# Value Creation & Value Capture in Indian Garment Sector Bargaining

Russell Morton*

November 2, 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 versus capture. While exporting can enhance growth, low prices paid by large exporters to their small suppliers raise concerns that gains accrue primarily to large buyers. However, when input markets, or complementary markets like insurance and credit, function poorly, 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—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 through threats to replace external suppliers with in-house production. Using novel data on a major Indian garment manufacturer's universe of fabric purchases from $\sim 500$ suppliers, I estimate a structural bargaining model to decompose how prices reflect these two features. Surplus creation from insurance-like contract terms explains low prices, with the buyer receiving $\sim 60\%$ of gains, rather than capture from replacement threats; difference-in-differences estimates yield consistent findings. Counterfactuals illustrate that greater buyer competition has limited effects on prices paid to small risk-averse suppliers, but 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 raise productivity in low- and middle-income countries (LMICs) (World Bank, 2020), policymakers raise concerns that exporters—the buyers domestically—pay low prices to small suppliers, capturing most gains (Boudreau et al., 2023). These concerns seem especially compelling in LMIC supply chains, where often a few large exporters purchase inputs from many small suppliers (Dominguez-Iino, 2023; Zavala, 2022). However, when input markets, or complementary markets like insurance and credit, function poorly, multifaceted vertical contracts govern these supply chain transactions. Importantly, prices in such vertical agreements reflect more than standard oligopsony markdowns.

Using proprietary data on the universe of fabric purchases by a large export-oriented Indian garment manufacturer, this paper examines how prices reflect two key vertical contract features that emerge when markets function poorly. First, agreements create surplus by mitigating frictions in input markets or complementary markets, such as insurance and credit. Fabric mills (the suppliers) face variable, lumpy orders, exposing them to quantity and profit risk. Agreements can specify stable quantities ("*quantity assurance*"), providing insurance against this risk.[1] When insurance markets are missing or frictional, suppliers may accept lower prices in exchange for quantity assurance. Prices may then 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;[2] then, prices bundle the transfer for the input with interest. Second, many large buyers respond to poor spot market performance by partially relying on in-house production, including the manufacturer studied in this paper.[3] Partial vertical integration could enable threats to replace external suppliers with in-house production, strengthening the buyer's bargaining position and facilitating surplus capture.[4]

This paper develops a framework to interpret vertical contract prices by decomposing them into surplus creation and capture. This decomposition clarifies that prices reflect more than markdowns from standard oligopsony models (*e.g.,* Cournot or differentiated-products

[1] Agreements between international buyers and exporters of Kenyan roses (Macchiavello and Morjaria, 2015) and between 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.

[4] Morton (2023) shows that vertical integration and agreements act as substitutes.

Bertrand), which feature only one source of market power: the wedge between supply and the buyer's marginal benefit of the input.

Distinguishing between these two mechanisms is necessary to understand the distributional effects of vertical relationships. Though both mechanisms can generate low prices, they have fundamentally different welfare implications. If low prices embody surplus created by mitigating frictions inside the vertical agreement, they represent the buyer's share of that surplus. Conversely, replacement threats reduce prices purely through rent extraction, which I refer to as "capture." Buyer market power simultaneously reflects both outside options—which incorporate surplus capture via replacement threats—and bargaining weights—which govern the buyer's share of surplus created. Crucially, only buyer power over surplus creation aligns incentives for agreement terms that mitigate frictions, such as insurance-like terms.

Empirically quantifying the separate effects of surplus creation and capture on prices requires overcoming two identification challenges. First, as agreements are informal (*i.e.*, no legal contract exists), rather than directly observing contract terms, I infer terms by comparing transactions across agreement and spot suppliers with identical price-relevant characteristics. Comparisons leverage the richness of the data, which cover ∼25,000 transactions with highly granular observables (*e.g.*, ∼6,200 product categories). Two stylized facts emerge: *i)* agreements reduce, without eliminating, quantity variability; and *ii)* agreements feature non-linear prices that stabilize profits,[5] while spot suppliers receive flat prices. Because 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: *i)* surplus from mitigating market frictions (the inside option), namely gains from insurance-like agreement terms in my setting; *ii)* the value of the next best alternative (the outside option), which incorporates replacement threats from in-house production; and *iii)* the surplus sharing rule (the buyer bargaining parameter). While the institutional context enables direct estimation of the outside option,[6] I develop a new approach to separate the *level* of surplus created from the *sharing rule* that divides it. My strategy leverages exogenous supplier-specific demand shocks over time.[7] As price adjustments to demand shocks provide insurance, their magnitude identifies supplier preferences for insurance-like terms.

[5]Prices are highest when profits would be lowest with constant prices, and *vice-versa*.

[6]The outside option, when the replacement threat is not exercised, is the spot market for both parties 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.

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

Conversely, average price levels identify buyer bargaining parameters.

I build a Nash bargaining model of the vertical contract that incorporates both surplus creation from insurance-like terms and rent extraction from threats to replace external suppliers with in-house production. The model formalizes how supplier preferences for smooth profits shape price effects of insurance-like agreement terms. Such preferences can reflect varied microfoundations, including borrowing and savings frictions that lead suppliers to fund frequent outflows (*e.g.*, payroll) from current profits.[8] For brevity, I hereafter refer to such preferences as "risk aversion." More risk-averse suppliers derive greater value from insurance-like terms, increasing the surplus generated by the agreement. Because the buyer receives its surplus share through discounts, greater surplus creation reduces prices. Therefore, surplus creation from insurance-like terms makes prices decrease in supplier risk aversion.

The model not only extends standard replacement-threat results (Ho and Lee, 2019; Ghili, 2022) by analyzing threats from in-house, rather than external, suppliers but also introduces a novel cutoff condition: threats only reduce prices when suppliers are sufficiently risk-averse. When risk aversion falls below this cutoff, replacement threats break agreements as insurance-like terms generate insufficient surplus. Because agreement and integrated quantities are inframarginal and cannot easily expand,[9] breakdown pushes sourcing to the expensive spot market rather than in-house production.[10] Therefore, the buyer "as if" endogenously chooses its outside option to preserve agreement discounts. It uses in-house supply as a replacement threat with more risk-averse suppliers while maintaining spot supply as the outside option otherwise.[11] Overall, the model advances a new explanation why replacement threats may not reduce prices, driven by the interaction of in-house, agreement, and spot supply.

I find that surplus creation alone, rather than surplus capture, generates all price reductions relative to spot prices, with the buyer retaining $\sim$60% of surplus created. 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

[8]Differences in financial access can generate heterogeneity in preferences for smooth profits. Alternatively, such preferences could arise from small firms inheriting the owner's risk aversion.

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

[10]Integrated supplier cost convexity, from capacity constraints, makes the buyer prefer increasing spot sourcing to expanding in-house production.

[11]I formalize that the buyer endogenously choosing its outside option as a reduced-form representation of either of two distinct games with both cooperative (*i.e.,* Nash bargaining) and non-cooperative stages in the spirit of Miller and Watson (2013). The first game has the buyer signal its residual in-house capacity, which is private information that varies stochastically and determines the actual outside option. The second considers the buyer committing to residual in-house capacity before bargaining, mirroring other agreement commitments (*e.g.,* paying agreement prices over spot prices at times).

for LMIC coffee mills (Blouin and Macchiavello, 2019).[12] On average, agreements provide a 1.5% discount relative to spot 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)* increased vertical integration—the buyer builds an integrated mill—and *ii)* heterogeneity in supplier exposure to it.[13] This validation complements the structural analysis because the two approaches rely on distinct variation: exogenous demand shocks identify risk-aversion parameters in the model, while the DiD uses a change in integration. Consistent with model predictions, as all supplier risk aversion estimates fall below the cutoff for surplus 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 (Ho and Lee, 2019) built for a setting without a spot market, highlighting how sourcing marginal quantities through spot markets shapes the effects of replacement threats.

Through the lens of the model, I quantify both *i)* the cutoff in 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 implausibly above other estimates.[14] In (hypothetical) agreements with suppliers sufficiently risk averse to be vulnerable to 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%).[15] Moreover, such highly 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).

Last, I use the estimated model to evaluate two counterfactual policies, focusing on the role of supplier risk aversion. First, I consider increased buyer competition (*e.g.*, from export promotion), modeled as the buyer awarding suppliers a higher 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

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

[13]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.

[14]Settings where agreements provide more quantity stability or costs are steeper feature a lower cutoff.

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

with risk aversion above the replacement threat cutoff, agreement surplus relative to in-house production is so small that even transferring all to the supplier barely shifts prices. Ignoring replacement threats overstates the benefits of competition by considering suppliers' outside option as the spot market, where much more surplus is available for redistribution. With more surplus to allocate, reallocation can meaningfully affect prices.

Second, introducing the missing product for profit insurance against spot market quantity risk increases prices broadly, especially for risk-averse suppliers. Insurance enables suppliers to smooth profits outside the agreement, strengthening 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 support profit smoothing by improving 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 delivers benefits not only by equalizing utility across states but also by raising prices. Therefore, insurance delivers not only well-known efficiency gains through smoothing, but also yields an important equity benefit 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 advances three specific literatures. First, it contributes to empirical bargaining (*e.g.*, Grennan (2013); Backus et al. (2020); Keniston et al. (2022); Freyberger and Larsen (2023)) by studying how weak markets shape vertical agreements. I highlight how agreements mitigate frictions within the vertical contract, creating surplus that partially integrated firms might capture through replacement threats from in-house production. To quantify these mechanisms, I incorporate risk aversion—as the friction-mitigating terms in my setting provide insurance—into my empirical model and develop a new identification strategy that separates risk aversion from bargaining parameters by leveraging how demand shocks affect both price levels and curvature. This approach has broad value as the trade-off between incentives and insurance is central to vertical relationships (Lafontaine and Slade, 2007). Models that abstract from risk aversion and insurance (*e.g.*, Cuesta et al. (2019); Lee et al. (2021)) underestimate the more risk-averse party's bargaining parameter, potentially generating biased price estimates. To study capture, I extend replacement threat models (Ho and Lee, 2019; Ghili, 2022) to settings where spot, integration, and vertical contracts co-exist. I show that buyers sourcing marginal quantities from spot markets *refrain* from replacement threats that break agreements.

Second, this paper contributes to vertical contracting in supply chains, particularly relational contracts enforced by future surplus rather than formal mechanisms (*e.g.*, Macchiavello and Morjaria (2015, 2020); Ghani and Reed (2022); Harris and Nguyen (2023); Brugues (2025), *etc.*). Prior work emphasizes relational enforcement; I instead focus on the surplus-creating role of terms that address market frictions and relational surplus sharing, which remain understudied. I quantify how prices reflect surplus creation versus rent extraction—crucial for welfare analysis but not directly observable in the data. To do so, I estimate a bargaining model rather than relying on commonly used take-it-or-leave-it models of relational contracts. While such models are well-suited for studying enforcement, they mechanically attribute all surplus to one party, making them unable to address distributional questions. I also examine how relational contracts respond to changes in the *vertical* organization of production, as prior work generally considers the effects of changes to market structure—the *horizontal* organization of production. As vertical integration—unlike market structure—is a strategic firm choice that potentially facilitates rent extraction via replacement threats, it introduces novel strategic considerations. This co-existence of integration and relational contracts is empirically relevant, as both arise when spot markets are weak.

Third, I contribute to research on market power along LMIC supply chains (*e.g.*, Boudreau et al. (2023); Dominguez-Iino (2023); Zavala (2022)). I show that policymaker concerns about buyers capturing gains from exporting should account for value created by contracts addressing market failures. To support small suppliers, I emphasize that policies targeting missing insurance markets can increase prices by strengthening suppliers' bargaining positions.

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.

# 2 Data, Context, & Stylized Facts

My empirical setting centers on an Indian garment manufacturer sourcing fabric simultaneously from agreement suppliers, the spot market, and in-house production. First, I describe the data and institutional context on fabric procurement. Next, I document two stylized facts about bargained agreements: *i)* agreements provide quantity assurance; and *ii)* agreement prices have a U-shaped relationship with quantities and a discount, on average, relative to spot prices. I interpret these facts as agreements providing insurance to suppliers. I then describe the increase in vertical integration from constructing a new mill. Last, I show the

buyer's overall fabric sourcing, illustrating that *marginal* quantities are from spot suppliers.

## 2.1 Data

I primarily rely on detailed proprietary transaction-level data on the universe of the buyer's fabric purchases between September 1, 2016 and December 31, 2019. These data are well-suited to study vertical relationships for two key reasons. First, I observe not only the price, quantity, and date but also granular information about the specific fabric (6,176 codes) and supplier identity. The 34,681 transactions represent over one billion square meters of fabric (worth ∼$180 million) purchased from 622 distinct suppliers (471 pre-integration). By analyzing transactions from one large buyer, and being able to track suppliers over time, I can compare transactions across supplier holding both the primary product characteristic (*i.e.,* what fabric it is) and the buyer fixed. Second, the transaction data include estimated production costs for fabrics that can be produced internally, allowing me to directly observe markups for some fabrics.[16]

I supplement the transaction data with cost data from the integrated mills that show how costs relate to capacity.[17] The data include a broad support of capacity levels from 41% to 102% (capacity over 100% indicates production on Sunday and/or overnight).

## 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 fabric type and color. The buyer then selects a supplier (a mill) to provide fabric for a few prototype garments, called "samples".[18] 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 an exact fabric match, as subtle fabric features are hard to reproduce and even minor deviations risk rejection by the end client.[19] Short timelines prevent trial-and-error with a different supplier to match the fabric, as only 28 days typically separate fabric purchases for samples and the

[16]Cost estimates are for a mill that is not capacity constrained and represent average variable costs.

[17]Data are a division-by-month panel from March 2019 to April 2020, where divisions refer to specific processes for fabric production without substitution across each other, such as weaving, dyeing, and knitting.

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

[19]Fabric procurement managers report that the sampling supplier fulfills the entire order for $\sim$ 90% of orders.

full order, consistent with "fast fashion."[20]

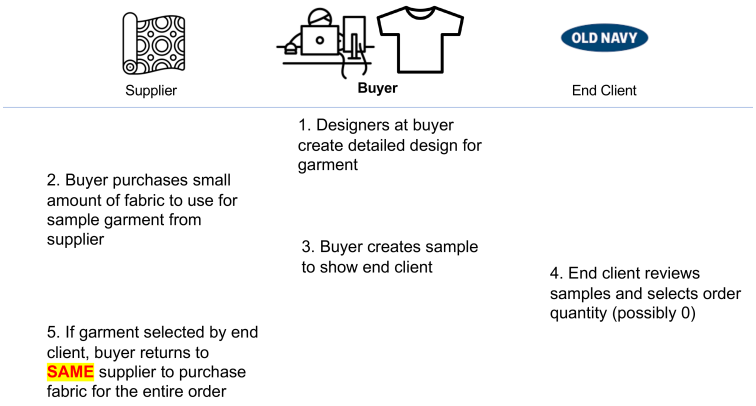



Figure 1: Garment Manufacturing Process

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 will fulfill the entire order but initially know neither which designs will be approved nor in what quantities. Second, because the buyer cannot switch suppliers after sampling, suppliers charge a high price, effectively acting as monopolists and holding up the buyer.[21] (Online Appendix OA1.4 discusses why prices for the full order are not set before 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 providing suppliers insurance in exchange for lower prices could benefit both parties. Consistent with this intuition, I document two stylized facts: *i)* agreements feature quantity stability, and *ii)* their prices are U-shaped in capacity utilization, with a discount on average for the buyer.

Because agreements are informal and unwritten, I infer their terms empirically by comparing transactions across agreement and spot suppliers. As the data do not directly indi-

[20] In the unlikely case that the supplier changes after sampling, production is typically brought in-house, facilitating stringent quality control, and is for an order with an unusually long time before delivery, permitting experimentation to match the fabric.

[21] The buyer is unlikely to hold up a supplier because fabric prices are set after the end client has ordered the garment. Then, while the specific fabric is very valuable to the buyer, the supplier could sell capacity to others. The buyer's primary opportunity to hold up a supplier would likely be after fabric production. However, unlike the agreement overall, individual transactions have simple explicit contracts that are presumably enforceable. Furthermore, reputational motives deter buyer hold-up.

cate supplier type, I classify suppliers by transaction frequency. This approach reflects that agreements are relationally enforced and relational contracts require repeated interaction to sustain the future surplus for self-enforcement—the standard in the literature.[22] I classify suppliers transacting every month as agreement suppliers.[23] These 19 suppliers account for ∼25% of pre-integration volume but only ∼4% of the 471 pre-integration external suppliers.

### 2.3.1 Fact 1: Agreements Provide Quantity Assurance

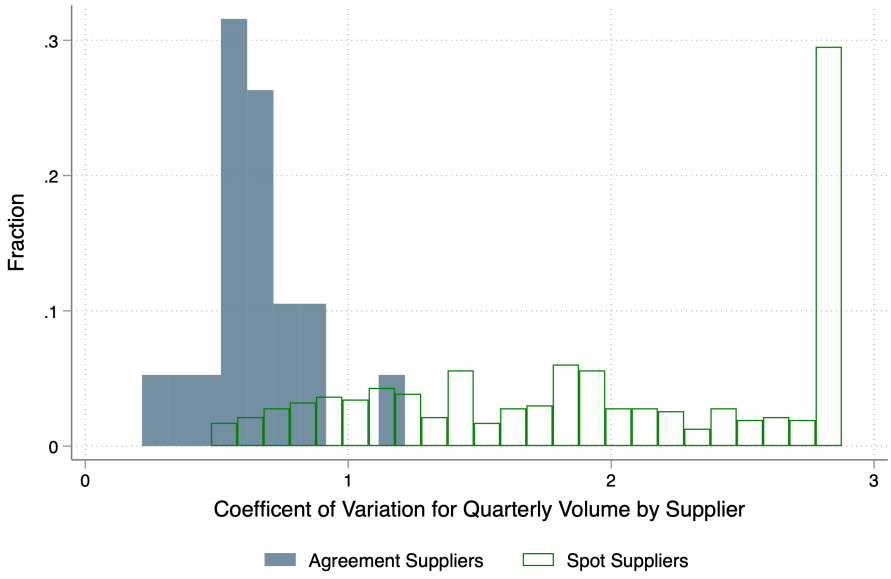



Figure 2: Coefficient of Variation and Supplier Type

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

First, I document that agreements provide quantity assurance (*i.e.,* reduce quantity variability). Figure 2 documents that the coefficients of variation for quarterly volume per supplier are noticeably lower for agreement than spot suppliers.[24] This result is not entirely mechanical: supplier classification depends on transaction frequency, not volume, so

[22]For example, Macchiavello and Morjaria (2015) and Macchiavello and Miquel-Florensa (2017) define relational contracts based on consistent trading every period ("at least three consecutive seasons" in the latter). 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).

[23]Classification only considers pre-integration months to ensure that analyses of the effects of vertical integration on agreement suppliers hold supplier type fixed. Appendix Figure A.1 shows the distribution of the percentage of months with transactions across suppliers.

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

agreement suppliers could have volatile quantities. Appendix Section A.1.2 provides additional evidence of quantity assurance by showing that agreement suppliers who are "owed" volume—those with below-typical recent quantities—receive additional volume (and vice-versa). This additional volume exceeds simple mean reversion. 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.

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

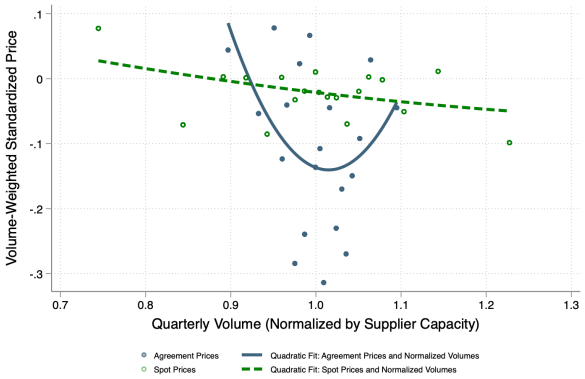



Figure 3: Pricing Function: Binned

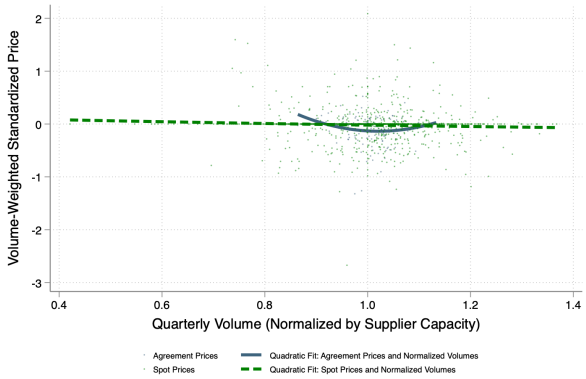



Figure 4: With Insurance

**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.

The second fact illustrates that agreements feature a U-shaped pricing function with respect to quantity (normalized by supplier capacity), while spot suppliers receive flat prices. To establish this fact, I build and analyze a supplier-by-quarter panel with each supplier's volume-weighted average standardized price and capacity utilization. Standardizing prices by month and fabric isolates the discount the buyer receives by removing cost differences from seasonal fluctuations[26] and input price differences.[27] Quantities are expressed as quarterly volumes relative to supplier productive capacity—effectively, quarterly supplier-specific

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

[26]While some seasons have systematically higher capacity utilization, such as the early fall for December holiday shopping, individual fabrics have large, unpredictable variation (see Online Appendix Figure OA1.1).

[27]To illustrate, consider two suppliers charging the same price for fabrics with different input costs. Then, the supplier of the high-cost fabric charges a lower markup.

*capacity utilization*—to account for heterogeneity in supplier scale.

Figures 3 and 4 show that agreement prices are U-shaped 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, but remains visible in the raw scatter despite a wider y-axis scale. Appendix Section A.1.3 confirms that only agreement, and not spot, prices statistically exhibit a U-shape. Furthermore, average agreement prices are below spot prices, highlighting that agreements include a discount (shown 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 generates profit variability.[28] Therefore, quantity assurance stabilizes profits (but agreements do not fully eliminate variability because orders are lumpy and cannot be split across suppliers).[29]

Second, the U-shaped pricing function also reduces profit variability relative to constant prices. When quantities are low, profits with constant prices would be relatively low because few units are sold; higher prices for low quantities smooth profits. When quantities are high, high 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. Higher prices at high capacity utilization offset these costs, further stabilizing profits.[31] Appendix A.1.3 illustrates that agreement prices are lowest when when profits would be highest with constant prices (at 100% capacity utilization), further supporting the insurance interpretation.[32] (Separately, I interpret constant spot prices as indicating that suppliers hold up the buyer. Therefore, spot prices are determined by the *buyer's* outside option, which is independent of a single supplier's capacity given their small size relative to the buyer.)

Last, I show that the stylized facts are unlikely to reflect two alternative explanations: *i)*

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

[29]Limited quantity assurance reflects that the buyer's primary instruments to stabilize quantities—frequency of sampling and the distribution of likely quantities for a given end client matched to the supplier—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: it would require extreme prices for some quantities or large fixed fees, either of which could induce buyer deviation. Additionally, quarterly supplier costs are unknown until the very end of the quarter, preventing an incentive-compatible spread of payments over time.

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

unobserved supplier heterogeneity between agreement and spot suppliers; and *ii)* misclassification of suppliers as agreement versus spot. Importantly, the U-shaped pricing function is difficult to rationalize besides insurance. Any alternative explanation must simultaneously account for: *i)* agreement prices are highest at both quantity extremes, *ii)* agreement prices are lowest at exactly 100% capacity utilization, *iii)* spot prices are flat for the same product, *iv)* agreements feature quantity assurance. For example, flat spot prices reject that the U-shape reflects a constant markup over U-shaped total cost curves.[33]

First, I consider unobserved supplier heterogeneity. For example, some suppliers may be more efficient (perhaps due to better management as technology is largely standardized among relevant suppliers). More efficient suppliers might pass on cost savings to the buyer, reducing average prices. However, efficiency alone does not explain why agreements feature more stable, rather than larger, quantities nor U-shaped prices. Regardless, to test this hypothesis, I replicate agreement versus spot supplier comparisons using high-volume spot suppliers ("placebo suppliers") that transact frequently but are plausibly not in agreements.[34] Importantly, plausible explanations of unobserved heterogeneity largely imply that supplier volume, rather than transaction frequency, should shape transactions. 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, implying they arise from agreements rather than correlates of size.[35]

Second, classifying suppliers as agreement or spot based on transaction frequency may introduce measurement error; frequent transactions are necessary but not sufficient for agreements. Such measurement error should generally bias results *against* finding meaningful differences across supplier types. Nevertheless, the placebo analysis suggests that misclassification is minor, as many placebo suppliers trade with the buyer almost every month but display different transaction patterns. Moreover, Online Appendix OA1.3 shows that both quantity variability and average prices—as the U-shaped pricing functions include a discount—shift by a statistically significant amount at the classification threshold.[36] Subsequent empirical analysis estimates supplier-specific parameters that rationalize observed

[33]Even a more convoluted story, such as U-shaped average *total* cost curve from fixed costs at the quarter level, would require the buyer to have more market power in agreements than in spot markets; however, unless agreement suppliers are benefiting from the agreement in some other way, agreement suppliers would simply deviate and hold up the buyer, charging the high spot price. In the agreements studied, I consider insurance to be the extra benefit in the agreement.

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

[35]This result does not imply that agreement suppliers are more or less efficient than other suppliers, but only that observed differences primarily reflect agreement terms.

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

agreements. Therefore, misclassifying spot suppliers as agreement suppliers is much more problematic than the opposite, further motivating the likely conservative classification standard of transactions every month.

This interpretation of the stylized facts raises two 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? Second, if the buyer only receives discounts from agreement suppliers, why not source exclusively from them? Online Appendices OA1.4 and OA1.5 address these questions.

## 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. (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 scale), contrasting with the characterization in Bloom et al. (2012) of large Indian garment manufacturers as messy and disorganized.

## 2.6 Context: Overall Sourcing

Figure 6 shows the buyer's overall volumes by quarter across all three supplier types: agreement; spot; and in-house. 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 fabric procurement costs) under $60,000, further emphasizing the buyer's large size.[37]

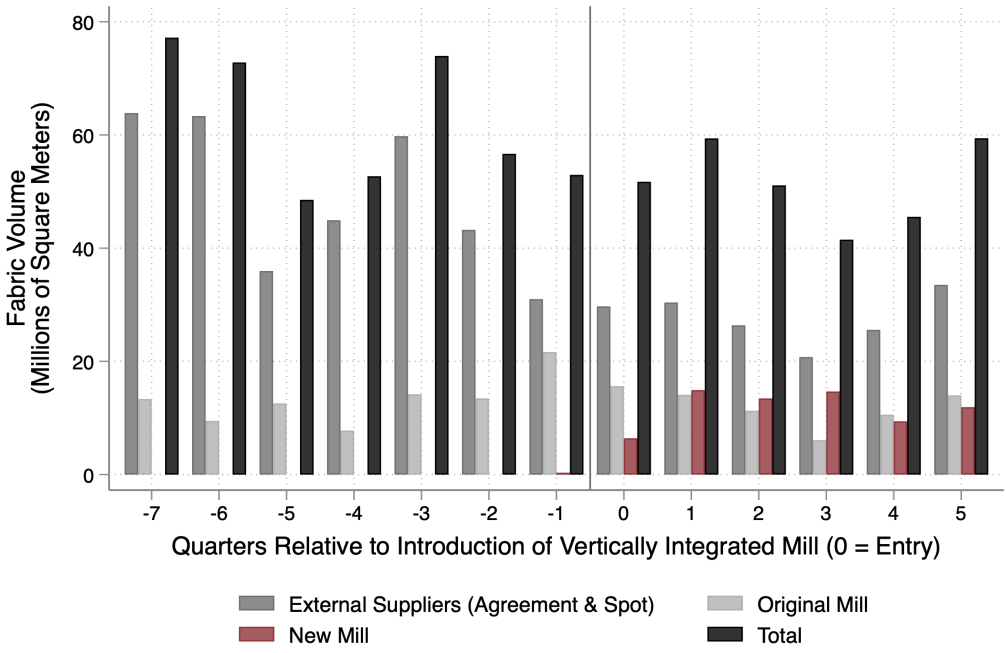



Figure 6: Fabric Sourcing Over Time

**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 make this setting particularly well-suited to analyze the price effects of insurance-like agreement terms and replacement threats from in-house supply. First, external suppliers remain empirically central post-integration, still providing most fabric. Therefore, the buyer has strong incentives to consider external supplier prices throughout. Second, the increase in vertical integration is sufficiently large to make replacement threats credible: the new mill produces ∼20% of the buyer's fabric, compared to only ∼7% from the largest external supplier. Third, end-client demand is stable, decreasing slightly from 250 to 215 million square meters per year. As a result, the buyer must strategically reduce the number of external suppliers, the average volume per external supplier, or both.[38]

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

[38] Increased integration is not likely a response to a supply-side change in production technology as a single cost function fits pooled data from both the old and new integrated suppliers well (see section 4.1.2).

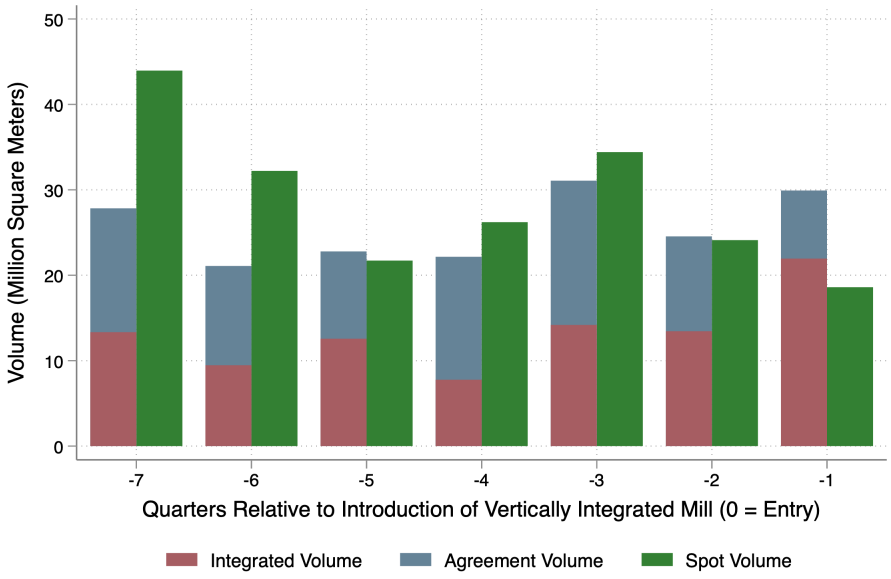



Figure 7: Pre-Integration: Fabric Sourcing by Supplier Type

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

Figure 7 shows that the buyer's overall sourcing strategy relies on integrated and agreement supply for stable quantities, while spot supply absorbs aggregate demand variability. Agreement and integrated suppliers receive stable volume (coefficient of variation for their aggregate volume is .152), while spot supply has much higher volatility (coefficient of variation is .324). Consistent with this result, spot volumes strongly correlate with total buyer demand ($R^2 = 0.83$), while agreement and integrated volumes show no meaningful response to total demand ($R^2 = 0.03$). These results suggest that the buyer uses agreement and integrated suppliers for *inframarginal* sourcing, while spot suppliers fulfill *marginal* orders.[39]

Two central institutional features rationalize the use of integrated and agreement supply for *inframarginal* quantities. First, increasing marginal costs of internal production create incentives to maintain stable in-house volumes, rather than increasing with overall demand. Empirically, integrated supply rarely exceeds 100% capacity (see Figure 10). Second, agreements specify stable quantities, preventing the buyer from adjusting agreement volumes to offset aggregate demand variability. (Relational contracts generally require time to establish trust, limiting scope to add new agreement suppliers (Macchiavello and Morjaria, 2015)).

[39]Online Appendix OA1.11 shows that these results are robust to evaluating the post-integration period instead, noting that post-integration quarters are analyzed separately to avoid conflating changes in sourcing strategy with changes in capacity by supplier type.

# 3 Theoretical Bargaining Model

This section develops a theoretical model that formalizes how prices reflect both *i)* surplus creation from insurance-like agreement terms and *ii)* surplus capture from threats to replace suppliers with in-house production. The model extends standard replacement-threat results by analyzing threats from in-house, rather than external, suppliers. It also reveals a new result: replacement threats *only* reduce prices when suppliers have risk aversion above a cutoff. In agreements with suppliers less risk averse than the cutoff, the buyer endogenously refrains from making replacements threats because such threats would break the agreement.

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. U-shaped prices emerge from the buyer and supplier bargaining over a pricing *function* (rather than a vector of unit prices, as common in Nash bargaining) that maps the realized quantity to a price. The model embeds quantity assurance through suppliers drawing quantities from different distributions in the agreement versus the spot market. Because the agreement provides *stable* but not necessarily *larger* quantities, both quantity distributions have the same mean quantities but agreement variance is smaller. The buyer's replacement threats appear through its outside option, with integration enabling the buyer to use in-house production as its outside option instead of the spot market.

While the model is static to transparently isolate surplus creation from surplus capture, I verify that incorporating dynamic incentive compatibility constraints (as in Macchiavello and Morjaria (2023)) yields qualitatively similar results for reasonable discount factors (Appendix A.3.1). The model represents bargaining over long-run surplus, with relational contracts enforcing the cooperative outcome modeled.

## 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 between the cost of purchasing quantity $q^A$ at price $p^A(q^A)$ in the agreement and the outside option value ($O_B$). I model the buyer as risk neutral given

its large size and sophistication.[40] Buyer surplus reflects only procurement costs, as revenue from end clients is determined before fabric procurement.

Supplier surplus is the utility of profit in the agreement less the value of the outside option ($O_S$). I model suppliers preferring smooth profits through the utility function, which presumably reflects that firms may need to make some payments every period (*e.g.,* payroll). Financial market frictions can make borrowing or saving difficult and costly, leading firms to use current cash inflows for these expenses.[41] For brevity, I refer to such preferences as risk aversion and use CARA utility.[42]

Insurance-like agreement terms generate surplus, thereby shaping agreement prices, through two channels. First, the U-shaped pricing function smooths profits, increasing supplier utility even for the same average profit level by reducing variability. Second, quantity assurance generates additional surplus beyond insurance by reducing average costs, which benefits even risk-neutral suppliers. With convex costs,[43] Jensen's inequality implies that reducing quantity variance lowers average costs: high costs at large quantities are not correspondingly offset by lower costs at small quantities. Cost convexity is documented in Figure 10.

Bargaining parameters determine surplus shares. The buyer bargaining parameter is $\alpha_B$, with higher values indicating that the buyer receives a larger share of surplus created by the agreement. As the buyer and supplier split the surplus, their bargaining parameters sum to one (*i.e.,* the supplier bargaining parameter is $1 - \alpha_B$).

The outside option for both the buyer (pre-integration) and supplier is transacting in the spot market. Because agreements are relational contracts, which require time to establish, neither the buyer nor supplier can switch to a new agreement partner 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 $p^S$, consistent with the

[40] 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.

[41] 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.

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

[43] 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 given frictions to hire and fire workers that prevent easily adjusting worker numbers.

flat spot prices in Figure 4.[44] Flat spot prices reflect supplier power from hold-up: suppliers extract surplus. Therefore, spot prices depend on the buyer's outside option, which does not vary with supplier capacity. Furthermore, even though the buyer is a large firm, the buyer is small relative to the Indian fabric market, with a market share under 0.5%.[45]

After integration, the buyer can use in-house supply as its outside option. I assume the integrated supplier uses the same production (*i.e.,* has the same cost function) as external suppliers, consistent with the buyer reporting that all relevant suppliers employ (near-)frontier technology. Morton (2023) empirically documents strong substitutability between integrated and agreement suppliers, making the integrated mill a credible threat to displace agreement suppliers given the meaningful, visible sunk cost of construction. Results are robust to cost heterogeneity between internal and external suppliers (Appendix OA1.5.2).[46]

Quantities reflect stochastic end-client demand rather than endogenous buyer or supplier choices. Agreements specify the *ex-ante distribution* of quantities, with quantity assurance leading to lower variance (but not affecting average volumes). Formally, I model spot quantities as a mean-preserving spread of agreement quantities. I set mean quantities for both to 100% capacity (which is the supplier's optimal capacity utilization given costs and prices). Quantities can be interpreted relative to capacity the supplier reserves for the buyer. Therefore, the importance of this agreement can affect supplier risk aversion (*i.e.,* suppliers with diversified buyers might be less risk averse over profits from any one buyer.).[47]

Formally, the outside options are:

[44]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.

[45]The buyer uses ∼250 million square meters of fabric per year, compared to 71 billion square meters of production in India in 2019. India is the sixth largest producer globally (India Brand Equity Foundation).

[46]Differences in production technology shift costs for both inside and outside options similarly, so payoffs change little for reasonable heterogeneity levels. Large cost heterogeneity seems implausible: if the integrated supplier were much more efficient, the buyer should fully integrate; if external suppliers were much more efficient, the buyer should not integrate.

[47]In principle, working with other buyers could affect cost convexity if suppliers balance demand shocks across buyers. However, seasonality limits this smoothing potential as demand shocks are correlated across buyers. Furthermore, results are robust to reasonable cost heterogeneity (see Online Appendix OA1.5.2).

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

The bargaining model incorporates two-sided market power. Market power depends on bargaining positions, which reflect both surplus shares (*i.e.,* the bargaining parameters) and outside options. Critically, I interpret these sources of market power differently: the buyer bargaining parameter represents the buyer's share of surplus created, aligning incentives for insurance provision. Conversely, replacement threats only extract rents, capturing surplus. The model also incorporates supplier market power in the spot market through hold-up.

Appendix A.4 details the model parameterization and solution. Parameters and functional forms are chosen to illustrate how surplus creation and capture affect prices (rather than to match any empirical moments). Section 4 builds and estimates a structural model closely tied to this framework for empirical quantification.

## 3.2 The Price Effects of Surplus Creation & Capture

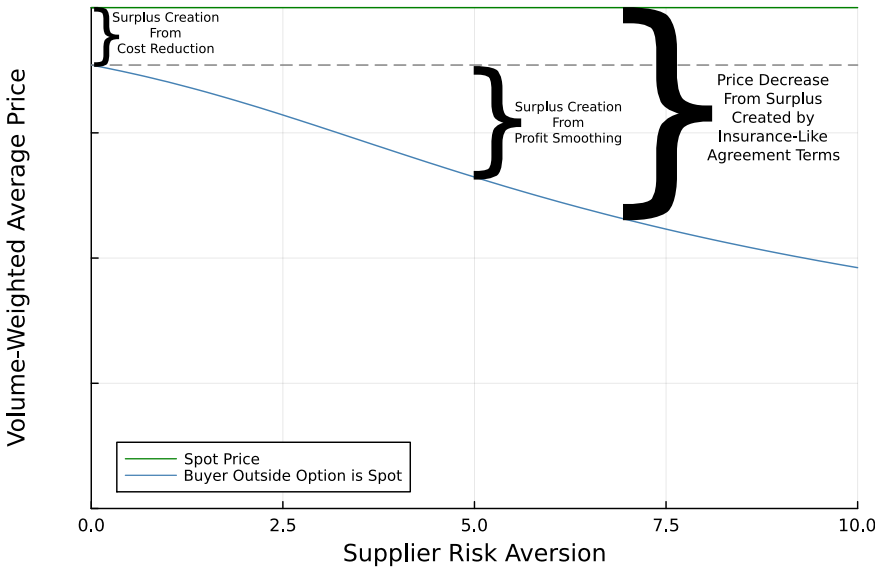



Figure 8: Surplus Creation From Insurance-Like Agreement Terms

**Note**: Costs represent volume-weighted average procurement costs, with prices from the parameterization in Appendix A.4.

Figure 8 illustrates how supplier risk aversion affects the buyer's expected price (*i.e.,* the volume-weighted average fabric price), holding the buyer bargaining parameter constant.[48] When the buyer purchases from the spot market, no bargaining occurs, so prices do not vary with supplier risk aversion. Instead, when the buyer's outside option is the spot market, surplus creation in the agreement generates discounts relative to spot prices.[49] Then, some surplus created in the agreement accrues to the buyer as a discount. The agreement discount with a risk-neutral supplier (*i.e.,* when risk aversion is 0) reflects surplus generated only by cost reductions from quantity assurance. As supplier risk aversion increases, the agreement creates additional surplus by smoothing profits through both insurance-like terms: *i)* the quantity assurance and *ii)* the U-shaped pricing function. Consequently, the discount grows.

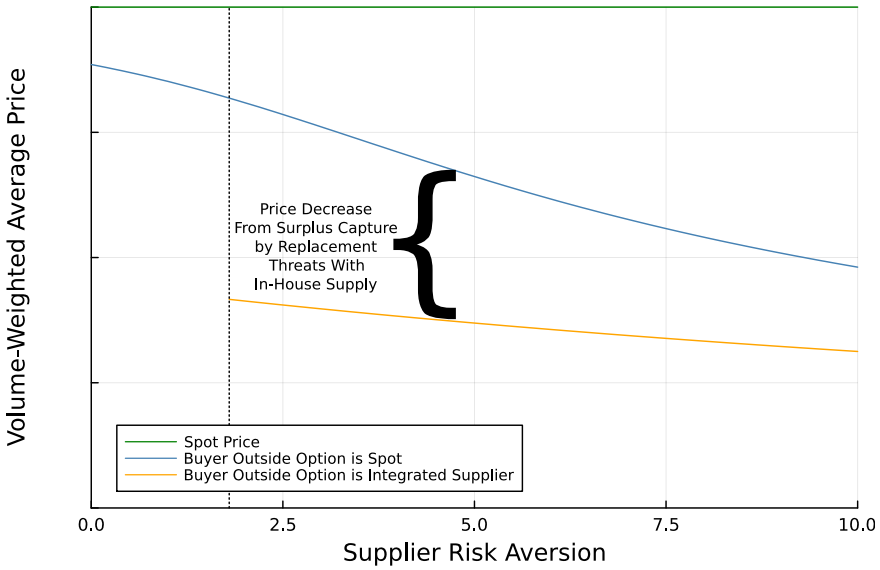



Figure 9: Surplus Capture From Replacement Threats

**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 surplus capture from the buyer threatening to replace suppliers with in-house production—when the buyer's outside option switches to its integrated supplier. The contracting environment constrains the benefits of replacement threats: there is a cutoff level of risk aversion below which the agreement breaks down in

[48] Online Appendix OA1.7 shows similar results when varying the bargaining parameter and holding supplier risk aversion constant. Although both influence prices, only risk aversion affects surplus creation from insurance-like agreement terms and capture from strategic replacement threats.

[49] Discounts emerge for any positive buyer bargaining parameter—when the buyer has some power over surplus creation.

response to replacement threats. This breakdown occurs because suppliers with low risk aversion receive insufficient surplus from insurance-like agreement terms, so no pricing function generates positive surplus for both parties.

When supplier risk aversion falls below this cutoff, the buyer would prefer to retain its pre-integration agreement because it includes a discount. However, in Nash bargaining, outside options are not choices but primitives that represent the value of the disagreement outcome. But in practice, parties may have incentives to strategically present specific outside options , especially when the credibility of the outside option is private information of one party. To formalize 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 the cooperative Nash bargain and a non-cooperative phase (Miller and Watson, 2013).

The first extension introduces a non-cooperative signaling game about residual capacity at the in-house supplier before cooperative bargaining occurs, consistent with signaling game models of incomplete information bargaining (*e.g.,* Backus et al. (2019)). Residual capacity determines whether the in-house supplier is the buyer's true outside option. It is also the buyer's private information, as only the buyer knows other orders at the in-house supplier, and varies stochastically with end-client demand. The signaling game incorporates beliefs about residual capacity in a principled way. (While Nash bargaining is axiomatic with perfect information, Lee et al. (2021) highlight that it has similarities with passive beliefs). 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 where the buyer commits to a level of residual capacity before bargaining. I similarly show 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

This section builds and estimates a structural model closely tied to the theoretical model to decompose agreement discounts, relative to spot prices, into surplus creation from insurance-like terms and surplus capture from replacement threats. To do so, I estimate the supplier risk aversion and buyer bargaining parameters that rationalize observed prices, allowing me to measure how much of the observed discount is from surplus creation, both from profit smoothing and cost reductions, and from surplus capture. I then use the model to examine how the decomposition of prices varies across supplier types, focusing on small suppliers, and evaluate policy counterfactuals to increase prices that small suppliers receive.

## 4.1 Estimation Approach

### 4.1.1 Buyer and Supplier Primitives

The structural model features three parameter groups: supplier-specific risk-aversion parameters; agreement-specific buyer-bargaining parameters; and a common tuning parameter. The tuning parameter penalizes deviations from risk neutrality, motivated by both economic intuition, as firms should not deviate extremely from risk neutrality, and econometric caution against overfitting given limited observations per supplier.[50]

Estimation proceeds in two steps. First, an "inner" optimization recovers supplier-specific risk-aversion and buyer-bargaining parameters that best fit prices, while the outer optimization selects a common tuning parameter using "out-of-sample" post-integration data.

The "inner" optimization has three steps. First, for each agreement, I derive the model-implied agreement pricing function $\hat{p}^A(q^A; \mathbf{\Omega}_s)$ that maximizes the Nash product in (1), where $\mathbf{\Omega}_s$ represents the parameter vector for each relationship. The agreement pricing function depends on these parameters because their values affect the Nash product and, consequently, the optimal pricing function.[51] Second, I estimate prices for observed agreement quantities using the model-implied pricing function. Last, I compute the objective function for each supplier—the root-mean-squared error (RMSE) between estimated and observed prices in pre-integration quarters. The objective function adds a LASSO-style penalty—the abso-

[50] Because the analysis focuses on long-run agreements, misclassifying a spot supplier as an agreement supplier would bias estimated risk aversion downward and understate surplus creation (given flat spot prices). This concern motivates my conservative classification approach. Regardless, prior analysis shows that the supplier classification performs well.

[51] I approximate $\hat{p}^A(q^A; \mathbf{\Omega}_s)$ with a second-order Taylor polynomial around full-capacity: $\hat{p}^A(q^A; \mathbf{\Omega}_s) = \beta_{0,s} + \beta_{1,s} q^A + \beta_{2,s}(q^A)^2$.

lute value of the risk-aversion parameter weighted by the tuning parameter—to discourage extreme deviations from risk neutrality (zero in CARA utility).

Intuitively, the "inner" optimization finds supplier-specific parameters that best match pre-integration prices. Formally, it solves the following objective function:

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

$S$ denotes the set of suppliers, $\{q^A_{s,t}\}$ are the observed quantities for agreement supplier $s$,[52] $T$ is the number of periods, $\hat{p}^A_s(q^A_{s,t};\mathbf{\Omega}_s)$ is the price from the agreement pricing function for supplier $s$ for quantity $q^A_{s,t}$, $p^A_s(q^A_{s,t})$ is the realized agreement price for supplier $s$ for quantity $q^A_{s,t}$, and $\theta_s$ is the supplier risk aversion parameter.

The "outer" optimization finds the tuning parameter that maximizes "out-of-sample" fit in the post-integration quarters, mitigating potential overfitting. The tuning parameter affects post-integration fit through the pricing function: a larger value pushes supplier risk aversion estimates toward risk neutrality. Formally, the "outer" optimization solves:

$$\min_{\lambda} \underbrace{\sum_{s\in S}}_{\substack{\text{Sum over}\\ \text{suppliers}}} \left[ \overbrace{\sum_{\{q^A_{s,t}\}}}^{\substack{\text{Sum over}\\ \text{post-integration}\\ \text{quarters}}} \underbrace{\sqrt{\frac{1}{T}\left(\hat{p}^A_s(q^A_{s,t};\mathbf{\Omega}_s) - p^A_{s,t}\right)^2}}_{\text{RMSE: Model \& Actual Prices}} \right] \quad (3)$$

### 4.1.2 Outside Estimation

Table 1 summarizes the additional moments for model estimation.

[52] While using monthly rather than quarterly data would increase the effective sample size, agreements empirically 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.

| 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 |

Table 1: Additional Moments for Estimating the Structural Model

I estimate the cost function using panel data on capacity and costs from both integrated suppliers. As the buyer's procurement staff indicate that agreement and integrated suppliers use similar technology, their cost functions should be comparable. I estimate the total cost function as $c(q) = \phi_0 + \phi_1 q + \phi_2 q^2$, where $\phi$ represents convexity. I do not impose convexity (*i.e.,* $\phi_2 > 0$), and estimates are 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 corresponds to a broad fabric group (*e.g.,* knits and striped fabrics).[53] Identification of the cost function leverages exogenous capacity shocks that that arise because sampling fixes the supplier before final order quantities are realized; thus, realized quantities reflect stochastic end-client demand rather than buyer choice.[54] As deviations for the quantity target are plausibly exogenous, OLS regressions provide causal estimates of the effect of capacity on costs.

Estimates demonstrate that the model fits the data well, shown in Figure 10. The $R^2$ is .795 and the RMSE is .07, which is small relative to overall cost variation. Additionally, the estimated cost function is convex, with the coefficient on $q^2$ at .7815 (p-value = .02).[55]

[53]Cost function estimations leverage the most temporally disaggregated data as aggregation can mask variation. For example, aggregated data would imply the same capacity utilization when a supplier operates at 100% for two months versus at 80% one month and at 120% the next.

[54]The buyer likely has some knowledge of the distribution of likely volumes from the end clients based on previous transactions. Although the buyer likely uses this knowledge when sampling to target capacity utilization at the integrated suppliers, the buyer cannot perfectly remove variation due to the lumpy, stochastic nature of orders. The data are consistent, as there is meaningful heterogeneity in capacity utilization.

[55]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 suggests that heterogeneity by fabric type is unlikely to be a central feature of the cost function.

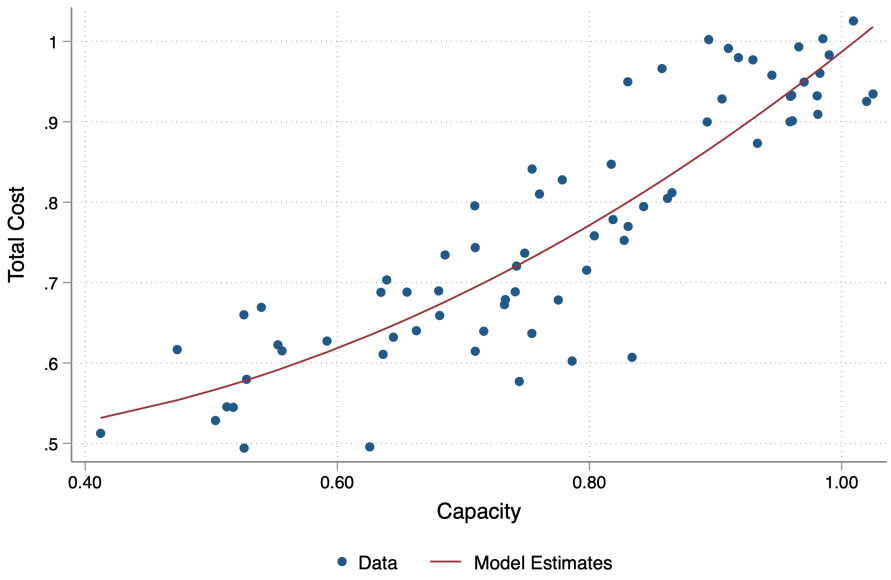



Figure 10: Cost Convexity at the Integrated Supplier

**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.

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

## 4.2 Identification of Risk Aversion & Bargaining Parameters

Average agreement prices alone cannot separately identify supplier risk aversion from the buyer bargaining parameter as increasing either parameter reduces prices. Higher supplier risk aversion increases surplus created by the agreement; since the buyer receives a share of that surplus, higher risk aversion lowers prices. Similarly, a higher buyer bargaining parameter increases the buyer's share of surplus, reducing prices.

I develop a novel identification strategy that leverages two distinct sources of variation in the data. First, the covariance of prices and capacity utilization identifies supplier risk aversion because greater risk aversion makes the pricing function more *convex*—more U-shaped contracts are associated with higher risk aversion because the U-shape smoothes profit. Second, variation in average agreement prices identifies the buyer bargaining parameter, as increasing the bargaining parameter reduces prices across capacity realizations.

With only one price-capacity tuple, it is impossible to separate supplier risk aversion from the buyer bargaining parameter. For any buyer bargaining parameter, there is a risk

aversion parameter that can rationalize any 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 with a buyer bargaining parameter of 0.72 result in the same price when capacity utilization is 1.

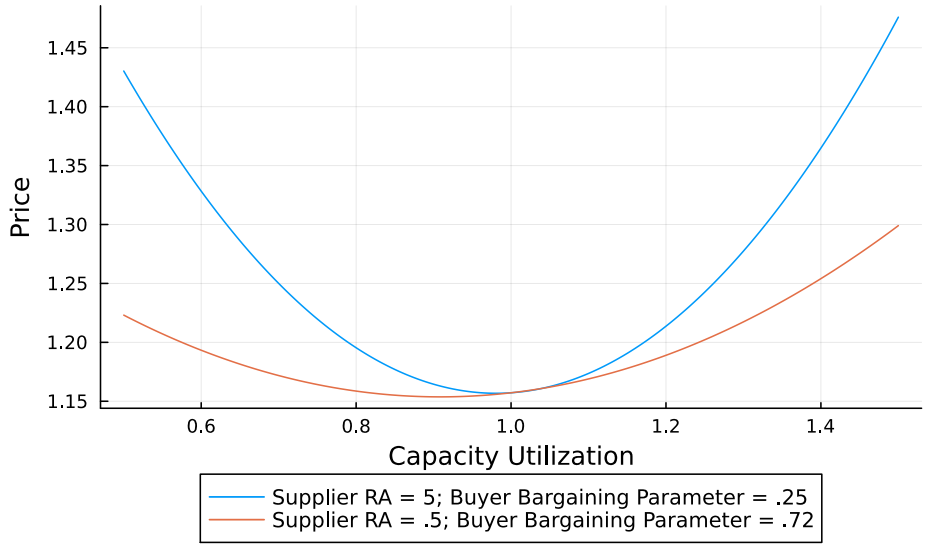



Figure 11: Convex Pricing Function for Different Parameters

However, a *second* price-capacity tuple (for a different capacity) enables identification. Lower risk aversion yields a flatter, less U-shaped contract, while a higher bargaining parameter results shifts prices downwards.[56] Figures 12 and 13 illustrate this identification: shifting the risk aversion primarily affects the shape of the pricing function, whereas shifting the buyer bargaining parameter mainly changes the average price level.

[56] A third price-capacity tuple identifies asymmetry in the U-shaped function that enables prices to smooth profit more effectively.

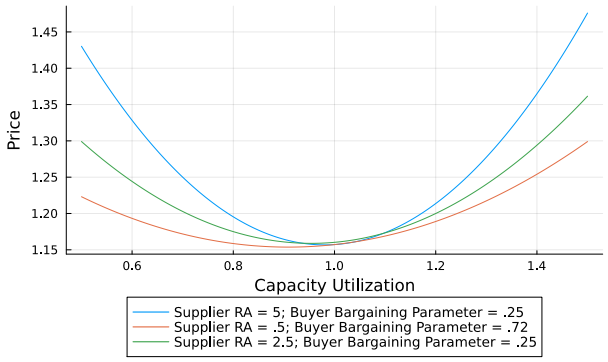



Figure 12: Shift Supplier Risk Aversion

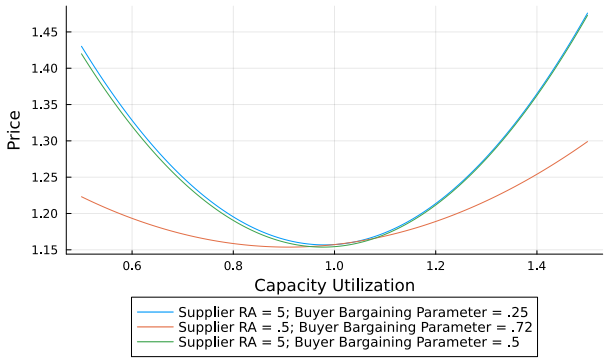



Figure 13: Shift Buyer Bargaining Parameter

To bring this identification argument to the data, I leverage quasi-exogenous shocks in capacity utilization. Garment production—where neither which samples will be ordered nor their quantities are known at the time of sampling—limits quantity assurance in the agreement to only reduce, and not eliminate, quantity variance. I interpret deviations from targeted agreement volumes as exogenous because neither the buyer nor the supplier choose which garments are ordered (nor their quantities); instead, orders reflect demand at end clients. Appendix A.8.1 shows results inconsistent with the main threat to exogeneity: that agreements include 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 for agreement suppliers resembles LMIC coffee mills in Blouin and Macchiavello (2019) (.0085 for the median agreement supplier and .0068 for coffee mills).[57] These coffee mills, described as "large firms by developing country standards," have ∼$3.5 million in annual sales and ∼$2 million in total assets. Given similarities in (likely) 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. Although these results suggest that agreement suppliers are large firms (especially by LMIC standards), the buyer studied is much larger: it purchases

[57]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.

∼$180 million yearly on fabric, and such expenses must be smaller than revenue .[58]

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 |

**Note**: Estimates from structural model at the supplier-level. The correlation between supplier risk aversion and buyer bargaining parameters is 0.28.

Estimated bargaining parameters imply that the buyer typically receives slightly more than half of surplus created by the agreement.[59] This result aligns with the aversion estimates as both suggest that agreements involve relatively shopshisticated, large suppliers (with the exception of one highly risk-averse supplier).[60] The large size of agreement suppliers is consistent with the necessity of sophisticated firm capabilitites to maintain complex contractual agreements (Macchiavello and Morjaria, 2023).

Because estimation includes outside moments, I prioritize model validation in Online Appendix OA1.8. First, the model matches an untargeted moment. Second, the model outperforms two alternatives in post-integration quarters; this comparison serves as an out-of-sample test because estimation primarily relies on pre-integration quarters.[61] (As the DiD directly estimates the price effects of surplus capture from replacement threats with in-house production, I present the DiD validation after the analagous model estimate.)

## 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. To do so, I first determine whether each supplier's risk aversion surpasses the cutoff where the buyer uses

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

[59]Section A.13 provides evidence of one microfoundation for heterogeneity in buyer bargaining parameters.

[60]One quarter with high leverage drives the outlier results.

[61]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.

replacement threats. Formally, I test whether any pricing function yields weakly positive surplus for both parties when the buyer's outside option is the integrated supplier. 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 the additional discount from surplus captured through replacement threats.

Estimates illustrate that agreement discounts are entirely attributable to buyer power over surplus creation, with no surplus capture from replacement threats. This result reflects the low, although still present, risk aversion of agreement suppliers. Agreement prices, on average, fall 1.5% below spot prices, which is a meaningful reduction given the nearly 7% markup in the spot market. These savings decrease the buyer's annual fabric procurement costs by ∼$0.87 million. I further decompose agreement discounts from surplus creation into cost reductions from quantity assurance and profit smoothing (*i.e.,* insurance). To isolate the role of cost reductions alone, I compute discounts for 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 account for roughly half of the total discount.

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% |

**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.

To assess robustness to the one outlier supplier with high risk aversion, I compute discounts with the *i)* median and *ii)* mean values of the supplier risk aversion and buyer

bargaining parameters.[62] The main finding is robust: all discounts arise entirely from surplus creation. Because all suppliers fall below the risk-aversion cutoff at which the buyer captures surplus through replacement threats, both the median and mean parameters also lie below that threshold. However, the level and composition of the discount differ. Discounts decrease overall because parameters affect prices non-linearly: doubling both the risk aversion and bargaining parameter more than doubles the discount.[63] This change reflects that decreasing the influence of one highly risk-averse supplier reduces the relative contribution of profit smoothing (insurance) to the discount.[64] Discounts are slightly higher with the median rather than mean parameter values; the larger median than mean buyer bargaining parameter more than offsets the lower median than mean risk aversion.

## 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.[65] The first difference is time: before versus after the buyer adds in-house capacity by building a new fabric mill. The second difference is exposure to the new vertical integration, defined by pre-integration fabric overlap with the integrated supplier: suppliers with overlap are "treated" (as the buyer can credibly threaten to switch to in-house production) while those without overlap are "controls." Importantly, all comparisons are within supplier type, so agreement suppliers are only compared to each other. I estimate a DiD with a propensity-score adjustment (Abadie, 2005) that relaxes parallel trends to *conditional* parallel trends; in my setting, comparisons condition on supplier exposure to fabric-demand trends that reflect evolving fashion tastes (detailed in Appendix $A.11$). The adjustment attains pre-treatment balance and, as outcomes are *not* matched pre-integration, pre-trends remain informative about parallel trends. Because neither difference identifies structural parameters, the validation exercise compares the model-implied replacement-threat effect with the DiD estimate.

The model generates sharp predictions for the effects of vertical integration on both prices

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

[63] This non-linearity has multiple sources. Two central reasons are: *i)* the Nash product is multiplicative, not additive; and *ii)* higher parameter values increase both the surplus size and the buyer's share.

[64] 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).

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

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.

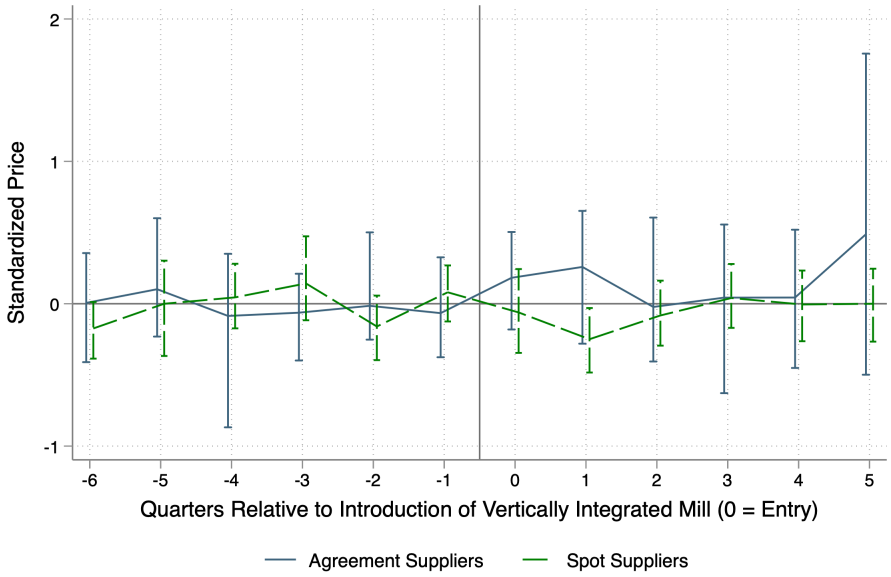



Figure 14: DiD Estimates: Price Effects

**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 formally in the theoretical model, the buyer sourcing *marginal* quantities from spot supply suggests that vertical integration should decrease spot rather than agreement volumes.[66] Appendix Figure A.9 confirms that agreement quantities do not change, but spot quantities do. Furthermore, the significant decline in spot quantities suggest that the DiD design has power to statsically detect economically meaningful effects.

I also examine average treatment effects pooled across periods in Appendix Table A.7.[67] Pooled results confirm that agreement suppliers do not experience any statistically significant effect on quantities or prices. (Point estimates are near zero, although not always precise,)

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

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

However, vertical integration causes a statistically significant and economically meaningful decrease in quantities for spot suppliers. The decline averages .4 standard deviations (∼120,000 square meters of fabric per quarter on average).[68] In Online Appendix OA1.10.2, I show robustness to the synthetic DiD approach in Arkhangelsky et al. (2021).

Online Appendix OA1.10.3 provides an additional validation exercise that compares DiD event study coefficents using actual prices with DiD estimates using model-implied prices for actual quantities. 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 DiD Estimates Are Inconsistent with NNTR

I compare the DiD with prices from a standard model of replacement threats: the Nash-in-Nash with Threat of Replacement (NNTR) model from Ho and Lee (2019). Ho and Lee (2019) study a bargaining environment—agreements between employers and health insurers on California—with neither a spot market nor integrated supply. Therefore, NNTR does not incorporate marginal sourcing from spot suppliers, 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 general critique of NNTR).

Figure 15 illustrates that price decreases predicted by the NNTR model are inconsistent with the DiD estimates of price changes from replacement threats.[69]

[68] 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).

[69] 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.[70] 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.

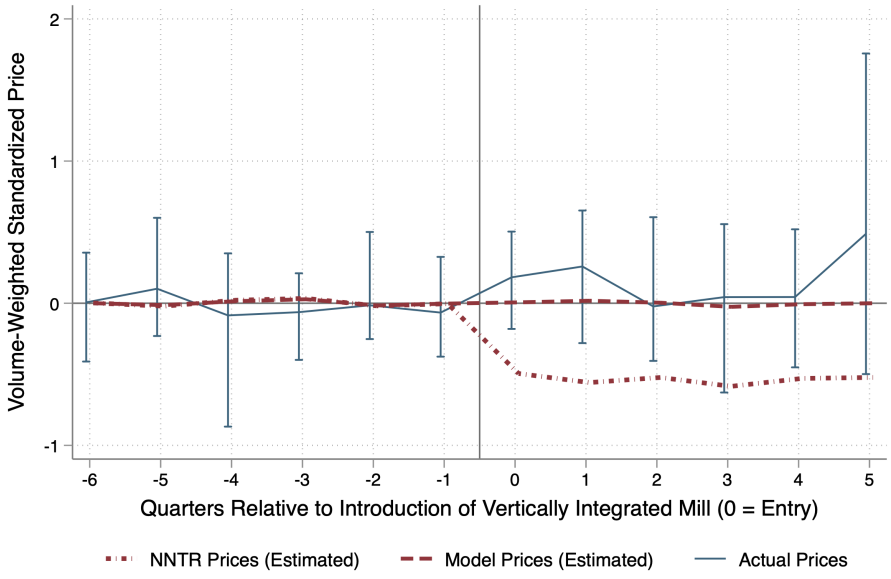



Figure 15: Comparison with NNTR Model

**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-Implied Prices for Small Suppliers

To study how suplus creation and capture affect prices for policy-relevant small suppliers—smaller than observed agreement suppliers—I simulate agreement prices using the model across a range of risk aversion and buyer bargaining parameters.

Figure 16 reports how agreement discounts, as a percent of spot prices, vary with supplier risk aversion and the buyer bargaining parameter. While discounts increase in both, the price change from higher risk aversion is modest. 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. For all risk aversion parameters shown, discounts reflect surplus creation, rather than surplus capture; the buyer always uses the spot market as its outside option with risk aversion below the cutoff. Because discounts do not increase much in risk aversion, cost reductions—rather than profit smoothing—drive price changes.

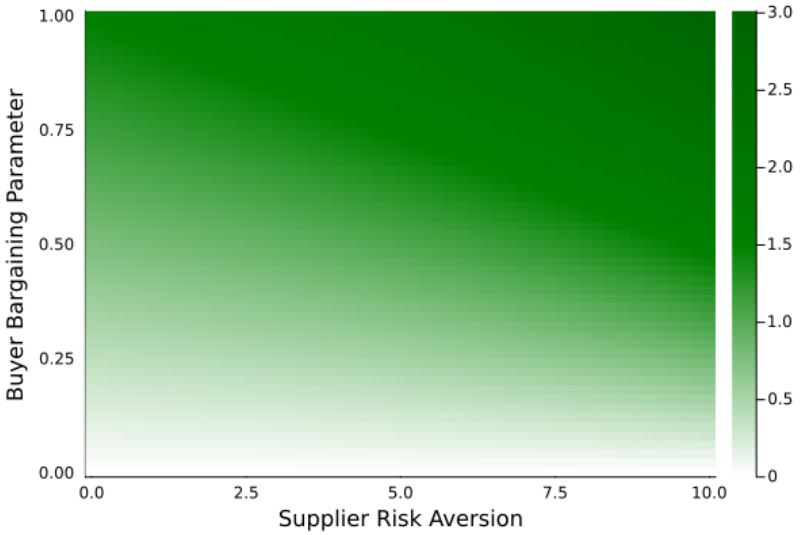



Figure 16: Agreement Discounts (% of market price)

**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 shows discounts for suppliers sufficiently risk averse for the buyer to capture surplus through replacement threats with in-house production. The cutoff level of risk aversion for replacement threats is 16.6, slightly exceeding estimates from other settings (12.9 in Gandelman and Hernandez-Murillo (2014)). This finding emphasizes how the contracting environment constrains replacement threats. However, as the risk-aversion threshold is not a structural parameter but depends upon them, the relevance of replacement threats should vary,. The cutoff level of risk aversion decreases when cost functions are more convex or agreements provide more insurance. Agreements can provide more insurance when they attenuate risk in both demand and supply; for example, agricultural workers facing productivity risk can face both supply-side and demand-side risk (Jayachandran, 2006).[71] LMIC smallholder farmers may satisfy these conditions and rely on bargained agreements (Macchiavello and Morjaria, 2020); they also could be risk averse and have highly convex production costs as a key input is fixed (land).

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

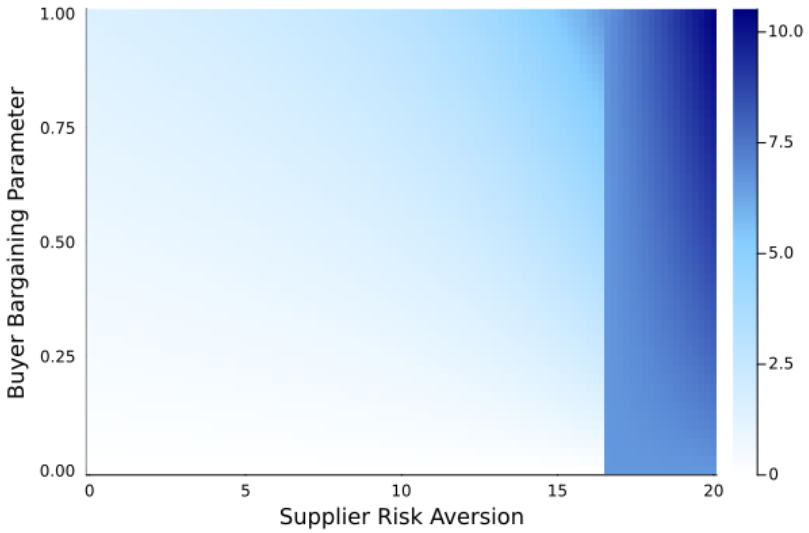



Figure 17: Agreement Discounts (% of market price)

**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 cutoff 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 agreement discounts average 1.5% of spot prices. The 6.7% discount exceeds 5% Indian VAT for apparel during the study period (5%)[72] and is large relative to markups in the spot market from hold-up (7.36%).

Online Appendix OA1.11 incorporates relational enforcement for a conservative calibration of the discount factor, documenting negligible differences in estimated prices.[73] 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 con-

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

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

sistent with low estimated risk aversion: both suggest agreement suppliers are large firms.[74]

# 5 Counterfactual Policy Evaluation

I use the estimated model to analyze counterfactual policies to increase prices paid to small suppliers. I focus on small suppliers becuase of their empirical importance as the majority of LMIC firms and employers (Hsieh and Olken, 2014), as well as distributional motiviations related to policy maker concerns that the gains from exporting do not reach up the supply chain. Consistent with this focus on small firms, ominpresent policy initiatives target small and medium enterprises (SMEs)[75] Furthermore, discounts from surplus creation and capture increase in risk aversion, and small suppliers are likely the most risk averse.

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 the importance of missing markets in LMIC settings, I create the missing insurance product so risk-averse firms can smooth profit outside the agreement.

## 5.1 Increase Downstream Buyer Competition

I analyze how increased downstream buyer competition affects prices. 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 through reducing the buyer bargaining parameter.[76] While competition from additional buyers should improve suppliers' outside options in the long-run, competition does not change suppliers' short-run outside option (the spot market) because suppliers cannot quickly establish a new agreement with relational

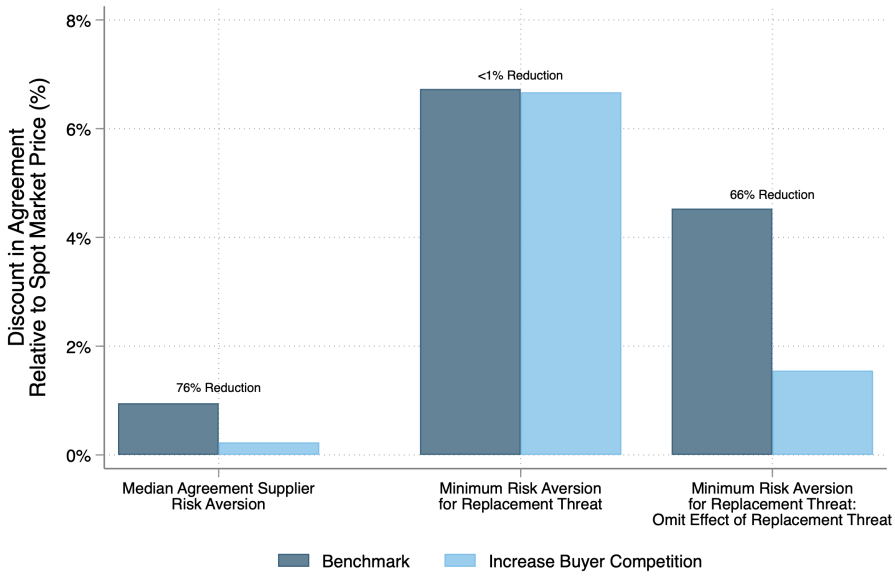
[74]The model features a cutoff level of the discount factor above which dynamic incentive compatibility constraints are slack in this model. Higher discount factors above this cutoff do not affect prices. When dynamic incentive compatibility constraints are not slack, they "flatten" the agreement pricing 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.

[75]Examples include the World Bank's policy advising on finance tools for SMEs and the IMF's 2019 conference on SME financing.

[76]I halve the buyer bargaining parameter.

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 guarantee 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: more valuable suppliers receive larger surplus shares (*i.e.,* lower estimated buyer bargaining parameters).[77]

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.

[77] Agreement suppliers producing fabrics with higher concentration, compared to other agreement suppliers, are likely more valuable to the buyer becauser supplier-fabric concentration affects the ease of providing quantity assurance. To illustrate, consider an extreme case: all agreement suppliers sell the same fabric. Then, when aggregate quantity for that fabric is low, the buyer cannot provide quantity assurance to all suppliers. Therefore, providing quantity assurance is relatively easier to suppliers in relatively concentrated fabric markets (relative to other agreement suppliers). Appendix Section A.13 describes the empirical 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 illustrates that increased downstream buyer competition meaningfully reduces the discount (as a percent of spot price) from a supplier as risk-averse as the median agreement supplier: discounts decrease by 76%, from .95% to 0.23%. However, for suppliers with risk aversion above the cutoff where the buyer uses 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 it, increased competition reduces 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.[78]

## 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 spot market profits affects agreement discounts. This insurance product mitigates how supplier risk aversion increases discounts by eliminating profit risk in the spot market.[79]

Although governments are unlikely to offer such insurance given moral hazard concerns, policymakers have many tools to provide limited insurance.[80] 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

[78]Online Appendix OA1.12 evaluates horizontal supplier mergers. Although supplier consolidation both reduces surplus created by the agreement, as the merged entity is less risk averse, and restrains surplus capture from replacement threats, such conslidation is inconsistent with the importance of small firms in LMIC settings and associated distributional concerns.

[79]Blouin and Macchiavello (2019) also explore how missing insurance markets interact with contracting behaviors, highlighting how suppliers' incentives for stategic default lead 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 distributional implications.

[80]Policies that provide some insurance are ubiquitous. For example, common price floors for agricultural goods insure producers against adverse prices.

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.

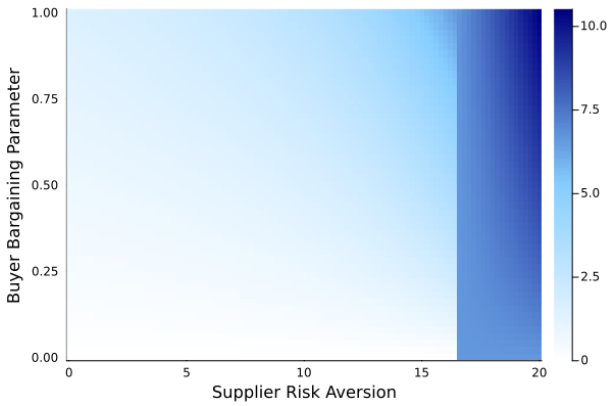



Figure 19: Without Insurance

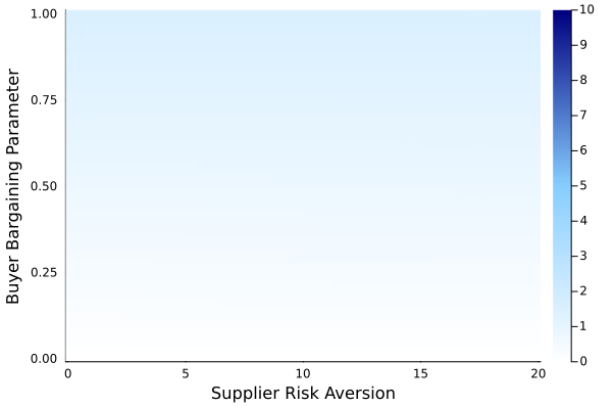



Figure 20: With Insurance

**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 highlights how low prices in vertical relationships can arise from two distinct forces with different welfare implications: surplus creation inside the agreement from mitigating market frictions and surplus capture via replacement threats from in-house supply. In the context of the Indian garment supply chain, I show stylized facts that agreements

include non-linear prices and quantity assurance. I interpret these facts as indicating that agreements smooth profits for suppliers. A theoretical model formalizes intuition about the price effects of surplus creation and capture: both make (average) prices decrease in supplier risk aversion, with a cutoff level in risk aversion above which capture affects prices. Estimates from a structural model, identified by plausibly exogenous demand shocks, show that all observed agreement suppliers have risk aversion below this cutoff, so the entire agreement discount ($\sim$1.5% of spot prices) is due to surplus creation. Becuase of buyer power over surplus creation, the buyer receives roughly 60% of the gains from insurance-like terms. Overall, the decomposition emphasizes how low prices can reflect the buyer's share of gains from addressing market frictions within the vertical relationship rather than pure rent extraction. Correspondingly, policy makers trying to increase prices small suppliers receive should consider targeting the frictions that lead small firms to rely on agreements rather than just addressing market power directly.

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.

**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, Bradley Larsen, Shengwu Li, J.J. Prescott, Bernardo Silveira, and Chuan Yu**, "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.

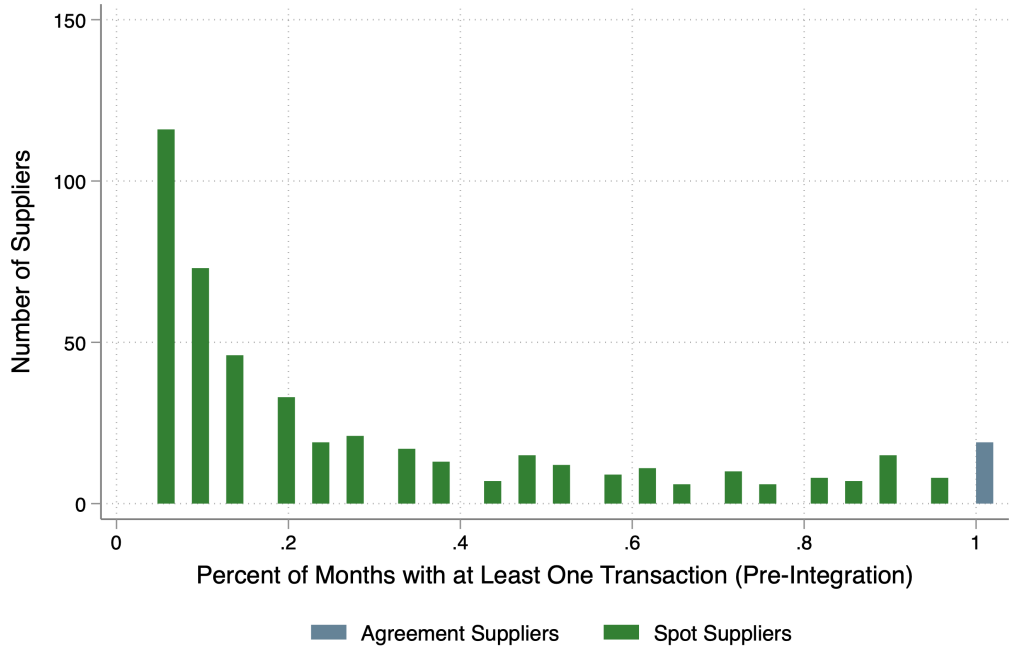
**Zavala, Lucas**, "Unfair Trade? Monopsony Power in Agricultural Value Chains," *Working Paper*, 2022.

# 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.

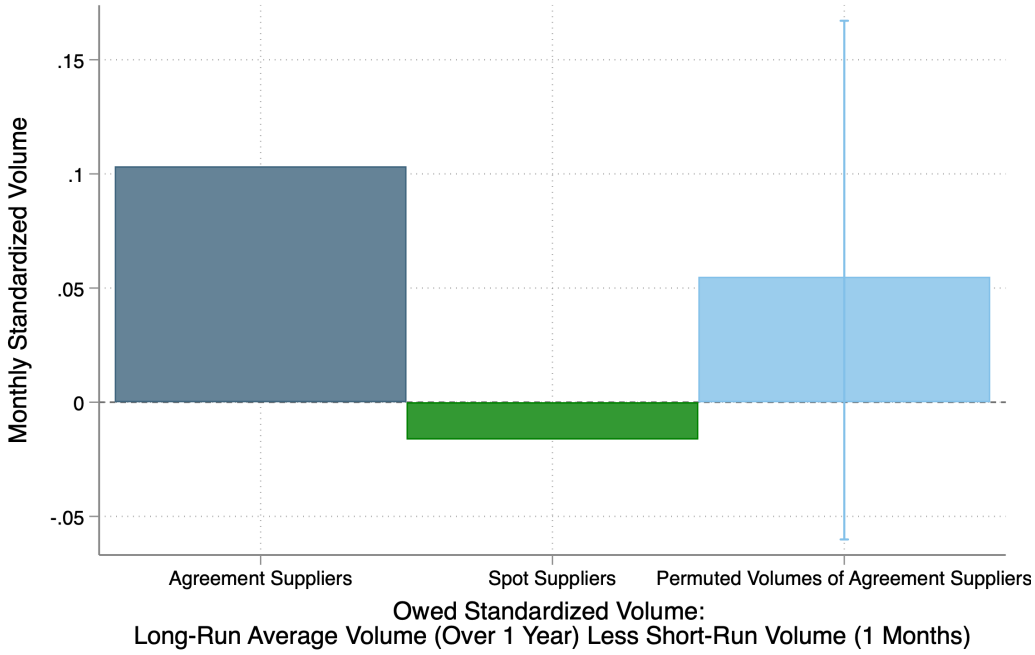
### A.1.2 Quantity Assurance: Owed Quantity Guarantees

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 measured as the difference between the proxy for the long-run quantity guarantee and the quantity fulfilled recently (*i.e.,* in the

prior month). Conversely, as there is no agreement with spot 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.[81] 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,

[81] 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).

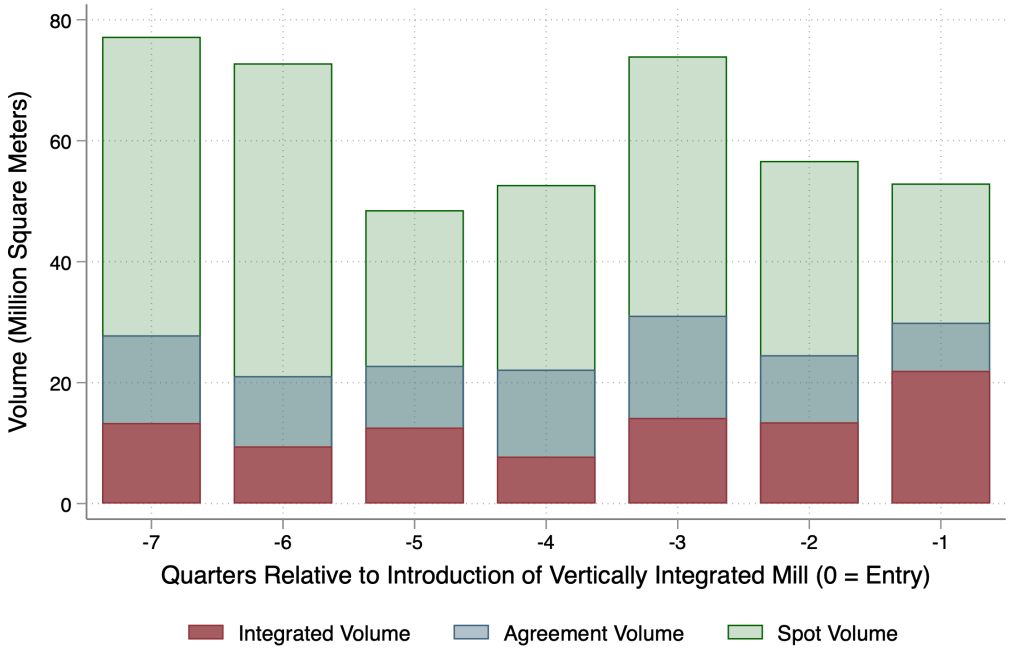
### A.1.3 U-Shaped Pricing Function

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 |

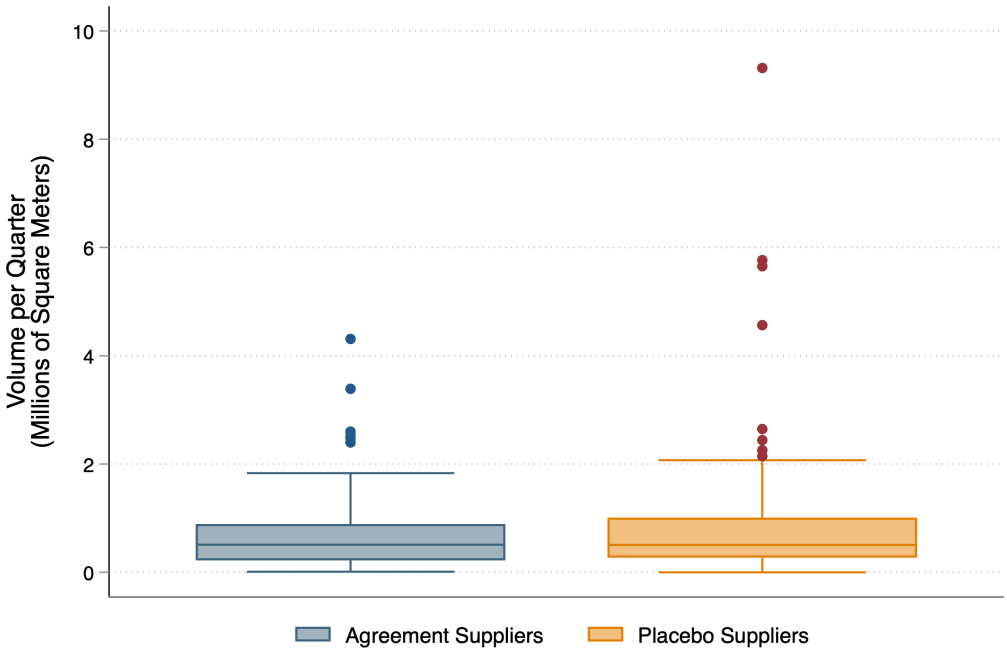
Figure A.3: Volume by Supplier Type (Pre-Integration)




**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, which is the highest price any supplier can charge since it is pinned down by the buyer's outside option. Therefore, I search over deviations and find the best static deviation, which occurs when $p^S q^{RC} - p^{RC} q^{RC}$ is lowest:

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

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:

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

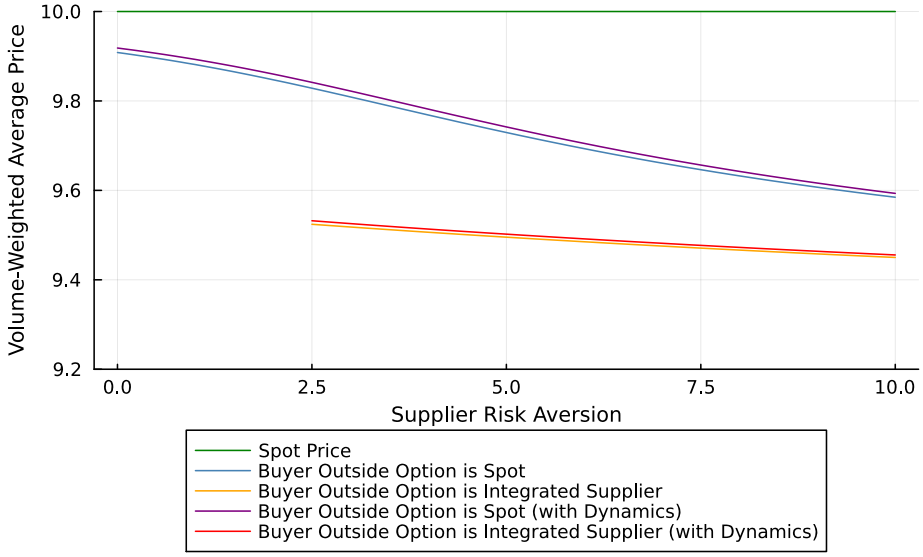
$$O_B = (\delta^{Bu})\Big(-\mathbb{E}_{q^{RC}}[C(q^{RC})]\Big) + (1-\delta^{Bu})\Big(-C(q^{RC})\Big) \text{ if integrated}$$

Note that for $\delta^S, \delta^B$ 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.[82]

[82] For this analysis, I set $\delta^{Su} = \delta^{Bu} = .9$. 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 & Supplier Risk Aversion with RC 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:

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

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, $\theta$ 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, $\theta$ should be higher for suppliers with higher risk aversion. I also normalize $\theta$ by subtracting off the expected payoff from selling to the spot market, such that $\theta$ 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 $\theta$, 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 $\gamma$ 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, $\gamma$ 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 $\gamma$ being a transfer; although not shown, results are qualitatively identical if $\gamma$ is exclusively a penalty.[83]

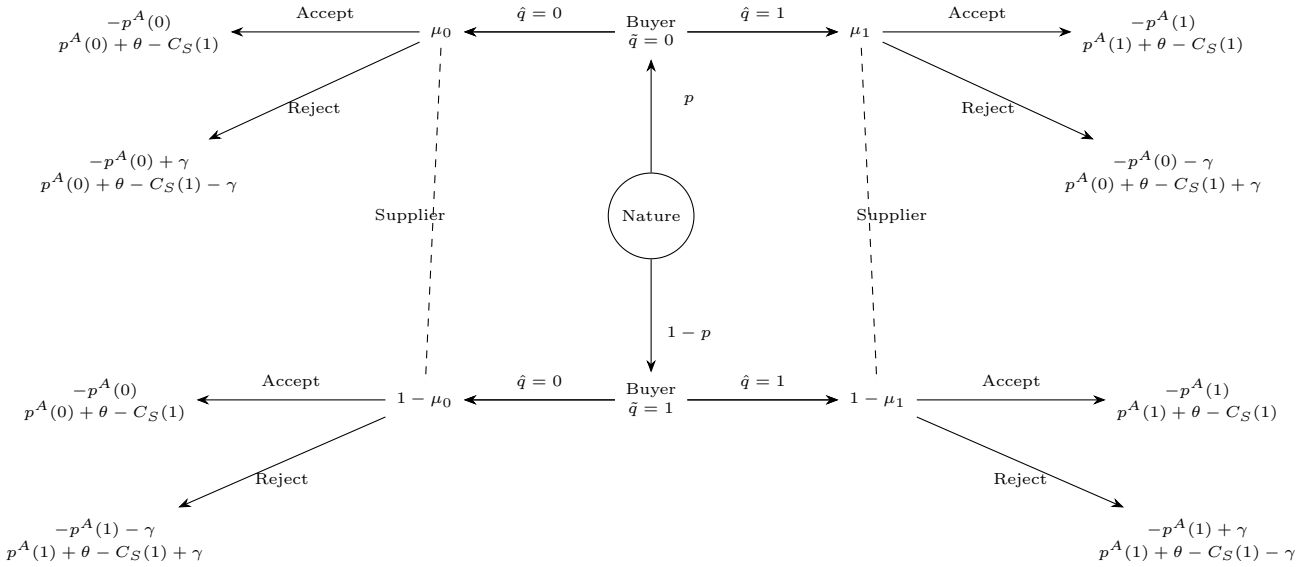
8. Nash bargaining prices are computed assuming that the buyer and supplier equally split the surplus (in other words, bargaining parameters are .5). I assume that the transfer $\gamma$ does not occur if the agreement dissolves.

Below, I show that there are two equilibria. First, there is an equilibrium where, with a sufficiently risk-averse supplier (*i.e.,* high enough $\theta$), the buyer pools and communicates $\hat{q} = 0$. However, this equilibrium no longer holds for low $\theta$. Second, there is an equilibrium where, with low $\theta$, 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 $\mu_0$ after the buyer plays $\hat{q} = 0$ and as $\mu_1$ 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 $\theta$ 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:

[83] 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.

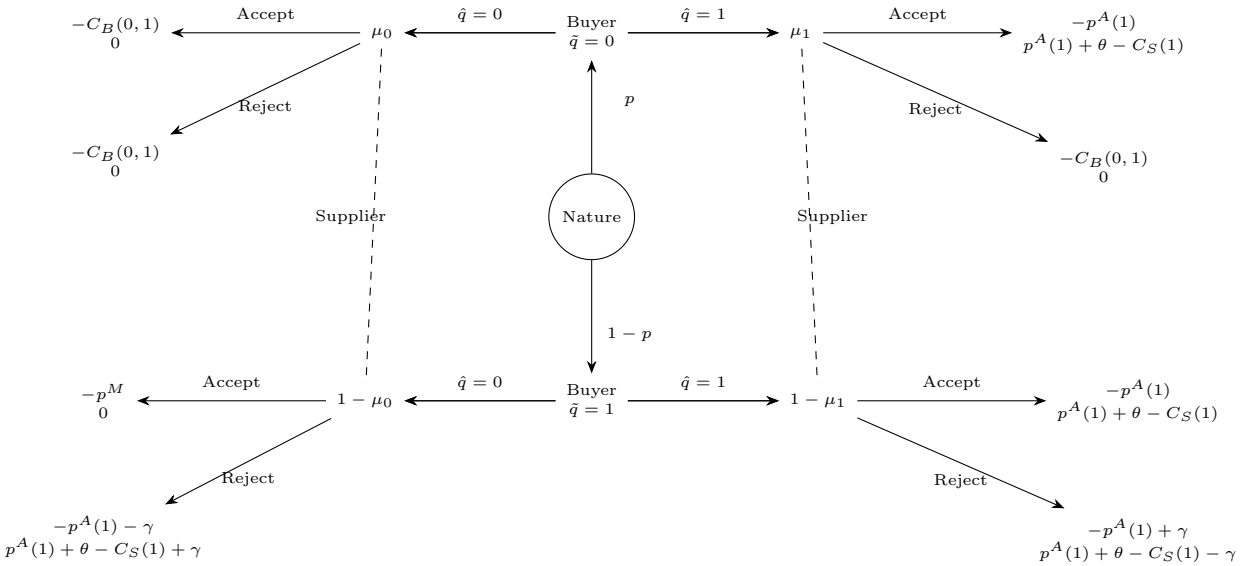
Formally, the equilibrium is that the buyer plays $\hat{q} = 0$ for both types, the supplier

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 $\gamma$ 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 $\gamma$, 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 $\theta$ 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 $\theta$ 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 $\theta$ 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 $\theta$ is low, then the incentive to reject after $\hat{q} = 1$ to earn $\gamma$, 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 $\mu_0$ 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 $\mu_0$.)

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 $\theta$ 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, $\theta$ 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, $\theta$ could represent (in a reduced-form way) higher risk aversion. I also normalize $\theta$ by subtracting off the expected payoff from selling to the spot market, such that $\theta$ 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 $\theta$ 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 $\theta$ 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 $\theta$, 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 $\theta$ 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 $\theta$ (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 $\theta$ 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.[84]

[84] 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.[85] 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.[86]

## 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.[87] 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.

[85] 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.

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

[87] 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.[88] Placebo suppliers are the closest counterfactual for agreement suppliers overall given their similar volumes.[89] 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.[90] 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

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

[89] 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.

[90] 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).[91]

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.

[91]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.

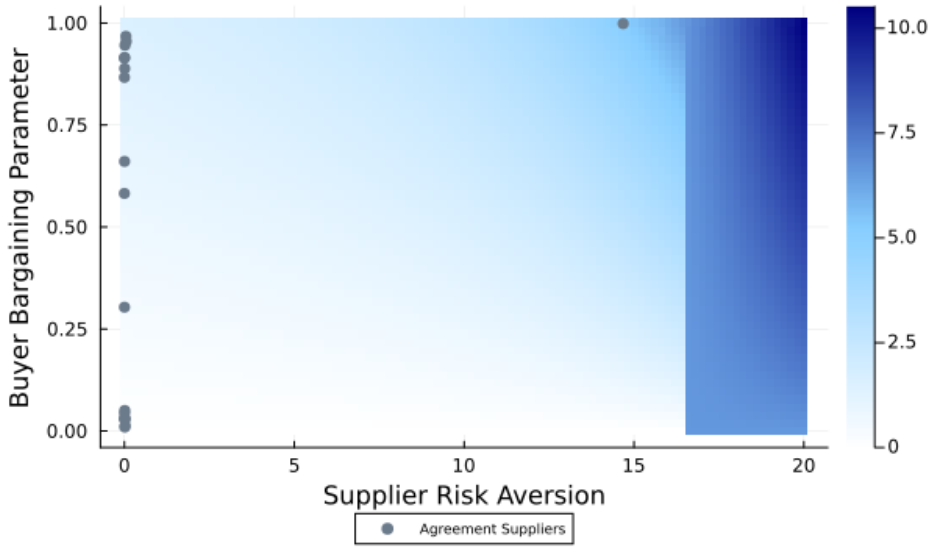
Table A.5: Buyer Bargaining Parameter and Supplier Risk Aversion: 95% CI

| *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

In this setting, a standard DiD design relies on parallel trends: absent integration, exposed suppliers' counterfactual outcomes would have evolved in parallel to those of control suppliers. Identification does not require random assignment of suppliers to treatment. Morton (2023) shows that *fabrics* (not suppliers) were endogenously selected: fabrics with more agreement suppliers are less likely to be brought in-house because vertical integration and agreements serve as substitutes. Conversely, suppliers are not selected for agreements based on individual fabrics (as discussed in Morton (2023)).[92] The data show that selection on fabrics to be brought in-house reflects trends in fabric demand: fabrics with increasing demand prior to the mill's construction are more likely to be produced at the new mill.

The endogenous choice of fabrics to produce in-house induces non-random supplier-level exposure to vertical integration via heterogeneity in suppliers's fabric portfolios. Therefore, I adopt a DiD design with propensity-score adjustment that accommodates non-random treatment (Abadie, 2005), detailed in Appendix A.11.1. This approach relaxes parallel trends to *conditional* parallel trends; in my setting, comparisons condition on supplier exposure to fabric-demand trends that reflect evolving fashion tastes. Falsification tests verify that propensity score adjustment achieves balance pre-treatment.[93] Furthermore, as outcomes are *not* matched pre-integration, pre-trends remain informative about parallel trends.[94]

### 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

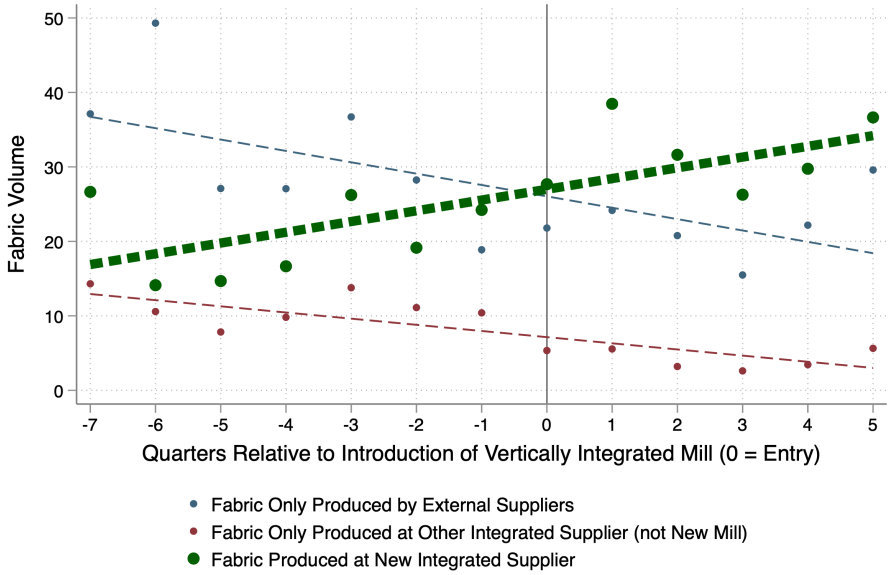
[92] There are two primary reasons why a supplier would not be selected for any individual fabric. First, the garment industry experiences short-run trends but agreements are long-run (for relational enforcement); changing trends makes selecting a supplier for an individual fabric a poor long-term strategy. 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. In fact, the median number of distinct fabrics sold by agreement suppliers in the 6 months prior to integration is large (42).

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

[94] 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.

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:

$$\mathbb{E}[Y_i^{post}(0) - Y_i^{pre}(0)|D_i = 1, \boldsymbol{X_i}] = \mathbb{E}[Y_i^{post}(0) - Y_i^{pre}(0)|D_i = 0, \boldsymbol{X_i}] \tag{4}$$

[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.

In equation 4, 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]

Table A.6: Balance on Untargeted Outcomes

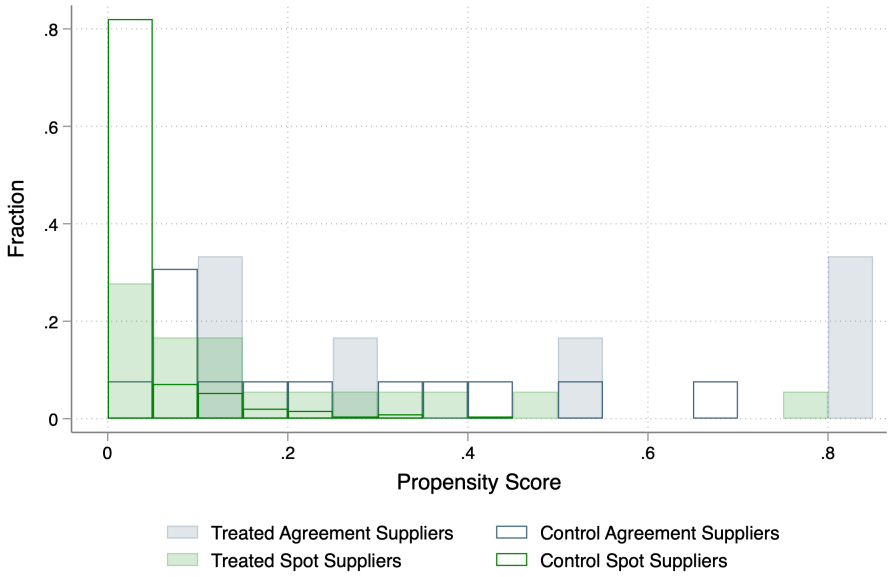
| | Coefficient on Treatment Indicator: | |
|---|---|---|
| Outcome Variable | Without IPW Adjustment | With IPW Adjustment |
| Pre-Integration Volume | 3033768 | 1085459 |
| | [1189072, 5489585] | [−1004042, 3383491] |
| Pre-Integration Transaction Count | 72.93 | −3.35 |
| | [44.49, 106.45] | [−46.3,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

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

[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 $\mathbb{E}[D_i] > 0$.

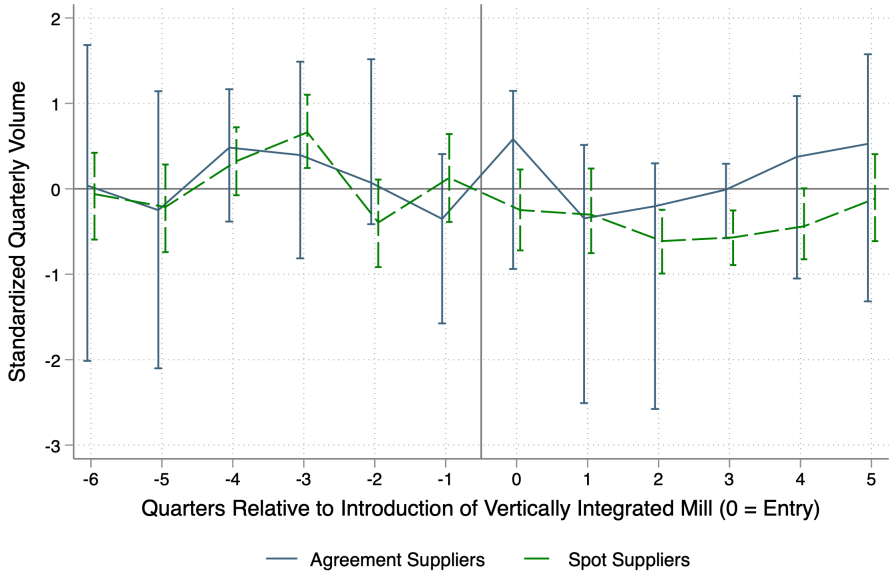
Figure A.8: Distribution of Propensity Scores by Supplier Type




**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] |
| Spot Suppliers Pre | .09 | 1.16 | −.02 | −.01 | .69 |
| | [−.05, .22] | [.01, 2.18] | [−.16, .07] | [−.24, .03] | [−.55, 1.83] |
| Spot 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

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

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).

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

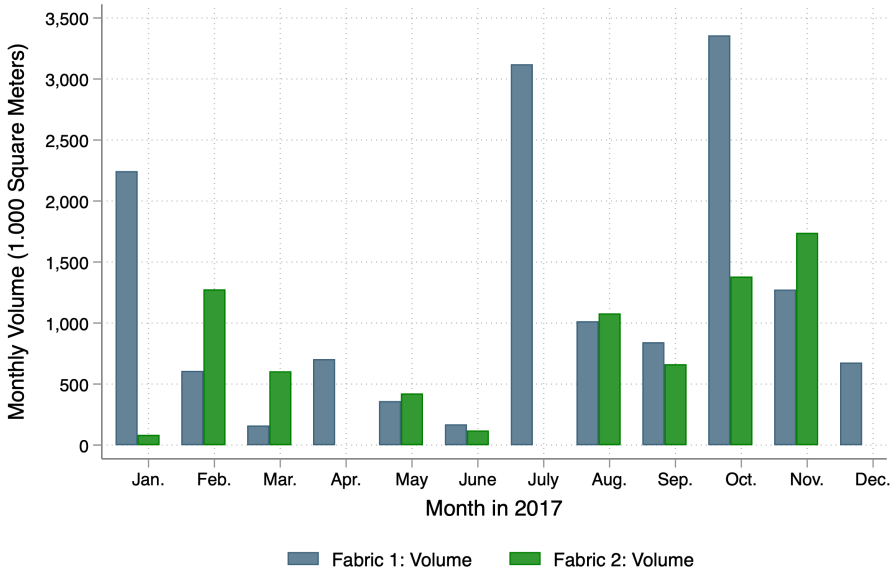
| Supplier Risk Aversion | | |
|---|---|---|
| | | No Outlier |
| Supplier Concentration | −20.417 | −.015 |
| (Volume-Weighted Supplier-Fabric HHI) | (13.89) | (.061) |
| N | 18 | 17 |
| $R^2$ | .355 | .006 |

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

# OA1 Online Appendix

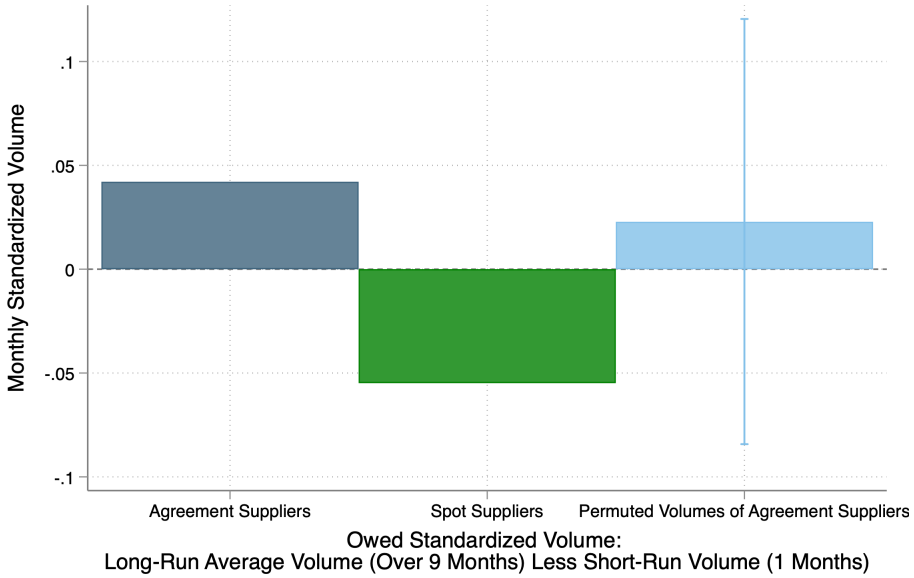
## OA1.1 Additional Descriptives About Buyer-Supplier Agreements

Figure OA1.1: Variation in Monthly Quantity by Fabric

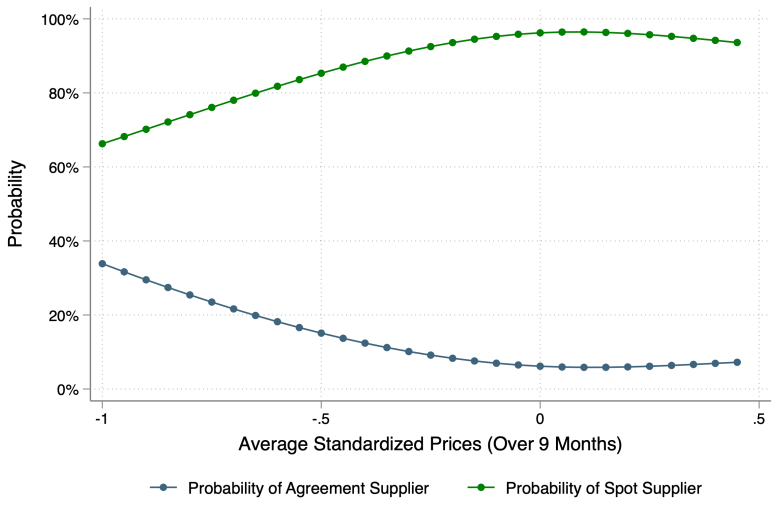



**Note**: Data from universe of fabric transactions by the buyer. Bars represent the quantity of fabric purchased during the month in 2017 for two selected fabrics. Fabrics were selected based on data availability to illustrate variation.

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

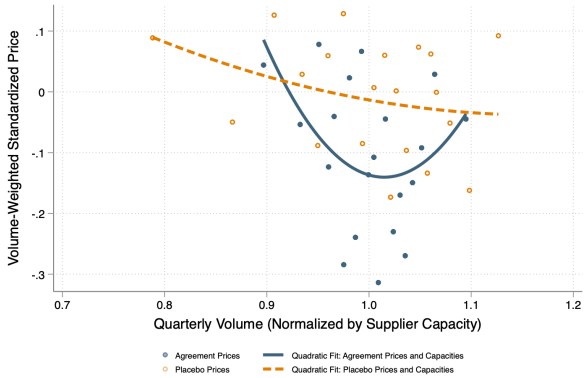
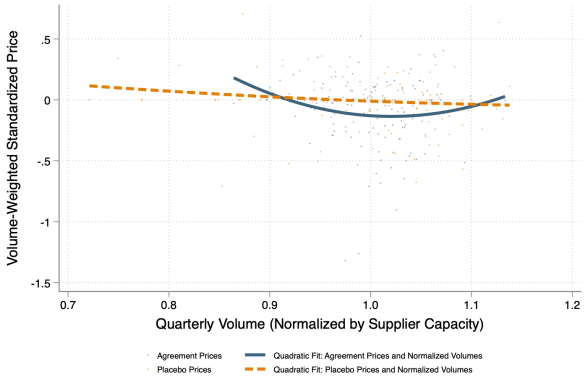



Figure OA1.4: Pricing Function (Placebo Comparison): Raw



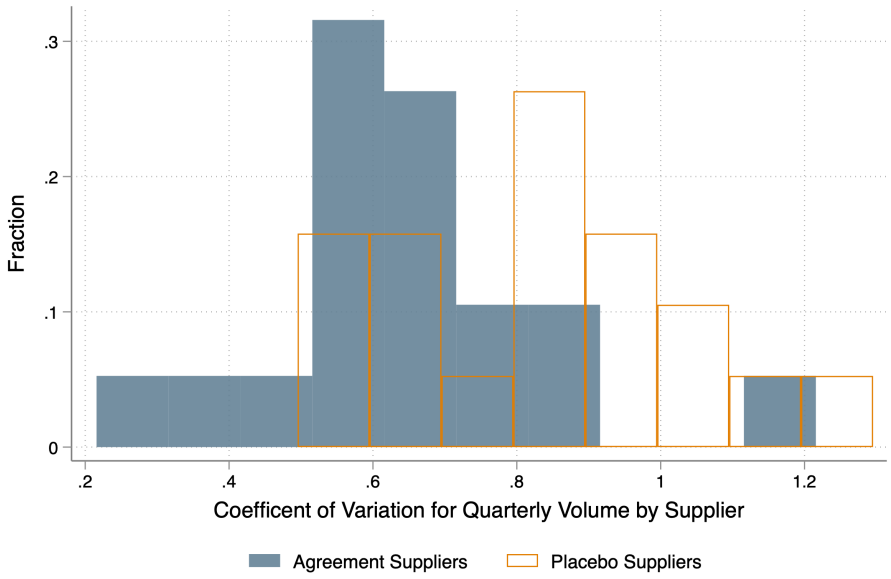

**Note**: Data from universe of fabric transactions by the buyer, transformed into supplier-by-quarter panel data. Quadratic fit is from a regression of volume-weighted standardized price on capacity and capacity squared. Figure OA1.3 bins prices and capacities using 20 quantiles. The mapping from quantities to capacities is as discussed in 4.1.2.

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

| | | | | | | |
|---|---|---|---|---|---|---|
| Normalized Volume | $-26.438^{**}$ | $-1.345$ | $-36.723^{**}$ | $-6.454$ | $-32.61^{**}$ | $-7.361$ |
| | (10.986) | (3.682) | (14.455) | (4.495) | (14.646) | (4.643) |
| Normalized Volume$^2$ | $12.95^{**}$ | .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 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 Normalized Volume | $-4.055$ | $-.596$ | $-5.415$ | $-1.864$ | $-4.662$ | $-2.079$ |
| 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 Normalized Volume | 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 |

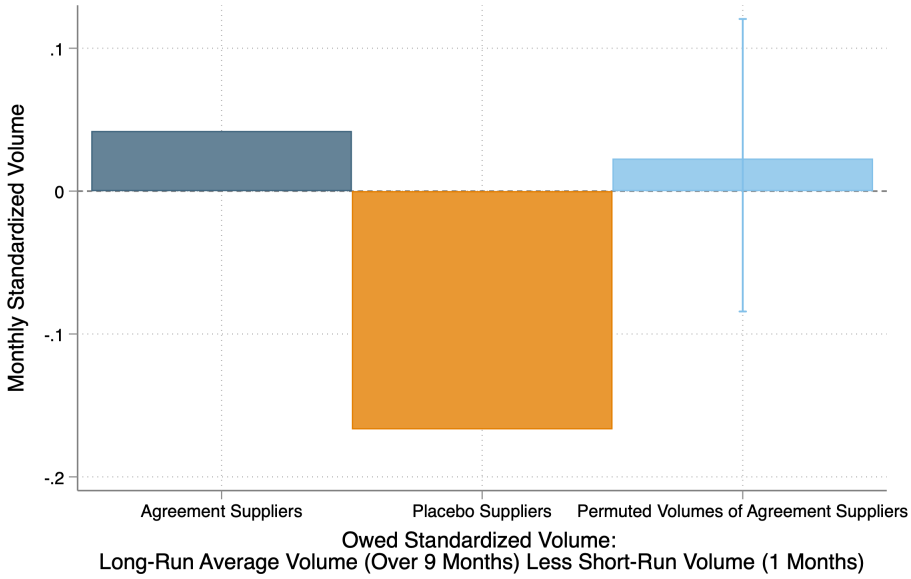
Table **??** illustrates prices with placebo suppliers never meet more than one of the four conditions for being U-shaped.

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.

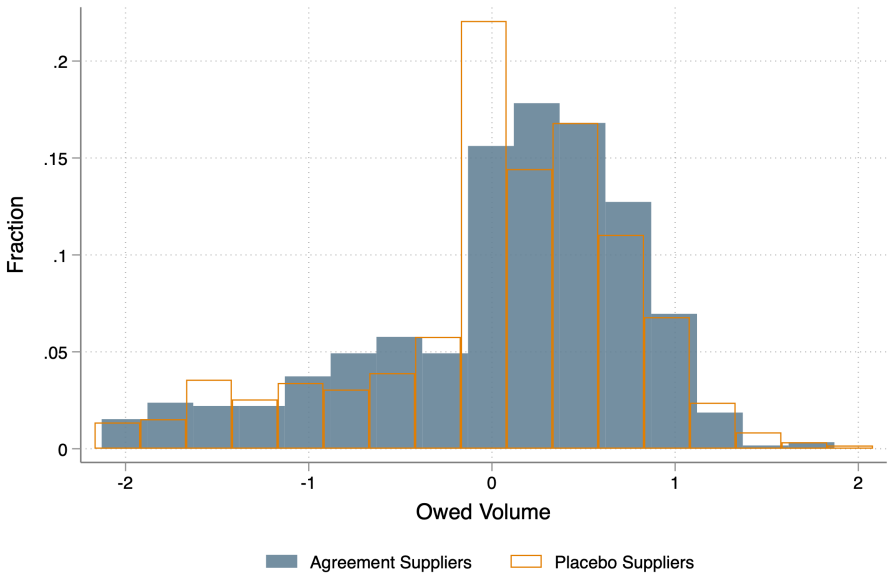
Figure OA1.6: Coefficient on Owed Quantity by Supplier Type (Placebo Comparison)



**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)



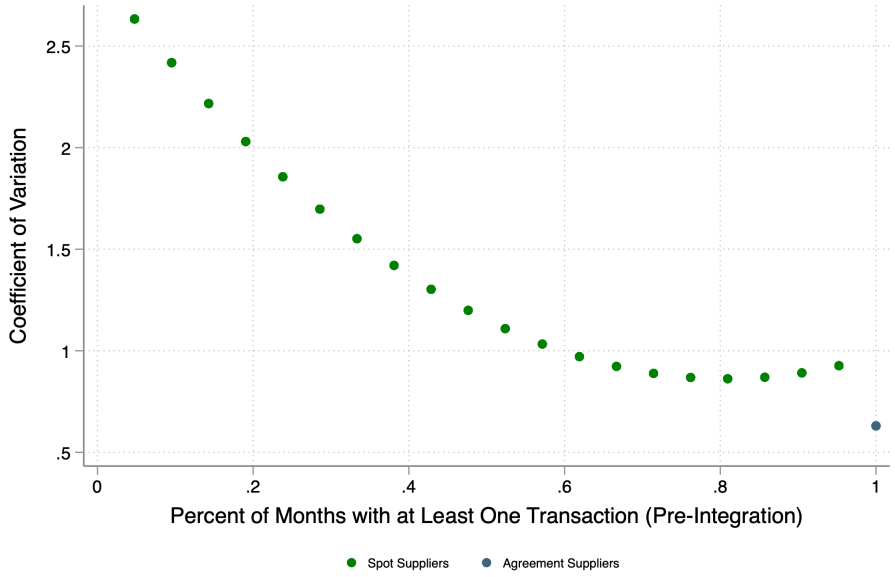
**Note**: Data from universe of fabric transactions by the buyer. Bars show the coefficient for owed standardized volume from a regression of current volume on owed volume and month fixed effects. Owed standardized volume is calculated for each supplier quarter as the standardized volume in the prior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

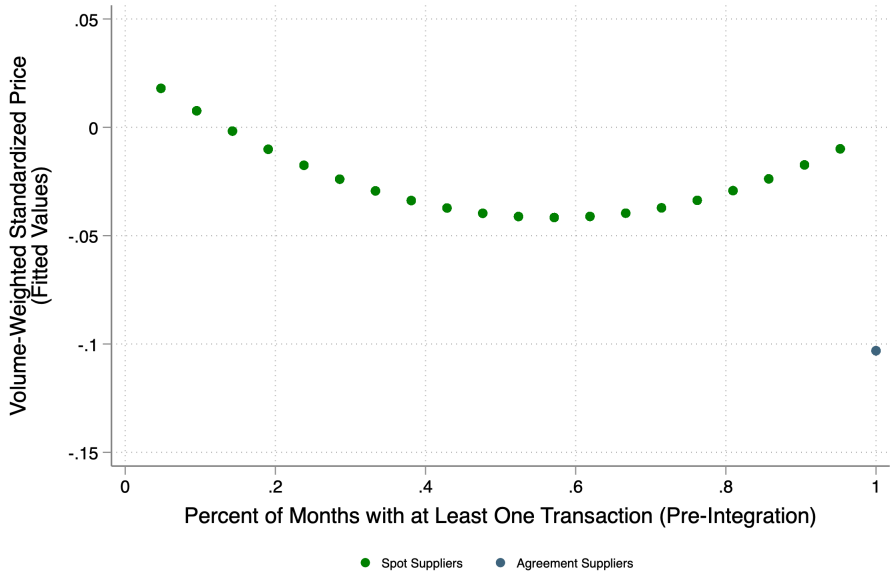
| | Coefficient of Variation for Supplier Quarterly Volume | Standardized Price |
|---|---|---|
| Agreement Supplier Indicator | −.346*** | −.102*** |
| | (.082) | (.031) |
| Series Controls ($k = 2$) for Transaction Frequency | Y | Y |
| Quarter FE | N | Y |
| N | 315 | 1603 |

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

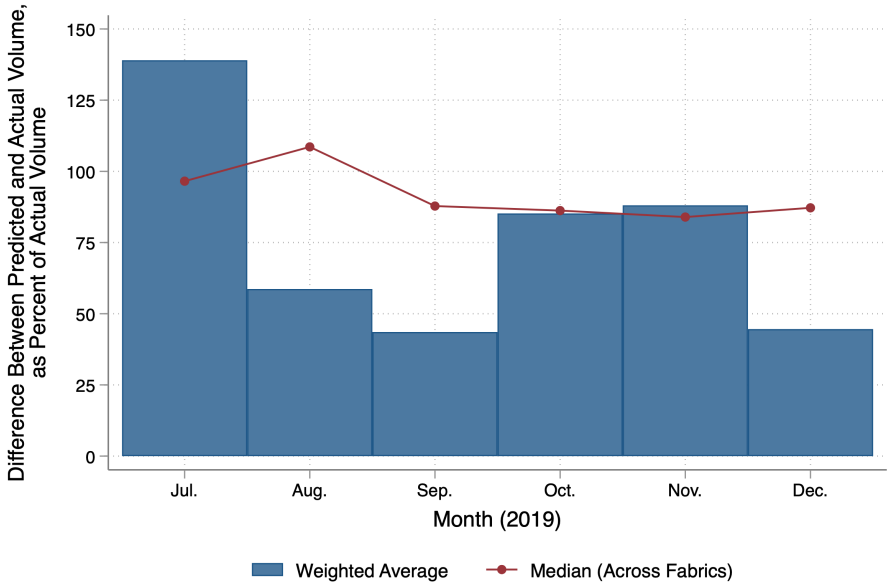
capacity, such a contract reduces supplier incentives to reject sampling opportunities (and, thereby orders), from other buyers to save capacity as the supplier would receive similar markups for all capacity states. It follows that suppliers would pass on costs to the buyer. Note that this concern is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

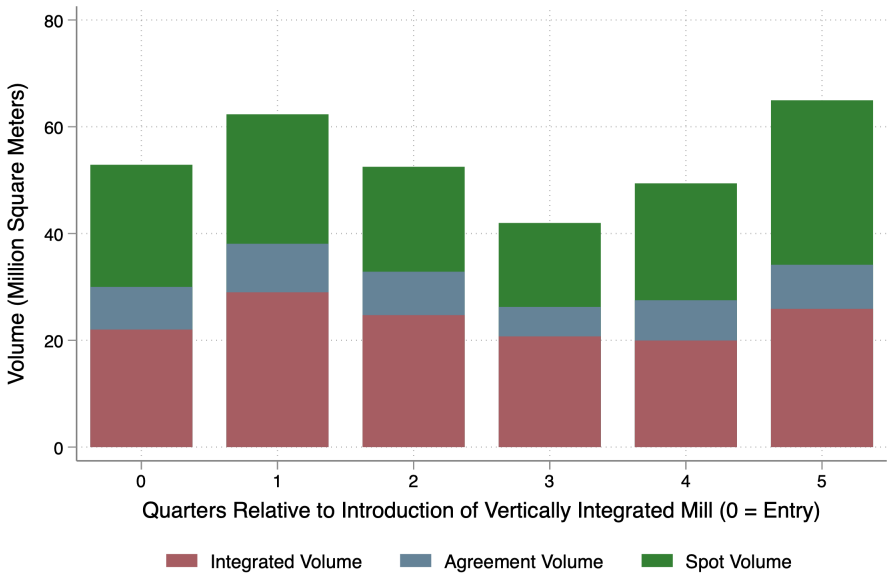
Figure OA1.10: Forecast Error in Volume Predictions




**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

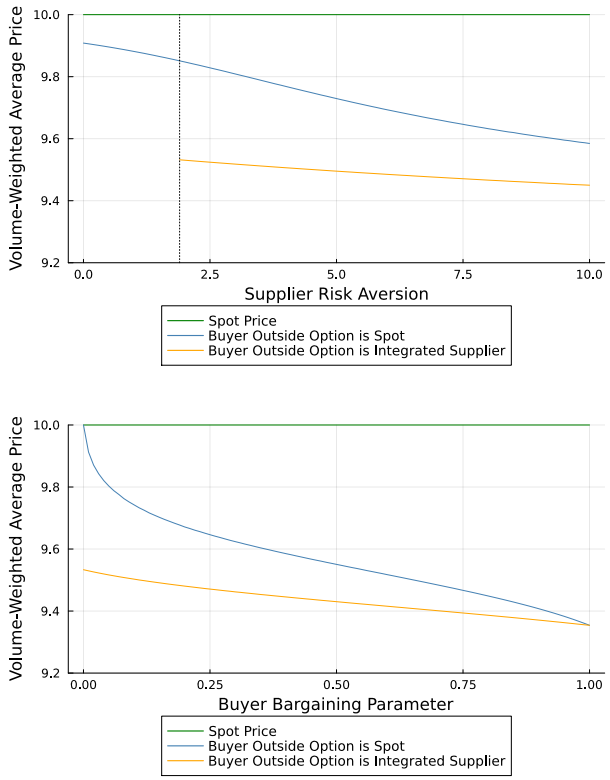



**Note**: Data from universe of fabric transactions by buyer. Bars represent the quantity of fabric purchased during the quarter from each supplier type.

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

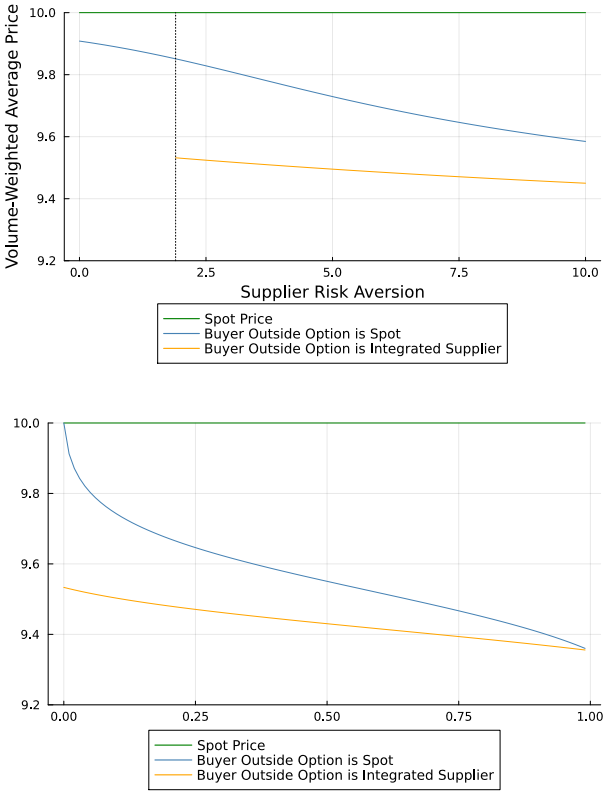
Results are robust to the external supplier having lower or higher costs than the integrated supplier, shifting costs by 10%. To ensure that optimal quantity is still $q = 1$ without changing market prices, I only shift the fixed part of costs by 10%, keeping convexity the same. Results (available upon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

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

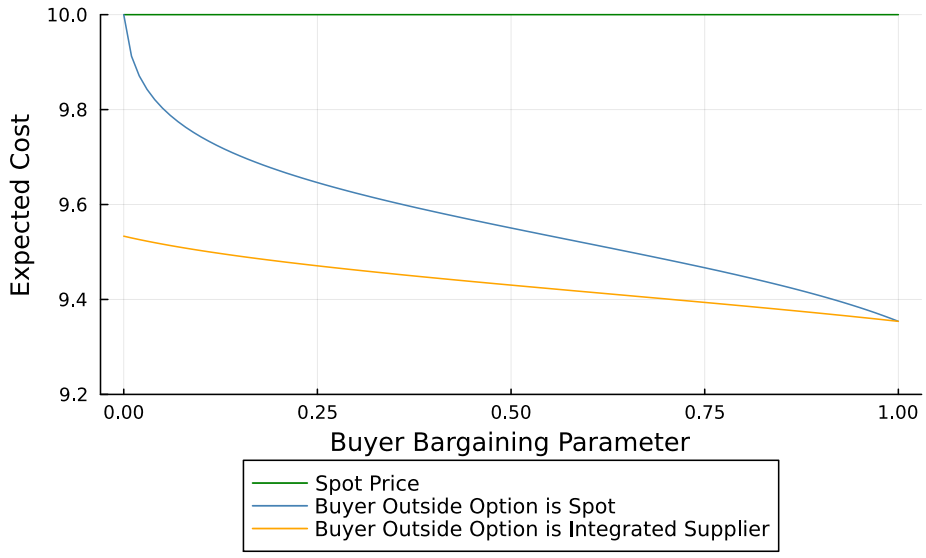
Although I do not observe $p$ directly from the data, note that $\frac{2p-1}{1-p}$ is increasing in $p \forall p \in [0, 1)$; therefore, for $p, \gamma$ sufficiently high, this inequality holds. Formally, as long as the inequality holds (meaning $p, \gamma$ are sufficiently high), then the equilibrium 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 $\theta$ 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

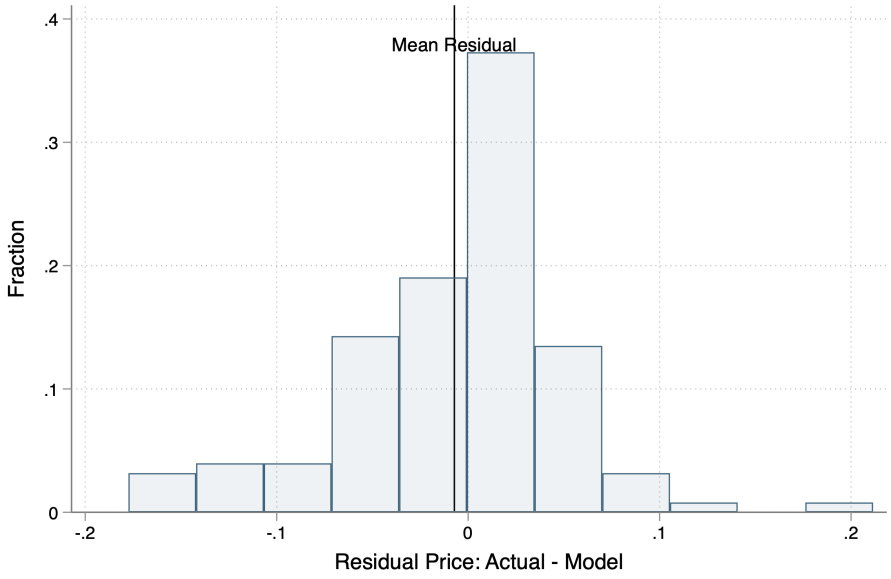
### OA1.8.1 Model Validation: Untargeted Moment

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.

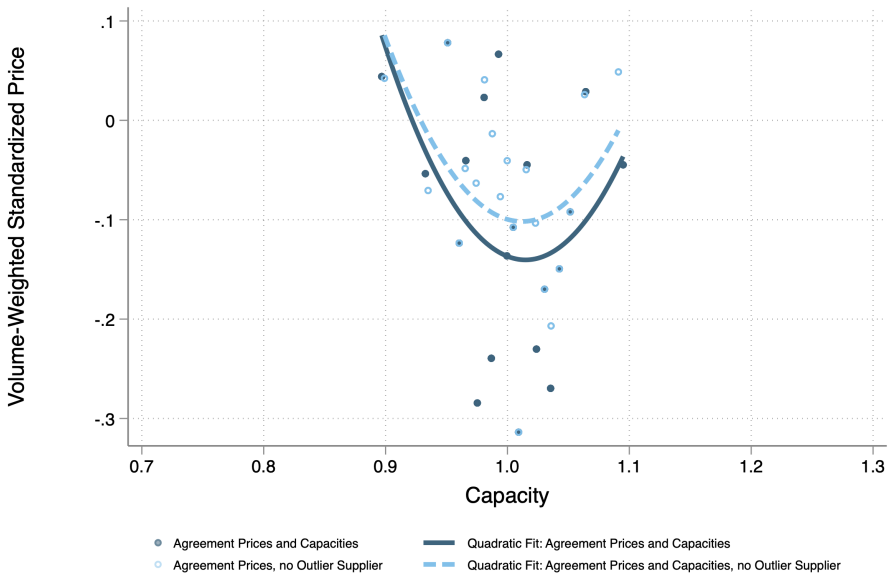
Table OA1.5: Out-of-Sample Fit Comparison (OLS with Tuning)

| | 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

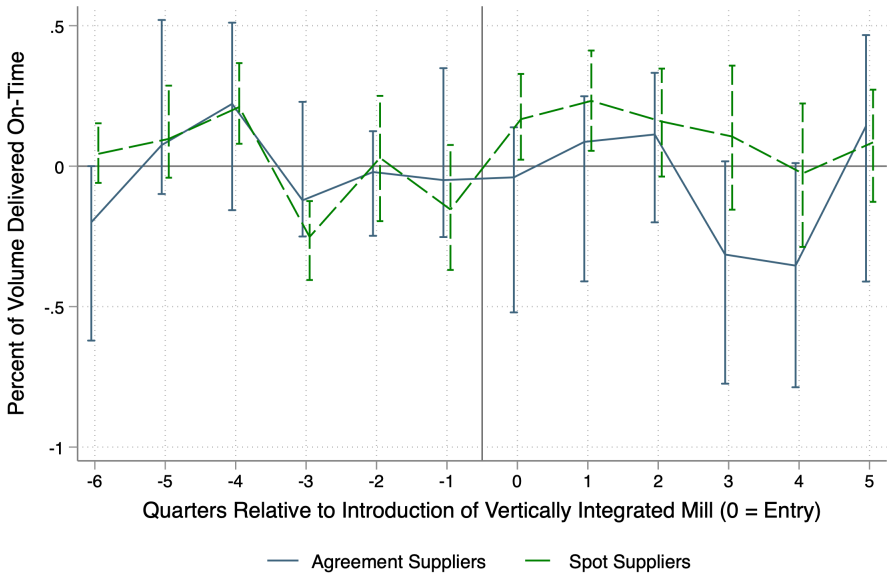



**Note**: Data from universe of fabric transactions by the buyer with agreement suppliers, transformed into supplier-by-quarter panel data. Quadratic fit is from a regression of volume-weighted standardized price on capacity and capacity squared. The figure bins prices and capacities using 20 quantiles.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

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

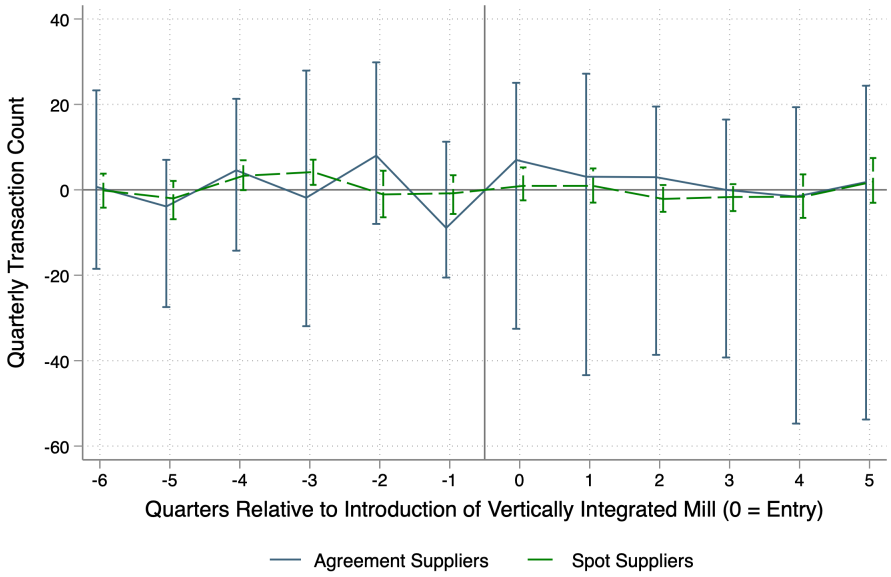
### OA1.10.1 Difference-in-Differences Additional Event Study Plots

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.

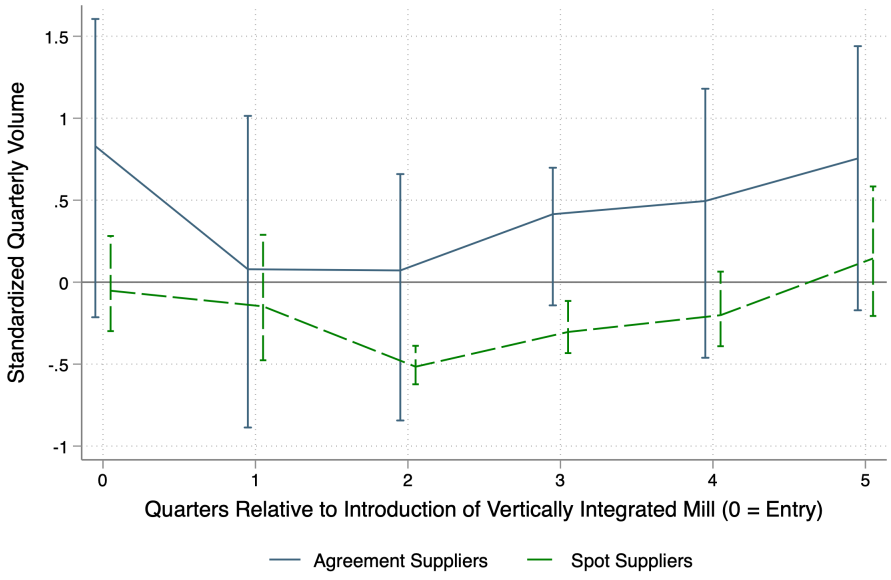
Figure OA1.18: Difference-in-Difference Estimates for Effect of Vertical Integration on Transaction Count




**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.

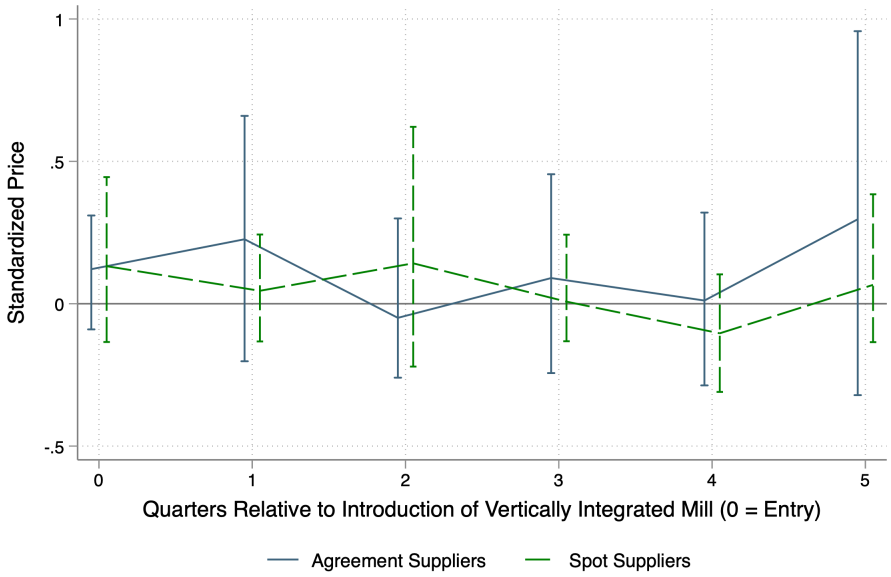
### OA1.10.2 Synthetic Difference-in-Differences

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



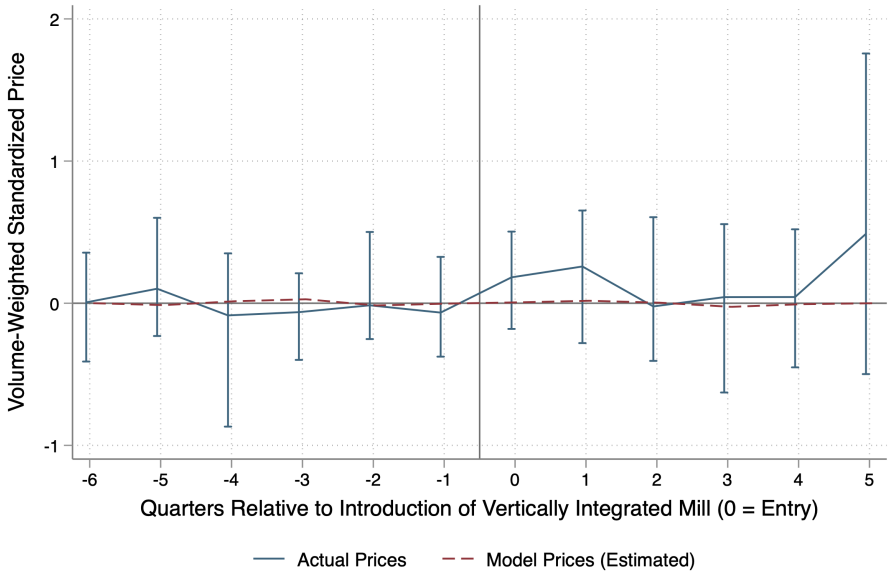

**Note**: Data from universe of fabric transactions by the buyer. Semiparametric difference-in-differences event study estimates with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

### 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).

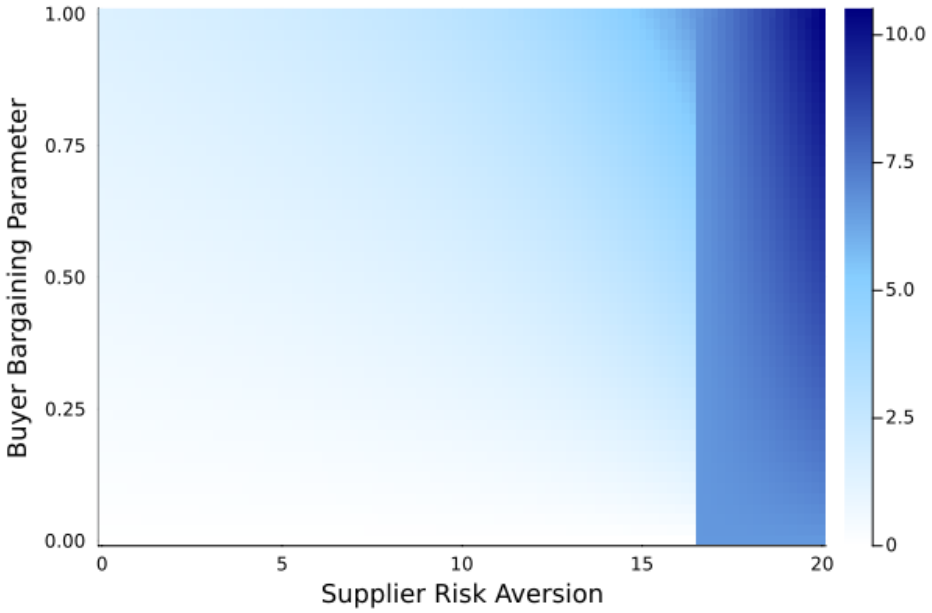
Figure OA1.21: Untargeted Moments: Difference-in-Differences




**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.