-neutral given its large size and sophistication, consistent with agreements shifting quantity risk onto the buyer so suppliers can receive stable quantities. It follows that buyer surplus is the discount it receives under the agreement. Because I model suppliers as risk averse, their surplus from the agreement reflects both the *level* and *variance* of agreement profits versus their outside option. Therefore, insurance-like agreement terms can create supplier surplus by reducing profit variance, even with lower average profits.

The model incorporates the insurance-like agreement features—both U-shaped prices and quantity assurance—as well as the buyer's replacement threats from in-house production. The model includes the U-shaped prices through the buyer and supplier bargaining over a pricing *function* (rather than a vector of unit prices, as common in Nash bargaining) that maps the ("future") realized quantity to a price. The model embeds quantity assurance by having suppliers face different distributions of quantities in the agreement versus the outside option of the spot market. Because the agreement provides *stable* but not necessarily *larger* quantities, both agreement and spot suppliers have the same mean quantities but

[39] Brugues (2025) studies how relational contracts begin and evolve over time to reach steady-state in the context of Ecuadorian manufacturing firms and their domestic buyers.

quantity variance is smaller in the agreement. The buyer's replacement threats appear in its outside option, as I compute model-implied prices both pre-integration—when the buyer's outside option is the spot market—and post-integration—when the buyer can use in-house production as its outside option.

While the static model presented abstracts from relational enforcement in order to highlight creation and capture of surplus, Appendix A.3.1 illustrates that extending the model to incorporate static incentives for parties to deviate from the agreement[40] yields similar qualitative conclusions for reasonable discount factors.

3.1 Model Set-Up

Equation 1 specifies the bargaining problem.

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

Buyer surplus is the difference in profit between buying fabric in the agreement—purchasing quantity q^A at price $p^A(q^A)$—and the outside option—with value O_B. I model the buyer as risk neutral given its large size and sophistication.[41] Buyer surplus omits revenue from sales to the end client because revenue is independent of fabric procurement costs as bargaining with the end client precedes fabric procurement. Therefore, the profit difference collapses to the cost difference between agreement purchases and the outside option.

The supplier's surplus is the utility of profit in the agreement—with CARA utility $U(\cdot)$ and production costs varying with capacity $C(q)$—less the value of the outside option O_S. I assume that suppliers have preferences for smooth profits, which presumably reflects that 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

[40] Specifically, I model the dynamic incentive compatibility constraint from Macchiavello and Morjaria (2023).

[41] 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.

expenses.[42] For brevity, I refer to such preferences as risk aversion.[43]

Therefore, insurance-like agreement terms can generate surplus for suppliers that affects agreement prices through two channels. First, the U-shaped pricing function smooths profits, increasing supplier utility even for the same average level of profit by reducing profit variability. Second, quantity assurance generates additional surplus beyond insurance by reducing average costs, which benefits even risk-neutral suppliers. Quantity assurance reduces average costs in the presence of increasing marginal costs (*i.e.,* convex costs) due to Jensen's inequality. Intuitively, this result reflects the non-linearity in the cost function implied by increasing marginal costs; the high costs associated with large quantities are not correspondingly offset by lower costs for low quantities. Increasing marginal costs are consistent with analysis of integrated supplier costs and capacity in Figure 10. (Marginal costs can increase because machine inputs are fixed in the short run; therefore, scaling up production can generate costly congestion. Additionally, labor costs could increase dramatically with larger quantities—the Factories Act of 1948 mandates double wages for overtime work.[44])

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 surplus 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 $1 - \alpha_B$.

The outside option for both the buyer prior to vertical integration and the supplier is to transact in the spot market. Because agreements are relational contracts, which require time and effort to establish, neither the buyer nor the supplier can switch to a new agreement trading partner, at least in the short run. Furthermore, because agreements feature quantity assurance, which reduces variability without necessarily increasing volumes, (average) agreement quantities cannot increase.

I model the spot market as perfectly elastic with constant price p^M which exceeds production costs due to hold-up. This assumption is consistent with the flat spot prices in

[42] Alternatively, firms may inherit owners' preferences over profits because small supplier profits are likely the primary income source for owners. I do not take a stance as to the specific source of these preferences.

[43] Blouin and Macchiavello (2019) and Karlan et al. (2014) also show that risk aversion shapes LMIC firm behavior.

[44] 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.

Figure 4.[45] Flat spot prices presumably reflect that the hold-up price is pinned down by the buyer's outside option, as suppliers extract surplus, and the buyer's outside option does not vary systematically with supplier capacity. (In other words, spot market prices reflect supplier, rather than buyer, market power). Furthermore, even though the buyer is a large firm, the buyer (and suppliers) are small relative to the Indian fabric market; the buyer's share of the domestic fabric market is less than 0.5%.[46]

After integration, the buyer's outside option could change to use the same production technology as the external supplier. This assumption reflects conversations with the buyer emphasizing that the buyer and agreement suppliers use similar technology.[47] The 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. Allowing for some cost heterogeneity between the buyer and supplier does not change results meaningfully (see Appendix OA1.5.2).[48]

Quantities in the model reflect the stochastic nature of demand in this setting, as end clients determine quantities rather than the buyer and supplier choosing it (although the agreement still specifies the *ex-ante distribution* of quantities). As the 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

[45] Incorporating buyer market power when purchasing inputs in the spot market leads to underestimates of surplus creation but has no effect on replacement threats. If the buyer has market power in the spot market, then the agreement creates additional buyer surplus by reducing spot prices because the buyer would face an upward-sloping supply curve. Buyer power does not affect analysis of replacement threats (given prices when the buyer's outside option is the spot market) because the spot market is not present in either the buyer's inside or outside option when in-house supply serves as the buyer's outside option.

[46] 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).

[47] Broadly, these fabric mills 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.

[48] 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). Furthermore, large cost heterogeneity 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.

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 diversified income streams might be less risk averse over production for any individual buyer.)[49] Formally, the outside options are:

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

Overall, I explicitly model two-sided market power within agreements through bargaining over prices. Market power in the agreement is determined by bargaining positions, which depend upon both surplus shares (*i.e.,* the bargaining parameters) and outside options. However, I interpret these sources of market power differently: the buyer bargaining parameter represents the buyer's share of surplus generated by insurance-like terms, whereas replacement threats are only rent extraction and do not depend upon surplus generated in the agreement. Therefore, discounts resulting from a positive buyer bargaining parameter align incentives for the buyer to provide insurance, whereas discounts from displacement threats only transfer surplus. The model also incorporates supplier market power in the spot market through their hold-up of the buyer.

Appendix A.4 provides details on the exact parameterization of the model and computation of bargained prices. Parameters and functional forms are chosen to formalize intution regarding key economics of surplus creation from insurance-like agreement terms and surplus capture from replacement threats using in-house production, rather than calibrated to match data or parameter estimates from other studies. Section 4 builds and estimate a structural model closely tied to this theoretical model for quantification.

[49] 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 cost heterogeneity does not affect surplus creation or capture, as shown in Online Appendix OA1.5.2.

3.2 Insurance-like Agreement Terms, Replacement Threats, & Supplier Risk Aversion

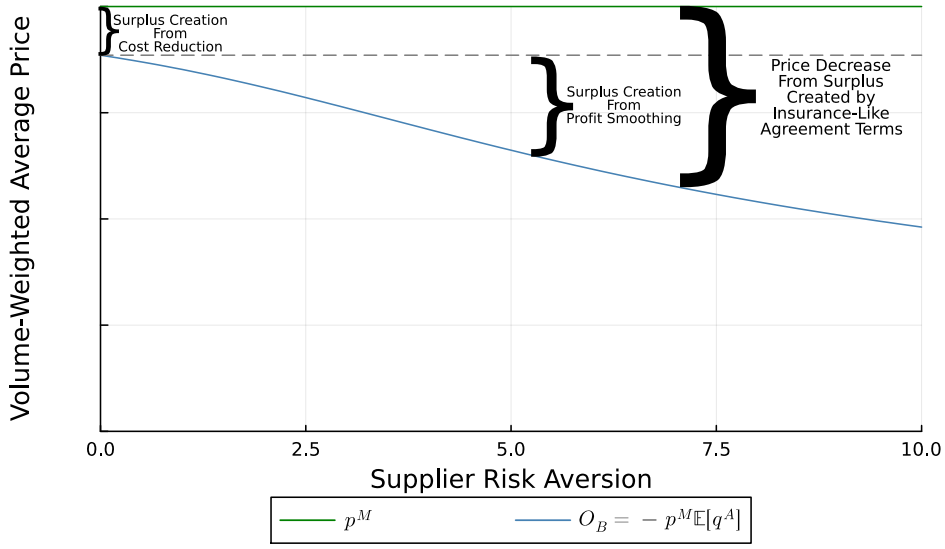


Figure 8: Surplus Creation From Insurance-Like Agreement Terms

Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.4.

Figures 8 illustrates how the expected cost to the buyer (*i.e.,* the volume-weighted average fabric price) varies with supplier risk aversion, holding the buyer bargaining parameter constant.[50] 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, surplus creation in the agreement results in price discounts for the buyer relative to spot prices as long as the buyer bargaining parameter is positive. Then, some of the surplus created in the agreement is shared with the buyer as a price discount. The discount in agreements with risk-neutral suppliers (*i.e.,* when risk aversion is 0) represents the buyer's share of surplus generated by cost reduction from quantity assurance given cost convexity. As supplier risk aversion increases, the agreement creates additional surplus

[50]For completeness, Online Appendix OA1.7 shows results from a similar analysis that holds supplier risk aversion constant but varys the bargaining parameter. Note that only supplier risk aversion, and not the buyer bargaining parameter, directly affects surplus creation from insurance-like agreement terms and capture from strategic replacement threats.

by smoothing profits through both insurance-like agreement terms: *i)* the quantity assurance and *ii)* the U-shaped pricing function. Consequently, the agreement discount grows.

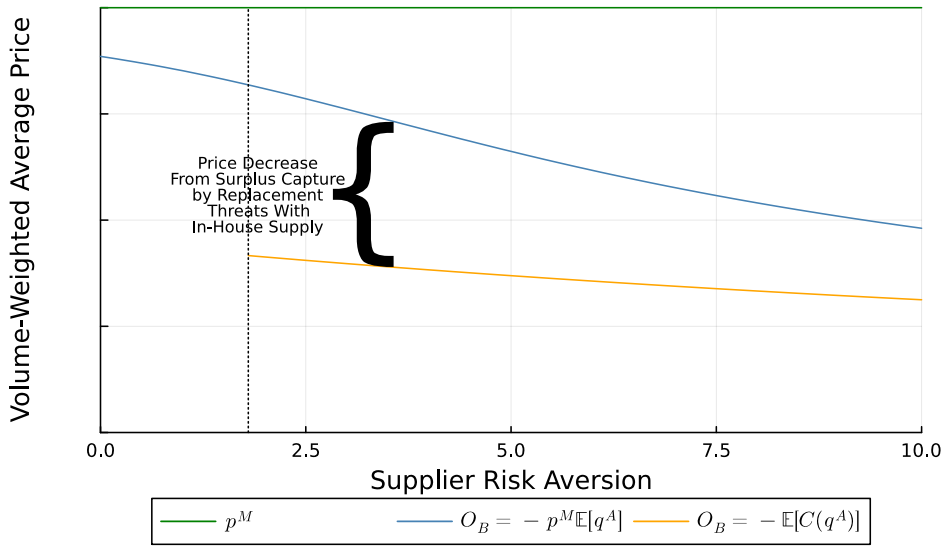


Figure 9: Surplus Capture From Replacement Threats

Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.4.

Figure 9 shows the price effects of the buyer's outside option switching to its integrated supplier. The additional discount represents surplus capture from threatening to replace suppliers with in-house production. 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 reflects that suppliers with low risk aversion receive little surplus from insurance-like agreement terms. Therefore, with such suppliers, no pricing function generates surplus for both parties.

What happens when supplier risk aversion falls below the cutoff, meaning the agreement would break 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).

The first extension begins with a non-cooperative signaling game about residual capacity at the in-house supplier before cooperative 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. Information asymmetry about residual capacity—the "state of the world" in this game—is central to the model. Residual capacity as the buyer's private information reflects that the buyer alone has knowledge about both orders from other end clients and other commitments in-house. (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 Nash bargaining has similarities with passive beliefs. Therefore, the non-cooperative phase incorporates beliefs in a principled way.)

I establish the existence of a weak Perfect Bayesian Equilibrium where the buyer claims that its in-house supplier is not capacity constrained to capture surplus, but only when bargaining with suppliers that value the non-price agreement terms above a cutoff (shown in Appendix A.5). This cutoff value for the valuation of non-price agreement terms corresponds to the cutoff level of supplier risk aversion for the buyer to capture surplus from replacement threats, as risk aversion determines surplus from insurance-like terms that need not affect average prices. The second extension, in Appendix A.6, considers a model with a simpler non-cooperative game that relies on buyer commitment to a level of residual capacity before bargaining, and similarly shows the existence of an equilibrium where the buyer only makes replacement threats with suppliers above a cutoff level for surplus from non-price terms.

Appendix A.6.4 discusses extensions from bilateral to multiateral bargaining. Results do not change under reasonable assumptions about information and timing.

4 Structural Model

I estimate a structural model closely tied to the theoretical model to quantify the price effects of surplus creation from insurance-like agreement terms and surplus capture from replacement threats with in-house production. The structural model also enables analyzing prices for any level of supplier risk aversion; computing prices for out-of-sample levels

of supplier risk aversion is valuable because only a limited, selected sample of agreement suppliers are likely observed. However, even very small firms in LMIC export sectors rely on bargained agreements, such as smallholder Rwandan coffee farmers (Macchiavello and Morjaria, 2020). But, given the large size of the buyer I study, such small suppliers may not be among the buyer's agreement suppliers. Last, I use the structural model to evaluate counterfactual policies to increase prices paid to suppliers and how policy effectiveness relates to risk aversion.

First, I discuss estimation and identification of the model. Then, I present risk aversion and bargaining parameter estimates and decompose agreement discounts into surplus creation and capture. Next, I examine prices for heterogeneous buyer and supplier (*i.e.*, relationship) types to recover both *i)* the cutoff level of supplier risk aversion such that the buyer captures surplus through strategic replacement threats with in-house production and *ii)* the price response to such threats. Subsequently, I validate the model by comparing the model-implied effect of surplus capture from replacement threats with an estimate using a difference-in-differences (DiD) approach. Notably, the DiD leverages variation not used for identifying structural parameters.

Last, I highlight how the bargaining environment restrains the buyer's ability to capture surplus through strategic replacement threats. I estimate a Nash-in-Nash with Threat of Replacement (NNTR) model in the spirit of Ho and Lee (2019). I demonstrate that prices from the NNTR model are inconsistent with DiD event study estimates, reflecting that the NNTR model considers an environment *i)* without spot and in-house co-existing supply alongside agreements and *ii)* where new agreements can form easily.

4.1 Estimation Approach

4.1.1 Buyer and Supplier Primitives

Estimating the structural model recovers three parameter groups: 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, as firms should not deviate too much from risk neutrality, and econometric concerns of overfitting given the limited sample size per

supplier.[51]

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 in (1). In the expression $\hat{p}^A(q^A; \mathbf{\Omega}_s)$, $\mathbf{\Omega_s}$ 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.[52] Second, I estimate prices for observed agreement quantities using the model-implied pricing function. Last, I compute the objective function for each supplier, which is the root mean squared error of price estimates from the pricing function using only pre-integration quarters.

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.

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-

[51] As I study prices in stable agreements, miscategorization of spot as agreement suppliers would underestimate risk aversion for miscategorized suppliers. I would then incorrectly conclude that surplus creation is not important. This concern motivates my conservative classification of agreement suppliers; as parameter estimates are supplier-specific, excluding some agreement suppliers has minimal implications for parameter estimates. Regardless, prior analysis suggests that the approach to supplier classification performs well.

[52] 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 4.2. 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 ($\beta_{1,s}$).

supplier objective functions:

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

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,[53] T is the number of periods, $\hat{p}^A_s(q^A_{s,t};\boldsymbol{\Omega}_s)$ 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 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:

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

4.1.2 Outside Estimation

Table 1 describes the additional moments necessary for model estimation.

[53] While using monthly rather than quarterly data would increase the effective sample size, agreements seem to operate at the quarterly rather than the monthly level. Specifically, reduced-form patterns are most consistent with the agreement at the quarterly level and, anecdotally, the buyer's fabric procurement staff emphasize that meetings with key suppliers occur quarterly.

Table 1: Additional Moments for Estimating the Structural Model

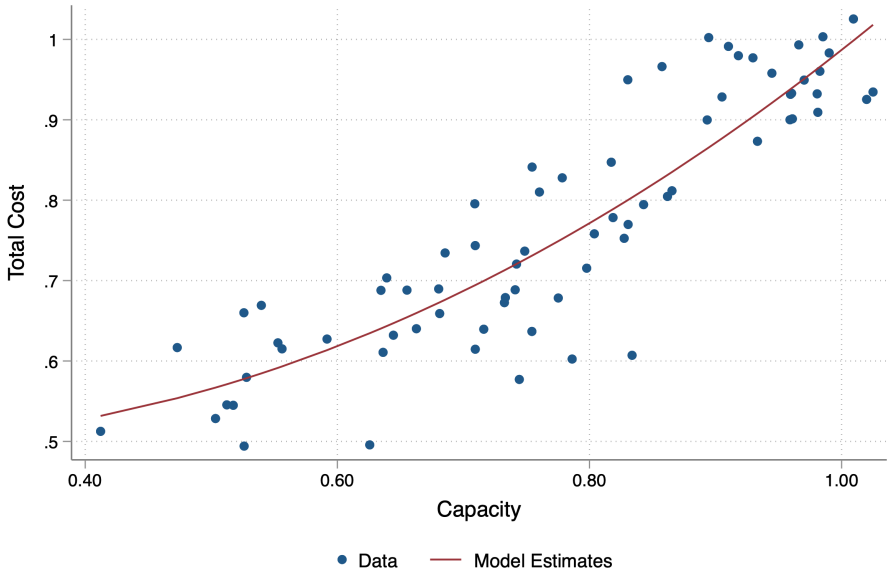
Parameter(s) in Model	Moment(s) to Match	Data
Cost Function Convexity	Actual Costs and Capacity	Cost and Capacity Data from Integrated Supplier
Mean Capacity	Average Log Quantity	Transaction Data
$\frac{\mathbb{V}[q^M]}{\mathbb{V}[q^{RC}]}$	Variance in Quantities for Spot Suppliers Compared to Variance 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 division maps broadly to fabric type (*e.g.,* knits and striped fabrics).[54] Identification of cost function parameters relies on exogenous capacity shocks. Capacity shock exogeneity derives from 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 end client demand rather than endogenous buyer decisions.[55] As quantities are (at least locally) exogenous, OLS regressions provide causal estimates of the effect of capacity on costs.

[54]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.

[55]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.

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 division. Capacity is reported directly in the data.

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.

Appendix A.7 provides details for remaining outside estimation as listed in Table 1.

4.2 Separate Identification of Supplier Risk Aversion from Bargaining Parameters

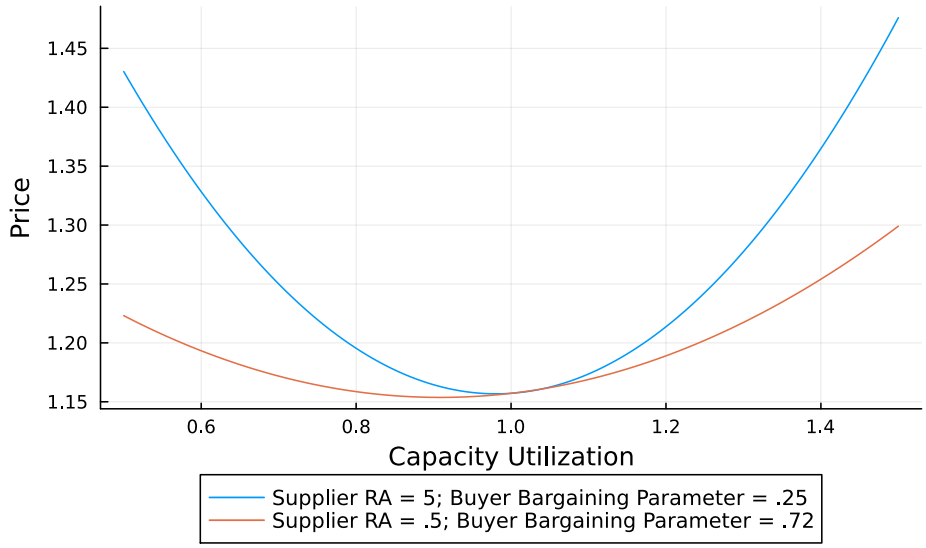
Separately identifying supplier risk aversion from the buyer bargaining parameter can be challenging because increasing either parameter alone reduces average agreement prices. When supplier risk aversion increases, holding the buyer bargaining parameter fixed, surplus created by the agreement increases. As some of this surplus created is shared with the buyer (for any positive buyer bargaining parameter), higher risk aversion results in lower prices.

Similarly, when the buyer bargaining parameter is high, the buyer receives a larger share of the surplus created by the agreement through lower prices.

I develop a novel identification strategy that leverages two key sources of variation in the data. First, 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 U-shaped contracts are associated with more risk-averse suppliers because price adjustments reduce profit variance across capacity fluctuations. Second, there is variation in the level of average agreement prices relative to spot prices. This variation in the level of average prices is essential for identifying the buyer bargaining parameter, as increasing this parameter reduces prices across capacity realizations.

When only one price-capacity tuple is observed, it is impossible to separate supplier risk aversion from the buyer bargaining parameter. For any given buyer bargaining parameter, adjusting supplier risk aversion can rationalize any observed price-capacity tuple (and vice-versa). Figure 11 provides an illustrative example: 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. As shown in Figure 11, lower risk aversion results in a less U-shaped contract and a higher bargaining parameter

results in lower average prices. (A third price-capacity tuple identifies asymmetry in the U-shaped function that enables prices to smooth profit more effectively.)

To 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 pricing function while shifting the buyer bargaining parameter mostly shifts the average price level.

Figure 12: Shift Supplier Risk Aversion

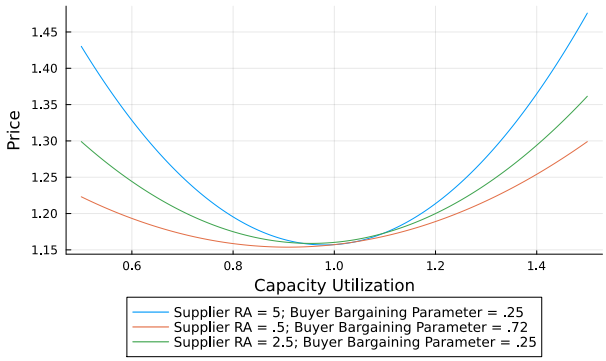
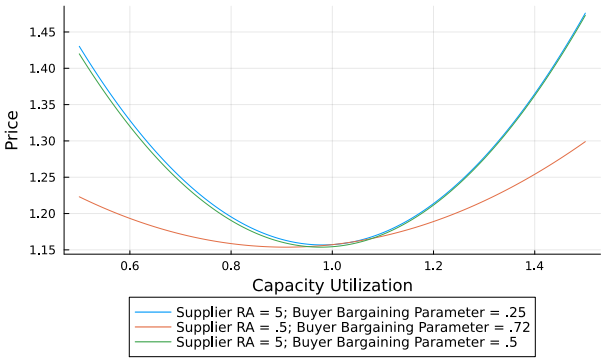


Figure 13: Shift Buyer Bargaining Parameter



To bring this identification argument to the data, I leverage supplier-specific capacity shocks from stochastic end client demand. Although quantity assurance implies the agreement has a target capacity, the garment production process—namely, that neither which samples will be ordered nor what quantities would be are known at the time of sampling—means the agreement reduces, but does not eliminate, quantity variance. Appendix A.8.1 provides evidence supporting the interpretation of deviations from target capacities as plausibly exogenous demand shocks. Specifically, I show results inconsistent with the most plausible threat to exogeneity: the agreement includes additional *dynamic* quantity assurance across quarters (*i.e.,* quarters following an extreme capacity are closer to the target).

4.3 Risk Aversion and Bargaining Parameter Estimates

Parameter estimates in Table 2 (confidence intervals in Appendix Table A.10) show that risk aversion estimates for agreement suppliers resemble coffee mills studied in Blouin and Macchiavello (2019) (.0085 for the median agreement supplier and .0068 for coffee mills).[56]

[56]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.4 shows (converted) CARA risk aversion parameters.

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 suppliers I study, the similarity of risk aversion estimates seems reasonable. Note that although these results suggest that agreement suppliers are large firms (especially by LMIC standards), the buyer studied is much larger with fabric purchases, which are likely smaller than revenue, reaching ∼$180 million per year.[57]

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.

Estimated buyer bargaining parameters imply that the buyer typically receives slightly more than half of surplus created by the agreement.[58] This result aligns with the risk aversion estimates, as both suggest that agreement suppliers are large firms (with the exception of one outlier supplier with high risk aversion and a high buyer bargaining parameter). The large size of agreement suppliers is consistent with the notion that creating and maintaining agreements, especially complex ones with highly non-linear prices, requires sophisticated firm capabilities (Macchiavello and Morjaria, 2023).

As the estimation routine includes outside estimates, I prioritize model validation in Online Appendix OA1.8. 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 primarily uses pre-integration quar-

[57]Note that the estimated tuning parameter (λ) is small at .002, ensuring that low estimated supplier risk aversion parameters are not primarily due to a large tuning parameter.

[58]I explore one possible microfoundation for heterogeneity in buyer bargaining parameters in A.13.

ters.[59] (Model validation using a DiD approach follows the decomposition of discounts into surplus creation and capture. As the DiD provides a direct estimate of surplus capture from replacement threats with in-house production, I discuss this validation after showing model estimates of surplus capture.)

4.4 Decomposing Prices into Surplus Creation and Capture

Table 3 decomposes agreement discounts into surplus creation from insurance-like agreement terms and capture from replacement threats with in-house production. 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 surplus 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 surplus creation. Then, the discount when the buyer's outside option is the spot market reflects the buyer's share of agreement surplus created, while surplus capture generates an additional discount when the buyer's outside option is in-house production.

Estimates illustrate that agreement discounts are entirely attributable to surplus creation, with no surplus 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 surplus. 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 agreement discounts 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. Importantly, as discounts reflect the buyer's share of agreement surplus, even discounts from surplus creation mean that the buyer extracts some share of the surplus based on its bargain-

[59] 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.

ing parameter. Estimated bargaining parameters indicate that the buyer receives a slight majority of agreement surplus.

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.[60] The primary finding—that surplus creation alone accounts for all observed discounts—remains unchanged. As all suppliers fall below the risk-aversion cutoff for the buyer to capture surplus through replacement threats, both the median and mean risk aversion also fall below the cutoff. However, both the level of the discount and its 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.[61] 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 insurance to the discount.[62] Discounts are slightly higher, on average, with the median rather than mean parameter values, reflecting that the larger median than mean buyer bargaining parameter more than offsets the lower median than mean risk aversion.

Table 3: Discounts from Agreement Suppliers

Supplier Type	Total Discount	Discount Decomposition: *Surplus Creation From Cost Reduction*	*Surplus Creation From Insurance*	*Surplus Capture From Replacement Threats*
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%

[60] 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).

[61] 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.

[62] 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 OA1.16).

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 surplus creation and capture. The discount from surplus creation computes prices when the buyer's outside option is the spot market, while the discount from surplus 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 surplus creation into cost reduction finds the discount for a risk-neutral version of the supplier, holding all else constant. Any remaining discount from surplus creation reflects profit smoothing from insurance-like agreement terms.

4.5 Model Validation using a Difference-in-Differences Design

I validate the structural model using a DiD design that directly estimates the effects of surplus capture from threats to replace suppliers with in-house production.[63] 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 in the structural model. Therefore, I compare the model-implied effects of replacement threats with DiD estimates.

4.5.1 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. Morton (2023) shows that *fabrics* (not suppliers) 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 serve as substitutes. However, suppliers are not

[63] Additional validation exercises are in Online Appendix $OA1.8$.

selected for agreements based on individual fabrics (as discussed in Morton (2023)).[64] (Instead, the data suggest that fabrics were selected to be brought in-house based on trends in fabric demand over time: fabrics with increasing aggregate demand prior to the mill's construction are more likely to be produced at the new mill.) However, suppliers are unlikely to be intentionally selected to be exposed to integration (even if fabrics are).

Even so, the buyer's endogenous choice of which fabric to produce in-house creates non-random assignment in exposure to vertical integration at the supplier level based on each supplier's fabric mix that correlates with volumes. Therefore, I adopt a DiD design, detailed in Appendix A.11.1, that accommodates non-random exposure to treatment through a propensity score adjustment (Abadie, 2005). This approach weakens the parallel assumptions trend to conditional parallel trends; in my empirical setting, comparisons are conditional on supplier exposure to trends in aggregate fabric demand that reflect demand shifters from evolving fashion tastes for specific fabrics. Falsification tests show that the propensity score adjustment based on these aggregate fabric demand trends achieves balance pre-treatment.[65] Furthermore, because outcome variables are *not* explicitly matched pre-integration, pre-trends still provide diagnostic information about the validity of the parallel trends assumption.[66]

4.5.2 Validation: Effects of Replacement Threats

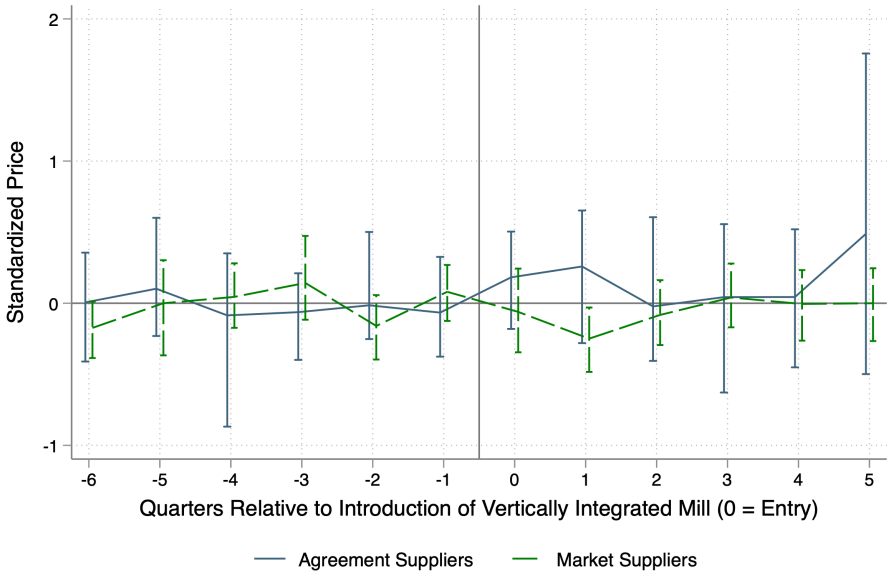
The model generates sharp predictions for the effects of vertical integration on both prices and quantities, given supplier risk aversion estimates. As even the most risk-averse agreement supplier has risk aversion below the cutoff for the buyer to make replacement threats (see Appendix Figure A.6), the model does not predict any price change for agreement suppliers. Figure 14 confirms that agreement supplier prices do not respond to integration, and that there are no pre-trends of concern.

[64] 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 large (42).

[65] Propensity scores are not calculated using supplier-level fabric demand. Therefore, these falsification tests do not simply guarantee balance by construction.

[66] 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.

Figure 14: DiD Estimates: Price Effects



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.

Although not explicitly in the theoretical model, the buyer sourcing *marginal* quantities from spot supply suggests that the quantity decline from external suppliers should come from spot rather than agreement supply;[67] marginal sourcing from spot supply reflects that spot suppliers charge the highest (average) prices. Appendix Figure A.9 confirms that agreement quantities do not change, but spot quantities do. Furthermore, the significant results for quantities suggest that the null price effects do not primarily reflect limited statistical power of the DiD design.

As the small number of agreement suppliers limits statistical power, I also examine average treatment effects pooled across all pre- and post-integration periods in Appendix Table A.7.[68] The pooled results confirm that agreement suppliers do not experience any statistically significant effect on quantities or prices. (Point estimates are economically close to zero, although not always precisely estimated.) However, vertical integration causes spot suppliers to experience a statistically significant and economically meaningful decrease in quantities.

[67]Recall that the new in-house supplier was not associated with increased end client demand, meaning that exernal sourcing mechnically should decrease.

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

The decline averages .4 standard deviations, amounting to approximately 120,000 square meters of fabric per quarter.[69]

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.[70] In Online Appendix OA1.10.2, 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.

Online Appendix OA1.10.3 provides an additional validation exercise by comparing DiD event study coefficents using actual prices with DiD estimates using model-implied prices for actual quantities. This validation exercise leverages that the DiD and structural model rely on different sources of identification. Event study estimates using model-implied prices in Figure OA1.21 are all within the 95% confidence intervals of event study coefficients using actual prices.

4.6 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 marginal sourcing from spot supply, 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 external validity of the effects of replacement threats (rather than a specific critique of NNTR).

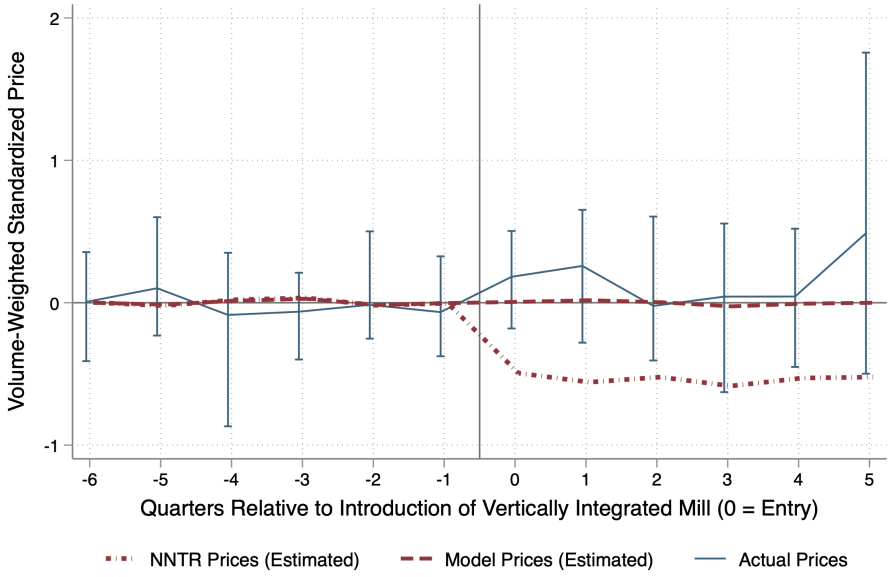
Figure 15 illustrates that price decreases predicted by the NNTR model are inconsistent

[69] 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).

[70] Full event study plots for these outcome variables are available in Online Appendix OA1.10.1.

with the DiD estimates of price changes from replacement threats.[71]

Figure 15: 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.12.

4.7 Model Estimates: Agreement Prices by Supplier Risk Aversion and Buyer Bargaining Parameter

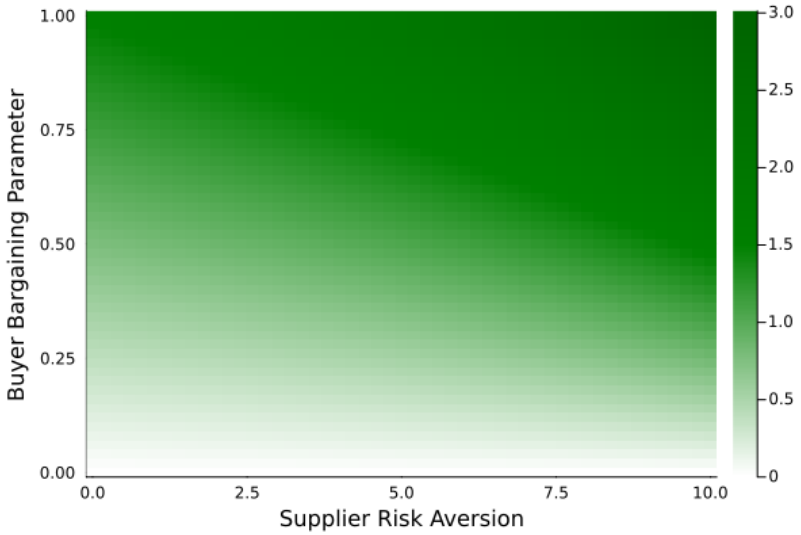
Taking the model seriously, as validation exercises suggest the model reasonably represents agreements, I compute agreement prices for a wide array of supplier risk aversion and buyer bargaining parameters. This exercise facilitates studying surplus creation and capture for policy-relevant small suppliers more risk averse than observed agreement suppliers.

Figure 16 reports agreement discounts, as a percent of spot prices, for different supplier risk aversion and buyer bargaining parameters. Discounts increase in both supplier risk

[71]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.[72] 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.12.

aversion and the buyer bargaining parameter. However, the additional discount from higher risk aversion is modest. For example, at the median estimated buyer bargaining parameter, the discount for a risk-neutral supplier is less than one percentage point smaller than the discount from a supplier at the median estimated risk aversion. Similarly, increasing risk aversion from the estimated median to the 75th percentile increases the discount by less than one percentage point. For all parameters included in this heatmap, discounts reflect surplus creation, rather than surplus capture. This result occurs because the buyer always uses the spot market as its outside option with suppliers with risk aversion parameters shown. Furthermore, as discounts do not increase much in risk aversion, cost reductions—rather than profit smoothing—-drive most discounts.

Figure 16: Agreement Discounts (% of market price)

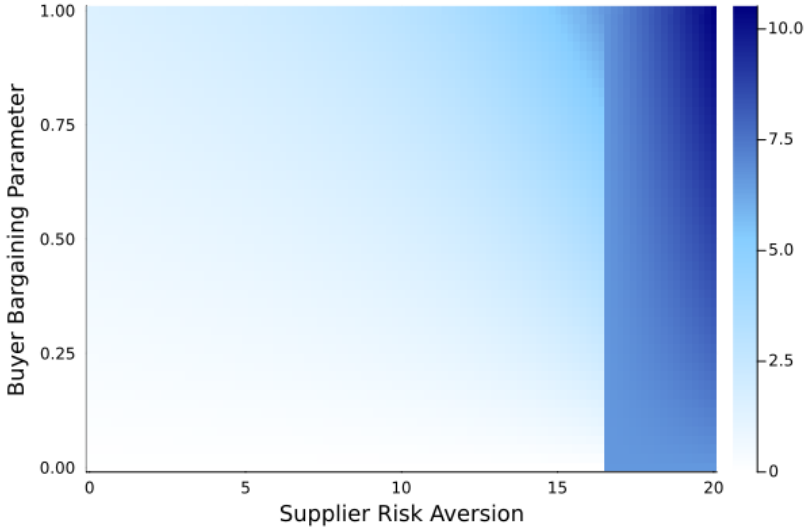


Note: Model estimates for the discount in the agreement relative to the spot market as a percent of the spot price. The heatmap legend is to the right of the heatmap and is in percentage units. Model estimates use the average capacity and quantity assurance across in-sample agreement suppliers.

Figure 17 replicates the analysis above, but for a larger set of risk supplier aversion parameters to show discounts for suppliers sufficiently risk averse for the buyer to capture surplus through replacement threats with in-house production. 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 depends upon 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 can be larger when it mitigates risk in both demand and supply; for example, risk on both sides of the market occurs for agricultural workers facing productivity risk (Jayachandran, 2006).[73] 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.

Figure 17: Agreement Discounts (% 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 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 surplus 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 discounts for agreement suppliers average only 1.5% of spot prices.

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

At 6.7%, the discount exceeds the 5% Indian VAT for apparel during the study period.[74] 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.

Online Appendix OA1.11 incorporates relational enforcement for a conservative calibration of the discount factor, documenting negligible differences in estimated prices.[75] 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 estimated risk aversion: both suggest agreement suppliers are large firms.[76]

5 Counterfactual Policy Evaluation

Both surplus creation from non-price agreement terms and surplus capture by strategic threats to replace suppliers with in-house production can result in low average agreement prices paid to small, risk-averse suppliers. Low prices raise distributional concerns as the most risk-averse firms are also likely the smallest. Distributional implications motivate policy remedies, given omnipresent policy initiatives targeting small and medium enterprises (SMEs), such as the World Bank's policy advising on finance tools for SMEs[77] and the IMF's 2019 conference on SME financing.[78] 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.

I evaluate two policies to increase prices paid 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

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

[75] Relational enforcement constrain how U-shaped prices are, influencing the level of surplus agreements can create and the discounts they can support.

[76] 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.

[77] World Bank SME Finance

[78] IMF SME Financing

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

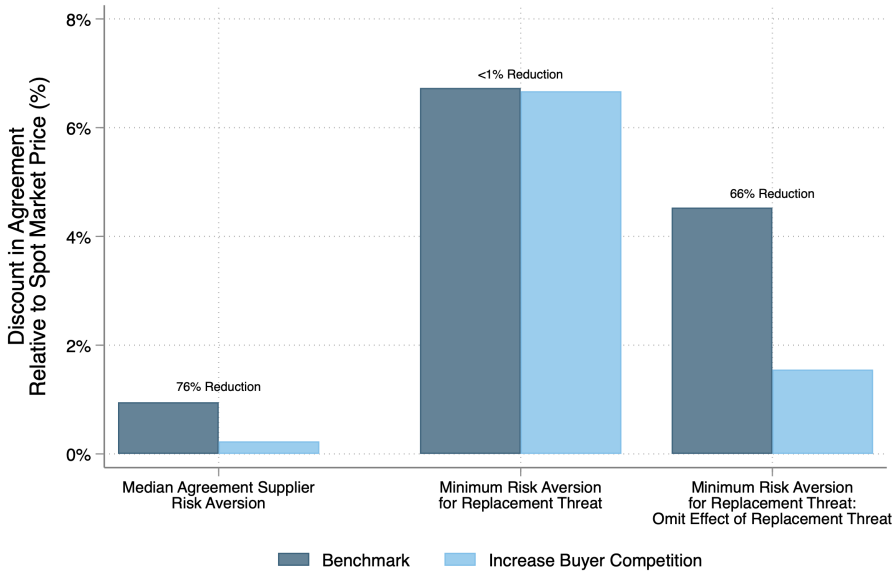
I first study the effect of increasing downstream buyer competition, estimating agreement prices when additional buyers compete to form agreements with suppliers. Policymakers can support the development of large exporting firms through many initiatives, including reducing red tape and other administrative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

I model increased downstream buyer competition as reducing the buyer bargaining parameter.[79] 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, 2025).) Therefore, the buyer responds to competition by offering suppliers a larger share of surplus to ensure their loyalty. This interpretation mirrors efficiency wages, as both incentivize loyalty through a higher surplus share (Katz, 1986). Empirical patterns support modeling buyer competition as a buyer bargaining parameter reduction; specifically, more valuable suppliers receive larger surplus shares (*i.e.,* have lower estimated buyer bargaining parameters).[80]

[79] I halve the buyer bargaining parameter.

[80] 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). Appendix Section A.13 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 18: 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 18 illustrates that increased downstream buyer competition meaningfully reduces the discount from a supplier as risk-averse as the median agreement supplier: discounts decrease 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 surplus capture from replacement threats *overestimates* the benefits of increased competition. Without surplus 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 uses in-house production as a replacement threat, agreement surplus is small relative to the buyer's participation constraint. When the apparent surplus is small, changing the sharing rule has minimal effects on prices. This intuition clarifies why abstracting from surplus capture generates incorrect policy conclusions for small risk-averse

firms, but also why this policy is effective at increasing prices paid to suppliers with low risk aversion when the buyer's threat point is the spot market. Then, there is enough agreement surplus that shifting sharing meaningfully affects prices.[81]

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.[82]

Although governments are unlikely to offer such insurance given moral hazard concerns, policymakers have many tools to provide limited insurance.[83] 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 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.

[81] Online Appendix OA1.12 evaluates 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.

[82] 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. While Blouin and Macchiavello (2019) focus on the efficiency implications of the missing insurance markets, I focus on effects on surplus creation and distribution.

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

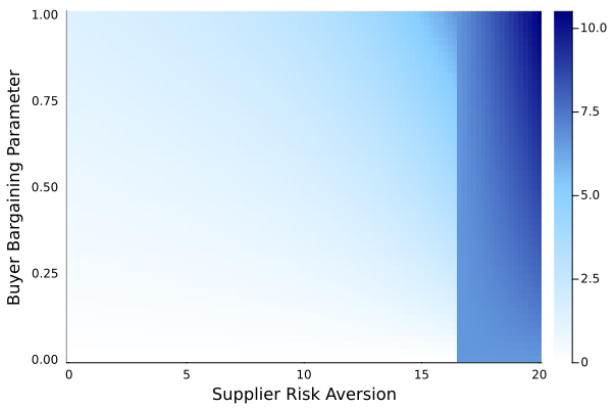


Figure 19: Without Insurance

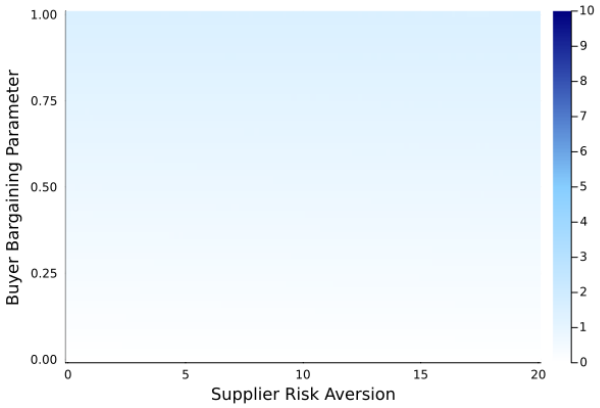


Figure 20: With Insurance

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 surplus through replacement threats, meaning that 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 examines how surplus is created and distributed in the Indian garment manufacturing supply chain, addressing concerns that low prices paid to small suppliers mean gains from exporting concentrate among the few buyers. A theoretical model shows that prices in buyer-supplier agreements reflect both surplus creation—via insurance-like terms that mitigate spot market frictions—and surplus capture by buyers threatening to replace external suppliers with in-house production. Marginal sourcing from the spot market restrains buyers' ability to use replacement threats to improve its bargaining position; the buyer only makes such threats in agreements with risk-averse suppliers who derive large benefits from insurance-like terms. Therefore, the buyer's ability to extract surplus is constrained by supplier risk aversion and contracting frictions.

Using detailed transaction data on the univerise of fabric purcahses by a major garment manufacturer, I estimate a structural model that decomposes observed discounts into

components of surplus creation and capture. Identification of risk aversion relies on how prices respond to exogenous demand shocks. The results indicate that all discounts result from surplus creation, with the buyer capturing slightly over half of the surplus. Risk aversion among suppliers is too low for buyers to credibly issue replacement threats, preventing capture by replacement threats. Counterfactual policy simulations illustrate that increased buyer competition has negligible effects for small firms, but access to the missing insurance market reduces discounts by removing reliance on insurance-like agreement terms for profit smoothing.

More broadly, this paper provides additional evidence that challenges to smoothing, especially relevant in LMIC settings, not only decrease welfare by generating different marginal utility across states, but also shape prices. 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 lead agricultural wage workers 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 because 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 least able to bear risk.

References

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

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.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts, "Does Management Matter? Evidence from India *," *The Quarterly Journal of Economics*, 11 2012, *128* (1), 1–51.

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*, 2025.

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.

Chen, Yanyou and Daniel Yi Xu, "A Structural Empirical Model of R&D, Firm Heterogeneity, and Industry Evolution," *The Journal of Industrial Economics*, 2023, *71* (2), 323–353.

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

Darmouni, Olivier, Simon Essig Aberg, and Juha Tolvanen, "Pulp Friction: The Value of Quantity Contracts in Decentralized Markets," *Working Paper*, 2023.

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.

Gaineddenova, Renata, "Pricing and Efficiency in a Decentralized Ride-Hailing Platform," *Working Paper*, 2022.

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.

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.

Hsiao, Allan, "Coordination and Commitment in International Climate Action: Evidence from Palm Oil," *Working Paper*, 2024.

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

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

Jimenez-Hernandez, Diego and Enrique Siera, "Should the Government Sell You Goods? Evidence from the Milk Market in Mexico," *Working Paper*, 2022.

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.

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 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.

Rajagopalan, Shruti and Kadambari Shah, "Why Indian Firms Don't Scale: Labor Edition," *India Political Economy Program: Research Paper*, 2024.

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.

Simone, Rebeeca De and Felipe Brugues, "Taxation when markets are not competitive: Evidence from a loan tax," *Working Paper*, 2024.

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

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

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

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

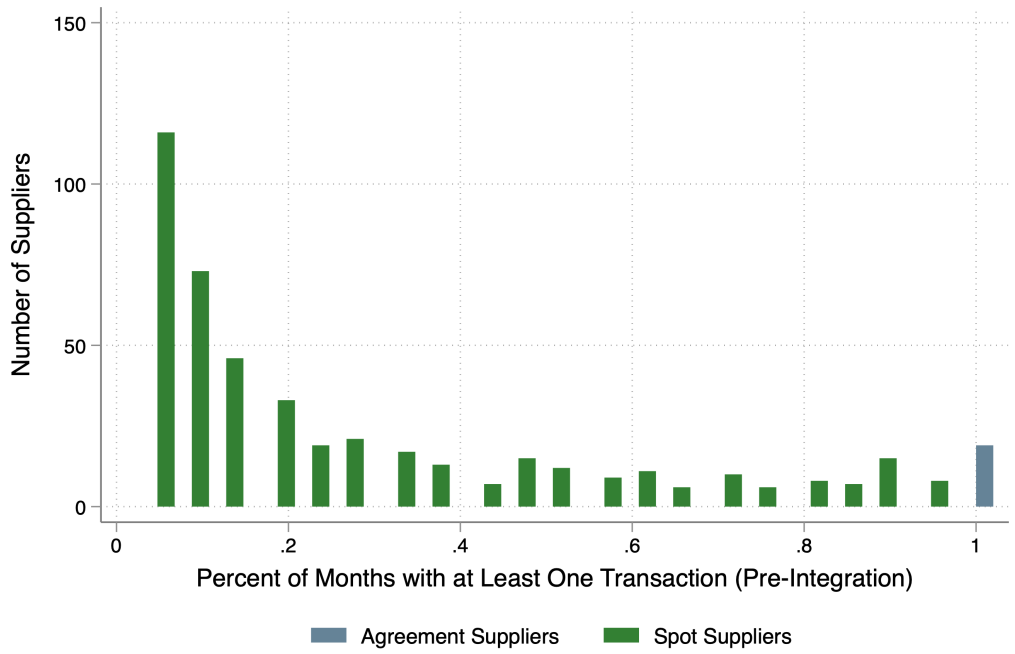
Zavala, Lucas, "Unfair Trade? Monopsony Power in Agricultural Value Chains," *Working Paper*, 2022.

A Appendix

A.1 Fabric Procurement Agreements

A.1.1 Supplier Classification: Agreement versus Spot

Figure A.1: 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.

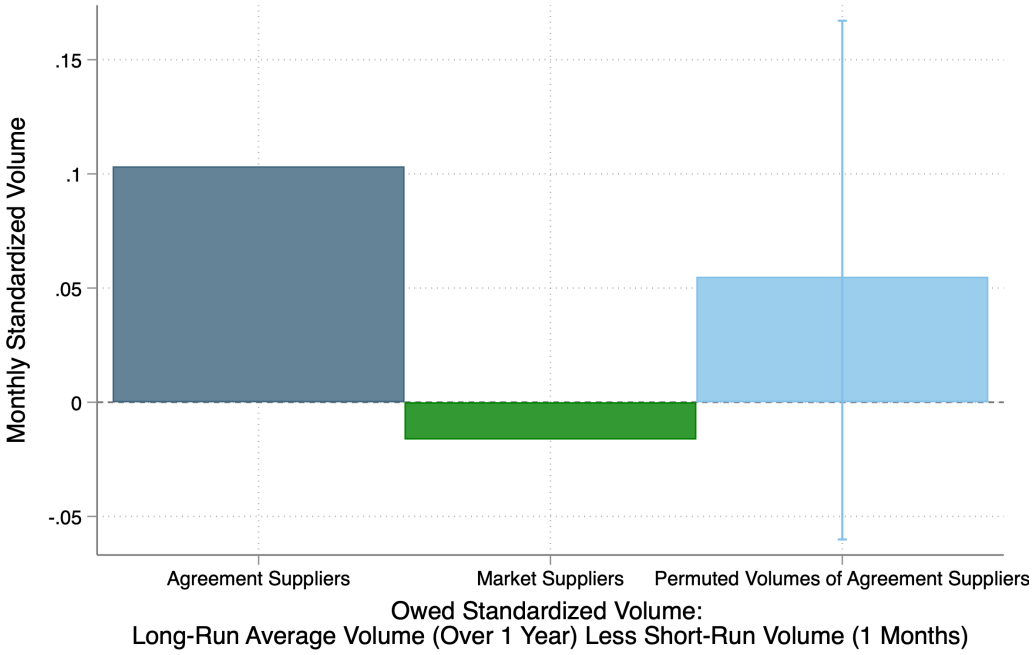
A.1.2 Quantity Assurance: Owed Quantity Guarantees

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.2 illustrates that a one standard deviation increase in owed volume for an agreement supplier is associated with a $\sim .1$ standard deviation increase in current volume.[84] Importantly, the same result does not hold for market suppliers, where an increase in owed quantity is associated with a small *decrease* in volume.

Figure A.2: 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. Specifically, I randomly permutes volume across months for the same supplier,

[84]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.

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, and find that the coefficient on owed volume 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.

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 could potentially pass on some cost savings to the buyer. However, if this contracting narrative were correct, the point estimate for the shorter definition of long-run (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 Online Appendix Figure OA1.2).

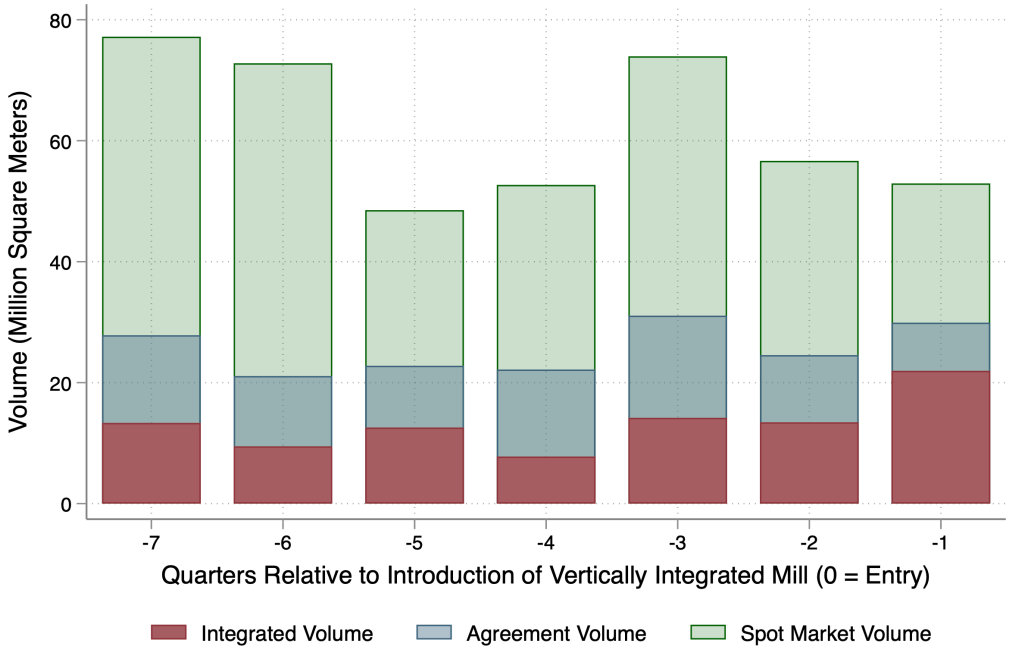
A.1.3 U-Shaped Pricing Function

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					
Normalized Volume	-26.438^{**}	$-.229$	-36.723^{**}	$-.755$	-32.61^{**}	$-.79$
	(10.986)	(.546)	(14.455)	(.717)	(14.646)	(.721)
Normalized Volume^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 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 Normalized Volume	-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 Normalized Volume	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

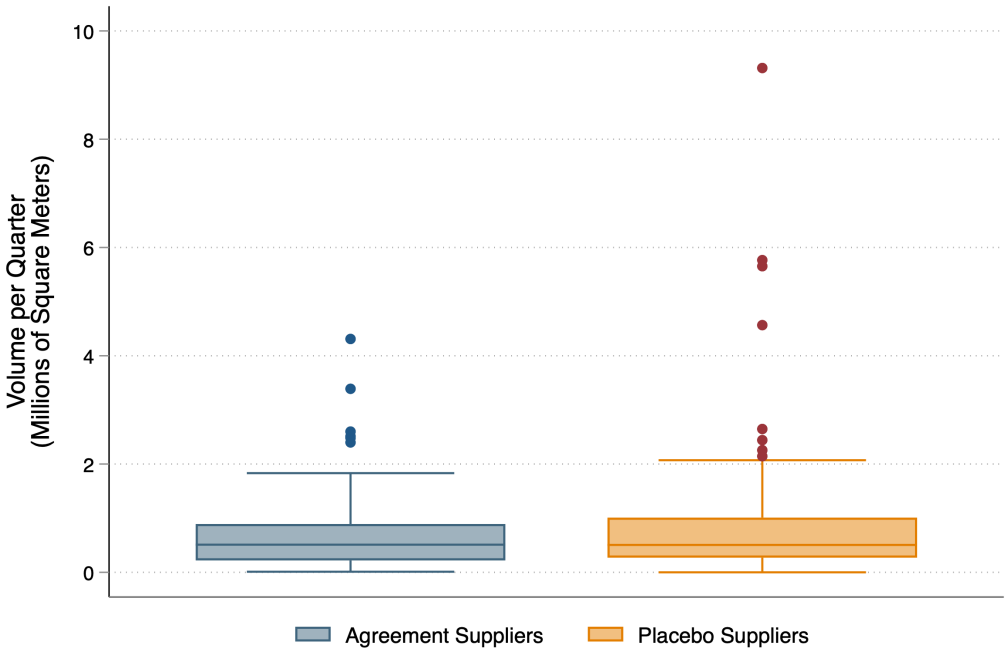
Figure A.3: Volume by Supplier Type (Pre-Integration)



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

A.2 Placebo Supplier Construction

Figure A.4: 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.

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 A.4 shows the similarity of agreement and placebo supplier quarterly volumes. 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).

A.3 Theoretical Bargaining Model Robustness

A.3.1 Relational Contracting 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 has a static 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:

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

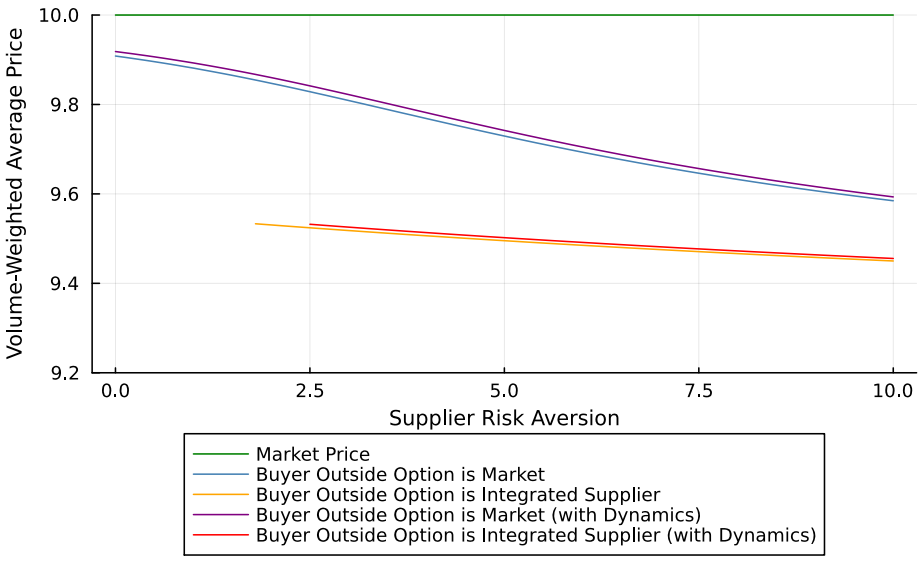
$$O_B = (\delta^B)\Big(-\mathbb{E}_{q^{RC}}[C(q^{RC})]\Big) + (1-\delta^B)\Big(-C(q^{RC})\Big) \text{ if integrated}$$

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.[85]

[85] 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.5: Agreement Prices and Supplier Risk Aversion with Relational Contract Dynamics



Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix A.4.

A.4 Parameterization and Solution of the Theoretical Model

The model is parameterized as:

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

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

A.5 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 procures 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.[86]

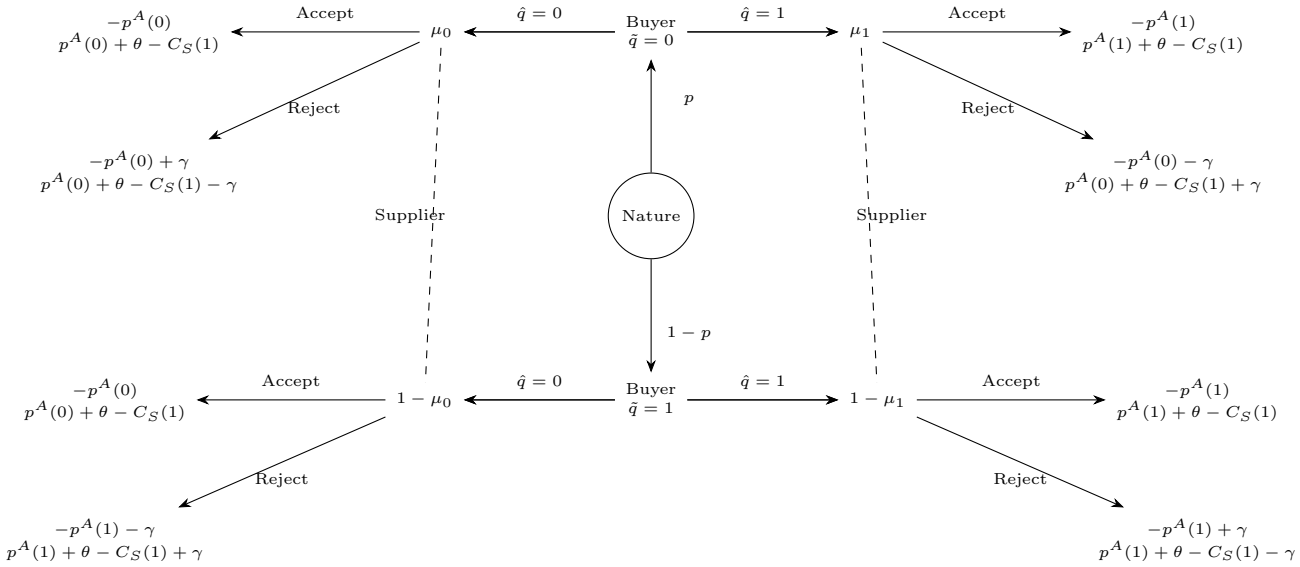
8. Nash bargaining prices are computed assuming that the buyer and supplier equally split the surplus (in other words, bargaining parameters are .5). I assume that the transfer γ does not occur if the agreement dissolves.

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 $\hat{q} = 1$. Note that replacing cooperative Nash bargaining with an infinitely repeated non-cooperative game yields similar qualitative results.

A.5.1 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:



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

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

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

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.

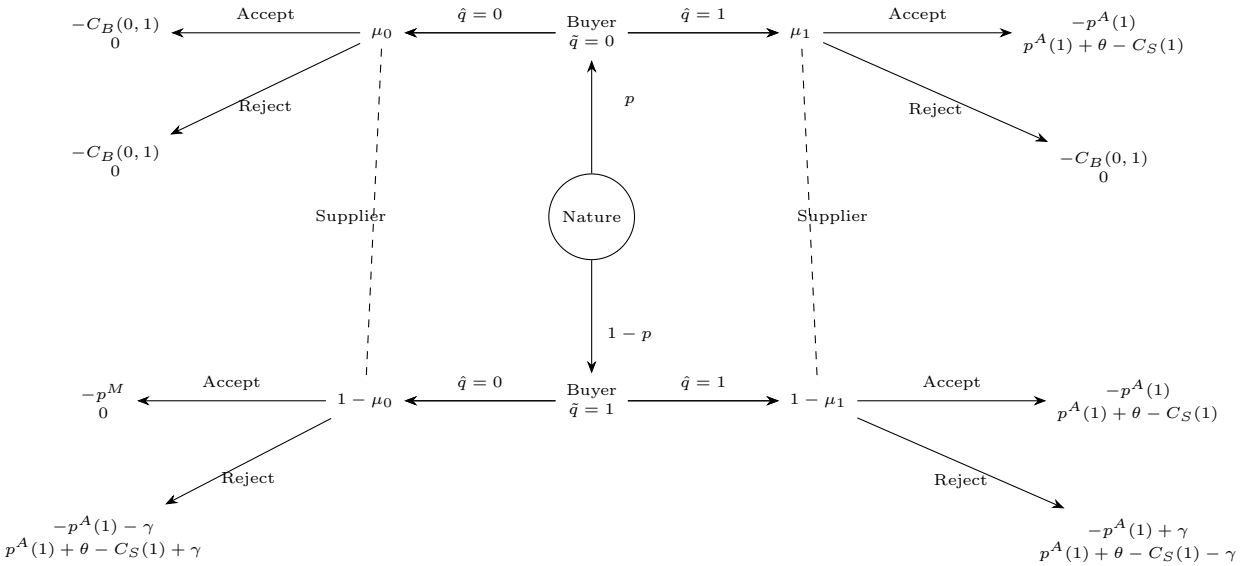
Formally, the equilibrium is that the buyer plays $\hat{q} = 0$ for both types, the supplier

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.5.2 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 on-path, so neither belief consistency nor sequential rationality discipline μ_0.)

Online Appendix OA1.6 provides a worked example of the signaling game to provide further intuition.

A.5.3 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 a risk-averse supplier. Although not shown, the pooling equilibrium at $\hat{q} = 1$ also exists when θ is high. However, if the buyer–as the first mover (plausibly consistent with being the 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.6 Replacement Threats in a Game with Commitment

Consider a game with commitment that proceeds as follows, using notation as in A.5.

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.6.1 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 occurs, 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.6.2 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.6.3 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. Specifically, the buyer has static incentives to deviate from the agreement when paying agreement prices over spot market prices. Instead, relational contract dynamics enforce commitment in this setting. As the relational contract already embeds commitment about prices, the relational contract could plausibly enforce additional comittment of outside options.

A.6.4 Multilaterial Bargaining Extension

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.5, 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.6 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.[87]

[87] 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.

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.[88] 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.[89]

A.7 Outside Estimation Details

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 Online Appendix Table OA1.3). Expressing quantities in logs captures this pattern in a reduced-form manner without introducing additional parameters to govern how the distribution of quantities over time affects 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.[90] 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.

[88] 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.

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

[90] 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, I can also evaluate 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.

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' counterfactual 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.[91] Placebo suppliers are the closest counterfactual for agreement suppliers overall given their similar volumes.[92] 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.[93] 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.

A.8 Structural Bargaining Model: Evidence Supporting Model Assumptions

A.8.1 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

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

[92] 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 OA1.8 to mitigate concerns that errors in outside moments lead to poor model performance.

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

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).[94]

To mitigate these concerns, I directly test the most likely examples of such behavior in Table A.2. 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.2: 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 OA1.8.1 should not find that model prices are unbiased.

[94]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.

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.3 by regressing capacity on an indicator for the post-period and total quantity of fabric procured across all other suppliers. I do not find any evidence that capacities are significantly different post-integration.

Table A.3: 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 OA1.8.1 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 OA1.8.2.

A.9 Structural Bargaining Model: Comparison of Risk Aversion Estimates

Table A.4: Comparison with Other Risk Aversion Estimates

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

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

Table A.5 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.

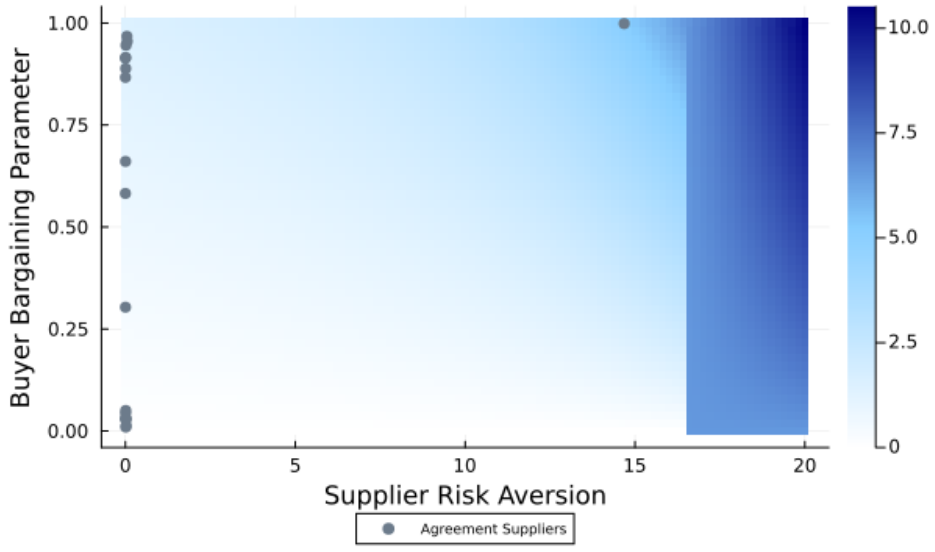
Table A.5: Buyer Bargaining Parameter and Supplier Risk Aversion: 95% CI

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.10.1 Model-Implied Surplus Capture from Agreement Suppliers

Figure A.6: Surplus Capture from Agreement Suppliers



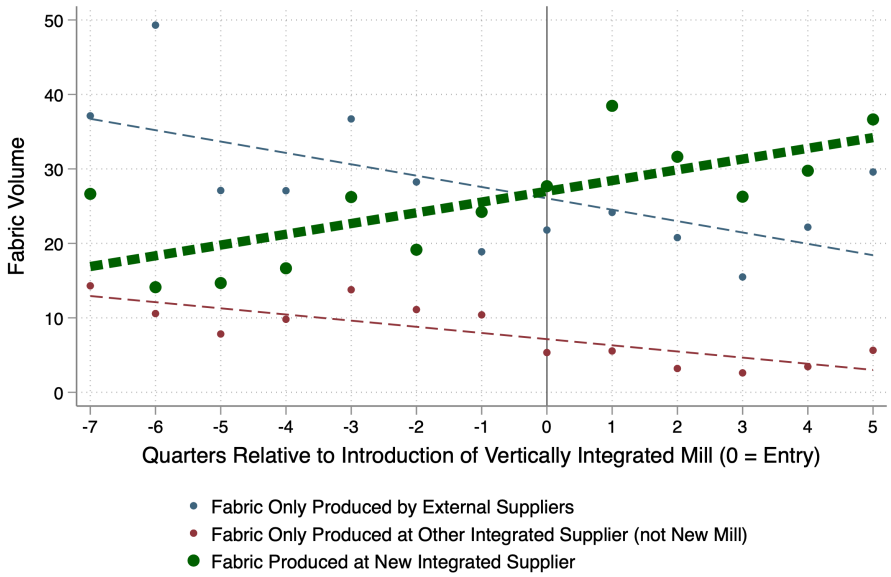
Note: Model estimates.

A.11 Validation using Difference-in-Differences Estimates

A.11.1 Semiparametric Difference-in-Differences Design and Identification

In addition to bringing less production in-house for fabrics with more agreement suppliers, Figure A.7 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.7: 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.[95] 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.[96] Formally, this assumption can be expressed as:

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

In equation 2, trends in fabric demand are included in $\boldsymbol{X_i}$,[97]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.6 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,[97] it is not mechanical that this propensity score approach would achieve balance. I also document in Figure A.8 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.[98]

[95]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.

[96]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.

[97]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.

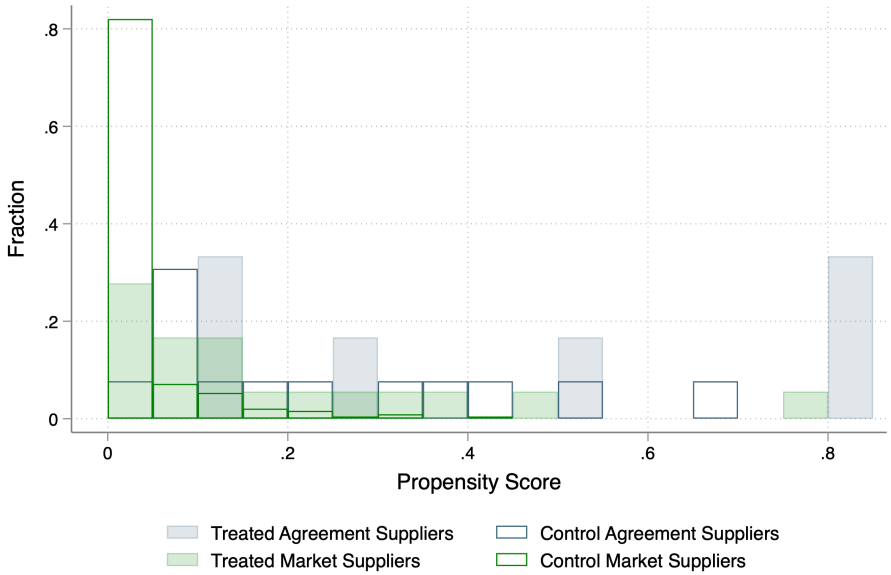
[98]Formally, as stated in Roth et al. (2022), "the conditional probability of belonging to the treatment group, given observed characteristics, is uniformly bounded away from one, and the proportion of treated units is bounded away from zero. That is, for some $\epsilon > 0, \mathbb{P}[D_i = 1|X_i]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

Table A.6: Balance on Untargeted Outcomes

Outcome Variable	Coefficient on Treatment Indicator: 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.

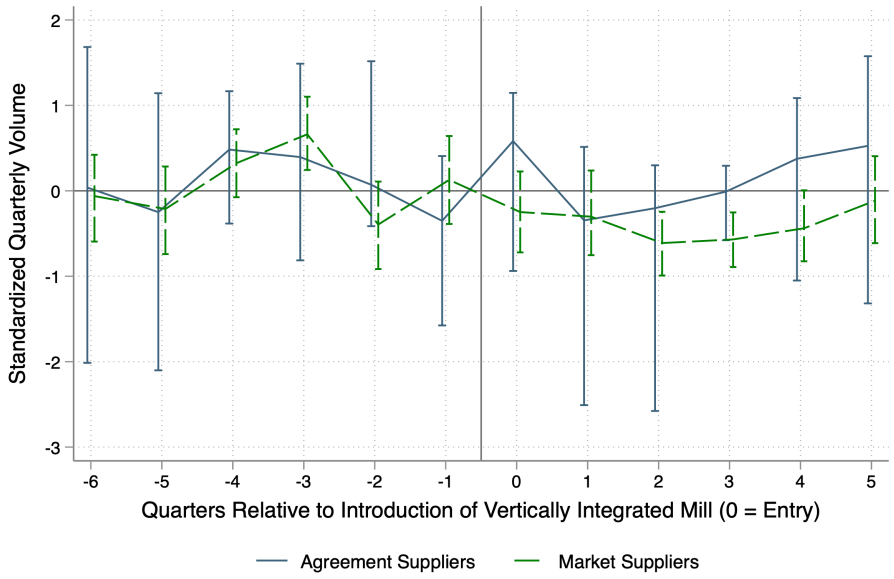
Figure A.8: Distribution of Propensity Scores by Supplier Type



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

A.11.2 Quantity Treatment Effects

Figure A.9: DiD Estimates: Quarterly Volume Effects



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.

A.11.3 Pooled DiD Estimates

Table A.7: 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.

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.

A.12 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).[99] 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 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 sole excluded supplier, as doing so brings down prices from included suppliers.

A.13 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).

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

Table A.8: Buyer Bargaining Parameter and Supplier Concentration

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

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

I find in Appendix Table A.8 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.9: Supplier Risk Aversion and Supplier Concentration

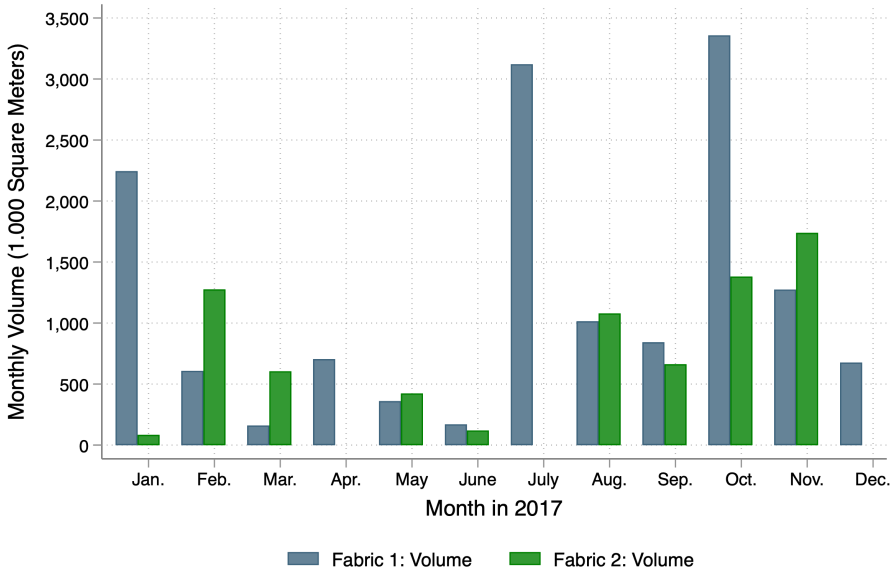
Supplier Risk Aversion		
		No Outlier
Supplier Concentration	−20.417	−.015
(Volume-Weighted Supplier-Fabric HHI)	(13.89)	(.061)
N	18	17
R^2	.355	.006

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

OA1 Online Appendix

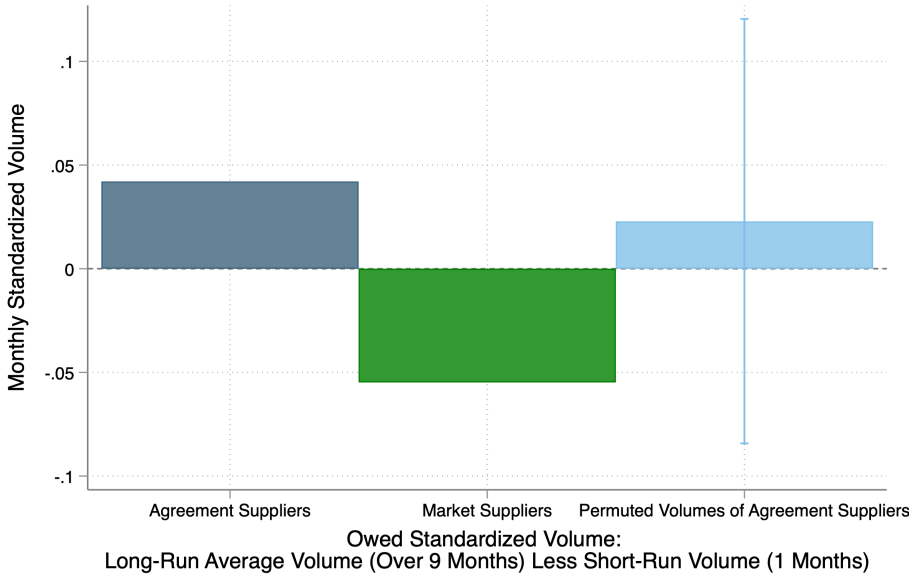
OA1.1 Additional Descriptives About Buyer-Supplier Agreements

Figure OA1.1: Variation in Monthly Quantity by Fabric

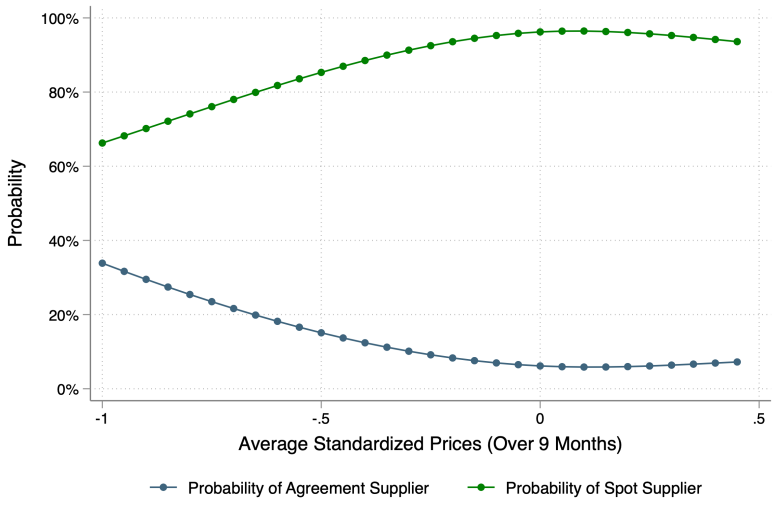


Note: Data from universe of fabric transactions by the buyer. Bars represent the quantity of fabric purchased during the month in 2017 for two selected fabrics. Fabrics were selected based on data availability to illustrate variation.

Figure OA1.2: 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.



Note: Data from universe of fabric transactions by the buyer. Estimated probabilities are from a kernel regression of an indicator for each supplier type on the average standardized markups in the three quarters prior to the construction of the new mill.

OA1.2 Placebo Supplier Analysis

Figure OA1.3: Pricing Function (Placebo Comparison): Binned

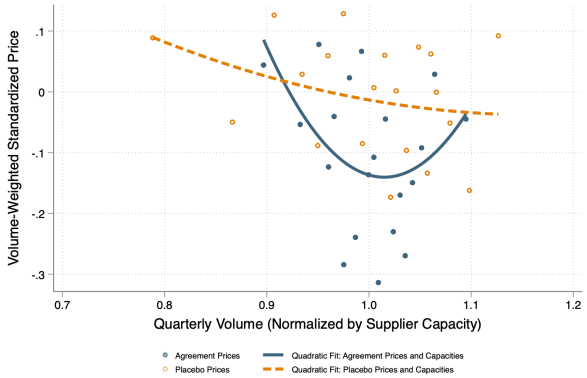
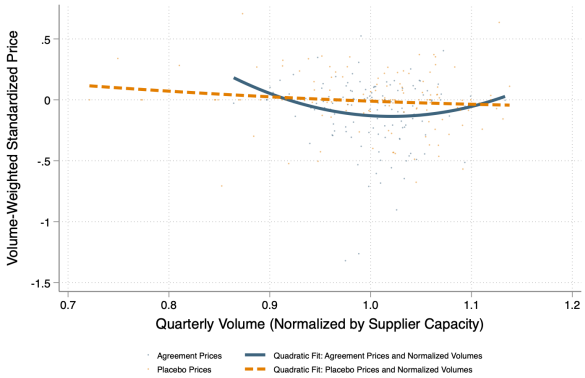


Figure OA1.4: Pricing Function (Placebo Comparison): Raw



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 OA1.3 bins prices and capacities using 20 quantiles. The mapping from quantities to capacities is as discussed in 4.1.2.

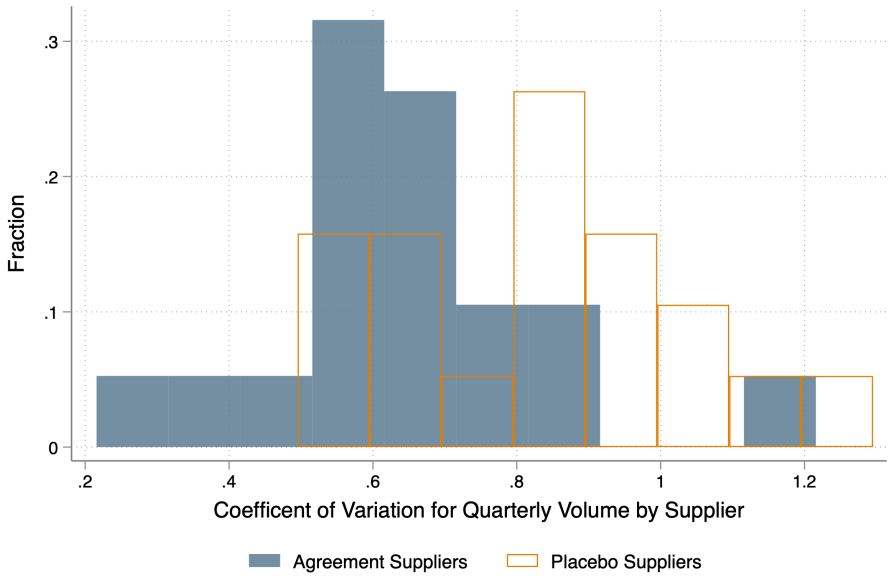
Table OA1.1: Pricing Function Convexity by Supplier Type (Placebo Comparison)

	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.

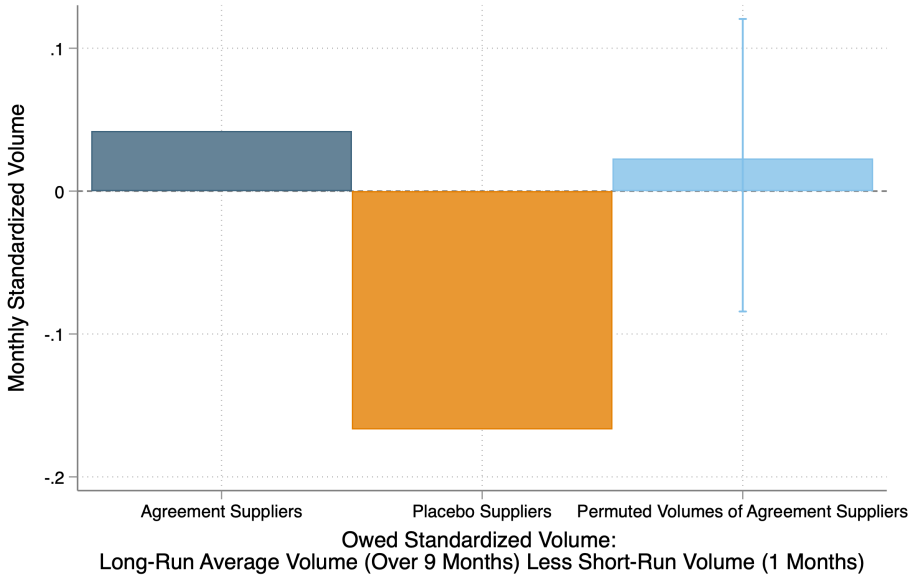
Table OA1.1 illustrates prices with placebo suppliers never meet more than one of the four conditions for being U-shaped.

Figure OA1.5: 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.

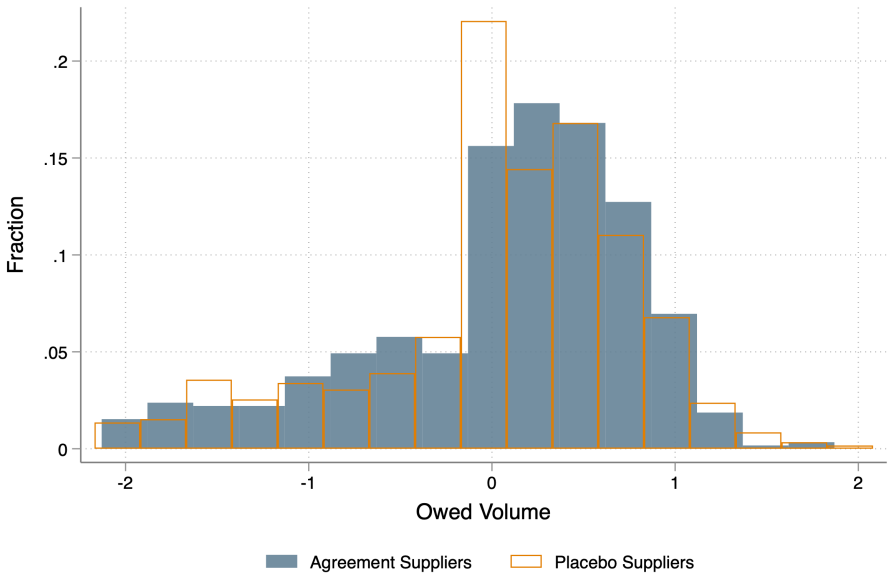
Figure OA1.6: Coefficient on Owed Quantity by Supplier Type (Placebo Comparison)



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 OA1.7: Distribution of Owed Quantity by Supplier Type (Placebo Comparison)



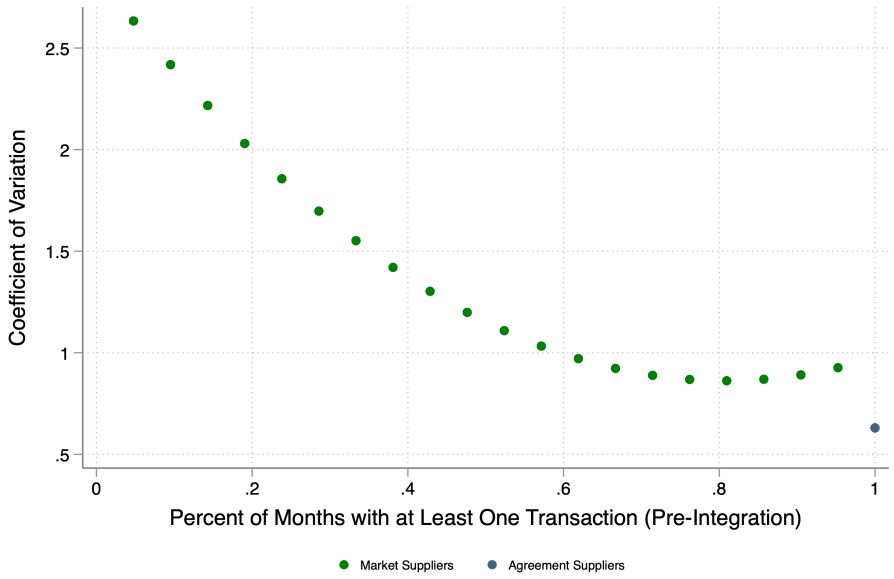
Note: Data from universe of fabric transactions by the buyer. Bars show the coefficient for owed standardized volume from a regression of current volume on owed volume and month fixed effects. Owed standardized volume is calculated for each supplier quarter as the standardized volume in the prior month 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.

OA1.3 Quantities & Prices by Transaction Frequency

Table OA1.2: Quantities, Prices, and Transaction Frequency

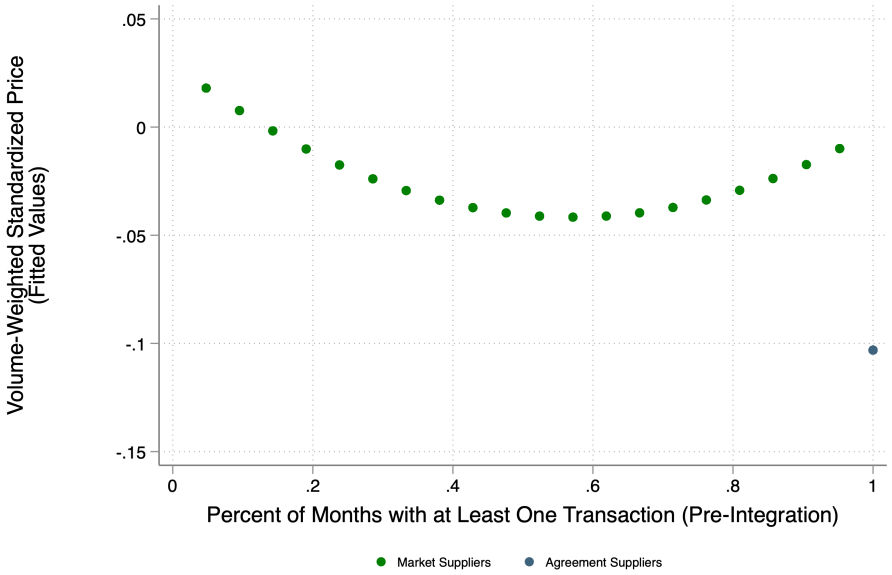
	Coefficient of Variation for Supplier Quarterly Volume	Standardized Price
Agreement Supplier Indicator	−.346***	−.102***
	(.082)	(.031)
Series Controls ($k = 2$) for Transaction Frequency	Y	Y
Quarter FE	N	Y
N	315	1603

Figure OA1.8: Distribution of Owed Quantity by Transaction Frequency



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

Figure OA1.9: Distribution of Average Prices by Transaction Frequency



Note: Fitted values from regression estimate of Data from universe of fabric transactions by the buyer.

OA1.4 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

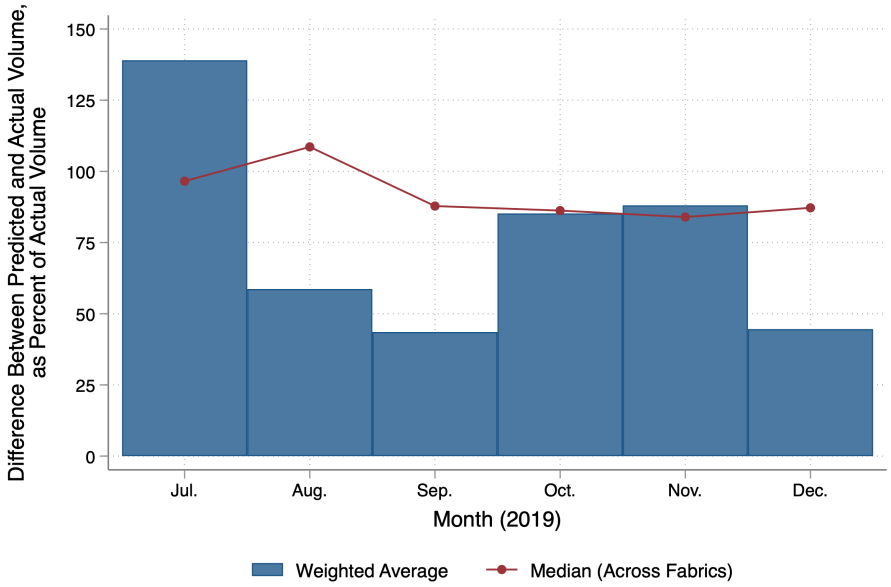
capacity, such a contract reduces supplier incentives to reject sampling opportunities (and, thereby orders), from other buyers to save capacity as the supplier would receive similar markups for all capacity states. It follows that suppliers would pass on costs to the buyer. Note that this concern is 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.

OA1.5 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 OA1.10 shows the average difference between predicted and actual volume, as a percent of actual volume, weighted by the fabric's share of total volume. These results highlight that errors are, on average, always more than 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.3.

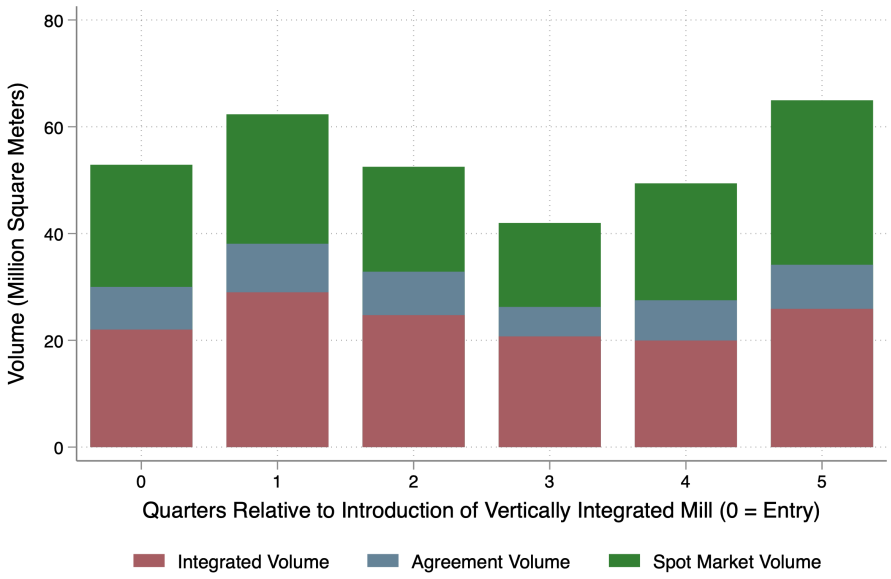
Figure OA1.10: Forecast Error in Volume Predictions



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.

OA1.5.1 Buffer Stock Post-Integration

Figure OA1.11: Post-Integration: Fabric Sourcing by Supplier Type

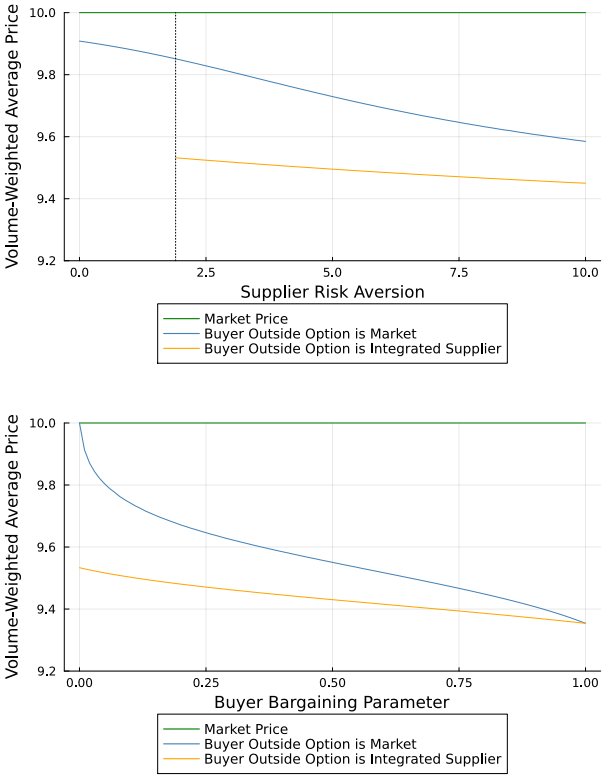


Note: Data from universe of fabric transactions by buyer. Bars represent the quantity of fabric purchased during the quarter from each supplier type.

OA1.5.2 Bargaining Model with Cost Heterogeneity Between External and Internal Supplier

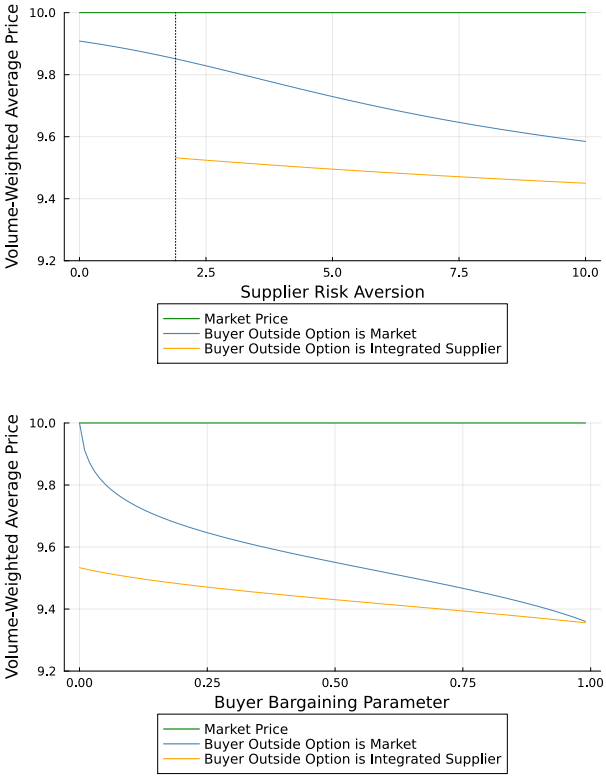
Results are robust to the external supplier having lower or higher costs than the integrated supplier, shifting costs by 10%. To ensure that optimal quantity is still $q = 1$ without changing market prices, I only shift the fixed part of costs by 10%, keeping convexity the same. Results (available upon 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 OA1.12: 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.4.

Figure OA1.13: 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.4.

OA1.6 Signaling Game: 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.

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

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

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

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

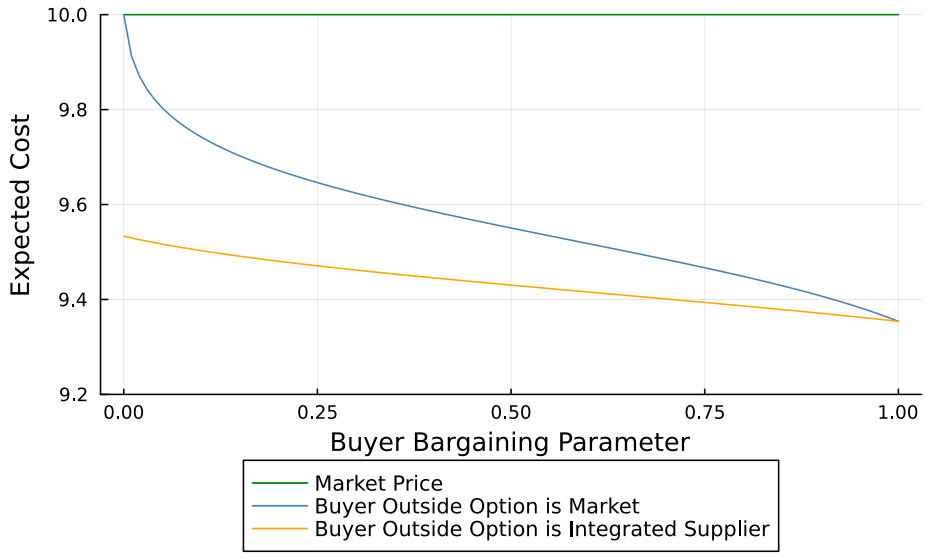
Although I do not observe p directly from the data, note that $\frac{2p-1}{1-p}$ is increasing in $p \forall p \in [0, 1)$; therefore, for p, γ sufficiently high, this inequality holds. Formally, as long as the inequality holds (meaning p, γ are sufficiently high), then the equilibrium features the buyer playing $\hat{q} = 0$ and 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 normalized 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 $\tilde{q} = 1$. It follows that the surplus generated by the agreement is $\theta - C_S(1) + p^M = -.5 - 1 + 2 = .5$. Then, the bargained price is $p^A(1) = 1.75$. However, when $\tilde{q} = 0$, the surplus from cooperation is $\theta - C_S(1) + C_B(1) = -.5 - 1 + 1 < 0$. In that case, the agreement breaks down.

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 $p(0) + (1-p)(p^A(1) + \theta - C_S(1) + \gamma) = (1-p)(.25 + \gamma) > 0 \forall p \in [0, 1)$. 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. Therefore, the buyer chooses between playing $\hat{q} = 0$, where the supplier rejects, and playing $\hat{q} = 1$, which the supplier accepts. The buyer prefers then to play $\hat{q} = 1$ for a cost savings of .25.

OA1.7 Surplus Creation and Capture: Buyer Bargaining Parameter

Figure OA1.14: 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.4.

Figure OA1.14 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.

OA1.7.1 Outside Estimation: Log Quantities

Table OA1.3: 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.

OA1.8 Model Validation and Fit

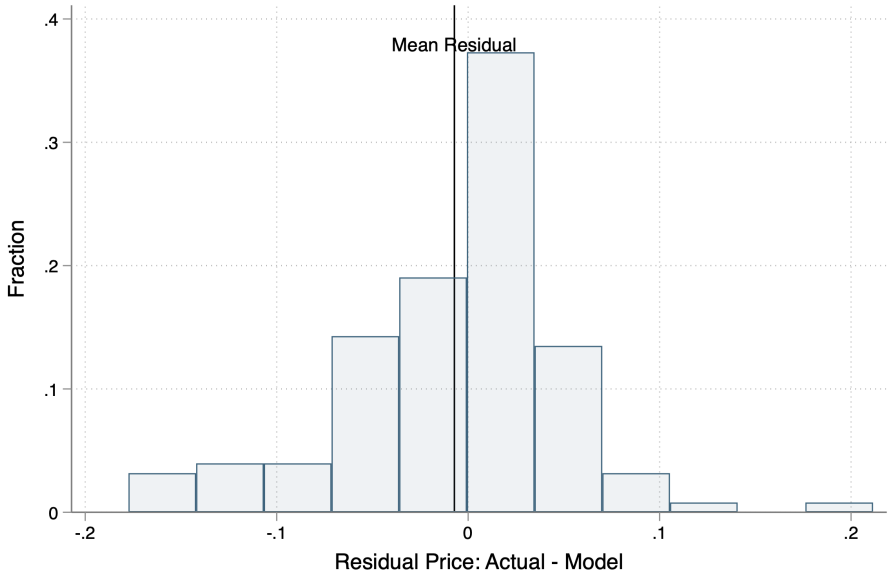
OA1.8.1 Model Validation: Untargeted Moment

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.[100] Figure OA1.15 demonstrates that the distribution of residuals is symmetric with the mean residual

[100]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.

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.

Figure OA1.15: Untargeted Moments: Bias



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

OA1.8.2 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 OA1.4.

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.

Table OA1.4: 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, Table OA1.5 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.

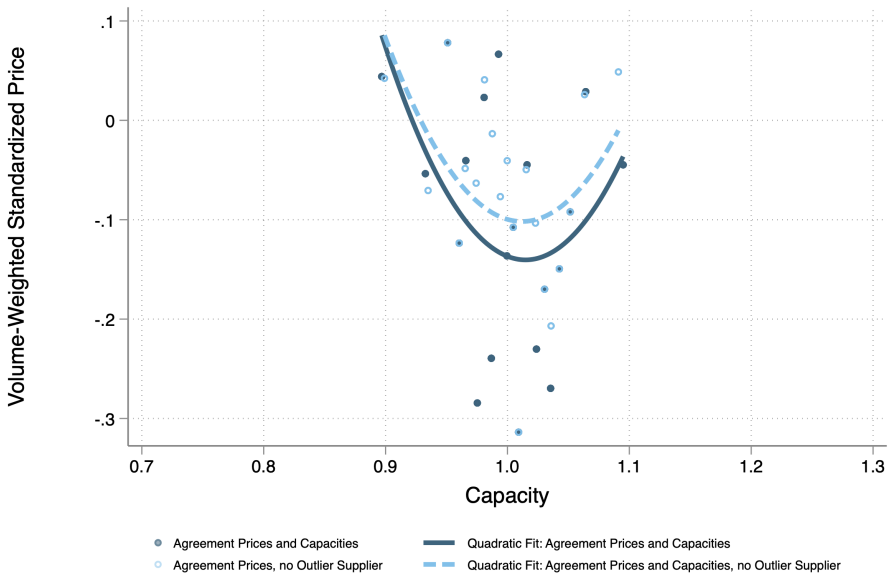
Table OA1.5: Out-of-Sample Fit Comparison (OLS with Tuning)

	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.

OA1.9 Pricing Function Robustness to Outlier Supplier

Figure OA1.16: Pricing Function & Outlier Supplier: Binned

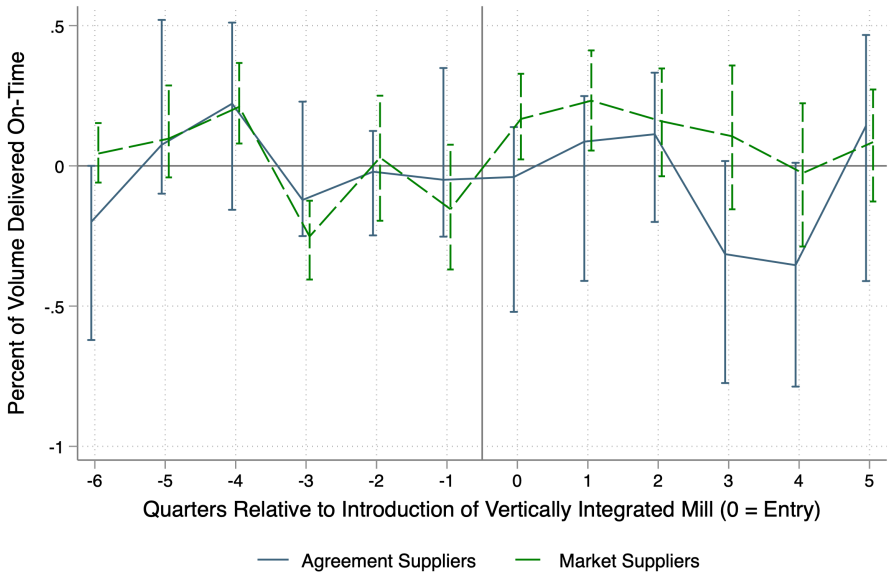


Note: Data from universe of fabric transactions by the buyer with agreement suppliers, transformed into supplier-by-quarter panel data. Quadratic fit is from a regression of volume-weighted standardized price on capacity and capacity squared. The figure bins prices and capacities using 20 quantiles. The mapping from quantities to capacities is as discussed in 4.1.2.

OA1.10 Difference-in-Differences Supplemental Analysis and Robustness Checks

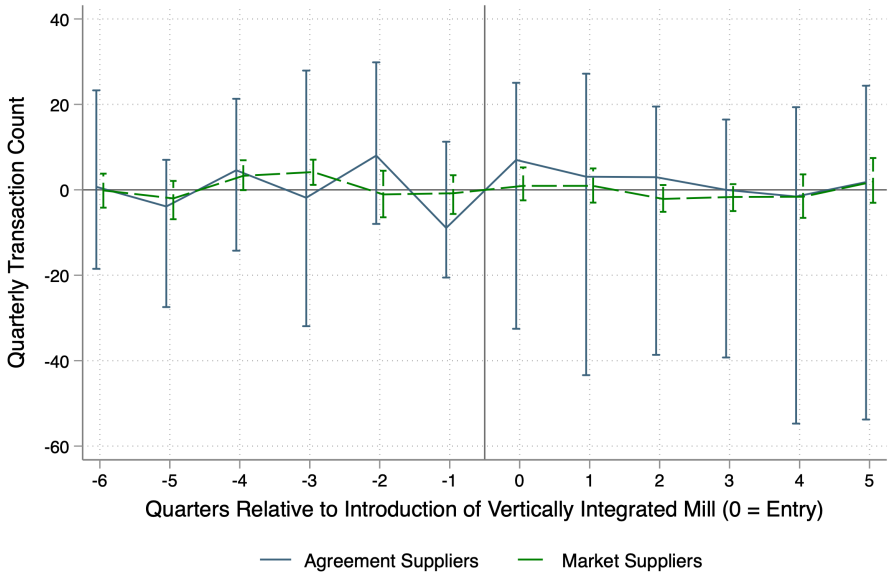
OA1.10.1 Difference-in-Differences Additional Event Study Plots

Figure OA1.17: 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.

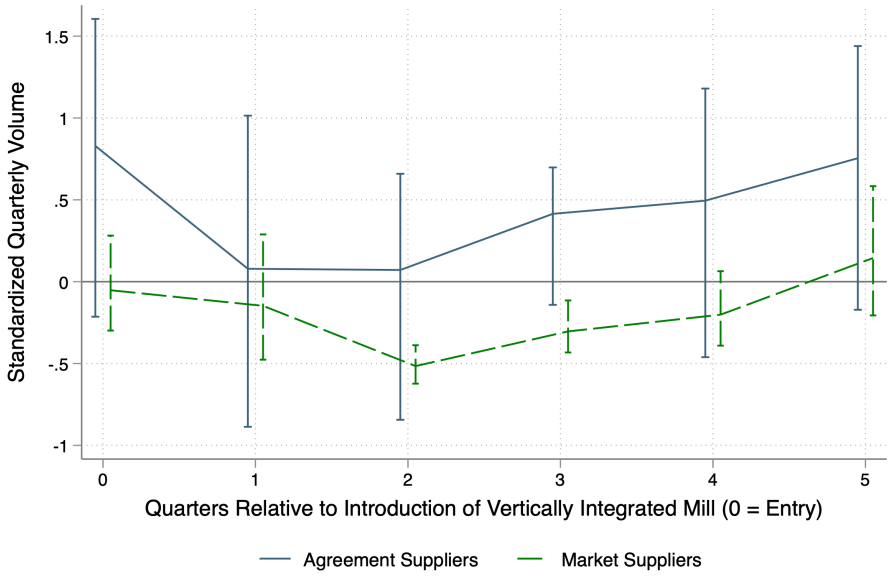
Figure OA1.18: Difference-in-Difference Estimates for Effect of Vertical Integration on Transaction Count



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.

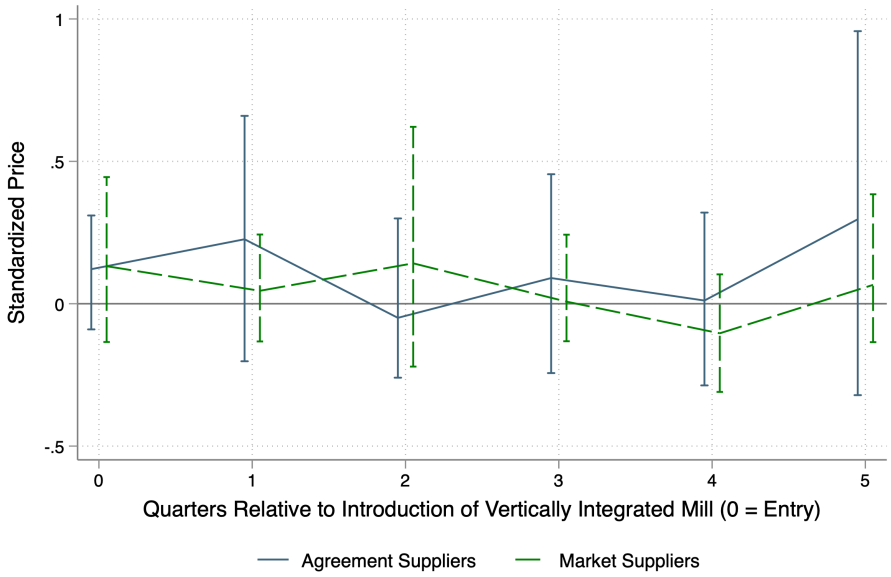
OA1.10.2 Synthetic Difference-in-Differences

Figure OA1.19: 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 OA1.20: Synthetic Difference-in-Differences Estimates for Effect of Vertical Integration on Standardized Prices



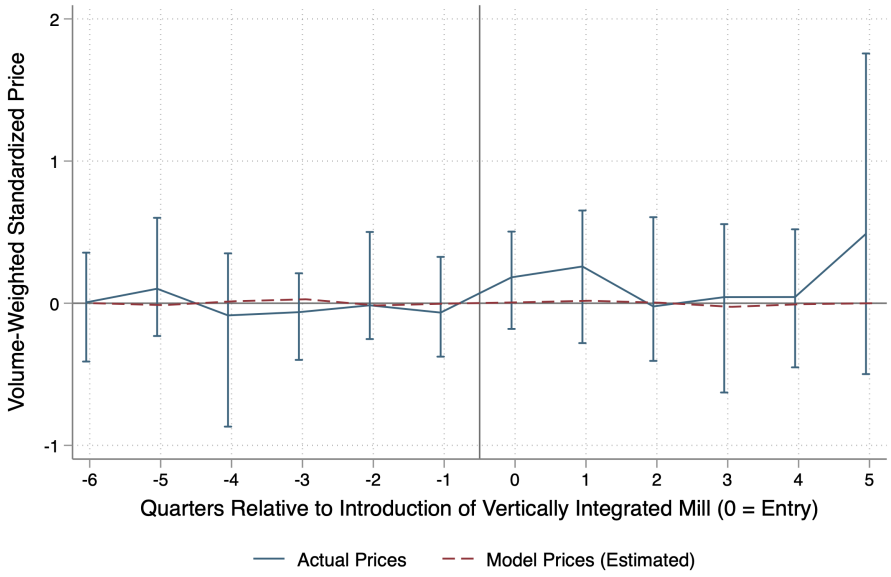
Note: Data from universe of fabric transactions by the buyer. Semiparametric difference-in-differences event study estimates with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

OA1.10.3 Difference-in-Difference Validation: Estimates as Untargeted Moments

As structural model estimation does not use the variation leveraged in the DiD design,[101] 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 OA1.21 are all within the 95% confidence intervals of event study coefficients using actual prices.

[101]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).

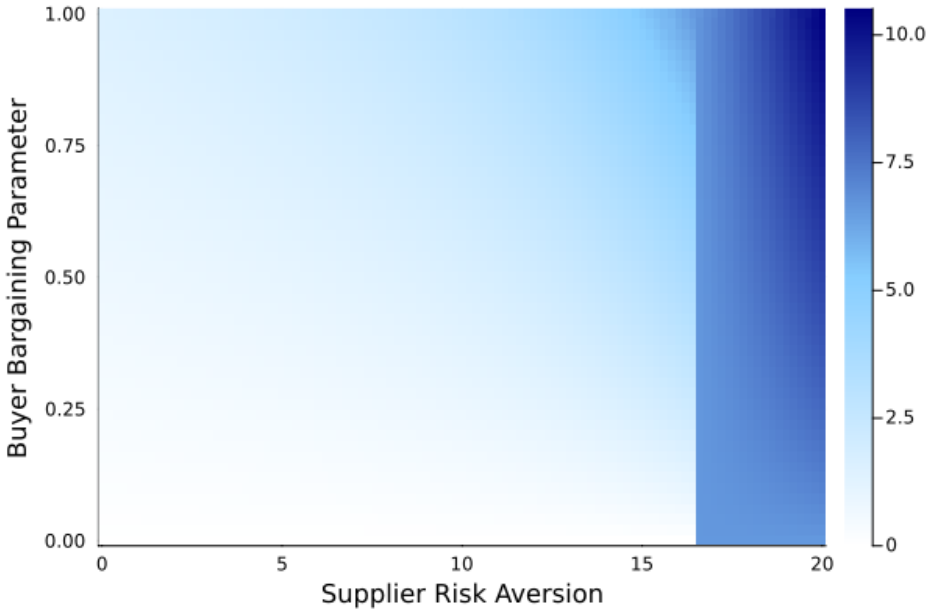
Figure OA1.21: Untargeted Moments: Difference-in-Differences



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.

OA1.11 Buyer Gains from Agreement with DICC

Figure OA1.22: 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 conservative, 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).

OA1.12 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%.[102]

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%.

[102]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.