Value Creation & Value Capture in Indian Garment Sector Bargaining

Russell Morton*

September 6, 2026

Click Here for Latest Version

Abstract

Low prices paid by large exporters to suppliers in low- and middle-income countries raise concerns that gains from exporting are unequally distributed. I study how input prices reflect two distinct forms of buyer power arising from buyers' responses to frictions in input and supporting markets. Buyer power over surplus creation reflects the buyer's retention of surplus created when vertical contracts mitigate frictions, such as by providing quantity assurance or trade credit when external insurance or credit markets function poorly. Buyer power from surplus capture reflects the buyer's appropriation of pre-existing rents through replacement threats enabled by partial vertical integration, shifting rents without increasing agreement surplus. Distinguishing between these two forms of buyer power matters because of their differing implications for efficiency, distribution, and policy. To quantify their effects, I build and estimate a structural bargaining model using proprietary data on the universe of fabric purchases by a large Indian garment manufacturer. Buyer power over surplus creation, not surplus capture, explains observed discounts; the buyer retains ∼60% of created surplus.

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, Soheil Ghili, 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, Luis Rojas, Meredith Startz, Yulu Tang, Leonard Treuren, and Eleanor Wiseman and seminar participants at Banco de la República – Colombia, Chicago DIL, CSAE, EEA, Firms, Labor Markets, & Development, IIOC, LSE IO-Devo Workshop, MWIEDC, NEUDC, Northwestern, OECD, PacDev, Paris School of Economics, PEDL Anniversary Conference, PEDL Early-Career Workshop, RIDGE IO, SIOE, SITE, STEG Annual Workshop, University of Michigan, University of Rochester (Simon Business School), USC, Wayne State University, and Yale University for their excellent feedback and comments. I appreciate Shahi Exports and Good Business Lab for sharing expertise 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), the low prices that exporters pay their suppliers raise concerns that powerful buyers extract much of the gains from trade (Boudreau et al., 2023; Goldberg, 2023). This concern is especially salient because many LMIC supply chains are funnel-shaped: a few large exporters purchase inputs from many smaller suppliers (Zavala, 2022; Dominguez-Iino, 2023). Yet, economically distinct forms of buyer power can generate the same low prices, with different implications for efficiency, distribution, and policy.

I quantify how two distinct forms of buyer power shape input prices in fabric procurement by a large Indian garment exporter. Buyer power *over surplus creation* reflects the buyer's retention of surplus created when a vertical contract mitigates a friction, increasing agreement surplus. For example, agreements can stabilize quantities or provide trade credit that suppliers cannot easily obtain elsewhere.[1] Because the agreement bundles the input purchase with the surplus-enhancing terms, suppliers may accept a lower average price—effectively paying for insurance or credit by offering the buyer a discount. Conversely, buyer power *from surplus capture* reflects the buyer's appropriation of pre-existing rents through a threat to replace the supplier with in-house production.[2] Such threats reallocate rents without increasing agreement surplus.

Distinguishing between these forms of buyer power is necessary to understand the implications of vertical contracts for efficiency, distribution, and policy. Surplus-creating agreements improve efficiency by mitigating frictions. Importantly, the buyer's retention of surplus created gives the buyer an incentive to incorporate insurance, credit, or other terms into the vertical contract. Otherwise, if all gains accrued to suppliers, the buyer would not benefit from providing them. By contrast, surplus capture shifts existing rents away from the supplier without increasing agreement surplus. Therefore, when low prices reflect surplus capture, the central challenge is buyer rent extraction. But when they reflect surplus creation, addressing the underlying market friction may best support suppliers.

These two forms of buyer power share a common origin: both arise from organizational responses to frictions in input and supporting markets. Vertical contracts mitigate such

[1] Agreements between exporters and buyers of Kenyan roses (Macchiavello and Morjaria, 2015) and between fishing firms and ice suppliers in Sierra Leone (Ghani and Reed, 2022) stabilize quantities. Rwandan coffee mills (Macchiavello and Morjaria, 2020) and Ecuadorian manufacturers (Brugues, 2025) provide trade credit to their suppliers.

[2] Many buyers, including the garment exporter studied, simultaneously rely on both in-house production and external suppliers (Breza and Liberman, 2017; Garg et al., 2023).

frictions directly; for example, agreement terms substitute for missing insurance markets. Vertical integration instead bypasses these frictions by bringing input production in-house rather than relying on external suppliers. Because the same frictions motivate both organizational responses, the coexistence of contracting and integration—and the distinct forms of buyer power each generates—is not incidental but a natural feature of LMIC supply chains.

The primary contribution of this paper is to decompose agreement prices into buyer power over surplus creation and buyer power from surplus capture. To do so, I build a structural bargaining model that incorporates insurance-like agreement terms and replacement options from both in-house production and coexisting spot supply. Estimation leverages proprietary data on the universe of fabric purchases by the garment exporter. I find that only buyer power over surplus creation reduces agreement prices. The buyer retains $\sim$60% of surplus created. The absence of surplus capture is striking because the buyer constructs a large, visible in-house mill during the sample period that supplies 20% of its fabric needs.

The combination of detailed transaction data, coexisting agreement and spot supply, and the opening of an in-house mill makes this setting unusually well suited to studying both forms of buyer power together. Comparing purchases across agreement and spot suppliers reveals the otherwise-unobserved agreement terms that create surplus, while the new mill shifts the buyer's scope for replacement, providing variation to test surplus capture. The data contain $\sim$35,000 transactions between September 2016 and December 2019, covering over 6,000 fabric categories. For each transaction, I observe supplier identity, price, quantity, and estimated production cost (for internally produced fabrics).

Because the agreements are informal rather than written, I infer their terms empirically by comparing purchases across agreement versus spot suppliers while holding fixed the buyer, time, and granular fabric category. Two stylized facts emerge. First, agreement suppliers receive more stable quantities than spot suppliers. However, agreements cannot eliminate quantity variability and associated profit risk because orders are lumpy and final quantities reflect demand shocks from end clients like Gap. Second, agreement prices are U-shaped in supplier capacity utilization, whereas spot prices are flat. At low utilization, higher agreement prices offset the effect of low quantities on supplier profits. At high utilization, higher prices compensate for elevated production costs. Because end-client demand shocks limit agreement quantity assurance, U-shaped agreement prices reduce remaining profit risk. Therefore, I interpret both vertical contract features as insurance-like agreement terms.

I build a Nash bargaining model that maps the two forms of buyer power to three economically distinct primitives and in which the buyer's outside option depends on in-house

production and spot supply. Buyer power implies a markdown—an input price below the buyer's marginal value of the input. In this setting, the buyer's marginal value is fixed, so primitives generate markdowns through price reductions. Risk aversion determines how much surplus is created by insurance-like agreement terms, while the buyer bargaining parameter determines how much of the surplus created is retained by the buyer. Therefore, a higher value of either primitive—supplier risk aversion or the buyer bargaining parameter—reduces prices, increasing buyer power over surplus creation. The buyer's outside option instead governs surplus capture through replacement threats. When bargaining with an agreement supplier fails, the buyer procures fabric through in-house or spot supply. Therefore, in-house costs and spot prices determine the buyer's disagreement payoff.[3] By contrast, in Nash bargains over a pure input transfer without surplus creation, a higher buyer bargaining parameter and a stronger outside option both simply shift rents to the buyer.

The coexisting spot market limits buyer bargaining leverage from replacement threats compared to canonical models without spot supply. In such models, the best price from a possible replacement disciplines the bargained price (Ho and Lee, 2019; Ghili, 2022). However, in my setting, spot suppliers absorb marginal demand fluctuations, while agreement and in-house suppliers provide stable base volumes. If an agreement dissolves, quantity assurance limits the buyer's ability to source more from other agreement suppliers. Furthermore, increasing in-house production is costly at high capacity utilization. Therefore, the buyer resorts to expensive spot supply if an agreement breaks.

The model shows that replacement threats reduce prices only when the agreement creates enough surplus to survive the threat. More risk-averse suppliers value profit smoothing more, making them willing to remain in the agreement even at a lower average price. Consequently, the model generates a cutoff in supplier risk aversion. Below the cutoff, a replacement threat would break the agreement, replacing the buyer's existing discount with expensive spot supply. Therefore, the buyer prefers to maintain the agreement over making a threat. Above the cutoff, the supplier values the insurance-like terms enough to remain in an agreement with the lower price due to the in-house replacement threat. The buyer can then obtain an additional discount through surplus capture.[4]

I develop a new strategy to empirically separate the surplus that agreement terms create,

[3]This specification of disagreement payoffs differs from canonical Nash-in-Nash models without a spot market where disagreement implies no bilateral trade, holding other agreements fixed.

[4]In Nash bargaining, outside options are primitives rather than strategic choices. I formalize this cutoff result through two distinct games in which, before Nash bargaining, the buyer either signals residual in-house capacity or commits to reserve in-house capacity. Either pre-bargaining extension yields this cutoff result without the buyer choosing its outside option in the Nash bargain.

the rule that divides surplus, and the rents that replacement threats reallocate. First, the curvature of agreement prices with respect to supplier capacity utilization identifies supplier risk aversion. More risk-averse suppliers value profit smoothing more; accordingly, the model implies a stronger U-shape to agreement prices as risk aversion increases. Second, conditional on this curvature, the average level of agreement prices identifies the buyer's bargaining parameter: lower average prices indicate higher buyer retention of surplus created by the insurance-like terms. Third, observed spot prices and internal production costs discipline the values of the buyer's spot and in-house outside options and, thereby, the scope for rent reallocation through replacement threats.

The timing of garment production provides the plausibly exogenous variation in supplier capacity utilization that identification requires. The buyer selects a fabric supplier before downstream clients determine which garments they will order and in what quantities. As a result, supplier-specific capacity utilization varies with downstream end-client demand unknown to the buyer when it selects a supplier. These end-client-induced demand shocks then trace out the mapping between capacity utilization and agreement prices.

All supplier risk-aversion estimates fall below the model-implied cutoff above which replacement threats reduce prices. This result explains why the model attributes agreement discounts entirely to buyer power over surplus creation. Risk aversion estimates align with those for export-oriented LMIC firms—coffee mills (Blouin and Macchiavello, 2019).

As additional evidence, I show that replacement threats do not affect agreements, using a difference-in-differences (DiD) design based on the construction of the in-house mill. Identification leverages this increase in vertical integration and heterogeneity in supplier exposure to it. The DiD complements structural estimates because the DiD relies on variation distinct from the price curvature and levels used to estimate the model. If in-house production strengthens the buyer's outside option against exposed agreement suppliers, agreement prices should fall. DiD estimates find no change in agreement prices or quantities. By contrast, exposed spot suppliers experience reductions in quantities but not prices. Therefore, internal production displaces marginal spot supply rather than disciplining agreement terms. These estimates are also inconsistent with the price responses predicted by a canonical replacement-threat model for a setting without a coexisting spot market. This contrast shows how sourcing marginal volumes from spot suppliers restrains bargaining leverage from vertical integration.

Surplus capture may matter more for small, financially constrained firms, where high risk aversion is particularly plausible. Small firms predominate empirically in India: 95% of

industrial units have 10 employees or fewer (Rajagopalan and Shah, 2024). Through the lens of the model, I quantify how agreement discounts change when supplier risk aversion exceeds the cutoff above which replacement threats reduce prices. The model implies that agreement discounts increase to 6.7% of spot prices—larger than Indian tax for apparel (5%).

I evaluate counterfactual policies to illustrate the different policy implications of the two forms of buyer power. First, I increase buyer competition. At the median estimated risk aversion, increased competition reduces the agreement discount relative to spot prices by 76%. However, when supplier risk aversion exceeds the cutoff above which replacement threats reduce prices, increased competition barely shifts prices. Once replacement threats discipline bargaining, little surplus remains to reallocate. Therefore, ignoring replacement threats overstates the benefits of buyer competition for highly risk-averse suppliers.

The second policy targets the underlying friction rather than buyer leverage directly. Insurance against spot-market profit risk increases prices for all suppliers, with the largest gains precisely where competition barely shifts prices—for suppliers above the risk-aversion cutoff. Insurance enables suppliers to smooth profits outside the agreement, both strengthening their outside option and eliminating buyer rent extraction through replacement threats. Intuitively, insurance reduces buyer power because the buyer no longer serves as the sole insurance provider. While moral hazard may limit the feasibility of profit insurance, policies that facilitate access to credit, savings, hedging, or cooperative risk sharing address the same friction. Therefore, insurance advances both efficiency and equity: it facilitates smoothing and counteracts buyer power.

Overall, the results show that low input prices by themselves need not imply extractive buyer power, although they may remain a distributional concern regardless. Accordingly, the appropriate policy response need not always focus on buyer leverage but should also consider the frictions that give rise to buyer power over surplus creation, an understudied form of buyer power (*e.g.*, Boehm and Oberfield (2020); Zavala (2022); Dominguez-Iino (2023); Bergquist et al. (2023); Startz (2025)).

This paper makes two related contributions to the literature on buyer power in vertical contracts (*e.g.*, Grennan (2013); Rubens (2023); Treuren (2025); Demirer and Rubens (2025)). First, I decompose observed price discounts into *i)* the buyer's retention of surplus created by friction-mitigating agreement terms and *ii)* the buyer's appropriation of pre-existing rents through replacement threats. To quantify the first component, I incorporate supplier risk aversion and insurance-like terms into an empirical bargaining model. To estimate the model, I develop an identification strategy that separates risk aversion from the

buyer bargaining parameter using the curvature and level of the pricing function. Although insurance plays an important role in vertical relationships (Lafontaine and Slade, 2007), empirical bargaining models often abstract from risk aversion and insurance (Cuesta et al., 2019; Lee et al., 2021), potentially attributing discounts generated by insurance-like terms to other sources of buyer power.

Second, I show how marginal spot supply constrains buyer power from replacement threats. Canonical replacement-threat models suggest substantial scope for buyer power through replacement threats: the best price available from excluded suppliers disciplines the buyer's outside option, reducing prices (Ho and Lee, 2019; Ghili, 2022). However, when spot supply fulfills marginal demand, in-house production may not generate bargaining leverage against agreement suppliers. This result connects buyer power to the vertical organization of production: the bargaining consequences of integration depend on how the buyer allocates stable and marginal sourcing across in-house, agreement, and spot supply.

Additionally, this paper advances research on contracting frictions and relational enforcement in LMIC vertical relationships (Macchiavello and Morjaria, 2015, 2020; Ghani and Reed, 2022; Brugues, 2025). Whereas much of this literature studies how informal agreements are sustained, I quantify how agreements create surplus by mitigating frictions and how bargaining divides that surplus. This focus motivates a bargaining model rather than a standard take-it-or-leave-it model that fixes the allocation of surplus by assumption.

The remainder of the paper proceeds as follows. Section 2 introduces the setting and data. Section 3 develops the model. Section 4 presents the 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 agreements: *i)* they provide quantity assurance; and *ii)* their prices feature a U-shaped relationship with quantities and a discount, on average, relative to spot prices. I interpret these facts as indicating that agreements insure suppliers. Then, I describe the buyer's construction of a new mill. Last, I depict the buyer's overall fabric sourcing, illustrating that spot suppliers provide *marginal* volume.

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 contracts for two key reasons. First, I observe not only transaction prices, quantities, and dates but also granular information about the specific fabric (6,176 codes) and supplier identity. The 34,681 transactions represent over 750 million square meters of fabric purchased (for ∼$180 million per year) from 622 suppliers (471 pre-integration). By analyzing transactions from one large buyer, and tracking suppliers over time, I can compare transactions across suppliers while 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 produced internally. As internal and external mills use similar technology, these cost estimates facilitate constructing markups directly from the data.[5]

I supplement the transaction data with cost data from the integrated mills that show how costs relate to capacity utilization.[6] 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

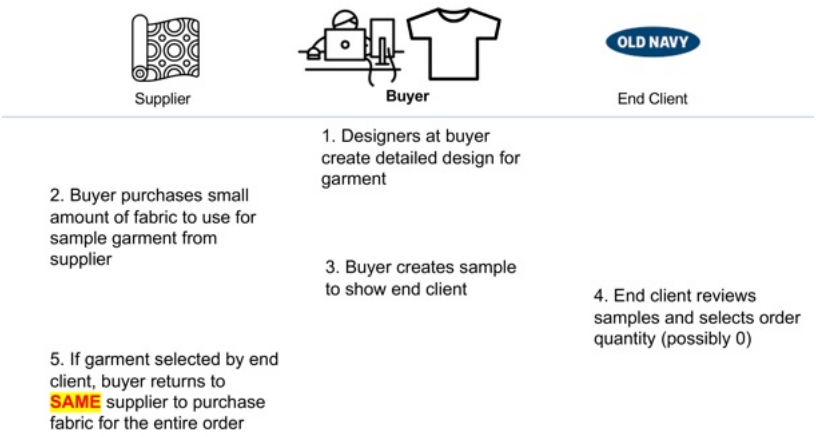
Garment manufacturing (in Figure 1) begins with buyer-employed designers sketching the shape and choosing a fabric type and color. The buyer then selects a supplier (a mill) to provide fabric for a few prototype garments ("samples"). The end client (*e.g.*, Old Navy) then reviews samples, and the end client and buyer agree on quantities and prices for a subset of garments. The buyer must return to the *same* supplier who produced fabric for the sample to ensure an exact fabric match; fabric features are difficult to reproduce and minor deviations create risk of garment rejection by the end client. Short timelines prevent trial-and-error with a different supplier to match the fabric, as only 28 days separate fabric purchases for the sample and the full order, consistent with the "fast fashion" context.[7]

[5]Cost estimates correspond to a mill that is not capacity constrained and represent average variable costs. Markup calculations assume common costs across in-house and external supply, reflecting that both use similar industry-standard technology.

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

[7]Fabric procurement reports that the sampling supplier fulfills the entire order for ∼ 90% of orders. In the rare cases when the buyer switches suppliers after sampling, production is typically brought in-house and only for orders with unusually distant delivery to facilitate experimentation and control to match the fabric.

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. Online Appendix OA1.4 discusses why prices for the full fabric order are not set before sampling, even though garment prices for sales to end clients are determined before the buyer finalizes fabric prices, and why the buyer cannot hold up suppliers.

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 which insure suppliers 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 indicate supplier type, I classify suppliers by transaction frequency, as is standard in the literature.[8]

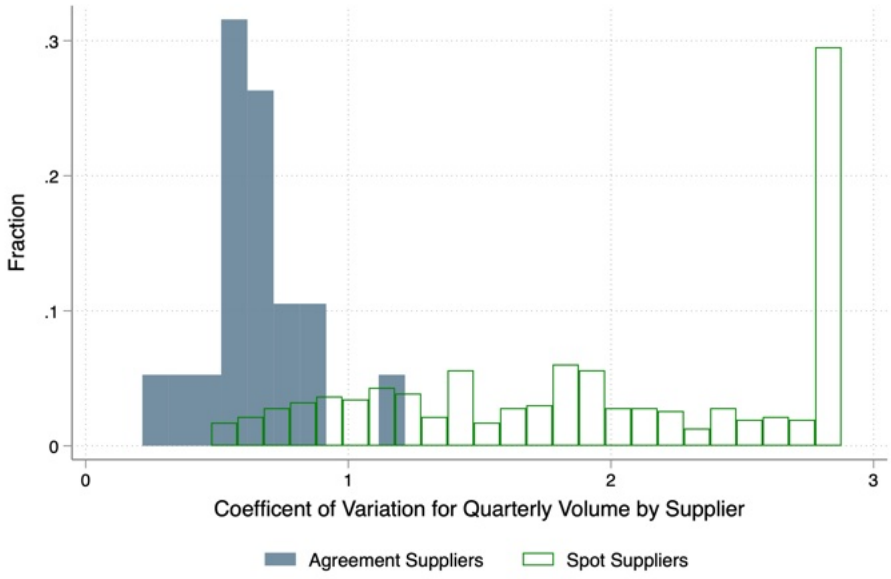
[8]For example, Macchiavello and Morjaria (2015) and Macchiavello and Miquel-Florensa (2017) define relational contracts based on consistent trading every period ("at least three consecutive seasons" in the latter). This relational contract is *not* a relational *incentive* contract (Levin, 2003), as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

This approach is motivated by the relational enforcement of agreements: repeated interactions generate the future surplus necessary for relational enforcement. I classify suppliers transacting every month as agreement suppliers.[9] These 19 suppliers account for ∼25% of pre-integration volume but only ∼4% of the 471 pre-integration external suppliers. Discussions with fabric procurement managers highlight the importance of preferred suppliers, with whom they meet in-person quarterly.

2.3.1 Fact 1: Agreements Provide Quantity Assurance

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 lower for agreement than for spot suppliers. The difference in average coefficient of variation by supplier type is statistically significant ($p < 0.001$). This result is not entirely mechanical: supplier classification depends on transaction frequency, not volume, so agreement suppliers could have volatile quantities.

Figure 2: Coefficient of Variation and Supplier Type



[9]Supplier classification only considers pre-integration months so analyses of the effects of integration on agreement suppliers hold supplier type fixed. Appendix Figure A.1 shows the distribution of the percentage of months with transactions across suppliers.

Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Only quarters before the buyer adds a new integrated supplier are included. Coefficients of variation are computed using the sum of all fabric volume in the quarter; each supplier is one observation ($N = 471$). The difference in average coefficient of variation by supplier type (agreement vs. spot) is statistically significant ($p < 0.001$).

Appendix Section A.1.2 provides additional evidence of quantity assurance: agreement suppliers who are "owed" volume—those with below-typical recent quantities—receive additional volume (and vice versa), even beyond mean reversion. Conversely, spot suppliers display the *opposite* pattern: *less* volume when they are "owed" (and vice versa). The buyer provides quantity assurance through both the number of samples and matching suppliers to end clients with different expected quantities.

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

Second, I illustrate that agreements feature a U-shaped pricing function with respect to capacity utilization (*i.e.,* 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 both seasonal fluctuations[10] and fabric input prices. Quantities are expressed as quarterly volumes relative to supplier productive capacity—effectively, quarterly supplier-specific *capacity utilization*—to account for heterogeneity in supplier scale.

Figure 3: Pricing Function: Binned

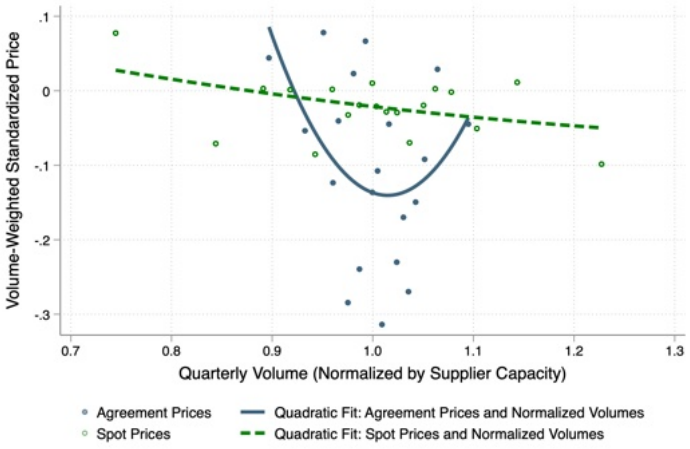
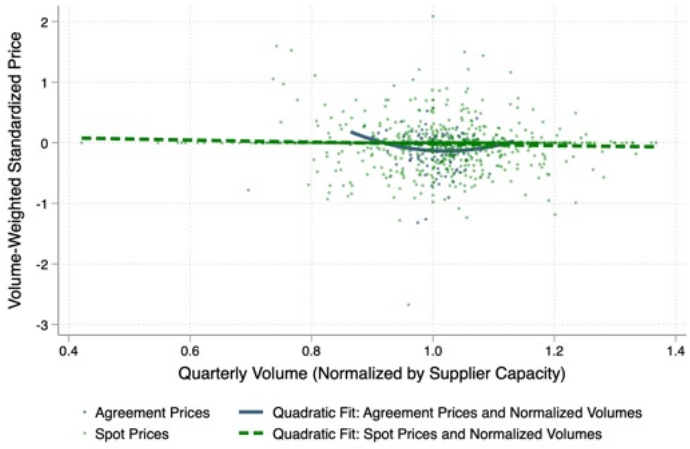


Figure 4: Pricing Function: Raw



[10]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).

Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Figures show both the quadratic fit and a scatter plot. The quadratic fit plots fitted values from a regression of volume-weighted standardized price on both normalized quarterly volume and normalized quarterly volume squared. Volume-weighted average prices are standardized by fabric and month. Normalized quarterly volume represents capacity utilization relative to supplier-specific capacity, where 1.0 = 100% capacity (details on measurement in 4.2). The scatter plot in Figure 3 shows binned prices and volumes using 20 quantiles, while the scatter plot in Figure 4 shows all raw data.

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 in the binned data, but remains visible in the raw scatter despite the wider y-axis scale. Because the buyer inspects all suppliers before sourcing from them, including spot suppliers, supplier capacity is observed by the buyer (and known to the supplier). Therefore, both parties can map realized quantities into capacity utilization when bargaining over prices.

Regressions confirm that only agreement—but not spot—prices statistically exhibit a U-shape (see Appendix A.1.3). Specifically, a regression of prices on capacity and capacity squared shows that only agreement prices exhibit both a negative, significant relationship with capacity utilization and a positive, significant relationship with squared capacity utilization, and fulfill additional criteria. The U-shape of agreement prices is more pronounced when regressions include supplier and quarter fixed effects, while spot prices never fulfill either slope criterion in any specification. Furthermore, average agreement prices are below spot prices, highlighting that agreements include a discount (see 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 in an environment with cost convexity (*e.g.,* due to capacity constraints).[11]

First, quantity variability generates profit variability. To illustrate why, consider a supplier who receives constant prices. When quantities are high, high costs erode margins. When quantities are low, even though margins may be high, the supplier's profits are low due to few units. Therefore, quantity assurance stabilizes profits (however, agreements do

[11] I empirically verify cost convexity using data from in-house suppliers in Figure 10. 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.

not fully eliminate variability because orders are lumpy and cannot be split across suppliers).

Second, as quantity assurance provides incomplete profit insurance, the U-shaped pricing function can further reduce profit variability relative to constant prices. Higher prices for low quantities smooth profits by compensating for the few units. When quantities are high, high costs erode margins under constant prices. Higher prices at high capacity utilization offset high costs, further stabilizing profits.[12] Appendix A.1.3 illustrates that agreement prices are lowest when profits would be highest with constant prices (at 100% capacity utilization), further supporting the insurance interpretation.

Last, I show that the stylized facts are unlikely to reflect two alternative explanations: *i)* unobserved 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 without insurance. Any alternative explanation should simultaneously account for the full set of empirical facts: *i)* agreements feature quantity assurance, *ii)* agreement prices are U-shaped: highest at both quantity extremes and lowest at 100% capacity utilization, *iii)* agreement prices are lower on average, and *iv)* spot prices are flat. For example, Nash bargaining over realized surplus with convex production costs could generate U-shaped agreement prices even with risk-neutral suppliers. But it is unclear why risk-neutral parties that trade repeatedly would adopt this complicated pricing rule rather than a constant price that yields the same expected profits—the only object relevant for risk-neutral parties. The fact that the U-shape appears only in repeated agreements, and not in spot transactions, further supports the profit-smoothing interpretation, as spot suppliers also have convex costs.

First, I consider unobserved supplier heterogeneity. For example, some suppliers may be more efficient (perhaps from better management as technology is largely standardized among relevant suppliers). Efficient suppliers might pass on cost savings to the buyer, reducing average prices, although efficiency alone does not explain why agreements feature more stable, rather than larger, quantities nor U-shaped prices. To investigate this hypothesis, I reproduce comparisons between agreement and spot suppliers, but restricted to spot suppliers with volumes comparable to agreement supplier volumes ("placebo suppliers" as in Appendix A.2.) because explanations related to unobserved heterogeneity typically imply that volume, rather than transaction frequency, should shape trade. Instead, Online Appendix OA1.2

[12] 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. Therefore, contracts could not specify an incentive-compatible spread of payments over time. The incompleteness of insurance in my empirical setting is consistent with a long literature on informal insurance in agreements with limited commitment (Ligon et al., 2002).

shows that differences in quantities and prices persist.

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. The subsequent empirical analysis estimates supplier-specific parameters that rationalize observed agreements. Therefore, misclassifying spot suppliers as agreement suppliers is much more problematic than the opposite, which further motivates my conservative approach.

This interpretation raises two questions addressed in Online Appendices OA1.4 and OA1.5. First, why does the buyer not resolve hold-up through supplier competition on prices *before* sampling (possibly with quantity-contingent prices given cost convexity)? Second, why does the buyer not source only through agreements to avoid high spot prices?

2.5 Context: Vertical Integration at a Large Buyer

Figure 5: Integrated Supplier Constructed in June 2018



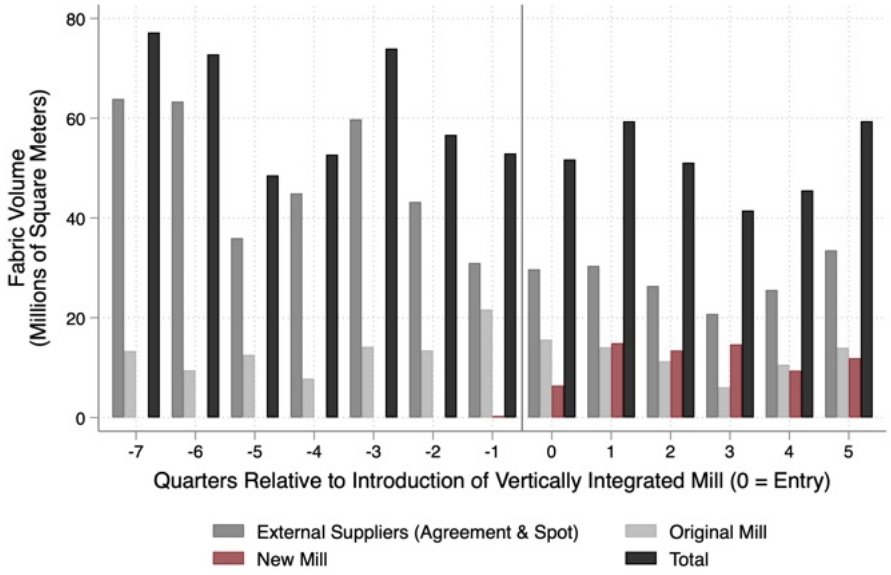
The buyer builds a new mill in the middle of the sample period. The buyer reports that high fabric prices motivate integration, consistent with spot suppliers holding up the buyer.

(Cost data from the integrated supplier suggest that spot markups are 7% over average variable costs for a supplier at 100% capacity utilization.) Figure 5 shows the new mill. The organization of the mill emphasizes the buyer's sophistication (and scale), contrasting with the messy, disorganized Indian garment manufacturers depicted in Bloom et al. (2012).

2.6 Context: Overall Sourcing

Figure 6 shows the buyer's overall volumes by quarter and supplier type. The large size of the buyer is apparent; on average, it purchases ∼$180 million (∼230 million square meters of fabric) per year. For comparison, 90% of Indian garment manufacturers report revenues (which presumably exceed their fabric procurement costs) under $60,000.[13]

Figure 6: Fabric Sourcing Over Time



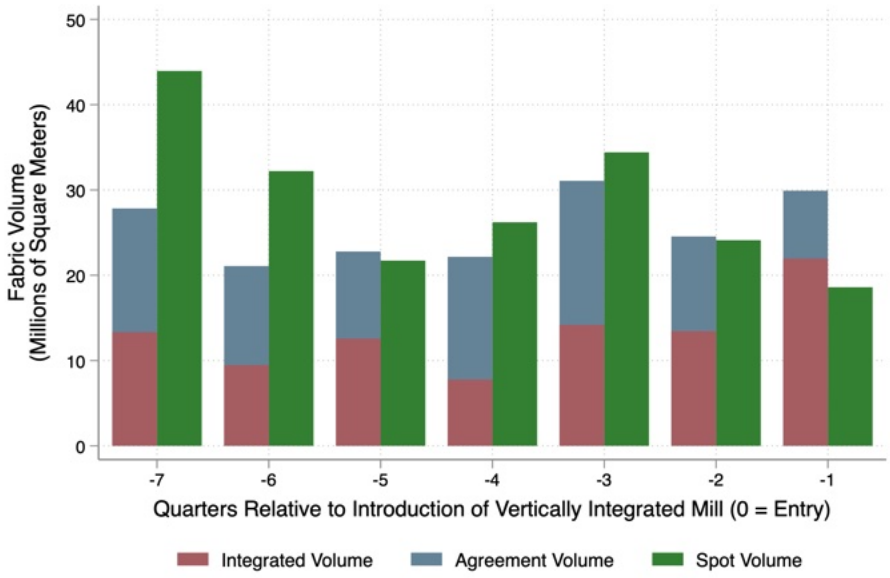
Note: Data from a panel constructed from the universe of fabric transactions by the buyer at the supplier-type-by-quarter level. Bars represent the quantity of fabric purchased during the quarter from each supplier type, summing across all suppliers of the type. Volume from external suppliers sums across all external suppliers, including all agreement and spot suppliers. Period 0 is the first full quarter of operation of the new mill (initial production began at the end of June, which is the end of period -1).

Three features make this setting particularly well suited to analyze the price effects of

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

insurance-like agreement terms and replacement threats from in-house supply. First, external suppliers remain empirically central post-integration, still providing most fabric. Therefore, external supplier prices remain economically relevant throughout the sample period. Second, the increase in vertical integration is sufficiently large to make replacement threats credible: the new mill produces ∼20% of the buyer's fabric, while its largest external supplier provides only ∼7%. The buyer also prominently displays the new integrated supplier on its website. 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.[14]

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



Note: Data from a panel constructed from the universe of fabric transactions by the buyer at the supplier-group-by-quarter level. Bars represent the quantity of fabric purchased during the quarter for *i)* integrated and agreement suppliers and *ii)* spot suppliers, summing across all suppliers in the group. Only quarters before the new mill is fully active are shown.

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 (the coefficient of variation for their aggregate volume is 0.155), while spot supply has much higher volatility (coefficient of

[14]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 4.2).

variation is 0.303). Consistent with this result, spot volumes strongly correlate with total buyer demand ($R^2 = 0.82$), while agreement and integrated volumes show no meaningful response to total demand ($R^2 = 0.09$).[15] These results reflect the buyer using agreement and integrated suppliers for *inframarginal* sourcing, while spot suppliers fulfill *marginal* orders.[16]

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 integrated production to fulfill additional orders. 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, which limits the buyer's scope to add new agreement suppliers (Macchiavello and Morjaria, 2015)).

3 Theoretical Model of the Vertical Contract

I develop a theoretical model to formalize how prices reflect the two forms of buyer power: buyer power over surplus creation from insurance-like agreement terms and buyer power from surplus capture through replacement threats. Two features distinguish my setting from both canonical Nash-in-Nash bargaining models and extensions that incorporate replacement-threats: replacement comes from in-house rather than excluded external supply, and spot suppliers fulfill marginal volumes. These features generate a key result: replacement threats reduce agreement prices only in agreements with suppliers who are sufficiently risk averse that the agreement creates enough surplus to survive the threat.

Recall that when agreements include insurance-like terms that create surplus beyond the input transfer, bargaining primitives have distinct economic meanings that map to the two forms of buyer power. The buyer bargaining parameter determines how much incremental surplus created by insurance-like terms accrues to the buyer. Therefore, it governs buyer

[15] Furthermore, a regression of log spot volumes on log overall volumes shows that spot volumes increase by 1.5% when overall volume increases by 1%, consistent with the buyer expanding spot sourcing to meet greater fabric needs. As spot volumes are only a fraction of total volume, they must increase by a larger percentage to deliver the same absolute increase in volume. Conversely, agreement volume is essentially unresponsive, increasing by only 0.2%.

[16] Online Appendix Figure OA1.9 shows that these results are robust to analyzing the post-integration quarters. I analyze post-integration quarters separately because the new in-house fabric mill mechanically increases integrated volumes. Therefore, even if the buyer continues to rely on spot suppliers for marginal volumes post-integration, combining pre- and post-integration quarters would risk conflating this mechanical capacity shock with a change in sourcing strategy.

power over surplus creation: a higher buyer bargaining parameter increases the buyer's return to incorporating surplus-generating terms into the agreement. When replacement threats improve the buyer's outside option, they shift rents to the buyer without increasing agreement surplus. Therefore, outside options shape buyer power from surplus capture through replacement threats. This mapping does not require that bargaining primitives affect prices only through these channels. Generally, when quantities are endogenous bargaining outcomes, bargaining primitives affect prices by changing the quantity traded and, in turn, gross surplus to divide. Importantly, Online Appendix OA1.6 relates each primitive to the resulting markdown wedge to illustrate that the connection between primitives and the distinct sources of buyer power remains even when primitives affect prices through quantity-driven surplus changes.

My empirical setting enables a clean separation between the two forms of buyer power because agreement quantities are inframarginal, with marginal quantities sourced from spot suppliers. Therefore, in-house supply need not reduce purchases from agreement suppliers, nor affect revenue from end clients associated with agreement quantities. It follows that gross surplus from agreement quantities—the gap between end-client revenue and total production costs—is unaffected by either bargaining primitive. Instead, replacement threats reduce prices on quantities the buyer purchases regardless of integration, meaning the buyer receives a higher share of unchanged gross surplus. Therefore, surplus capture through replacement threats is pure rent extraction. Analogously, because agreement quantity targets are fixed, the buyer bargaining parameter affects prices through the provision and distribution of surplus from insurance-like terms rather than changes in gross surplus.

The model embeds empirically observed agreement terms in a cooperative bargaining framework that facilitates decomposing prices into the two forms of buyer power. It links surplus created by quantity assurance and U-shaped pricing to supplier risk aversion, the buyer's retained share of that surplus to the buyer bargaining parameter, and surplus capture through replacement threats to the buyer's outside option. To keep this mapping parsimonious and transparent, I use a static model that I interpret as bargaining over long-run surplus in established agreements (with relational contracts sustaining the cooperative agreement). Appendix A.3.1 shows that the key results are robust to incorporating dynamic incentive compatibility constraints, following Macchiavello and Morjaria (2023).

3.1 Model Setup

The buyer and supplier bargain over a pricing function that maps realized agreement quantities into prices. 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 expected procurement-cost savings from purchasing fabric through the agreement rather than using the buyer's outside option. This focus on procurement costs reflects the timing of garment production: the buyer and end client agree on both garment quantities and prices *before* fabric procurement occurs. Therefore for a given garment, end-client revenue and non-fabric production costs are fixed and only fabric input costs vary with input sourcing.[17] I model the buyer as risk neutral given its large size and sophistication, so expected procurement spending enters linearly. (The buyer also mitigates its main risk—delayed fabric delivery, leading to expensive air transport—through financial penalties to suppliers when fabric arrives late.) Formally, the agreement payoff is $-\mathbb{E}_{q^A}[p^A(q^A)q^A]$; because O_B is also written as a payoff—the negative of procurement costs outside the agreement—buyer surplus is $-\mathbb{E}_{q^A}[p^A(q^A)q^A] - O_B$. Equivalently, buyer surplus is outside-option procurement costs minus agreement procurement costs.

Supplier surplus is expected utility from agreement profits less the value of the outside option (O_S). I allow suppliers to value smoother profits, motivated by financial frictions: firms must make recurring payments (*e.g.,* payroll), but borrowing or saving to smooth cash flows may be difficult or costly. Separately, firm-level risk aversion may reflect owners' preferences, as in owner-managed suppliers whose profits are their owners' primary income source. For brevity, I refer to these preferences as risk aversion and use constant absolute risk aversion (CARA) utility.[18]

Two insurance-like agreement terms create surplus, thereby shaping agreement prices. First, both quantity assurance and U-shaped prices smooth supplier profits, increasing the value of the agreement for risk-averse suppliers. Second, quantity assurance also creates

[17] The model abstracts from marginal insurance provision costs because agreement quantities are inframarginal: the buyer can provide quantity assurance by reallocating sourcing across suppliers, and U-shaped prices need not generate additional sourcing costs. Allowing nondecreasing marginal provision costs would strengthen the role of buyer power over surplus creation because a higher buyer bargaining parameter increases the buyer's marginal return to surplus-generating agreement terms.

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

surplus through cost reduction. With convex costs (see Figure 10),[19] Jensen's inequality implies that reducing quantity variance lowers average costs because high costs at large quantities are not fully offset by low costs at small quantities. Therefore, agreements create surplus even for risk-neutral suppliers by reducing costs.

Bargaining parameters determine surplus shares. The buyer bargaining parameter is α_B. Higher values indicate that the buyer receives a larger share of agreement surplus. As the buyer and supplier split the surplus, 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 the supplier can switch to a new agreement partner in the short-run.

I model spot prices (p^S) as constant, motivated empirically by flat spot prices (Figure 4).[20] Even though the buyer is large relative to other garment exporters (and the typical Indian firm), it is small relative to the Indian fabric market: its market share is less than 0.5%.[21,22] The decomposition of agreement discounts uses p^S as an empirical benchmark both to help interpret the size of the discount and because it constitutes the price the buyer would pay absent the agreement. If the buyer truly takes spot prices as given and chooses quantities optimally, then p^S equates to the buyer's marginal value of the input (perhaps only approximately with lumpy orders). The agreement discount is then not only a comparison but also, under this stronger interpretation, exactly the markdown,

After integration, the buyer's outside option shifts to in-house supply. I model the integrated supplier using the same productive technology as external suppliers (*i.e.,* the same cost function) because the buyer believes that all relevant suppliers employ (near-)frontier technology. Morton (2023) empirically documents strong substitutability between integrated

[19] Marginal costs can increase because machine inputs are fixed in the short run and scaling up production can generate costly congestion. Additionally, labor costs could increase dramatically at larger quantities—hiring and firing frictions prevent firms from easily adjusting employment and the Factories Act of 1948 mandates double wages for overtime work.

[20] Incorporating buyer market power in the spot market would increase surplus creation through additional buyer gains from trade because agreements would reduce spot prices. Buyer power in the spot market does not affect buyer power from surplus capture because the spot market does not impact the buyer's gains from trade when in-house supply serves as the buyer's outside option.

[21] The buyer uses ∼230 million square meters of fabric per year, whereas India produced 71 billion square meters of fabric in 2019. India is the sixth largest fabric producer globally (India Brand Equity Foundation).

[22] Flat spot prices are consistent with the hold-up interpretation in Section 2.2: once the garment is selected by the end client, the supplier from sampling is the only viable option. Accordingly, it raises prices until the buyer's participation constraint binds. Because this constraint does not vary with supplier capacity utilization, hold-up could generate flat spot prices.

and agreement suppliers, making displacement by the integrated mill a credible threat to agreement suppliers given the visible sunk cost of mill construction. Results are robust to reasonable cost differences between internal and external suppliers (see Online Appendix OA1.7).[23]

Formally, the outside options are:

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

The quantities in the bargaining problem reflect stochastic realizations of end-client demand rather than endogenous buyer or supplier choices in the fabric market. Because the end client and buyer bargain over garment quantities and prices before the buyer purchases fabric, realized fabric demand is determined by client orders rather than by the buyer's fabric-sourcing decision. Agreements specify the *ex ante distribution* of quantities: quantity assurance reduces the variance of realized orders without necessarily increasing average volume. Formally, I model spot quantities as a mean-preserving spread of agreement quantities, with both distributions centered at 100 percent capacity utilization, the supplier's optimal utilization rate. Agreement quantities do not respond to spot prices directly because agreements are inframarginal (and provide quantity assurance). Instead, higher spot prices increase the buyer's marginal input costs, as spot suppliers fulfill marginal orders, thereby reducing the buyer's total fabric demanded.

Appendix A.4 details the model parameterization and solution. I choose parameters and functional forms to illustrate how surplus creation and capture affect prices. Section 4 builds and estimates a structural model closely tied to this framework for quantification.

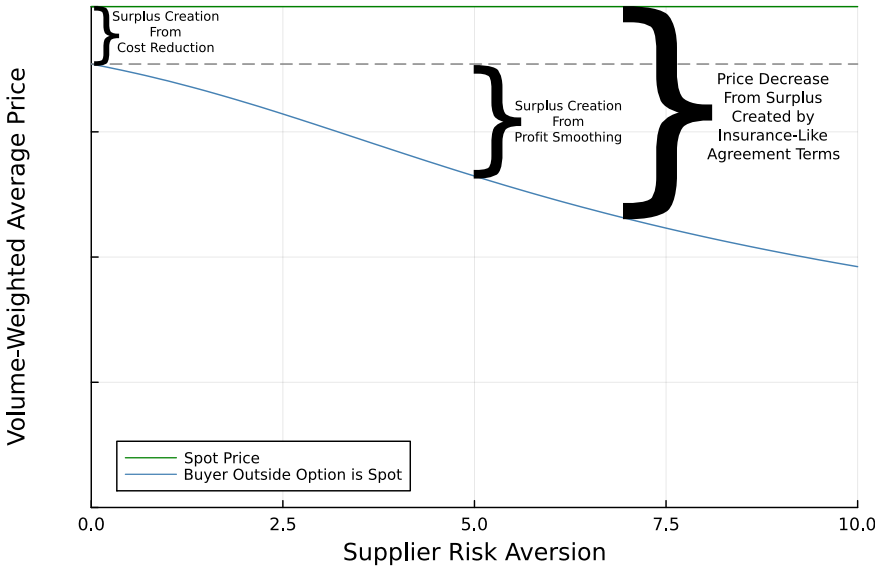
3.2 Price Effects of Buyer Power Over Surplus Creation

The price effects of buyer power over surplus creation depend on two objects: how much surplus agreement terms create (beyond the input transfer) and how much of that surplus the buyer retains. Created surplus increases in supplier risk aversion because risk-averse suppliers value profit smoothing more. A higher buyer bargaining parameter lowers prices

[23] Large cost heterogeneity seems implausible: if integrated supply were much more efficient, the buyer would integrate aggressively; if external suppliers were much more efficient, observed integration would be difficult to rationalize.

by increasing the buyer's retained share of surplus created. Therefore, agreement prices fall when either primitive increases. Because changes in supplier risk aversion and the buyer bargaining parameter do not affect supplier production costs or buyer revenue from end clients, resulting price reductions reflect buyer power over surplus creation.

Figure 8: Price Effects of Surplus Creation from Insurance-Like Agreement Terms



Note: Costs represent volume-weighted average procurement costs that reflect prices from solving the model as parameterized in Appendix A.4. Spot prices are flat because there is no bargaining on the spot market; therefore, supplier risk aversion does not affect spot prices. The price decrease from surplus created by insurance-like agreement terms is the discount relative to spot prices in the agreement. I decompose this discount from surplus creation. First, the discount from surplus created by cost reduction is the discount in an agreement with a risk-neutral supplier (*i.e.,* supplier risk aversion is 0.0). Second, the additional discount associated with increased supplier risk aversion is from surplus created by profit smoothing.

First, Figure 8 illustrates that the discount from buyer power over surplus creation increases in supplier risk aversion, holding all other parameters fixed. The discount is the difference between the spot price and the buyer's volume-weighted expected price. The discount with a risk-neutral supplier (*i.e.,* when the CARA coefficient is 0) reflects surplus generated only through cost savings from quantity assurance. As supplier risk aversion increases, quantity assurance and U-shaped prices create additional surplus by smoothing profits, so the discount grows.

Second, Online Appendix OA1.9 shows the analogous exercise for the buyer bargaining

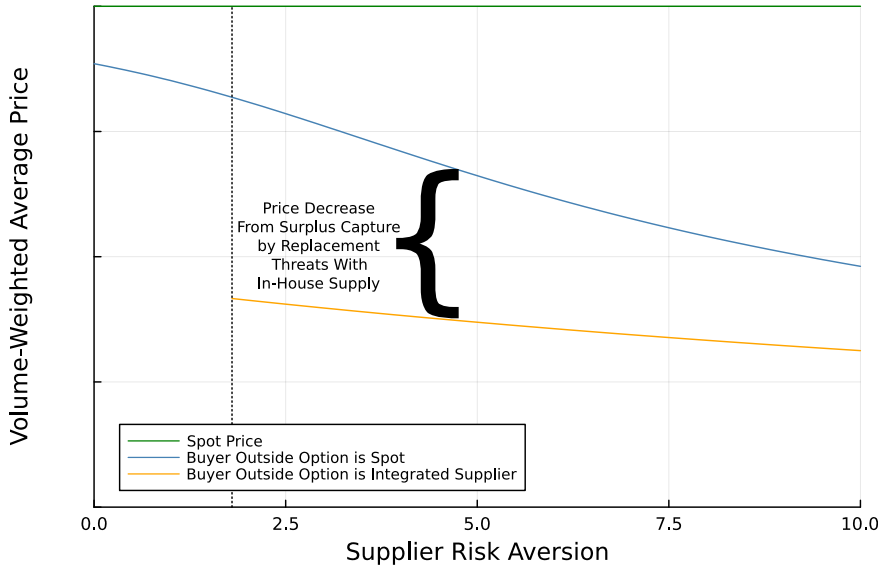
parameter. Agreement prices fall as the buyer bargaining parameter increases because the buyer retains a larger share of the surplus created by agreement terms.

3.3 Price Effects of Buyer Power From Surplus Capture

Figure 9 illustrates the price effects of buyer power from surplus capture through replacement threats from in-house production. These price effects operate by shifting the buyer's outside option to its integrated supplier. The model reproduces canonical replacement-threat logic when agreements survive the outside option change: a stronger buyer outside option reduces prices. Because this price reduction occurs without changing supplier production costs or buyer revenue from end clients, it reflects buyer power from surplus capture through replacement threats.

However, my empirical setting—specifically, that spot suppliers fulfill marginal volumes—limits when the canonical logic applies. Below a cutoff level of supplier risk aversion, insurance-like agreement terms create too little surplus for the agreement to survive the in-house replacement threat. If the agreement breaks, the buyer loses the agreement discount and must replace inframarginal agreement volume through the higher-priced spot market: quantity assurance limits expansion by other agreement suppliers, while capacity constraints limit expansion of in-house production. Therefore, replacement threats do not reduce agreement prices when supplier risk aversion falls below the cutoff. When supplier risk aversion exceeds the cutoff, the agreement creates enough surplus to survive the stronger outside option, and replacement threats reduce prices by shifting rents toward the buyer.

Figure 9: Price Effects of Surplus Capture through Replacement Threats



Note: Costs represent volume-weighted average procurement costs that reflect prices from solving the model as parameterized in Appendix A.4. Spot prices are flat because there is no bargaining on the spot market; therefore, supplier risk aversion does not affect spot prices. The price decrease from surplus capture is the additional discount when the buyer's outside option becomes its integrated supplier over the discount in the agreement purely from surplus creation. The vertical dotted line depicts the cutoff value of risk aversion below which the agreement breaks down in response to a replacement threat (*i.e.,* there is no surplus in the agreement for both the buyer and the supplier).

The cutoff result should not be interpreted as the buyer literally choosing a different Nash outside option for each supplier. In Nash bargaining, outside options are primitives: they represent disagreement payoffs, not endogenous choices. Instead, I establish that the cutoff emerges as the equilibrium outcome of two richer bargaining environments in which the buyer can take pre-bargaining actions—making a low-price replacement-threat offer or committing to reserve residual in-house capacity—that make in-house replacement credible. In equilibrium, the buyer takes such actions only when the supplier values agreement terms beyond the input transfer enough for the agreement to survive the threat. Generally, terms beyond the input could refer to insurance-like terms or trade credit.

The first extension, in Appendix A.5, introduces a signaling game in which the buyer privately observes a forecast of residual in-house capacity and can make a low-price replacement threat. The supplier may accept or reject the threat. Rejection reveals whether residual in-house production is capacity constrained, after which bargaining uses the buyer's true outside option. Rejection disrupts procurement, so the supplier provides compensation to the buyer.

The resulting equilibrium features a supplier-valuation cutoff: accepted replacement threats reduce prices only when the supplier values the agreement sufficiently highly. The second extension, in Appendix A.6, considers a game where the buyer commits to reserve residual in-house capacity before cooperative bargaining, in the spirit of capacity-commitment games (Dixit, 1980).[24] Both extensions generate a cutoff result: the supplier accepts the replacement threat, or the buyer reserves capacity, only when the supplier's valuation is high enough for the agreement to survive the resulting threat. Therefore, replacement threats reduce prices only when supplier valuation of agreement terms exceeds the cutoff.

Online Appendix OA1.10 highlights theoretical and empirical issues with an alternative model of replacement threats where vertical integration enables the buyer to make a take-it-or-leave-it offer to extract surplus.

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 buyer power over surplus creation and buyer power from surplus capture. Estimation recovers the supplier risk aversion and buyer bargaining parameters that rationalize observed agreement prices, empirically linking the primitives to buyer power.

4.1 Estimation Approach: 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. Estimation proceeds in two steps. First, an inner optimization recovers the risk aversion and bargaining parameters that best fit prices, given the tuning parameter. Second, the outer optimization selects a tuning parameter. The tuning parameter penalizes deviations from risk neutrality, both to rule out implausibly large risk aversion estimates and to limit overfitting given the small number of observations per supplier.

The inner optimization maps candidate bargaining primitives into predicted agreement prices. For each agreement, a candidate parameter vector $\mathbf{\Omega}_s$—consisting of supplier risk aversion and the buyer bargaining parameter—implies an agreement pricing function $\hat{p}^A(q^A; \mathbf{\Omega}_s)$

[24] Partial vertical integration is a costly, largely irreversible, and visible investment. Buyer commitment also appears elsewhere in the agreement, such as when the buyer pays agreement prices above spot prices in some periods.

that maximizes the Nash product in (1) over a quadratic pricing function.[25] The inner optimization chooses parameters to minimize the RMSE between observed pre-integration agreement prices and prices implied by $\hat{p}^A(q^A; \mathbf{\Omega}_s)$, plus a LASSO-style penalty on risk aversion. Formally, the inner optimization solves:

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

S denotes the set of suppliers, $\{q^A_{s,t}\}$ are the observed quantities for agreement supplier s, 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 θ_s is the supplier risk aversion parameter. I estimate the model at the quarterly level for two reasons: meetings with agreement suppliers occur quarterly, and reduced-form patterns are clearest at that frequency.

The outer optimization finds the tuning parameter that maximizes out-of-sample fit in post-integration quarters, akin to a holdout sample in regularization. The tuning parameter affects post-integration fit through the pricing function: a larger value pushes supplier risk aversion toward zero (*i.e.,* risk neutrality). Formally, the outer optimization solves:

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

The outer optimization uses post-integration prices because the construction of the new mill provides uniquely informative variation for distinguishing parameterizations that fit pre-integration prices similarly but imply different responses to the buyer's changed replacement

[25] I estimate $\hat{p}^A(q^A; \mathbf{\Omega}_s)$ with a second-order Taylor polynomial around full capacity: $\hat{p}^A(q^A; \mathbf{\Omega}_s) = \beta_{0,s} + \beta_{1,s} q^A + \beta_{2,s}(q^A)^2$. This functional form maintains flexibility while reflecting the parsimonious U-shaped schedules observed in the data. It excludes unrestricted complete-insurance schedules that are inconsistent with the informal agreements studied, as discussed in Section 2.4.

opportunity.

4.2 Estimation Approach: Outside Estimation

Table 1 summarizes the additional moments for model estimation.

Table 1: Additional Moments for Estimating the Structural Model

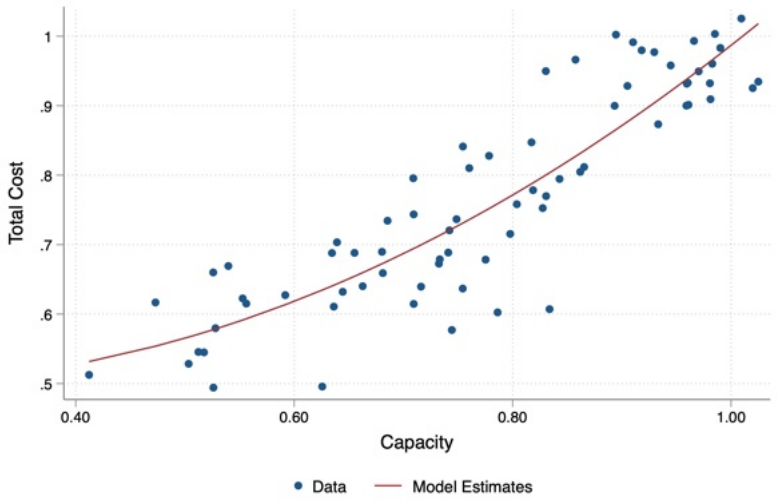
Parameter(s) in Model	Moment(s) to Match	Data
Cost Function Convexity	Actual Costs and Capacity	Cost and Capacity Data from Integrated Suppliers
Mean Capacity	Average Log Quantity	Transaction Data
$\frac{\mathbb{V}[q^S]}{\mathbb{V}[q^A]}$	Variance in Quantities for Spot Suppliers Compared to Variance in Quantities for Agreement Suppliers	Transaction Data
p^S	Ratio of Prices from Spot Suppliers to Estimates from Integrated Supplier	Transaction Data

I discuss estimation of the cost function below; Appendix A.8 provides details for remaining outside estimates in Table 1.

I estimate the cost function using panel data at the division level on in-house capacities and costs from April 2019 to March 2020, where division corresponds to a broad fabric group (*e.g.,* knits and striped fabrics). The cost function, $c(q) = \phi_0 + \phi_1 q + \phi_2 q^2$, allows for convexity when $\phi_2 > 0$ without mechanically imposing it. Because procurement staff indicate that agreement and integrated suppliers use similar technology, costs should be comparable.

Identification leverages plausibly exogenous capacity shocks that arise because sampling fixes the supplier before final order quantities are realized. Therefore, realized quantities reflect stochastic end-client demand rather than buyer choice; Section 4.3 discusses this exogeneity assumption and provides supporting evidence. When deviations from the quantity target are exogenous, regressions directly estimate how capacity affects costs.

Figure 10: Cost Convexity of Integrated Supply



Note: Data from a division-by-month panel of proprietary cost and capacity data from both integrated suppliers, where a division is a distinct production process without meaningful substitutability across divisions (*e.g.,* dyeing and weaving). Costs are normalized total costs, constructed by dividing unit costs by the division's mean unit cost (for comparability across divisions) and multiplying by capacity utilization. Capacity utilization for that division-month is reported directly in the data as a percentage of capacity used.

The estimated cost function in Figure 10 fits the data well: the R^2 is .795 and the RMSE is .07. Estimates show cost convexity: the coefficient on q^2 is .7815 (p-value = 0.02).[26]

4.3 Identification of Risk Aversion & Bargaining Parameters

Average agreement prices alone cannot separately identify supplier risk aversion from the buyer bargaining parameter because increasing either parameter individually reduces prices. When supplier risk aversion increases, the agreement creates more surplus. As the buyer retains some surplus created through a discount, higher risk aversion reduces prices. Similarly, a higher buyer bargaining parameter lowers prices by increasing the buyer's surplus share.

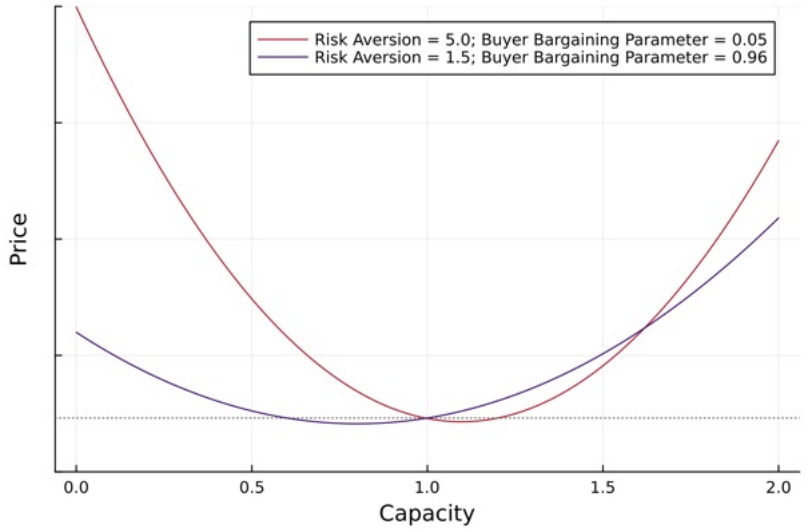
I develop a novel identification strategy that leverages two distinct sources of variation in the data. First, price curvature identifies supplier risk aversion. More risk-averse suppliers value price adjustments to capacity to smooth profit more, so the model predicts a more pronounced U-shaped pricing function. Second, variation in average agreement prices identifies the buyer bargaining parameter, as increasing the bargaining parameter reduces prices

[26] I express costs as percentages relative to mean costs within fabric group, allowing estimates to use cost data across fabric groups. The strong fit supports this pooled specification.

across capacity realizations.

With only one price-capacity tuple, it is impossible to separate supplier risk aversion from the buyer bargaining parameter. For any value of the 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 0.05 and a supplier risk aversion of 1.5 with a buyer bargaining parameter of 0.96 result in the same price when capacity utilization is 1.

Figure 11: Convex Pricing Function for Different Parameters



Note: Prices computed from a simulation of the model with parameters chosen to illustrate the identification argument. First, I compute the pricing function when supplier risk aversion is 5.0 and the buyer bargaining parameter is 0.05. Then, I shift the risk aversion to 1.5 and find the buyer bargaining parameter that results in the same price when capacity utilization is 1.0, which is 0.96. This figure highlights that multiple parameters can result in the same prices for any given price-capacity tuple; however, different parameters generate different prices at other capacity utilizations.

However, *additional* price-capacity tuples at *different* capacity realizations enable identification. Risk aversion primarily determines the *shape* of the pricing function: lower risk aversion makes prices flatter and less U-shaped. The buyer bargaining parameter largely shifts the average price *level*. Figures 12 and 13 illustrate this logic: reducing supplier risk aversion changes curvature, while reducing the buyer bargaining parameter raises prices across capacities.

Figure 12: Shift Supplier Risk Aversion

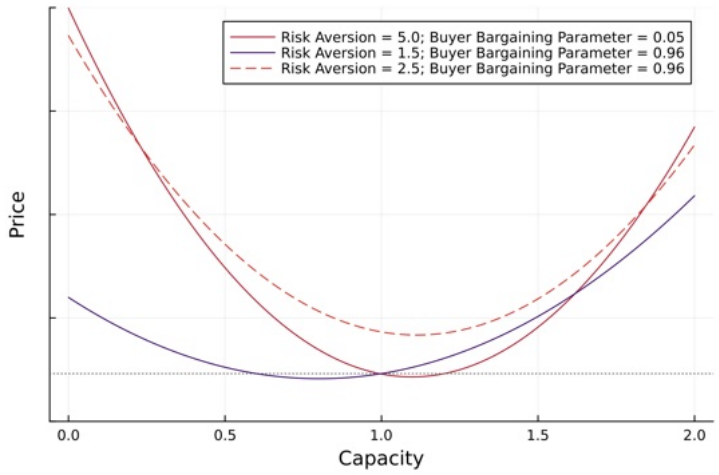
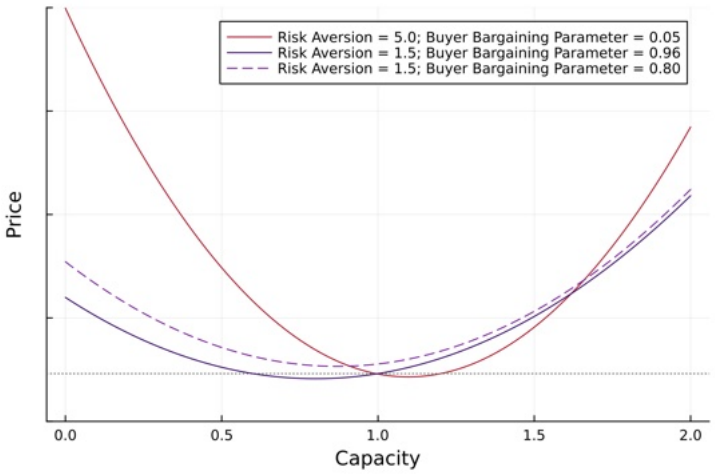


Figure 13: Shift Buyer Bargaining Parameter



Note: Prices computed from a simulation of the model with parameters chosen to illustrate the identification argument. Prices are shown for the same parameters as in Figure 11, and for an additional pricing function that reduces supplier risk aversion (Figure 12) or the buyer bargaining parameter (Figure 13).

To take this identification argument to the data, I leverage quasi-exogenous shocks in capacity utilization that reflect the institutional details of garment production. Because the supplier is chosen before the end client determines which garments to order and in what (lumpy) quantities, agreements reduce, but cannot eliminate, quantity variance. Deviations from targeted agreement volumes reflect, at least in part, end-client demand rather than strategic buyer or supplier choices. Importantly, quantity targeting *within* a quarter does *not* invalidate the identification argument. Instead, identification leverages residual quantity variation that remains after within-quarter quantity targeting.

Appendix A.9.1 provides supporting evidence that residual quantity variation is plausibly exogenous. First, supplier-fabric volumes respond strongly to aggregate fabric-period demand, including leave-out demand that excludes the supplier's own volume, but not to lagged supplier-fabric volumes. Therefore, quantity deviations are consistent with changes in end-client demand but not supplier-fabric targeting. Second, 63% of average agreement volume per quarter is of a fabric type not produced in the prior quarter.[27] Therefore, suppliers are not simply allocated the same fabric every quarter. Third, I find negligible evidence of dynamic quantity assurance across quarters: quarters following extreme capacities are not systematically closer to the target.

[27]Similarly, on average, 11.7 of the 16.9 fabrics produced per agreement supplier in a quarter were not made by that supplier in the prior quarter.

4.4 Risk Aversion and Bargaining Parameter Estimates

Table 2 (confidence intervals in Appendix A.11) shows that the estimated supplier risk aversion resembles estimates for LMIC coffee mills in Blouin and Macchiavello (2019): .0085 for the median agreement supplier versus .0068 for coffee mills.[28] These coffee mills—"large firms by developing country standards"—have ∼$3.5 million in annual sales and ∼$2 million in total assets. As these coffee mills are large firms by LMIC standards and similarly positioned as firms working in a commoditized export setting, the similarity of the risk aversion estimates seems reasonable. Although these results suggest that agreement suppliers are relatively large, the buyer is presumably much larger: it purchases ∼$180 million of fabric yearly, with revenue exceeding fabric expenses.[29] The estimated tuning parameter (λ) is only .002; therefore, low estimated supplier risk aversion is not primarily driven by regularization. Additionally, one high-leverage quarter drives the outlier risk-aversion estimate.

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: The table shows summary statistics for parameters estimated separately for each agreement (*i.e.,* one risk aversion and one buyer bargaining parameter for each supplier) from the structural model. Risk aversion parameters are CARA coefficients, where 0 represents risk neutrality. The buyer bargaining parameter is a Nash bargaining weight between 0 and 1 that represents the buyer's surplus share. The tuning parameter (λ) estimate is .002. I omit one agreement supplier from estimation because it is missing price data for one quarter. The correlation between supplier risk aversion and buyer bargaining parameters is 0.28.

Estimated bargaining parameters imply that the buyer typically retains slightly more than half of agreement surplus.[30] Together with the risk-aversion estimates, these estimates

[28] Blouin and Macchiavello (2019) and other studies estimate risk aversion using different utility functions. For comparability, I convert reported risk aversion parameters to CARA coefficients. Appendix Table A.4 shows the converted estimates.

[29] Low estimates may partly reflect that suppliers smooth demand across other buyers. Then, estimates measure residual preferences for smooth profits. Notably, seasonality limits cross-buyer smoothing through correlated buyer demand.

[30] Section A.14 provides evidence of one microfoundation for heterogeneity in buyer bargaining parameters.

suggest that agreement suppliers are relatively large, sophisticated firms, consistent with the capabilities needed to maintain the complex contractual agreement studied (Macchiavello and Morjaria, 2023).

Online Appendix OA1.12 provides two model validation exercises. First, the model matches an untargeted moment. Second, the model outperforms two alternatives in post-integration quarters; this comparison serves as a largely out-of-sample test because estimation primarily relies on pre-integration quarters. As the DiD directly estimates the price effects of surplus capture through replacement threats, DiD validation follows presentation of the analogous model estimate.

4.5 Decomposition: the Price Effects of Buyer Power

Table 3 decomposes agreement discounts into buyer power over surplus creation from insurance-like agreement terms and buyer power from surplus capture through replacement threats. The decomposition first calculates prices when the buyer's outside option is the spot market: buyer power over surplus creation generates the discount relative to spot prices. Buyer power from surplus capture is any additional discount when the buyer's outside option switches to in-house production. If supplier risk aversion falls below the cutoff above which the agreement survives replacement threats, there is no buyer power from surplus capture.

Table 3: Discounts from Agreement Suppliers

Supplier Type	Total Discount	Discount Decomposition: *Surplus Creation from Insurance-Like Agreement Terms*		*Surplus Capture*
		From Cost Reduction	*From Profit Smoothing*	*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: All discounts are the difference between the relevant model-implied price and the spot price, expressed as a percentage of the spot price. The total discount considers the volume-weighted average agreement price implied by the model. In the decomposition of the discount, the portion from surplus creation reflects the price when the buyer's outside option is the spot market. The decomposition of the discount from surplus creation into the share from cost reduction is the discount when supplier risk aversion is shifted to 0.0 (*i.e.,* when the supplier is risk neutral), holding all else constant. Profit smoothing from insurance-like agreement terms generates the remaining discount from surplus creation. Figures 8 and 9 visualize the decomposition (for parameters chosen to illustrate the model rather than the estimated relationship types).

Estimates show that agreement discounts are entirely attributable to buyer power over surplus creation, without any surplus capture through replacement threats. Average agreement prices are 1.5% below spot prices, reducing annual fabric procurement costs by ∼$0.87 million. This discount is meaningful relative to spot markups near 7%, where the markup is calculated for a spot supplier at 100% capacity utilization. I further decompose discounts from buyer power over surplus creation into surplus from either cost reductions from quantity assurance or profit smoothing. Discounts for a risk-neutral version of each supplier, holding all else fixed, isolate the cost-reduction component; surplus from profit smoothing explains the remainder. I find that cost reductions and profit smoothing each account for approximately half of the total discount.

To assess robustness to the one outlier supplier with high risk aversion, I compute discounts using the *i)* median and *ii)* mean values of supplier risk aversion and the buyer bargaining parameter.[31] The central conclusion holds: all discounts arise entirely from surplus creation. However, the level and composition of discounts differ. Discounts decrease overall because prices are nonlinear in supplier risk aversion and the buyer bargaining parameter.[32] Reducing the influence of the high-risk-aversion supplier decreases the relative contribution of profit smoothing to the discount.[33]

[31] I consider two hypothetical suppliers, each with the average capacity variance of agreement suppliers. Calculating average discounts across agreement suppliers after replacing each supplier's risk aversion and bargaining parameter with the median or mean value yields nearly identical results (less than .0001 p.p. different).

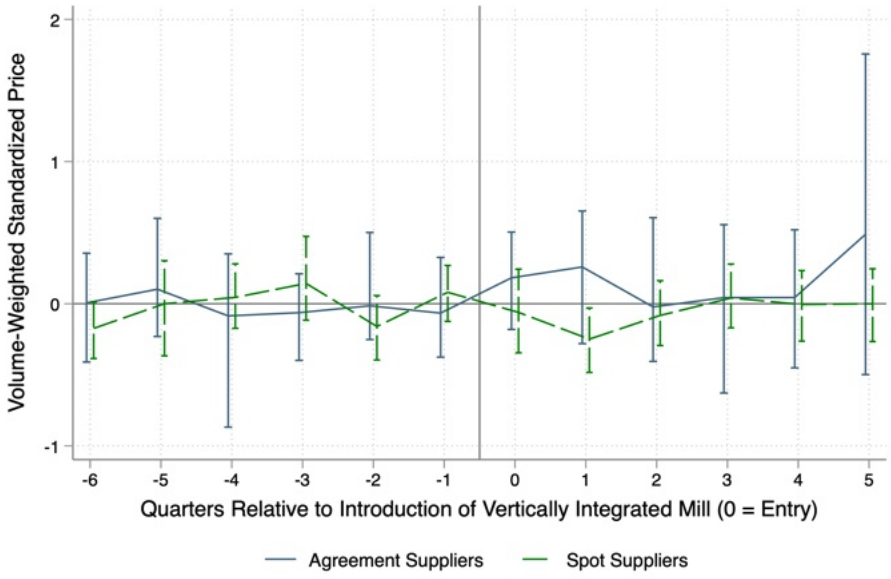
[32] This nonlinearity has multiple sources. Two central reasons are that *i)* the Nash product is multiplicative, not additive, and *ii)* higher parameter values increase both total surplus and the buyer's share of that surplus. Discounts are slightly higher with median than mean parameter values because the larger median buyer bargaining parameter more than offsets the lower median risk aversion.

[33] Given that identification leverages the shape of the agreement pricing function, Online Appendix Figure OA1.14 confirms that the U-shape is robust to removing the most risk-averse supplier.

4.6 Model Validation: Difference-in-Differences Estimates

I validate the structural model by comparing results with a DiD design that directly estimates the effects of buyer power from surplus capture through replacement threats with in-house production. The first difference is time: I compare quarters before versus after the buyer builds the new fabric mill that enables replacement threats. The second difference is exposure to the new mill, defined by pre-integration product overlap with the new integrated supplier. I classify suppliers producing fabric subsequently made by the new mill as treated because the buyer can more credibly replace them with in-house production; remaining suppliers are controls. Importantly, all comparisons are *within* supplier type (*i.e.,* agreement suppliers are compared to each other). DiD estimates incorporate a propensity-score adjustment (Abadie, 2005) that relaxes parallel trends to *conditional* parallel trends. Here, estimates condition on suppliers' exposure to trends in fabric demand that reflect evolving fashion tastes, as detailed in Appendix A.12.[34] The DiD validates the result using a distinct source of identifying variation: the mill opening and differences in supplier exposure to it, rather than the downstream demand shocks that identify parameters in the structural model.

Figure 14: DiD Estimates: Effects on Prices



[34] I confirm that propensity-score adjustment attains pre-treatment balance, which remains diagnostic because DiD outcome variables are not matched.

Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Event study estimates use a semiparametric difference-in-differences design (Abadie, 2005) detailed in Appendix *A*.12 and are estimated separately for agreement and spot suppliers. Suppliers are considered treated if they produce a fabric also made at the newly built internal fabric mill. Confidence intervals are bootstrapped to account for the two-stage estimation and clustered by supplier. Volume-weighted standardized prices are standardized by fabric and month (as described in Section 2.3.2).

The theoretical model generates sharp predictions for buyer power from surplus capture given supplier risk aversion estimates. Because even the most risk-averse agreement supplier has risk aversion below the cutoff where agreements survive replacement threats (see Appendix Figure A.6), the model implies no change in agreement prices. Figure 14 confirms that prices paid to agreement suppliers do not respond to integration. It also shows no concerning pretrends.

Because the buyer sources *marginal* quantities from spot suppliers, and end-client demand did not increase contemporaneously with the new mill, the model implies that vertical integration should reduce spot rather than agreement volumes. Appendix Figure A.9 supports this model implication empirically: agreement quantities do not change, while spot quantities decline. The significant decline in spot quantities also suggests that the DiD design has sufficient statistical power to detect economically meaningful effects.

Appendix Table A.7 shows average treatment effects pooled across periods.[35] Pooled results show no statistically significant effect on agreement prices or quantities; the point estimates are near zero, although not always precisely estimated. However, vertical integration causes a statistically significant and economically meaningful decrease in quantities for spot suppliers. The quantity decline averages .4 standard deviations—∼120,000 square meters of fabric per quarter. Online Appendix OA1.14.2 confirms robustness to the synthetic DiD approach from Arkhangelsky et al. (2021).

Online Appendix OA1.14.3 provides an additional validation exercise that compares DiD event study coefficients using actual prices with equivalent estimates using model-implied prices. Event study estimates using model-implied prices all fall within the 90% confidence intervals of estimates using actual prices.

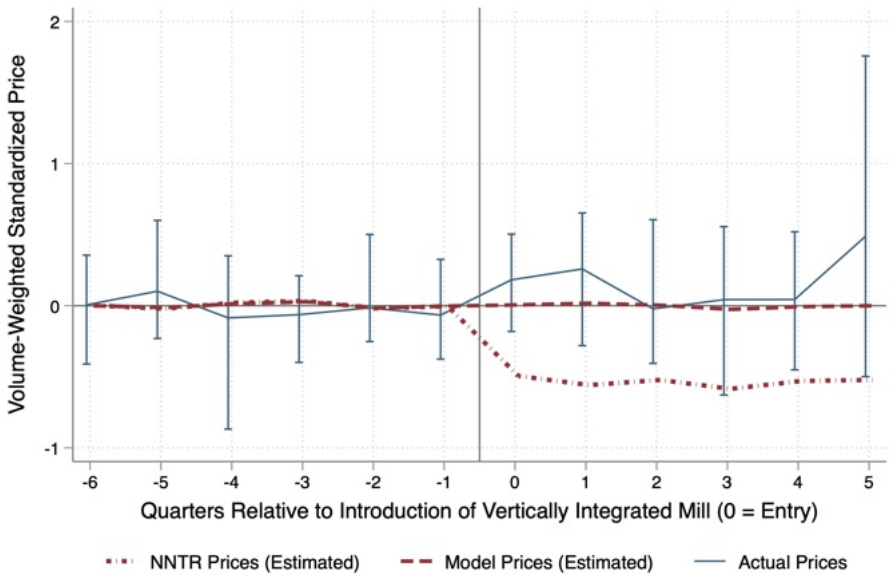
[35]Bootstrapped confidence intervals account for both clustering at the supplier level, as highlighted in Bertrand et al. (2004), and uncertainty in the propensity score.

4.7 DiD Estimates: Spot Sourcing Constrains Replacement Threats

To illustrate how spot markets constrain replacement threats, I compare my DiD results with prices implied by a canonical replacement-threat model: the Nash-in-Nash with Threat of Replacement (NNTR) model from Ho and Lee (2019). Because Ho and Lee (2019) study agreements between hospitals and health insurers in California, the NNTR bargaining environment differs from my setting in three key ways: my empirical setting features *i)* marginal sourcing from spot suppliers, *ii)* constraints on the formation of new agreements due to relational enforcement, and *iii)* replacement threats from in-house production. These differences make NNTR a useful benchmark for showing how the contracting environment shapes the effects of replacement threats. Appendix A.13 details my NNTR implementation.

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

Figure 15: DiD Estimates: Comparison with NNTR Model



Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Event study estimates use a semiparametric difference-in-differences design (Abadie, 2005) detailed in Appendix *A*.12. NNTR prices are estimated following Ho and Lee (2019) (see Appendix A.13). Suppliers are considered treated if they produce a fabric that is also made at the newly built integrated fabric mill. Confidence intervals are bootstrapped to account for the two-stage estimation and clustered by supplier. Volume-weighted standardized prices are standardized by fabric and month (see Section 2.3.2).

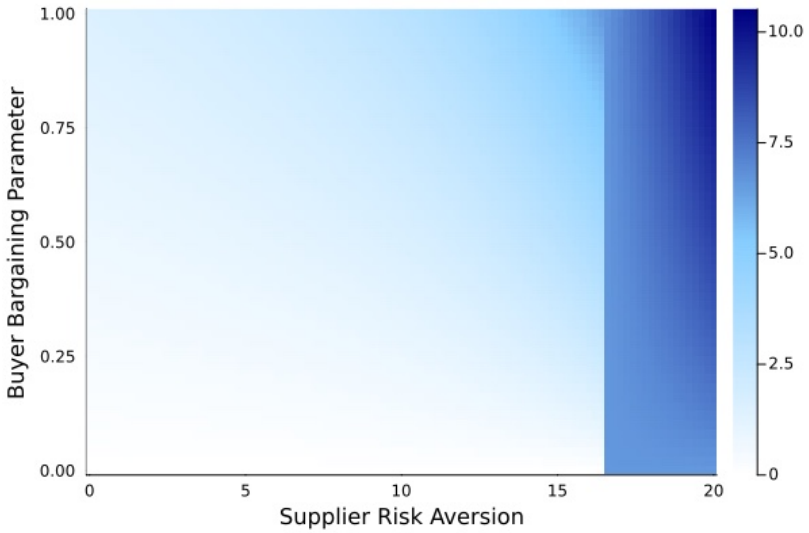
The DiD result that agreement prices do not respond to partial vertical integration is

inconsistent with an alternative bargaining model in which integration enables TIOLI offers because that model implies integration reduces prices (see Online Appendix OA1.10).

4.8 Surplus Creation & Capture for Highly Risk-Averse Suppliers

I use the model to characterize discounts for suppliers more risk averse than observed agreement suppliers. This exercise is motivated by small suppliers' policy relevance: smaller firms are likely more exposed to risk and the frictions that make smoothing profits challenging, and observed agreement suppliers are likely relatively large firms. Therefore, to study how buyer power—especially from surplus capture—affects such suppliers, I find both the cutoff level of supplier risk aversion above which replacement threats reduce prices and the price effects of such threats.

Figure 16: Agreement Discounts (% of Spot Price)



Note: The heatmap shows the model-implied discount in the agreement relative to the spot price, expressed as a percentage of the spot price, for different values of the supplier risk aversion and buyer bargaining parameters. The heatmap legend is to the right of the heatmap and units are percentages of spot prices. Model estimates consider the average capacity and quantity assurance across agreement suppliers.

To illustrate how buyer power affects smaller, more risk-exposed suppliers, Figure 16 plots model-implied agreement discounts, as a percent of spot prices, for higher levels of supplier risk aversion. This analysis reveals that the cutoff level of risk aversion above which replacement threats reduce prices is 16.6, exceeding estimates from other settings (12.9 in Gandelman and Hernandez-Murillo (2014)), but the same order of magnitude. This finding

emphasizes how the contracting environment constrains replacement threats. The buyer's discount increases meaningfully once supplier risk aversion surpasses the cutoff. For example, when the buyer bargaining parameter splits agreement surplus equally, 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 Indian goods and services tax for apparel in the sample period (5%)[36] and is large relative to markups in the spot market from hold-up (7.36%), where costs are computed for a supplier at 100% capacity utilization.

As the risk aversion threshold is not a structural parameter itself, the relevance of replacement threats varies across settings. 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).[37] LMIC smallholder farmers may satisfy these conditions and rely on bargained agreements (Macchiavello and Morjaria, 2020); they could be highly risk averse and have convex production costs as a key input is fixed (land).

Online Appendix OA1.15 incorporates relational enforcement for a conservative calibration of the discount factor, documenting negligible differences in estimated prices.[38] This result suggests that suppliers are at least as patient as the calibrated discount factor, which is based on microfinance interest rates. A high discount factor for agreement suppliers is consistent with low estimated risk aversion—both suggest agreement suppliers are large firms.

5 Counterfactual Policy Evaluation

I use the model to evaluate policies aimed at increasing prices paid to small suppliers. I focus on small suppliers because they account for most firms and employment in LMICs (Hsieh and Olken, 2014) and are central to distributional concerns that gains from exporting do not reach upstream producers; accordingly, myriad policy initiatives explicitly target small

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

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

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

firms.[39] Furthermore, limited access to borrowing and savings markets can make profit volatility especially costly for these firms; these frictions mean that such firms likely have stronger preferences for smooth profits (higher supplier risk aversion in the language of the model). This focus is especially relevant because discounts from both buyer power over surplus creation and buyer power from surplus capture increase with supplier risk aversion.

I consider two counterfactuals that reflect distinct policy responses to low supplier prices. First, I increase downstream buyer competition, motivated by concerns about buyer concentration. Second, I introduce insurance over spot-market profits, motivated by missing smoothing markets that make insurance-like agreement terms valuable, especially for highly risk-averse suppliers. Comparing these policies highlights both how they address different forms of buyer power and why their effects vary with supplier risk aversion.

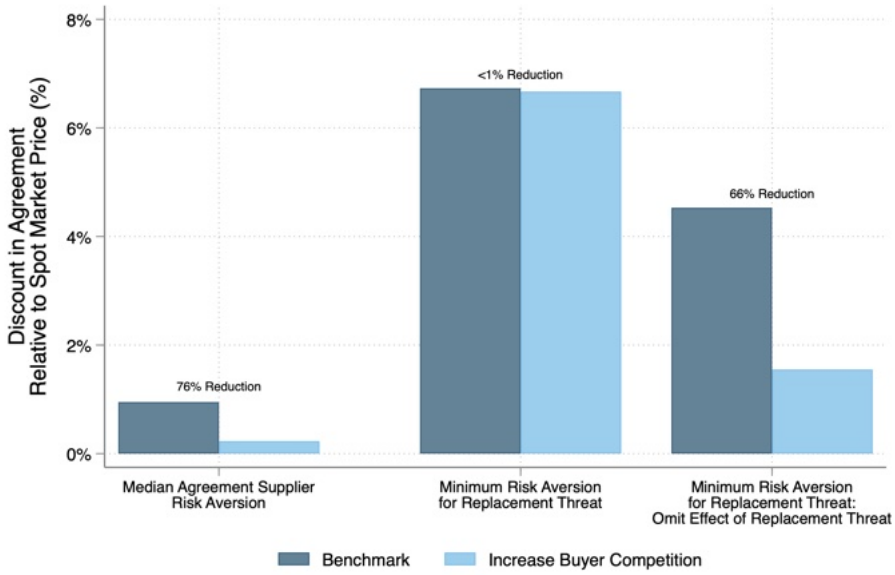
5.1 Increase Downstream Buyer Competition

First, I evaluate increased buyer competition, given concerns about buyer concentration in export-oriented supply chains where market structure is often "funnel-shaped"—few buyers for many suppliers. Policies can increase such competition by encouraging the development of additional exporting firms, such as by reducing red tape and other administrative costs, or facilitating connections to international end clients ("export promotion").

I model the connection between buyer competition in agreements and prices through the buyer bargaining parameter. Competition from additional buyers should improve suppliers' outside options in the long run, but not in the short run: suppliers cannot quickly establish new agreements because they are relationally enforced, and the trust underpinning relational contracts requires time and effort to develop (Cajal-Grossi et al., 2023; Brugues, 2025). Therefore, the counterfactual considers the buyer responding to competition by offering suppliers a larger share of surplus to maintain supplier loyalty. This interpretation mirrors efficiency wages, as both incentivize loyalty through a higher surplus share (Katz, 1986). Empirical patterns in Appendix A.14 support modeling buyer competition as a reduction in the buyer bargaining parameter: more valuable suppliers receive larger surplus shares (*i.e.,* have lower estimated buyer bargaining parameters). Specifically, I halve the buyer bargaining parameter.

[39] Examples include the World Bank's policy advice on financial tools for SMEs and the IMF's 2019 conference on SME financing. Such policies also exist in high-income countries, such as preferences for small business in public procurement and favorable credit access.

Figure 17: Downstream Buyer Competition and Agreement Prices



Note: Bars show the discount in the agreement relative to the spot market, expressed as a percentage of the spot price, at the benchmark (on the left) and with increased buyer competition (on the right) at two different levels of supplier risk aversion: the median of agreement suppliers and the cutoff in risk aversion above which replacement threats reduce prices. The benchmark uses the median estimated buyer bargaining parameter among agreement suppliers, while the scenario with increased buyer competition halves the buyer bargaining parameter. Model estimates consider the capacity and quantity assurance for the median agreement supplier.

Figure 17 illustrates that the effects of increased buyer competition vary sharply with supplier risk aversion. For a supplier with the median estimated risk aversion, increased competition meaningfully reduces the discount: it falls by 76%, from 0.95% to 0.23%. However, when supplier risk aversion exceeds the cutoff where replacement threats reduce prices, the discount barely changes, falling from 6.73% to 6.67%. Omitting surplus capture through replacement threats *overestimates* the benefits of increased competition for highly risk-averse suppliers; without it, increased competition reduces the discount by 66%—4.5% to 1.5%.

Intuitively, increased buyer competition has negligible effects on prices paid to highly risk-averse suppliers because, once replacement threats from in-house production discipline bargaining, the agreement generates limited gains from trade relative to the buyer's outside option. Prices are then governed by the buyer's participation constraint, leaving little scope for price adjustments. This logic explains why abstracting from surplus capture can generate misguided policy conclusions for small risk-averse firms. It further explains why increased

competition meaningfully raises prices for suppliers with low risk aversion: when replacement threats do not discipline prices, discounts reflect the buyer's share of surplus created inside the agreement. Then, changes in surplus sharing pass through to prices.

5.2 Insure Spot-Market Profits

The importance of missing markets in LMIC settings, together with the limited effect of increased buyer competition for highly risk-averse suppliers, motivates a second counterfactual: access to insurance over spot-market profits. With actuarially fair insurance, suppliers fully insure spot-market profits, making utility constant across states. Formally, each supplier's outside option becomes $O_S = U(\mathbb{E}_{q^S}[p^S q^S - C(q^S)])$. Insurance eliminates profit risk in the spot market, effectively removing the profit-smoothing component of agreement surplus. Agreement surplus then only reflects cost reductions generated by quantity assurance. This counterfactual highlights the benefits of many policies that help firms smooth profits without relying on the agreement, such as improving access to financial markets so firms can borrow and save to buffer against variability, even though moral-hazard concerns likely limit the viability of profit insurance as an actual policy.[40]

Figure 18: Agreement Discounts Without Insurance

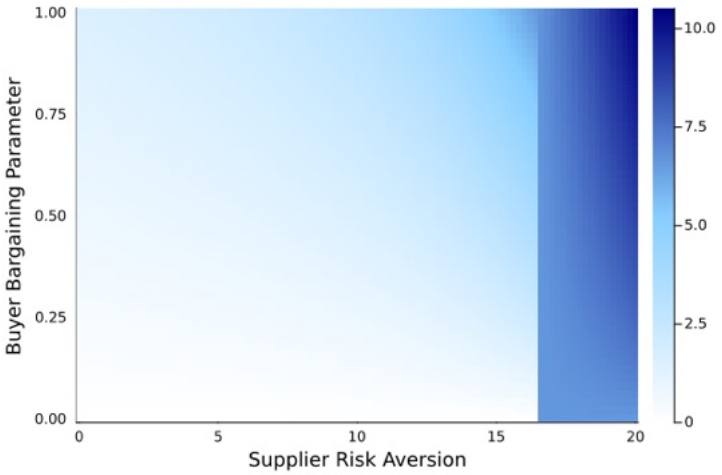
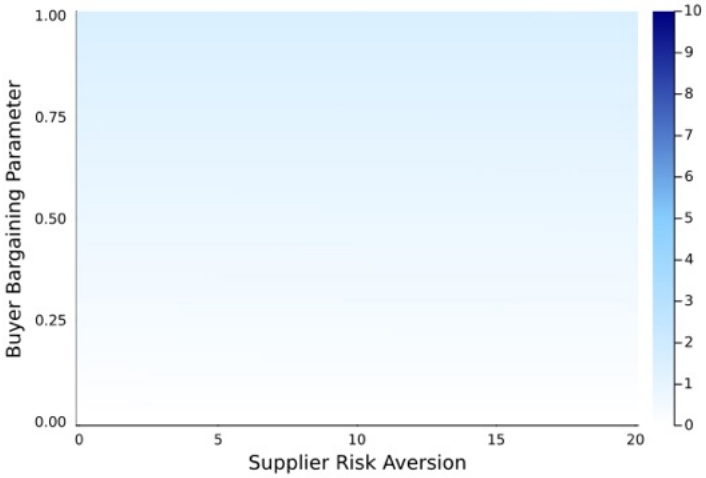


Figure 19: Agreement Discounts With Insurance



[40] Policies that provide some insurance are common. For example, price floors for agricultural goods insure producers against adverse prices. Similarly, subsidies or regulations that expand access to credit and savings for small firms can help them self-insure. Policies can also facilitate access to derivatives and other financial products that directly hedge risks, though financial literacy and transaction costs may limit take-up.

Note: The heatmap shows the model-implied discount in the agreement relative to the spot price, expressed as a percentage of the spot price, for different supplier risk aversion and buyer bargaining parameters. The heatmap legend is to the right of the heatmap and units are percentages of spot prices. Model estimates consider the average capacity and quantity assurance across agreement suppliers. Figure 18 reproduces Figure 16, while Figure 19 shows discounts when suppliers have access to actuarially fair insurance over profits when selling to the spot market.

Figure 19 highlights that allowing suppliers to insure against spot-market profit risk reduces discounts: buyer power from surplus capture no longer lowers agreement prices. This result arises because surplus capture derives from the buyer serving as the de facto monopoly provider of profit smoothing. When smoothing is unavailable outside the agreement, the supplier's outside option is less valuable because it exposes the supplier to risk. Insurance-like agreement terms therefore create sufficient surplus for agreements to survive replacement threats, enabling those threats to reduce prices. Actuarially fair insurance replaces this monopoly provision of profit smoothing with the perfectly competitive benchmark. More broadly, this counterfactual highlights how low input prices in rich vertical contracts can reflect not only concentration in the input market, but also limited or absent competition in supporting markets whose services are bundled into the vertical agreement.

6 Conclusion

Low vertical contract prices can arise from two distinct forms of buyer power with different implications for efficiency, distribution, and policy. Buyer power over surplus creation reflects the buyer's retention of surplus created by terms that mitigate market frictions; buyer power from surplus capture reflects the buyer's appropriation of pre-existing rents through replacement threats from in-house supply. Studying fabric purchases by a large Indian garment exporter, I show that agreements include two insurance-like terms: quantity assurance and U-shaped prices. A bargaining model formalizes the price effects of the two forms of buyer power: replacement threats reduce prices only when insurance-like terms create sufficient surplus for the agreement to survive the threat. As agreement surplus increases in supplier risk aversion, buyer power from surplus capture only affects prices when supplier risk aversion exceeds a cutoff. I estimate a structural bargaining model to decompose discounts into the two forms of buyer power. Identification leverages price responsiveness to exogenous demand shocks. As all agreement suppliers have estimated risk aversion below the cutoff, observed discounts reflect buyer power over surplus creation rather than surplus

capture. The model implies that the buyer retains ∼60% of surplus created.

These results emphasize that low input prices in vertical contracts need not reflect pure rent extraction in input markets. The appropriate policy response depends on which form of buyer power most strongly shapes prices. Increasing buyer competition has negligible effects in agreements with highly risk-averse suppliers, as prices are disciplined by replacement threats. Instead, strengthening the supporting markets that agreement terms substitute for gives suppliers access outside the agreement, reducing the discount the buyer can obtain from serving as their sole provider.

More broadly, this paper provides additional evidence that challenges to smoothing, especially relevant in LMIC settings, shape both welfare and prices. Prior work shows that limited smoothing can depress prices received by LMIC producers and wages received by LMIC workers (Burke et al., 2018; Jayachandran, 2006). In my setting, missing smoothing markets lead small risk-averse firms to accept lower average prices in exchange for insurance-like agreement terms. Policies that reduce suppliers' exposure to risk can therefore improve both *efficiency*—by helping firms use scarce resources when they are most valuable—and *equity*—by strengthening the bargaining position of firms least able to bear volatility.

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.

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.

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

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

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.

Boehm, Johannes and Ezra Oberfield, "Misallocation in the Market for Inputs: Enforcement and the Organization of Production," *The Quarterly Journal of Economics*, 06 2020, *135* (4), 2007–2058.

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.

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

Demirer, Mert and Michael Rubens, "Welfare Effects of Buyer and Seller Power,"

Working Paper, 2025.

Dixit, Avinash, "The Role of Investment in Entry-Deterrence," *The Economic Journal*, 1980, *90* (357), 95–106.

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.

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.

Goldberg, Pinelopi Koujianou, *The Unequal Effects of Globalization* Ohlin Lectures, Cambridge, MA: The MIT Press, 2023. With Greg Larson.

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

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

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

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

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

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

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

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

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).

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall, "Informal Insurance Arrange-

ments with Limited Commitment: Theory and Evidence from Village Economies," *The Review of Economic Studies*, 01 2002, *69* (1), 209–244.

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.

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," *Working Paper*, 2022.

Rubens, Michael, "Market Structure, Oligopsony Power, and Productivity," *American Economic Review*, September 2023, *113* (9), 2382–2410.

Startz, Meredith, "The value of face-to-face: Search and contracting problems in Nigerian Trade," *Working Paper*, 2025.

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

Treuren, Leonard, "Wage markups and buyer power in intermediate input markets," *Working Paper*, 2025.

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

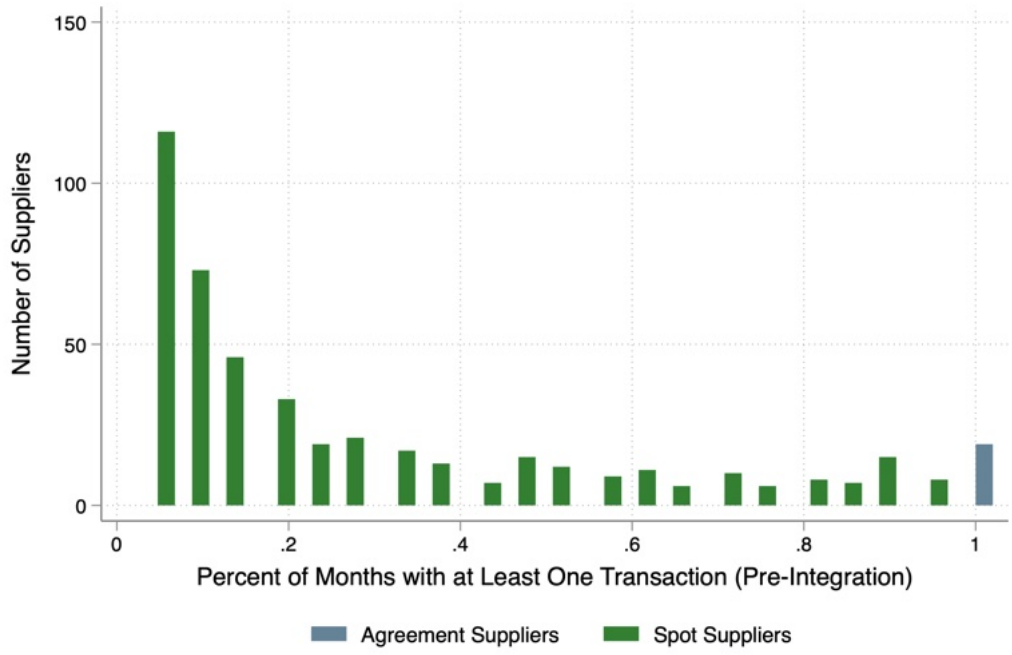
Zavala, Lucas, "Unfair Trade? Monopsony Power in Agricultural Value Chains," *Working Paper*, 2022.

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 the universe of fabric transactions by the buyer. For each supplier, I compute the percentage of months with at least one transaction, considering all pre-integration months. Bars represent the number of suppliers for each percentage of months with a transaction.

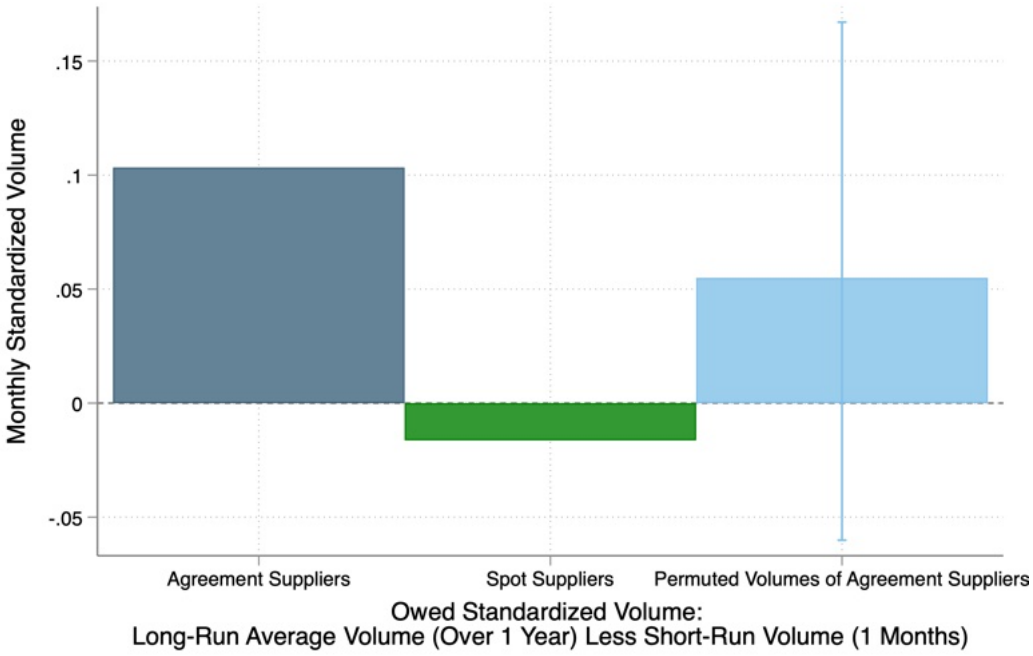
A.1.2 Quantity Assurance: Owed Quantity Guarantees

In addition to showing that agreement suppliers have more stable quantities, I show that the buyer directs quantities towards agreement suppliers to provide quantity assurance. To stabilize volumes, agreement suppliers should receive more volume when the agreement "owes" them volume, relative to the quarterly target. To test this hypothesis empirically, I regress volume in a month on a proxy for the volume owed in the agreement. Long-run monthly average volume serves as a proxy for the quantity target in the agreement. Owed quantity is measured as the difference between the proxy for the long-run quantity target and the quantity fulfilled recently (*i.e.,* in the prior month). I standardize both monthly volume and owed volume within a supplier for ease of interpretation. Conversely, there is no reason to

expect a relationship between owed volume and current volume for spot 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.[41] Importantly, the same result does not hold for spot 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 a supplier-by-month panel constructed from the universe of fabric transactions by the buyer. Bars show coefficient for owed standardized volume from a regression of current volume on owed standardized volume and month fixed effects. Owed standardized volume is calculated for each supplier-month as the average of monthly standardized volume in the prior year, not including the prior month, less the standardized volume in the prior month. Permuted volumes randomly reshuffle standardized volumes across months within a supplier. Regression estimates then consider volumes computed 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, its magnitude is larger than that of the buyer randomly redirecting the *same* quantities to agreement sup-

[41] The pass-through of less than one could reflect the challenges the buyer faces in providing the quantity guarantee given the stochastic, lumpy nature of garment orders.

pliers, which features mean reversion by construction.[42] Additionally, the most reasonable alternative mean-reverting contracting strategy distinct from the agreement described would presumably direct volumes to suppliers 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 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 be higher, as more recent volume should be 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 methodology of Lind and Mehlum (2010). The approach specifies four tests that all must be met for a relationship to be U-shaped. First, from a regression of the dependent variable on the linear and quadratic terms of the independent variable, the coefficient on the quadratic term must be positive and significant. Second, the implied 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 of the independent variable. Table A.1 shows that agreement prices satisfy the criteria for a U shape, while spot prices do not. Furthermore, as the minimum price should be at target capacity in this empirical setting, it is natural to expect that the minimum price would occur at a capacity of one. This additional test is not strictly necessary to satisfy the requirements for being U-shaped.

[42] Specifically, I repeat the regressions but using volumes randomly permuted across months for the supplier. Therefore, the set of volumes associated with each supplier is the same across all simulations but the mapping between month and volume changes. The coefficient on owed volume has a smaller point estimate on average across 500 randomizations than the observed association in the data. 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.

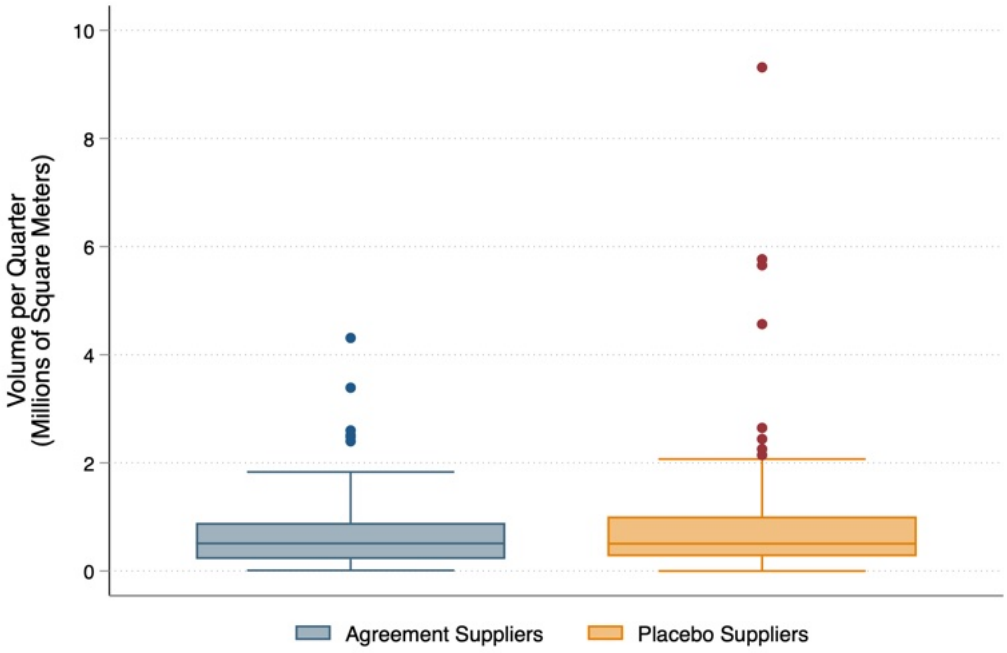
Table A.1: Pricing Function Convexity by Supplier Type

	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
Normalized Volume at 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

Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Coefficients are shown from a regression of volume-weighted standardized price on both normalized quarterly volume and normalized quarterly volume squared. Volume-weighted average prices are standardized by fabric and month. The construction of normalized quarterly volume is described in Section 4.2. Robust standard errors are shown in parentheses and $*, **, ***$ denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. The 95% confidence interval for the minimum price and the slopes at the minimum and maximum are computed from standard errors using the delta method.

A.2 Placebo Supplier Construction

Figure A.3: Volume Comparison: Agreement and Placebo Suppliers



Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Boxplots show volume per supplier per quarter using all quarters prior to integration. A supplier is considered a placebo supplier if it is among the top 19 spot suppliers with respect to average quarterly volume *regardless* of transaction frequency. The restriction to 19 suppliers is designed to match the number of 19 agreement suppliers.

I construct a set of placebo suppliers as the spot suppliers with the largest volumes pre-integration that are not classified as agreement suppliers. To make the set of placebo suppliers 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 of agreement and placebo suppliers are similar, noting that this approach does not mechanically ensure that volumes of placebo and agreement suppliers resemble each other. Figure A.3 confirms that agreement and placebo suppliers have similar quarterly volumes. In fact, placebo suppliers have slightly *larger* volumes than agreement suppliers, with higher 25th and 75th percentiles, higher mean, and almost identical medians (≈1% different).

A.3 Theoretical Bargaining Model Robustness

A.3.1 Relational Contracting Dynamics

The key model implications are robust to evaluating a dynamic model of relational contracts that incorporates dynamic incentive compatibility constraints (DICC), as in Macchiavello

and Morjaria (2023). Compared to the static model, only deviation payoffs change. On-path payoffs do not change because the relational contract specifies a pricing rule before quantities are realized. The model features two-sided moral hazard: the supplier wants to hold up the buyer when the realized relational contract price is low, and the buyer wants to deviate to the spot market when the realized relational contract price is higher than the spot market. Both parties accept spot prices because they should be indifferent from doing so, and then trade on the spot market (with each other or with other parties) in future periods (*i.e.*, grim-trigger punishment).[43]

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, the best static deviation from the supplier's perspective occurs when $p^S q^{RC} - p^{RC} q^{RC}$ is highest. Denoting the supplier discount factor as δ^{Su}, I write the contract as below and consider the best static deviation.

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

Similarly, the buyer can break the relational contract. The buyer's best deviation occurs when $p^{RC} q^{RC} - p^S q^{RC}$ is the highest. In this case, the buyer's outside option is:

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

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

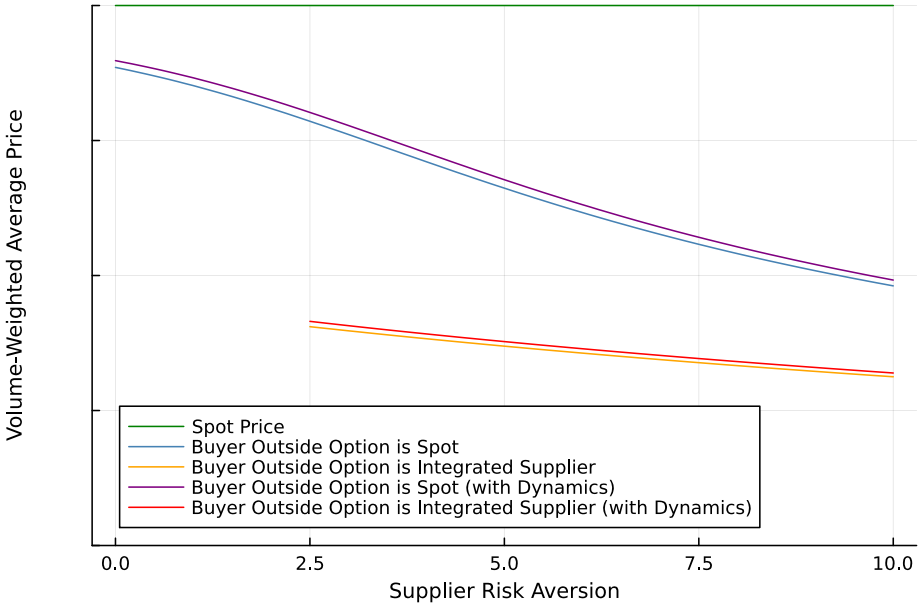
Note that for δ^S, δ^B sufficiently large, dynamic incentives are irrelevant as only long-run payoffs shape behavior.

The analysis below shows that results are robust to evaluating a dynamic relational contract model of the agreement.[44]

[43] If anything, suppliers slightly prefer the contract with the spot price because there are positive markups at spot prices.

[44] For this analysis, I set $\delta^{Su} = \delta^{Bu} = .9$. As empirical analysis of this model is at the quarterly 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.4: Agreement Prices & Supplier Risk Aversion with RC Dynamics (DICC)



Note: Costs represent volume-weighted average procurement costs that reflect prices from solving the model as parameterized in Appendix A.4. Spot prices are flat because there is no bargaining on the spot market; therefore, supplier risk aversion does not affect spot prices. Prices from both the static and dynamic model are shown. The dynamic model considers a quarterly discount factor of .9.

A.4 Parameterization and Solution of the Theoretical Model

The model is parameterized as:

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

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

I now formalize an incomplete-information bargaining game where the buyer privately observes a forecast of residual in-house capacity. The buyer then sends a signal about whether in-house capacity is relatively more likely to be constrained or not—effectively, whether to threaten replacement or not. The supplier then decides whether to accept or reject the offer. If the supplier rejects the offer, then the game proceeds with costly verification of realized residual in-house capacity. Bargaining then uses the buyer's true outside option, which depends on whether residual in-house production is capacity constrained. The supplier's rejection disrupts procurement planning and requires a costly supplier audit of the buyer's in-house supply. Therefore, if the agreement survives, the supplier provides compensation. This compensation can be interpreted as part of the relational agreement, and ensures suppliers do not always just challenge the buyer and demand costly verification. Importantly, the supplier benefits from non-average price terms in the agreement, such as insurance-like terms in my empirical setting but trade credit or other services in other contexts. When suppliers value these terms sufficiently highly, they accept the replacement threat; otherwise, suppliers reject the replacement threat and prefer selling to the spot market. Therefore, accepted replacement threats reduce prices only when supplier valuation is high.

Let c^S denote the external supplier's cost of producing the order, and let c^B denote the buyer's internal production cost when residual in-house capacity is slack. Assume $c^B < p^S$, where p^S is the buyer's procurement cost when residual in-house capacity is constrained. The supplier receives common-knowledge nonprice agreement value θ, relative to its normalized outside payoff of spot market sales. Given this normalization, the supplier's outside option has value zero.

Nature draws the buyer's private forecast type:

$$\rho_i \in \{\rho_L, \rho_H\}, \qquad 0 \leq \rho_L < \rho_H < 1$$

with $\pi = \mathbb{P}(\rho_i = \rho_H)$. Here, ρ_i is the probability that suitable residual in-house capacity will be slack if the replacement threat is rejected. I refer to ρ_L as the low-slack-probability type and ρ_H as the high-slack-probability type.

The buyer's in-house mill has median capacity utilization of 78%, making relatively low probabilities of slack residual in-house capacity empirically plausible. This should not be interpreted as a mechanical mapping from average capacity utilization to ρ_i.

The buyer chooses between a normal offer N, which treats spot supply as the relevant outside option, and a lower-price replacement-threat offer T, which signals slack residual in-house capacity. I refer to N as normal because, empirically, the buyer's in-house mill runs at high capacity utilization. Assuming equal bargaining weights, bargained prices are:

$$p_N(\theta) = \frac{p^S - \theta + c^S}{2}, \qquad p_T(\theta) = \frac{c^B - \theta + c^S}{2}$$

Because $c^B < p^S$, $p_T(\theta) < p_N(\theta)$.

The game proceeds as follows:

1. Nature draws ρ_i, and the buyer privately observes whether it is the low- or high-slack-probability type.

2. The buyer chooses the signal to send about residual in-house supply: either N or T.

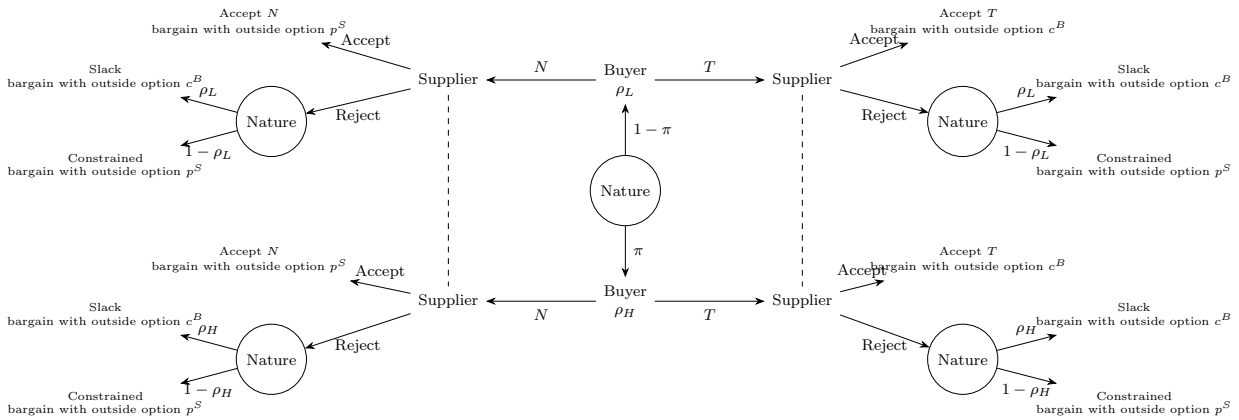
3. After observing N, the supplier chooses whether to accept or reject the normal offer. Acceptance leads to trade at $p_N(\theta)$.

4. After observing T, the supplier either accepts the replacement-threat offer or rejects the threat.

5. Following rejection of either offer, the parties learn whether residual in-house production is capacity constrained. Residual in-house capacity is slack with probability ρ_i and constrained with probability $1-\rho_i$. The parties then Nash bargain using the true buyer outside option: c^B if residual in-house production is not capacity constrained and p^S if residual in-house production is capacity constrained. If the post-rejection bargain survives, the supplier provides compensation $\kappa > 0$ to the buyer for rejection-induced disruption. The same κ applies regardless of whether residual in-house production is capacity constrained. If post-rejection bargaining breaks down, no compensation is paid. I impose one natural bound on κ: it must not eliminate the possibility of gains from threatening replacement. Therefore, $p_N(\theta)+\theta-c^S-\kappa \geq p_T(\theta)+\theta-c^S \geq 0 \implies \kappa \leq \frac{p^S-c^B}{2}$.

The Nash price is determined first, using the verified outside option. The compensation is then imposed as a separate adjustment under the bargaining protocol.

Figure A.5: Signaling Stage



Before continuing, first consider the supplier's response to N. Accepting the normal offer gives the supplier:

$$p_N(\theta) + \theta - c^S = \frac{c^S - \theta + p^S}{2} + \theta - c^S$$
$$= \frac{p^S + \theta - c^S}{2}$$

If the supplier rejects, verification occurs and any surviving post-rejection bargain gives the supplier its state-contingent Nash surplus less κ. Because the normal offer already gives the supplier the payoff associated with the more favorable spot-supply outside option, but without the compensation cost, rejection cannot improve its payoff. Therefore, the supplier accepts N whenever the normal agreement generates gains from trade relative to constrained in-house supply—the economically relevant and interesting setting. This setting prevails when $\theta + p^S - c^S \geq 0$

A.5.1 Pooling on T When θ is High

Here, I establish a weak Perfect Bayesian Equilibrium where the buyer threatens to replace a supplier with θ sufficiently high to survive the threat. Formally, a supplier with θ : $\theta + c^B - c^S \geq 0$ satisfies this condition; let $\theta^H \equiv c^S - c^B$ such that $\theta \geq \theta^H$ is a high θ, indicating that bargaining survives the replacement threat. If the supplier accepts the threat (signal T), the buyer has no incentive to deviate because it obtains the lowest procurement price when its outside option is not capacity constrained (formally, as $p_T(\theta) < p_N(\theta)$). The key constraint is then ensuring that it is incentive compatible for the supplier to accept in this case. I consider values of θ sufficiently high that both post-rejection bargains survive after compensation. Because the bargain with c^B as the buyer's outside option generates the lower supplier surplus, this requires $\theta \geq c^S - c^B + 2\kappa$.

It is incentive compatible for the supplier to accept if the payoff from accepting is weakly higher than the payoff of rejecting. The supplier's payoff from accepting is $p_T(\theta) + \theta - c^S$. The supplier's payoff from rejecting depends on its belief about the buyer type. When buyers pool, sequential rationality dictates that the supplier believes that the probability of the buyer having residual slack capacity in-house is $\bar{\rho}_T = \pi\rho_H + (1-\pi)\rho_L$. Therefore, the supplier's expected payoff from rejecting is $\bar{\rho}_T(p_T(\theta) + \theta - c^S) + (1 - \bar{\rho}_T)(p_N(\theta) + \theta - c^S) - \kappa$. It follows that the supplier accepts if:

$$p_T(\theta) + \theta - c^S \geq \bar{\rho}_T(p_T(\theta) + \theta - c^S) + (1 - \bar{\rho}_T)(p_N(\theta) + \theta - c^S) - \kappa$$
$$\kappa \geq (\bar{\rho}_T - 1)(\frac{c^B + \theta - c^S}{2}) + (1 - \bar{\rho}_T)(\frac{p^S + \theta - c^S}{2})$$
$$\kappa \geq \frac{1}{2}[(\bar{\rho}_T - 1)c^B + (1 - \bar{\rho}_T)p^S]$$
$$\kappa \geq \frac{1}{2}(1 - \bar{\rho}_T)(p^S - c^B) = (1 - \bar{\rho}_T)\frac{p^S - c^B}{2}$$

Therefore, whenever θ exceeds the cutoff $\theta \geq c^S - c^B + 2\kappa$ and κ satisfies the condition above, the wPBE exists. Note that the lower bound on κ highlights the role of this cost: when the buyer signals T and the agreement will survive, the supplier can only benefit from rejecting the signal in a world without this compensation. Namely, in such an environment, the supplier either receives the same low price or, if the buyer is revealed to be constrained, receives a higher price. It follows that, intuitively, κ helps discipline the supplier from always rejecting. Also, note that for any $\bar{\rho}_T > 0, \kappa$ does not violate the relevant bound that $\kappa \leq \frac{p^S - c^B}{2}$; however, for $\bar{\rho}_T$ very small, the window for κ is narrow.

A.5.2 Discount from T Breaks When θ is Low

It is simple to show that the wPBE above with a discount for the buyer does not persist for θ below the cutoff θ^H. If the supplier accepts the signal, because θ is below the cutoff, bargaining breaks because no surplus remains and the supplier has a payoff of 0 from spot sales. If the supplier rejects, it receives 0 when the buyer has slack capacity, again because the agreement breaks. When the buyer is capacity constrained, the supplier may or may not have a positive payoff given the natural bounds on κ and θ. If the bargaining does not dissolve, then the supplier prefers to deviate and obtain those positive profits with probability $1 - \bar{\rho}_T$. If bargaining dissolves, the supplier's payoff is 0 on path and off-path. While the wPBE does not technically break in this case, the equilibrium does not feature the buyer receiving a discount.

A.5.3 Pooling on N when θ is Low

Here, I establish a weak Perfect Bayesian Equilibrium where the buyer *never* threats to replace a supplier with θ below the cutoff θ^H to survive the threat. Recall that it is already established that the supplier accepts the signal N. Therefore, the remaining deviation to check is whether the buyer would prefer to make a replacement threat. Equivalently, here we check if the buyer is indeed constrained by bargaining breakdown. We consider the buyer type with the strongest temptation to deviate: ρ_H. This buyer's procurement cost on-path is $p_N(\theta) = \frac{p^S - \theta + c^S}{2}$.

If the buyer deviates to threaten replacement, its payoff depends on the realized capacity state and the supplier's strategy. Recall that when the buyer pools, the supplier always reject when θ is low. Therefore, the buyer's deviation procurement cost is c^B when the buyer has slack capacity. Here, I analyze games with suppliers who have low valuations of the agreement, which affects the buyer's procurement cost when the buyer is constrained. Specifically, when the buyer is capacity constrained, we consider supplier with a sufficiently low valuation of the agreement that surplus after rejection would be negative once accounting for κ. To characterize the lower-θ region in which rejection leads to breakdown in either capacity state, define $\theta_B \equiv c^S - p^S + 2\kappa$. When $\theta < \theta_B$, even the more favorable post-rejection bargain—the bargain when the buyer is revealed to be capacity constrained—does not leave the supplier enough surplus to cover the compensation: $p_N(\theta) + \theta - c^S - \kappa < 0$. Because the supplier receives still less surplus when residual in-house capacity is slack, post-

rejection bargaining also breaks down in that state. Therefore, no compensation is paid following rejection. A buyer that deviates to T incurs procurement cost c^B when residual in-house capacity is slack and p^S when it is constrained.

$$\begin{aligned}\frac{p^S - \theta + c^S}{2} &\leq \rho_H c^B + (1-\rho_H)(p^S) \\ \theta &\geq c^S - p^S + 2\rho_H(p^S - c^B)\end{aligned}$$

This result emphasizes another key piece of intuition for the effect of replacement threats: the buyer needs to receive *some* discount in the agreement to want to preserve it. This requires a minimum bound on the θ possible, although this bound is generally quite loose. For example, as long as the supplier makes spot profits, $c^S < p^S \implies c^S - p^S < 0$. Therefore, for small ρ_H, this bound need not even be positive.

A.5.4 Conditions Supporting Both Pooling Equilibria

Here, I establish that a range of κ satisfy the three conditions to support both pooling equilibria.

First, pooling on N requires bounding θ. The upper bound $\theta < \theta_B \equiv c^S - p^S + 2\kappa$ restricts attention to suppliers with sufficiently low valuation of agreement terms that a rejected replacement threat leads to breakdown in both capacity states. The lower bound $\theta \geq c^S - p^S + 2\rho_H(p^S - c^B)$ ensures that even the buyer type most tempted to deviate prefers the accepted normal agreement, because the agreement discount is large enough to make it worthwhile to preserve.

Second, pooling on T imposes the bound $\kappa \geq (1-\bar{\rho})\frac{p^S - c^B}{2}$, which prevents the supplier from rejecting the offer to seek a better price after verification. The third condition ensures that κ is not so large that it eliminates gains from threatening replacement: $\kappa \leq \frac{p^S - c^B}{2}$.

Combining these conditions, a single value of κ supports both pooling equilibria whenever:

$$\max\left\{\rho_H(p^S - c^B), (1-\bar{\rho})\frac{p^S - c^B}{2}\right\} < \kappa < \frac{p^S - c^B}{2}$$

where:

$$\bar{\rho} = \pi\rho_H + (1-\pi)\rho_L$$

This interval is nonempty when $\rho_H < 1/2$ and $\bar{\rho} > 0$. The first condition is empirically plausible because the buyer's in-house mill generally operates at high capacity utilization (median of 78%), making suitable residual capacity unlikely to be slack more than half the time. The second condition is vacuously true.

A.5.5 Discussion

This game formalizes how replacement threats can depend on supplier valuation of nonprice agreement terms. When θ is low, rejecting the signal of T leads to a breakdown of bargaining. The buyer then loses the agreement discount and must rely on either in-house or spot supply. Because suppliers with low θ reject signal of T, but not N, the buyer prefers to signal N to preserve its discount. When θ is sufficiently high that bargaining survives compensation, the supplier then accepts after a signal of T. Then, because the bargain after T generates a lower price than N, both buyer types pool on the replacement-threat offer once the supplier accepts it.

In my empirical setting, θ maps to the value of insurance-like agreement terms to the supplier. Supplier risk aversion increases this value, so the cutoff in θ corresponds to a cutoff in supplier risk aversion. Accepted replacement threats therefore reduce prices only for suppliers that value the agreement terms sufficiently highly.

A.6 Replacement Threats in a Game with Commitment

Consider a game with commitment that proceeds as follows, using the same cost and payoff normalization as above.

1. The buyer commits to a capacity level at its integrated supplier of $q^C \in \{0, 1\}$. Commitment to $q^C = 0$ means that the buyer commits to saving capacity at the in-house supplier (*i.e.,* leaving the in-house supplier unconstrained). Therefore, production is lower cost in-house than the spot price. Conversely, $q^C = 1$ means that the in-house supplier is capacity constrained, so production is more costly in-house than the spot price. Equivalently, $q^C = 1$ represents the buyer committing to use the integrated supplier to displace other spot supply, should an order arrive. Note that the model considers the opportunity cost of unused capacity.

2. Then, the buyer and supplier engage in Nash bargaining relative to the buyer's outside option. The agreement breaks if no surplus remains.

3. Last, the buyer receives an exogenous capacity shock, which represents a new order from an end client. With probability p, it does not receive an additional order. Therefore, with probability $1 - p$, the buyer receives an additional order.

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

5. The supplier's payoff inside an agreement is $p^A(q^C) + \theta - C_S(1)$ if there is surplus in the agreement. In this function, θ represents any net nonprice benefits of the agreement, such as profit smoothing through a U-shaped pricing function. Therefore, in the context of the agreement I study, θ could represent (in a reduced-form way) higher risk aversion. I also normalize θ by subtracting off the expected payoff from selling to the spot market, such that θ represents the net benefits of the relationship compared to spot market

without the revenue or production cost in the agreement (such that the supplier's outside option value is 0). $C_S(1)$ is the production cost for the order.

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

A.6.1 Replacement Threat: Committing to $q^C = 0$

When θ is sufficiently high for agreement surplus even when the buyer is *not* capacity constrained at its in-house supplier, the equilibrium choice of the buyer is to commit to $q^C = 0$ to improve its bargaining position.

Then, when the buyer commits to $q^C = 0$, its payoff is $-p^A(0)-(1-p)C_B(0,1)$. The first term references the cost paid to the agreement supplier and the second term is the internal production for (possible) additional volume. From the Nash bargaining problem with in-house replacement, $p^A(0) = \frac{C_S(1)-\theta+C_B(0,1)}{2}$. Conversely, if the buyer commits to $q^C = 1$, its payoff is $-p^A(1) - (1-p)C_B(0,1)$. Therefore, the buyer trivially picks $q^C = 0$ because it receives a better price from the supplier as $p^A(1) = \frac{C_S(1)-\theta+p^S}{2}$ (and, by construction, $C_B(0,1) < p^S$).

A.6.2 No Replacement Threat: Committing to $q^C = 1$

Now, consider the case when θ is lower, as occurs for a less risk-averse supplier, such that agreement surplus does not remain when $q^C = 0$. However, there are gains from trade (*i.e.,* there is surplus) relative to $q^C = 1$. It follows that $\theta : C_S(1) - C_B(0,1) > \theta > C_S(1) - p^S$.

Then, if the buyer selects $q^C = 0$, its payoff is $-C_B(0,1)-(1-p)p^S$ because the agreement dissolves. When the agreement dissolves, the buyer first uses in-house capacity and then sources the next volume through the spot market (the buyer only uses in-house supply for the first order given cost capacity). However, if the buyer selects $q^C = 1$, its payoff is $-p^A(1)-(1-p)C_B(0,1) = -\frac{C_S(1)-\theta+p^S}{2}-(1-p)C_B(0,1)$, as it has saved the internal capacity to use for a possible additional order. Therefore, the buyer prefers $q^C = 1$ as long as:

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

We can conclude that the right-hand side is indeed positive by noting that both the numerator and denominator are positive. The numerator is positive based on the definition of $\theta : \theta > C_S(1) - p^S \Rightarrow p^S + \theta - C_S(1) > p^S - C_S(1) + C_S(1) - p^S = 0$. Therefore, the numerator is strictly positive. The denominator is also positive as $p^S > C_B(0,1)$. Therefore, as long as p is sufficiently small, which means an additional order is likely (recall that additional orders occur with probability $1-p$), then the buyer plays $q^C = 1$. This result only applies to settings when supplier risk aversion is low (*i.e.,* $\theta : C_S(1) - C_B(0,1) > \theta$).

A.6.3 Discussion

A natural question is whether the buyer's commitment is credible to suppliers. Specifically, the buyer might want to renege after committing to $q^C = 0$ and use in-house capacity for the first order. In other words, it would 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. If the buyer can renege after bargaining—by producing the guaranteed order in-house—then it will no longer have slack capacity available for a potential second order. Anticipating this, suppliers may effectively bargain as if the buyer will be capacity constrained. The buyer has no incentive to renege when the expected cost of reneging is weakly higher than the expected cost of honoring the commitment:

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

Therefore, as long as the additional volume is sufficiently likely (*i.e.,* $1-p$ is large), the buyer does not face commitment problems. Furthermore, this result highlights that commitment concerns are not likely binding when spot markups are high ($p^S - C_B(0,1)$), as observed empirically in this setting.

Separately, commitment may also be credible in the context of the agreement given that 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 to pay agreement prices over spot prices, the relational contract could plausibly enforce additional commitment to specific outside options.

A.7 Multilateral Bargaining Extension

The model considers only a bilateral bargaining game, but in reality 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, relative to total demand from end customers, shape the buyer's private forecast of residual in-house capacity. 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 the two pooling equilibria emphasized for surviving relationships, regardless of the bargaining order with agreement suppliers.[45]

This equivalence reflects that, in the two pooling equilibria emphasized for surviving relationships, no rejection occurs on path: in the pooling-on-N equilibrium, no threat is made; in the pooling-on-accepted-T equilibrium, the supplier accepts without rejection. Therefore, prior agreements with other external suppliers do not change the credibility of in-house production as a replacement threat because they are incorporated into the buyer's private capacity forecast.[46] Similarly, a simultaneous (rather than sequential) multilateral game, such as Nash-in-Nash, would also produce the same multilateral equilibrium in these two pooling cases. The equivalence of the simultaneous game reflects that, in these two pooling equilibria, neither the signaling game nor the commitment game creates contracting externalities across negotiations.[47]

A.8 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 from quantities to capacity utilization. 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

[45] In the commitment game, the buyer would need to bargain 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 (*i.e.,* starting with agreement suppliers).

[46] Only private equilibria can be maintained in the multilateral sequential signaling game. For example, if suppliers have homogeneous costs and agreement prices are observable, then a supplier who sees agreement prices consistent with the pooling equilibrium where the buyer makes normal offers should update about the buyer's private capacity forecast. This condition is stronger than reality, as capacity forecasts 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 avoid revealing its willingness to make replacement-threat offers to less risk-averse suppliers.

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

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. Cost function estimations leverage the most temporally disaggregated data as aggregation can mask variation; therefore, estimation uses monthly data. 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.

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 draws from the supplier-specific gamma distribution that corresponds with the method of moments estimator for the observed agreement quantities. I choose a gamma distribution because it is a two-parameter distribution with positive support and fits the data well.[48] Furthermore, a gamma distribution nests many other distributions, including many commonly used distributions (*e.g.,* chi-squared and exponential distributions are special cases of a gamma distribution). Overall, this approach incorporates scale economies, suppliers with larger capacity experience the same quantity deviation as a smaller change in capacity.

Additionally, I must estimate the counterfactual capacity variance that agreement suppliers would experience with their outside option—the spot market. As the agreement includes quantity assurance, counterfactual spot market quantity variance should be larger than agreement quantity variance. Therefore, I assume that agreement suppliers' 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 would overestimate the variance that agreement suppliers would face outside the agreement. This overestimate reflects that any individual spot 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.[49] Placebo suppliers are the closest counterfactual for agreement suppliers overall given their similar volumes.[50] Encouragingly, the estimate of counterfactual variance using placebo suppliers is

[48]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.

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

[50]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 Online Appendix OA1.12 to mitigate concerns that errors in outside moments lead to poor model performance.

reasonable: quantity variance is 8.8 times larger outside the agreement. 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.[51] The spot price, which incorporates hold-up, has an estimated margin of 7 percent over costs, using costs for a supplier at 100% capacity utilization. This estimate also seems quite plausible: it is a large enough markup to create incentives for strategic sourcing through both vertical integration and agreements but not so large as to be implausible.

A.9 Structural Bargaining Model: Evidence Supporting Model Assumptions

A.9.1 Quantity Exogeneity

Although capacity utilization for any supplier is certainly not random, as the buyer provides quantity assurance in the agreement, deviations from the targeted agreement capacity in any one quarter are plausibly exogenous. Specifically, these deviations reflect stochastic end-client demand, as the buyer knows neither which samples will be selected nor their quantities. Importantly, strategic buyer behavior to reduce capacity variance within a period does not invalidate capacity deviation exogeneity. Instead, such strategic behavior is what drives the observed quantity assurance in the agreement.

I validate that these deviations are consistent with end-client demand shocks in Table A.2. Specifically, I build a supplier-by-fabric-by-quarter panel and regress supplier-fabric volume on lagged supplier-fabric volume and fabric-period demand. Supplier-by-fabric fixed effects mean that these regressions compare a supplier-fabric pair's volume in a quarter to its own typical volume for that fabric. I find that lagged supplier-fabric volume is small and statistically insignificant, inconsistent with persistent supplier-fabric-specific quantity adjustments. By contrast, contemporaneous fabric-period volume, including leave-out fabric-period volume (*i.e.,* omitting the supplier's quantity from the fabric-period volume), strongly predicts supplier-fabric volume. These results support interpreting the identifying variation as fabric-level demand shocks that are spread across suppliers, rather than strategic targeting of specific supplier-fabric pairs.

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

Table A.2: Exogeneity of Supplier-Fabric Deviations

	Supplier-Fabric Volume			
	(1)	(2)	(3)	(4)
Lagged Supplier-Fabric Volume	.0327	.0773	.0725	.0162
	(.0797)	(.0862)	(.0928)	(.0784)
Fabric-Period Volume		.1954***		
		(.0406)		
Leave-Out Fabric-Period Volume			.1525***	
			(.0218)	
Supplier-Fabric FE	Y	Y	Y	Y
Period FE	Y	Y	Y	N
Period-Fabric FE	N	N	N	Y
N	10157	10157	10157	1484
R^2	.4201	.669	.526	.7865

Note: Data are from a supplier-by-fabric-by-quarter panel dataset constructed from the universe of fabric transactions by the buyer. Coefficients reported are from a regression of the volume for the supplier-fabric-quarter on the lag of the same variable. All regressions include supplier-by-fabric fixed effects, such that the analysis considers deviations from typical volumes. In columns 2-3, regressions also control for measures of aggregate fabric demand, while column 4 replaces the aggregate fabric demand control with fabric-by-quarter fixed effects. Clustered standard errors are in parentheses, with two-way clustering at the supplier-fabric and quarter levels.

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).[52]

To mitigate such concerns, I directly test the most likely examples of such behavior in Table A.3. Specifically, if the buyer provides additional dynamic capacity variance reduction, it would presumably take the form of ensuring that extreme capacity shocks are unlikely to occur repeatedly. I directly test this hypothesis by examining the autocorrelation of extreme capacity utilization, which I define as deviations in absolute capacity utilization at or above the 90th percentile. I find that the association between prior and current extreme capacity realizations is small and not statistically significant. Furthermore, I even find that the association between lagged capacity utilization and current capacity utilization at all is small and statistically insignificant.

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

Table A.3: Capacity Utilization & Dynamic Insurance

	Extreme Capacity Utilization (≥90th Percentile)		Capacity Utilization	
	(1)	(2)	(3)	(4)
Lagged Extreme Capacity Utilization	−.043	−.056		
	(.059)	(.048)		
Lagged Capacity Utilization			.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 are from a supplier-by-quarter panel dataset constructed from the universe of fabric transactions by the buyer. In columns 1-2, coefficients reported are from a regression of an indicator for a quarter of extreme capacity utilization (*i.e.,* normalized volume, as described in Section 4.2) on the lag of the same variable, where extreme capacity utilization means a capacity deviation at the 90th percentile or higher. In columns 3-4, coefficients are from a regression of capacity utilization, in levels, on its lag. Robust standard errors are in parentheses and the 95% empirical confidence intervals from a cluster bootstrap by supplier are shown.

Separately, if parameter estimates are biased, then model prices should be biased. In this case, the untargeted moment validation exercise in OA1.12.1 should not find that model prices are unbiased.

A.10 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.11 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.

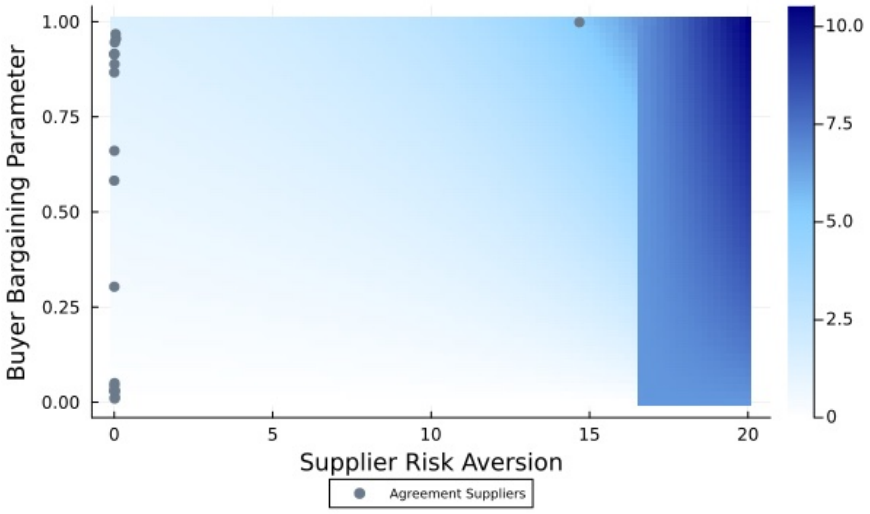
Table A.5: Buyer Bargaining Parameter and Supplier Risk Aversion: 95% CI

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: The table shows point estimates and confidence intervals for parameters estimated from the structural model, estimated separately for each agreement (*i.e.,* one risk aversion and one buyer bargaining parameter for each supplier). Risk aversion parameters are CARA coefficients, where 0 represents risk neutrality. The buyer bargaining parameter is a Nash bargaining weight between 0 and 1 that represents the buyer's surplus share. Confidence intervals use 100 bootstrap replications clustered by supplier; each replication re-samples normalized quantity and prices for each supplier (with replacement). I omit one agreement supplier because of missing price data.

A.11.1 Model-Implied Surplus Capture from Agreement Suppliers

Figure A.6: Surplus Capture from Agreement Suppliers



Note: This figure plots agreement supplier types over figure 16, showing that all agreement suppliers have risk aversion below the cutoff.

A.12 Validation using Difference-in-Differences Estimates

The standard parallel trends assumption in this setting would be that, absent integration, exposed (*i.e.,* treated) suppliers' outcomes would have evolved in parallel to those of unexposed (*i.e.,* control) suppliers. Identification does not require random assignment of suppliers to treatment. Morton (2023) shows that *fabrics* (*not* suppliers) are endogenously selected for integration: fabrics with more agreement suppliers are less likely to be brought in-house because vertical integration and agreements serve as substitutes. Importantly, suppliers are *not* selected into agreements based on individual fabrics (discussed in Morton (2023)).[53] The data show that selection on fabrics brought in-house reflects trends in overall 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' fabric portfolios. Therefore, I adopt a DiD design with propensity-score adjustment that accommodates non-random

[53]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 (necessary for relational enforcement); trends make choosing a supplier for an individual fabric a bad long-term strategy. Second, it is easier to fulfill quantity guarantees when suppliers sell a wide array of fabrics; therefore, the buyer selects agreement suppliers based on their fabric portfolios. Consistent with that approach, the median number of distinct fabrics sold by agreement suppliers in the 6 months prior to integration is large (42).

treatment (Abadie, 2005), detailed below in Appendix A.12.1. This approach relaxes parallel trends to *conditional* parallel trends; in my setting, comparisons condition on suppliers' exposure to trends in fabric demands that reflect evolving fashion tastes. Falsification tests verify that propensity score adjustment achieves balance pre-treatment.[54] Furthermore, as outcomes are *not* matched pre-integration, pre-trends support the validity of parallel trends.[55]

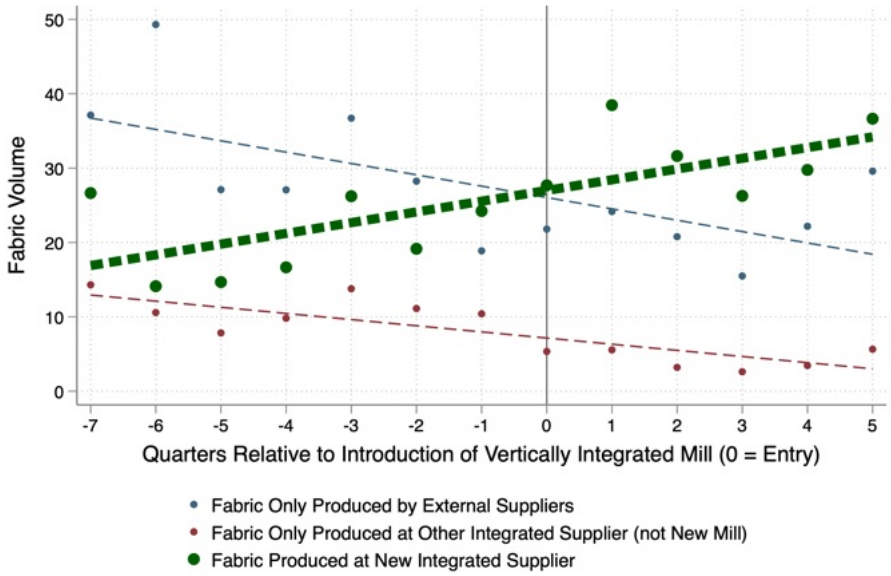
A.12.1 Semiparametric Difference-in-Differences Design and Identification

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, a preferred supplier might stop receiving orders after integration because the fabrics they produce lose 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 of long-run in-house supply, such as building an integrated mill, is consistent with evidence from other settings, such as initial electricity prices shaping long-run electricity consumption of manufacturing plants (Hawkins-Pierot and Wagner, 2022).

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

[55] 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 upon request.

Figure A.7: Trends in Fabric Demand and Vertical Integration



Note: Data from the universe of fabric transactions by the buyer, transformed into a panel at the fabric group-by-quarter level. There are three fabric groups: *i)* fabrics produced only by external suppliers; *ii)* fabrics produced only at the other integrated supplier; and *iii)* fabrics produced at the newly constructed integrated supplier (*i.e.,* fabric mill). Each observation represents the total volume summing across all fabrics within the fabric group for the quarter. Trend lines shown are linear with respect to time.

Trends in fabric demand motivate a semiparametric difference-in-differences approach (Abadie, 2005) that weakens the parallel trends assumption to *conditional* parallel trends.[56] Applied to 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 suppliers' exposure to vertical integration.[57] Formally, this assumption can be expressed as:

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

[56]The approach in Freyaldenhoven et al. (2019) would not work well in this setting because covariate adjustment relies on long-run trends rather than short-term contemporaneous movements. In other words, covariate adjustment is mostly at the supplier, rather than supplier-by-time, level, which is not the ideal variation for the covariate adjustment approach.

[57]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. Suppliers' 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 $\boldsymbol{X}_i$,[58] and $Y_i^{post}(0)$ denotes the untreated potential outcome in the post period. $D_i = 1$ indicates treatment units and $D_i = 0$ indicates 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. 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, 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. Because the propensity scores use fabric, and not supplier, trends,[58] 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.[59]

Table A.6: Balance on Untargeted Outcomes

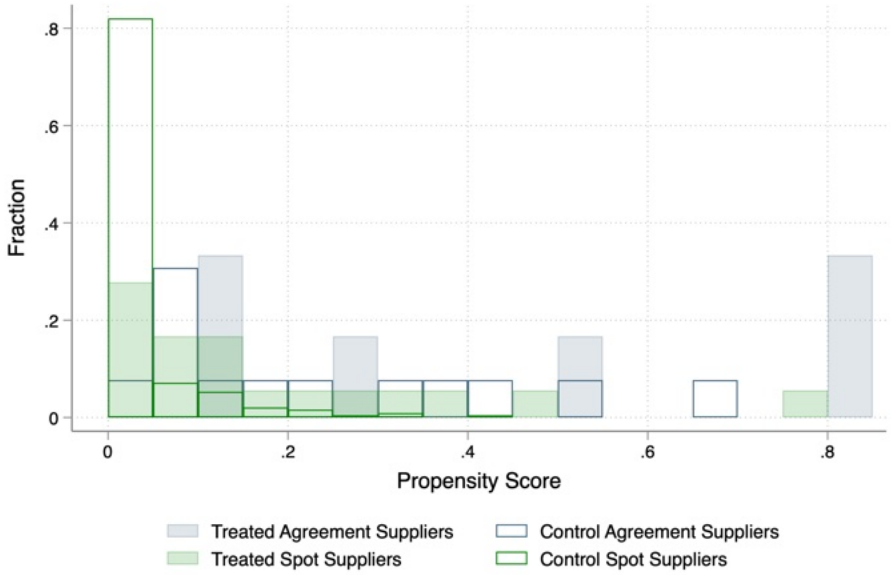
	Coefficient on Treatment Indicator:	
Outcome Variable	Without IPW Adjustment	With IPW Adjustment
Pre-Integration Volume	3033768	1085459
	[1189072, 5489585]	[−1004042, 3383491]
Pre-Integration Transaction Count	72.93	−3.35
	[44.49, 106.45]	[−46.3, 15.1]
Pre-Integration Fabric Count	26.71	−.98
	[17.08,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

[58]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, the interaction of agreement supplier indicator and the linear term for fabric growth, and the interaction of the agreement supplier indicator and the count of distinct fabrics.

[59]Formally, as 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 $\mathbb{E}[D_i] > 0$.

Note: Data from the universe of fabric transactions by the buyer, using only pre-integration quarters. Coefficients shown are from a supplier-level regression of the outcome variable, expressed in the first column, on a treatment indicator. Suppliers are considered treated if they produce a fabric that is also made at the newly built integrated fabric mill. The column "Without IPW Adjustment" does *not* incorporate the weighting procedure from Abadie (2005), while the column "With IPW Adjustment" does. The propensity score model in the IPW adjustment relies on exposure to fabric-level trends in demand across all end clients. Bootstrapped 90% confidence intervals in brackets, where the bootstrap accounts for the two-stage estimation (*i.e.,* estimating the propensity score). The propensity score is estimated using a logit model on the full set of variables in $\boldsymbol{X}_i$: 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, the interaction of agreement supplier indicator and the linear term for fabric growth, and the interaction of the agreement supplier indicator and the count of distinct fabrics.

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



Note: Data from the universe of fabric transactions by the buyer, using only pre-integration quarters. The propensity score model in the IPW adjustment relies on exposure to fabric-level trends in demand across all end clients and is estimated at the supplier level. The histogram shows the distribution of estimated propensity scores. Bars indicate the fraction of all suppliers with the associated propensity score and in the relevant group (treated or control and spot or agreement).

A.12.2 DiD: Effects on Quarterly Volume & Pooled Results

Figure A.9: DiD Estimates: Effects on Quarterly Volumes

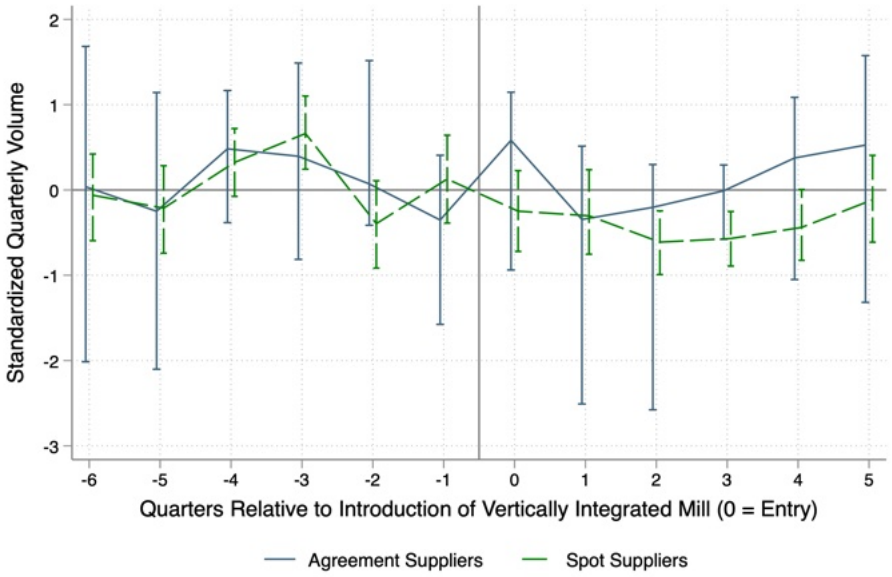


Table A.7: Pooled Pre- and Post-Treatment Effects

	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 a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Event study estimates use a semiparametric difference-in-differences design (Abadie, 2005) detailed in Appendix *A.12* and are estimated separately for agreement and spot suppliers. Suppliers are considered treated if they produce a fabric that is also made at the newly built integrated fabric mill. Standardized quarterly volume is standardized at the supplier level (*i.e.,* each observation is the z-score of volume for that supplier, where the mean and standard deviation are calculated using pre-integration data only). Volume-weighted average prices are standardized by fabric and month (as described in Section 2.3.2). Pooled estimates sum all pre- and post-period event study coefficients separately. Confidence intervals shown below the point estimates are bootstrapped to account for the two-stage estimation and clustered by supplier.

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.13 NNTR Implementation

My NNTR estimation approach aims to capture the strategic incentives for threats of replacement in Ho and Lee (2019). I incorporate their bargaining environment, where there is neither a spot market nor vertical integration. To focus on the buyer's strategic incentives, I hold the supply-side constant. However, I update my original buyer model in two 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 manufacturer using estimates from Cajal-Grossi et al. (2023).[60] Second, I model the integrated supplier as an additional external supplier, as the bargaining environment in Ho and Lee (2019) does not include any integrated production. I assume that the integrated supplier has zero risk aversion, as it is part of a large firm, and the buyer can extract all surplus from it. Then, I implement the bargaining model as in NNTR, finding the average price that would make the buyer indifferent between keeping a supplier and switching to the excluded supplier for each excluded supplier. Inside agreements, suppliers charge the minimum of the bargained price and the agreement price. I find that the buyer actually uses the integrated supplier as the sole excluded supplier, as doing so brings down prices from included suppliers.

A.14 Buyer Bargaining Parameter and Supplier Market Power

Compared to other agreement suppliers, suppliers producing fabrics in more concentrated fabric markets should be more valuable to the buyer because concentration affects the ease of providing quantity assurance. Consider an extreme case in which all agreement suppliers sell only one identical fabric. Then, when demand for that fabric is low, the buyer cannot provide quantity assurance to all suppliers. Therefore, it is easier to provide quantity assurance to suppliers in relatively concentrated fabric markets. Below, I show empirically that concentration in fabric markets is significantly negatively correlated with estimated buyer bargaining parameters, but not with estimated risk aversion parameters. The low correlation with risk aversion mitigates concerns that the connection between supplier value and buyer bargaining parameters is spurious.

I measure concentration in supplier's product markets through two steps. First, I compute fabric-specific HHIs as the HHI for the fabric considering volumes from agreement

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

suppliers one quarter prior to integration. Second, I aggregate the fabric-specific HHI to the supplier level as the volume-weighted average of the individual fabric-specific HHIs. I only use volumes from agreement suppliers because buyer market power only pertains to agreement suppliers,[61] and the trade-off of using volume from one order to fulfill quantity assurance across suppliers exclusively applies to agreement suppliers.

Table A.8: Buyer Bargaining Parameter and Supplier Concentration

	Buyer Bargaining Parameter
Supplier-Fabric 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. The coefficient shown is from a regression of the buyer bargaining parameter estimated from the structural model on supplier-fabric concentration, where supplier-fabric concentration is the supplier-fabric HHI across all agreement suppliers one quarter prior to integration. As the supplier bargaining parameter is one less the buyer bargaining parameter, a negative coefficient indicates less buyer market power and more supplier market power. Robust standard errors are shown in parentheses and $*, **, ***$ denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively.

Appendix Table A.8 illustrates that the proxy for supplier market power is strongly, and statistically significantly, correlated with the buyer bargaining parameter estimated from the structural bargaining model. 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 agreement. Results are robust to using all suppliers to compute the supplier-fabric HHI: the point estimate is -1.116 and is significant at the 5% level.

This correlation is unlikely to be spurious, as estimated supplier risk aversion is not meaningfully correlated with the market power proxy.

Table A.9: Supplier Risk Aversion and Supplier Concentration

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

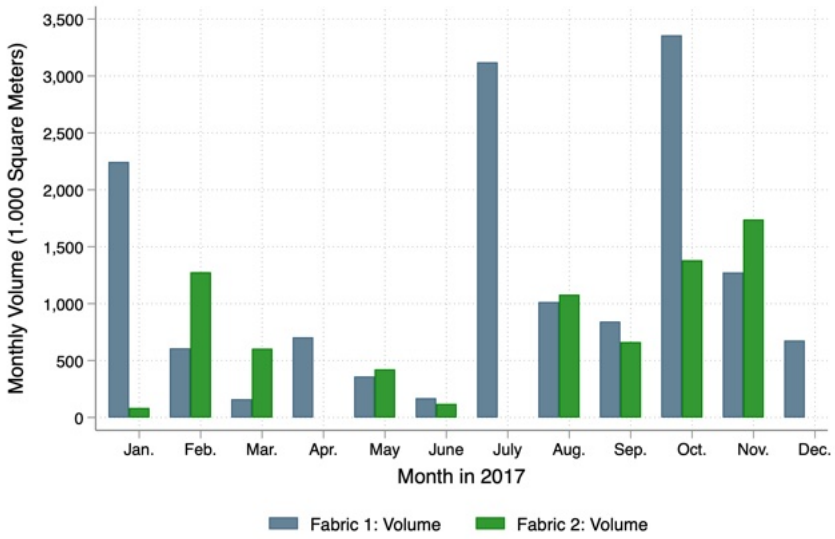
Note: Data from universe of fabric transactions by the buyer and model estimates. The coefficients shown are from a regression of the risk aversion parameters estimated from the structural model on supplier-fabric concentration, where supplier-fabric concentration is the supplier-fabric HHI across all agreement suppliers one quarter prior to integration. Robust standard errors are shown in parentheses and $*, **, ***$ denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively.

[61] Spot suppliers hold up the buyer, extracting all surplus.

OA1 Online Appendix

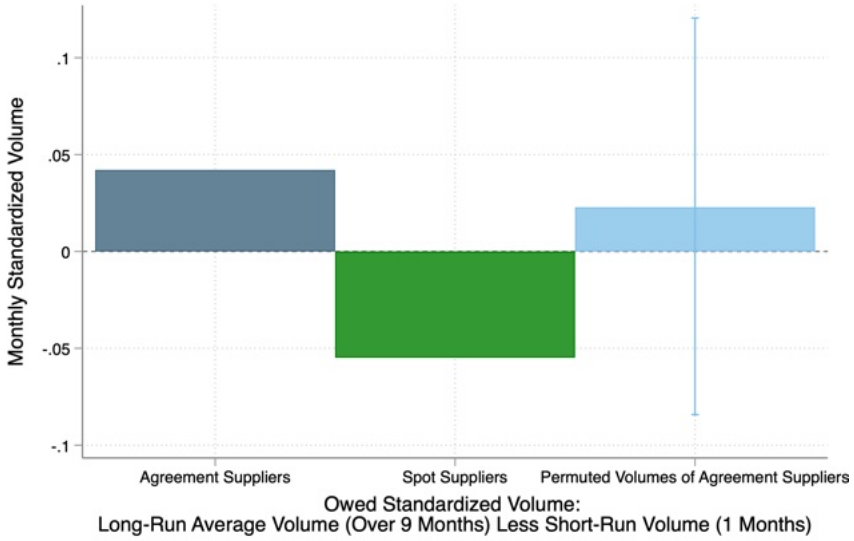
OA1.1 Additional Descriptives About Buyer-Supplier Agreements

Figure OA1.1: Variation in Monthly Quantity by Fabric



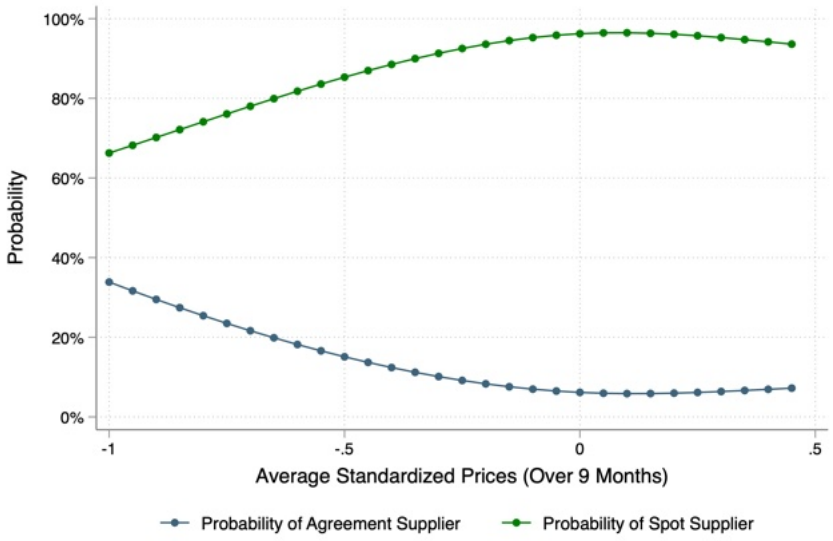
Note: Data from universe of fabric transactions by the buyer. Bars represent the quantity of fabric purchased during the month in 2017 for two selected fabrics. Fabrics were selected based on data availability to illustrate variation.

Figure OA1.2: Coefficient on Owed Quantity by Supplier Type



Note: Data from a supplier-by-month panel constructed from the universe of fabric transactions by the buyer. Bars show coefficient for owed standardized volume from a regression of current volume on owed standardized volume and month fixed effects. Owed standardized volume is calculated for each supplier-month as the average of monthly standardized volume in the prior three quarters, not including the prior month, less the standardized volume in the prior month. Permuted volumes randomly reshuffle standardized volumes across months within a supplier. Regression estimates then consider volumes computed using the permuted data across 500 permutations. The value shown for permuted volumes for agreement suppliers represents the mean across permutations; the 90% empirical CI is also shown.

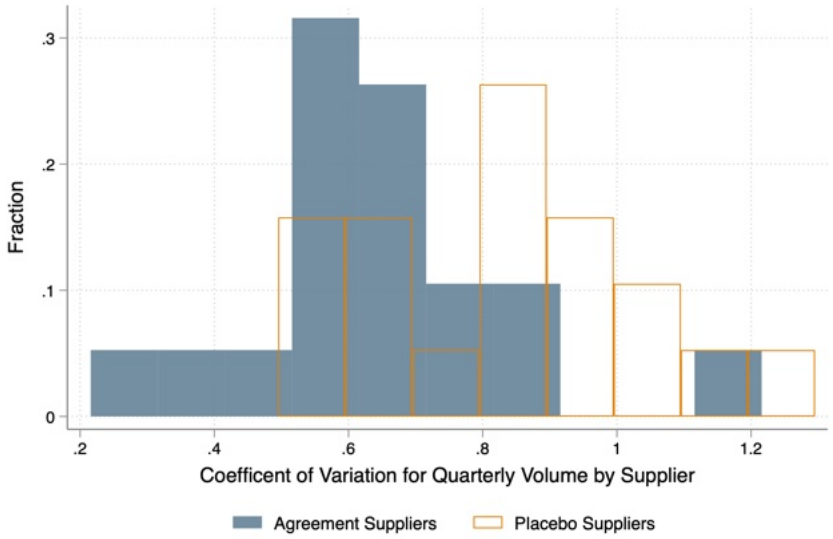
Figure OA1.3: Supplier Types and Prices



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.

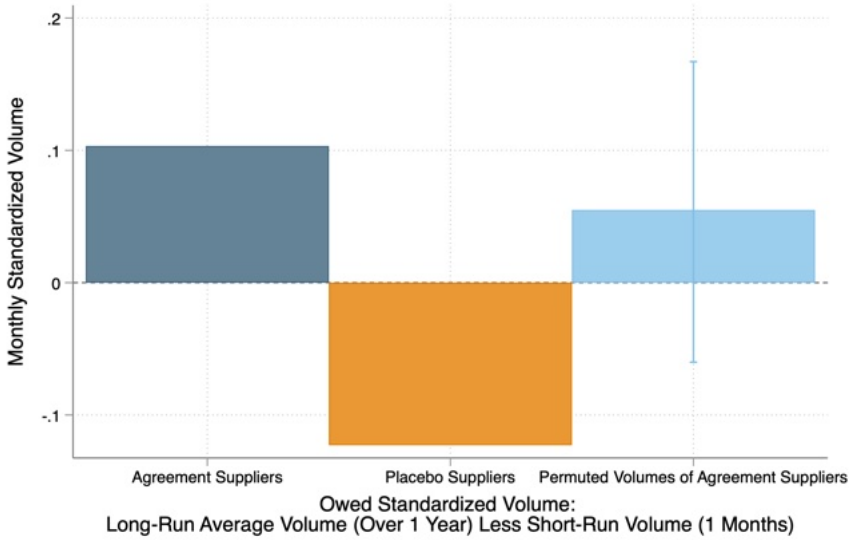
OA1.2 Placebo Supplier Analysis

Figure OA1.4: Coefficient of Variation by Supplier Type (Placebo)



Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Only quarters before the buyer adds a new integrated supplier are included. Coefficients of variation are computed using the sum of all fabric volume in the quarter; each supplier is one observation in the histogram ($N = 38$). The difference in average coefficient of variation by supplier type (agreement vs. placebo) is statistically significant ($p < 0.001$).

Figure OA1.5: Coefficient on Owed Quantity by Supplier Type (Placebo)



Note: Data from a supplier-by-month panel constructed from the universe of fabric transactions by the buyer. Bars show coefficient for owed standardized volume from a regression of current volume on owed standardized volume and month fixed effects. Owed standardized volume is calculated for each supplier-month as the average of monthly standardized volume in the prior year, not including the prior month, less the standardized volume in the prior month. Permuted volumes randomly reshuffle standardized volumes across months within a supplier. Regression estimates then consider volumes computed using the permuted data across 500 permutations. The value shown for permuted volumes for agreement suppliers represents the mean across permutations; the 90% empirical CI is also shown.

Figure OA1.6: Pricing Function (Placebo): Binned

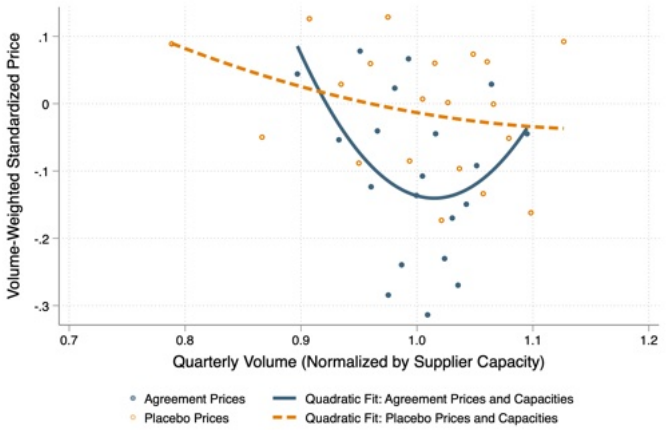
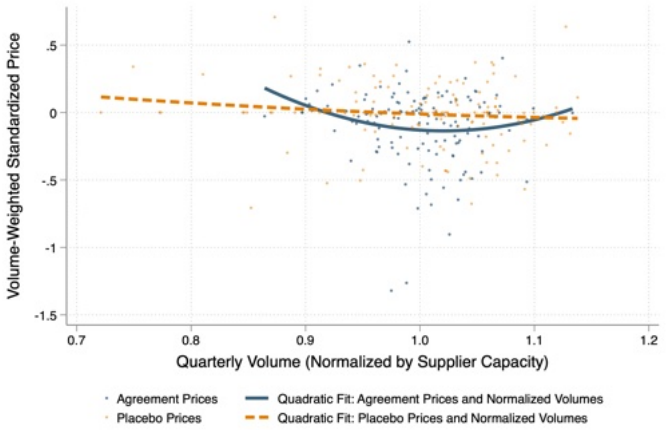


Figure OA1.7: Pricing Function (Placebo): Raw



Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Figures show both the quadratic fit and a scatter plot. The quadratic fit plots fitted values from a regression of volume-weighted standardized price on both normalized quarterly volume and normalized quarterly volume squared. Volume-weighted average prices are standardized by fabric and month. Normalized quarterly volume represents capacity utilization relative to supplier-specific capacity, where 1.0 = 100% capacity (details on measurement in 4.2). The scatter plot in Figure OA1.6 shows binned prices and volumes using 20 quantiles, while the scatter plot in Figure OA1.7 shows all raw data.

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

	Volume-Weighted Standardized Price					
Capacity	-26.438^{**}	-1.345	-36.723^{**}	-6.454	-32.61^{**}	-7.361
	(10.986)	(3.682)	(14.455)	(4.495)	(14.646)	(4.643)
Capacity^2	12.95^{**}	.519	18.113^{**}	3.182	16.169^{**}	3.663
	(5.562)	(1.952)	(7.16)	(2.391)	(7.252)	(2.471)
Supplier Type	Agreement	Placebo	Agreement	Placebo	Agreement	Placebo
Supplier FE	N	N	Y	Y	Y	Y
Quarter FE	N	N	N	N	Y	Y
Normalized Volume at 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 Capacity	-4.055	$-.596$	-5.415	-1.864	-4.662	-2.079
95% CI	$[-6.828, -1.282]$	$[-2.354, 1.162]$	$[-9.541, -1.29]$	$[-3.958, .23]$	$[-8.865, -.459]$	$[-4.236, .079]$
Slope at Maximum Capacity	2.914	$-.164$	4.333	.788	4.04	.973
95% CI	$[-.349, 6.178]$	$[-1.756, 1.429]$	$[.778, 7.887]$	$[-1.14, 2.715]$	$[.431, 7.648]$	$[-1.024, 2.97]$
R^2	.034	.015	.462	.25	.497	.281
N	133	130	133	130	133	130

Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Coefficients are shown from a regression of volume-weighted standardized price on both normalized quarterly volume and normalized quarterly volume squared. Volume-weighted average prices are standardized by fabric and month. The construction of normalized quarterly volume is described in Section 4.2. Robust standard errors are shown in parentheses and $*, **, ***$ denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. The 95% confidence interval for the minimum price and the slopes at the minimum and maximum are computed from standard errors using the delta method.

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

OA1.3 Quantities & Prices by Transaction Frequency

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

	Coefficient of Variation for Supplier Quarterly Volume	Volume-Weighted Standardized Price
	(1)	(2)
Agreement Supplier Indicator	$-.346^{***}$	$-.102^{***}$
	(.082)	(.031)
Series Controls ($k = 2$) for Transaction Frequency	Y	Y
Quarter FE	N	Y
N	315	1603

Note: Data from the universe of transactions by the buyer. Coefficients of variation—the outcome variable in column (1)—are computed from supplier-by-quarter panel data that sums all quantities within the quarter. Therefore, there is one observation per supplier. Volume-weighted standardized prices—the outcome variable in column (2)—are standardized by fabric and month. Therefore, there is one observation per supplier-quarter. Values reported are from a regression of the outcome variable on an indicator for agreement suppliers and nonparametric controls for transaction frequency (the percentage of months with a transaction, defined at the supplier level). The nonparametric controls are a series estimator (using both the level and the square). The regressions consider only quarters before the buyer adds an integrated supplier for construction of all outcome variables. Robust standard errors are shown in parentheses and $*, **, ***$ denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Only the regression for prices uses quarter fixed effects because the coefficient of variation data are defined only at the supplier, rather than supplier-by-quarter level. The regression with prices as the outcome variable is robust to omitting the fixed effect ($\hat{\beta} = -.102$, p-value $= .001$).

OA1.4 Why Not an Explicit Contract? Why no Buyer Hold-Up?

There are (at least) three 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 explicit contracts (Rao (2022); Banerjee and Duflo (2000)). Second, an effective contract would specify prices as a function of capacity utilization; a contract with constant prices 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 with constant prices (at least not without extremely high prices). As there are many possible future capacity states due to both the unpredictable stochastic demand from end clients, as well as unpredictable demand from other buyers, such a contract would likely be difficult (or at least expensive) to write because it would need to specify prices for

many different states, even abstracting from enforcement concerns. 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 prices as a function of the supplier's capacity, it may introduce supplier moral hazard that would undermine the effectiveness of the contract. Specifically, such a contract would reduce (or even eliminate) supplier incentives to accept or reject sampling opportunities from other buyers to safeguard capacity for the buyer. This concern seems relevant if the buyer believes that suppliers try to target their capacity to match the buyer's needs.

The buyer is unlikely to hold up a supplier because fabric prices are determined after the end client has ordered the garment. At that point, while the specific fabric is very valuable to the buyer, the supplier could sell capacity to other buyers. The buyer's primary opportunity to hold up a supplier would likely be after fabric production. However, in contrast to the agreement overall, individual transactions have simple explicit contracts that are presumably enforceable. Furthermore, reputational motives deter buyer hold-up.

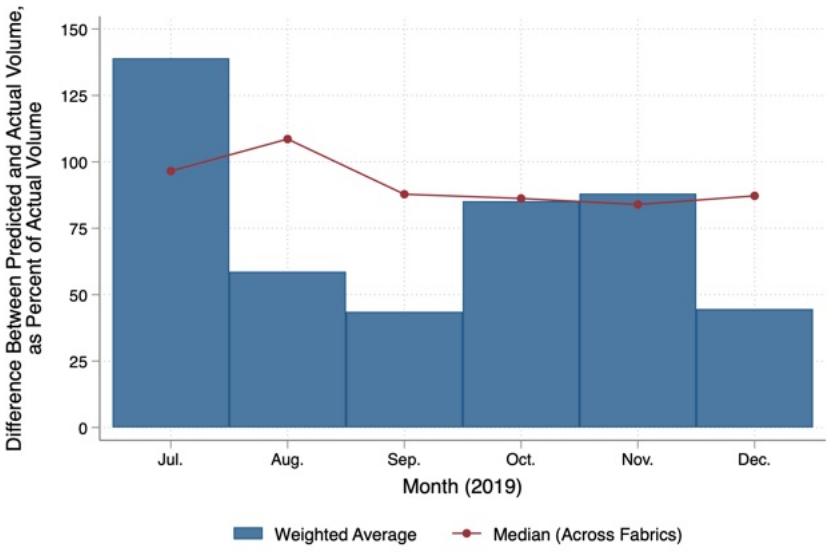
OA1.5 Why Not Exclusively Agreement Suppliers?

Given the evidence that agreement suppliers offer lower prices than spot 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. As 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 OA1.8 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 agreement and integrated suppliers appears to be quite stable pre-integration, with spot suppliers absorbing stochastic variation in quantities. This "buffer-stock" sourcing approach is visible in Figure

7.

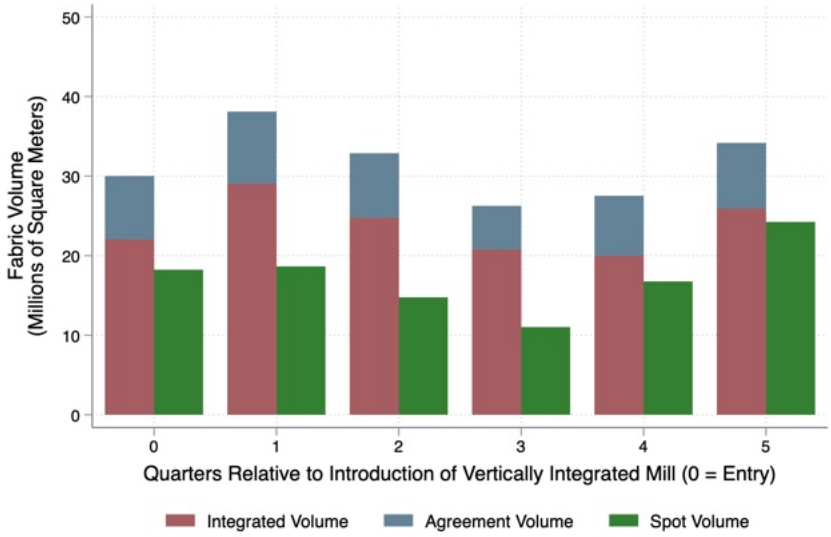
Figure OA1.8: Forecast Error in Volume Predictions



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 OA1.9: Post-Integration: Fabric Sourcing by Supplier Type



Note: Data from a panel constructed from the universe of fabric transactions by the buyer at the supplier group-by-quarter level. Bars represent the quantity of fabric purchased during the quarter for *i)* integrated and agreement suppliers and *ii)* spot suppliers, summing across all suppliers in the group. Only quarters after the new mill is fully active are shown. The coefficient of variation remains almost double for spot (0.255) relative to integrated and agreement supply (0.141).

OA1.6 Bargaining Primitives, Surplus Creation and Capture, & Markdowns

This appendix clarifies how the bargaining primitives map to buyer power measured as markdowns, even when quantities adjust, and how these mappings correspond to surplus creation and surplus capture. A markdown is the wedge between the buyer's marginal revenue product (MRP) from an input and the input price. MRP reflects downstream revenue from output produced using the input. The MRP schedule therefore does not mechanically shift with either the buyer's disagreement payoff or the buyer bargaining parameter. The arguments below state sufficient conditions under which the two bargaining primitives increase markdowns through the two channels emphasized in the paper.

The buyer bargaining parameter maps to markdowns through surplus creation. To see this, compare a higher buyer bargaining parameter to a lower one, holding fixed outside options and the technology for surplus-generating agreement terms. When the buyer bargaining parameter is higher, the buyer receives a larger share of incremental surplus created by contract terms that mitigate frictions, such as insurance or trade credit. These terms increase agreement surplus and lower the agreement price because the buyer retains part of the

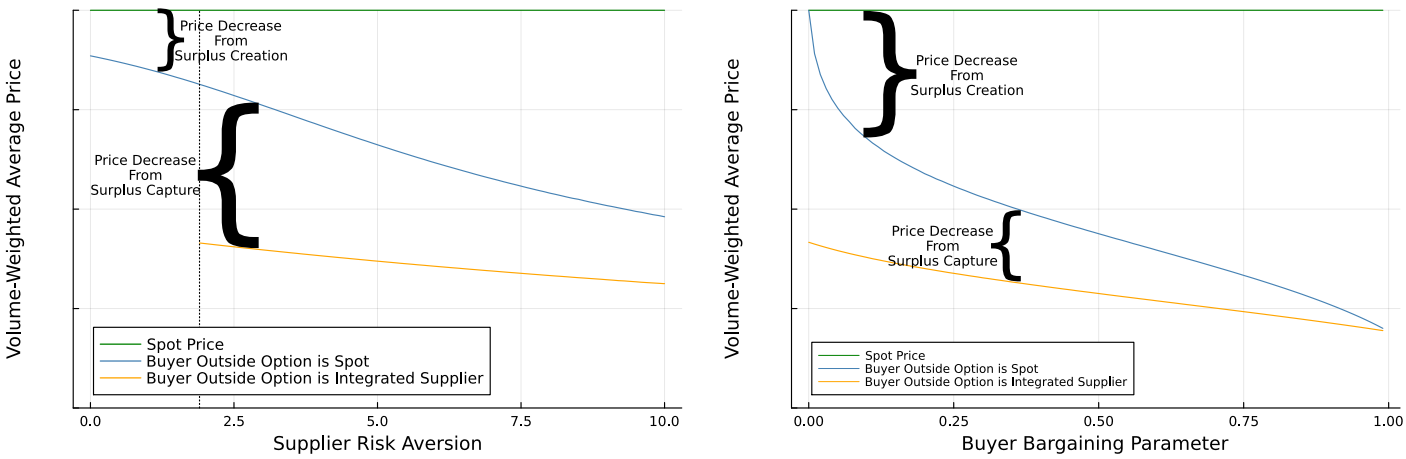
created surplus. When quantities adjust, the quantity response depends on how incremental surplus created by the agreement term varies with quantity. Increases in the buyer bargaining parameter weakly reduce quantities whenever the net incremental surplus from the surplus-generating contract terms weakly decreases in quantity, where the net incremental surplus is the supplier valuation net of supplier marginal input production costs and buyer costs of providing the terms. Under this condition, a higher buyer bargaining parameter increases the buyer's return to expanding surplus-generating terms without increasing the marginal gross surplus from additional quantity, so agreement quantities weakly fall. As long as buyer MRP is weakly downward-sloping—a standard monotonicity assumption—MRP at the marginal unit weakly increases. Therefore, under these conditions, the buyer bargaining parameter increases markdowns through surplus creation: prices fall because the agreement creates incremental surplus, part of which is retained by the buyer, while MRP at the marginal unit weakly increases.

Replacement threats map to markdowns through surplus capture. A stronger buyer outside option shifts rents toward the buyer by making disagreement less costly. Conditional on quantity, the stronger outside option lowers the negotiated price for any interior bargain with buyer bargaining parameter below one. With quantity adjustment, standard monotone-comparative-statics conditions—weakly upward-sloping supplier marginal cost and weakly downward-sloping buyer bilateral willingness-to-pay for sourcing through the agreement relative to the replacement option—ensure that replacement threats weakly reduce agreement quantities. This bilateral willingness-to-pay differs from MRP: it reflects the buyer's value of sourcing through the agreement relative to its replacement option, whereas MRP reflects downstream revenue from output produced using the input. Since replacement threats do not shift the MRP schedule, the weak reduction in quantity weakly increases MRP at the marginal unit. Replacement threats therefore increase markdowns: prices fall and MRP at the marginal unit weakly increases. The price-reduction component reflects surplus capture because it transfers rents to the buyer through a stronger disagreement payoff rather than by creating additional surplus inside the agreement.

OA1.7 Bargaining Model with Cost Heterogeneity Between External and Internal Suppliers

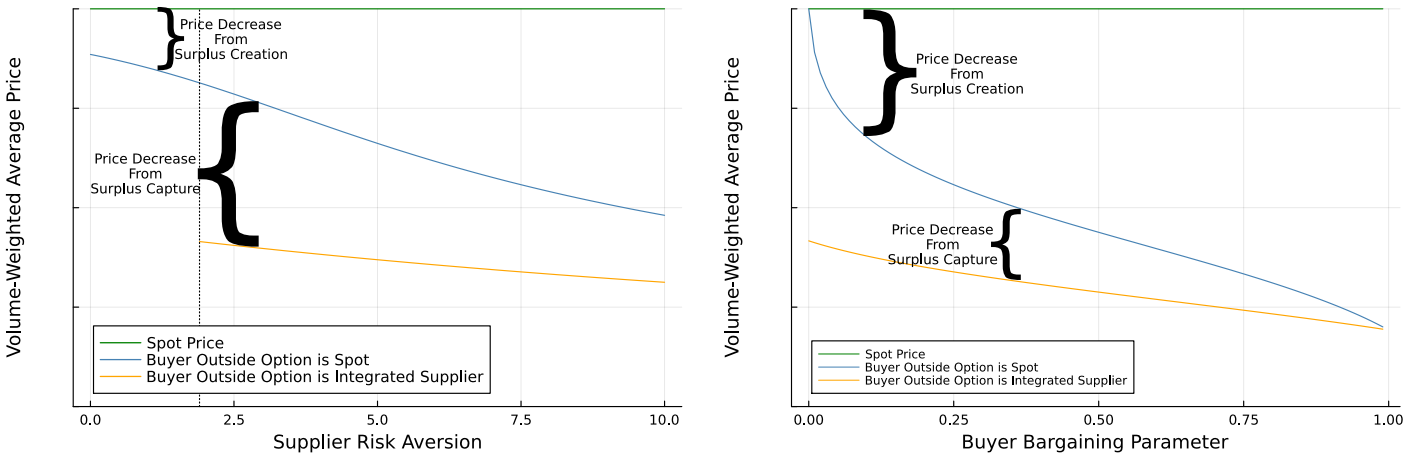
Results are robust to the external supplier having lower or higher costs than the integrated supplier, shifting costs by 10% in either direction. To ensure that optimal quantity is still $q = 1$ without changing market prices, I only shift the fixed part of costs by 10%, keeping convexity the same. Results (available upon request) shifting all costs are similar. Large heterogeneity in production costs between the buyer and supplier seems 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 its productivity).

Figure OA1.10: Agreement Prices, Supplier Risk Aversion, & Buyer Bargaining Parameter when External Supplier has Lower Costs



Note: Costs represent volume-weighted average procurement costs that reflect prices from solving the model as parameterized in Appendix A.4. Spot prices are flat because there is no bargaining on the spot market; therefore, supplier risk aversion does not affect spot prices. External supplier costs are 10% lower than integrated supplier costs for the same quantity.

Figure OA1.11: Agreement Prices, Supplier Risk Aversion, & Buyer Bargaining Parameter when External Supplier has Higher Costs



Note: Costs represent volume-weighted average procurement costs that reflect prices from solving the model as parameterized in Appendix A.4. Spot prices are flat because there is no bargaining on the spot market; therefore, supplier risk aversion does not affect spot prices. External supplier costs are 10% higher than integrated supplier costs for the same quantity.

OA1.8 Signaling Game: A Numerical Example

I now provide a numerical example for Appendix A.5. The example illustrates that a single value of κ can support both pooling equilibria, that the pooling-on-N region is nonempty, and that sufficiently high supplier valuation supports pooling on an accepted replacement-threat offer T. The example is illustrative and does not characterize equilibria outside these two pooling regions. Let $c^S = c^B = 1$ and $p^S = 2$. I consider two buyer capacity-forecast types: $\rho_L = 0.20, \rho_H = 0.30$, where each type is equally likely (*i.e.,* $\pi = \frac{1}{2}$). Consider $\kappa = .475$, $\theta^H = 1$, and $\theta^L = -.1$. Recall that negative θ is reasonable given that θ is the supplier valuation of non-price agreement components less spot profits, and spot prices can be positive.

Therefore, the bargained prices are $p_N(\theta) = 1.5 - \frac{\theta}{2}, p_T(\theta) = 1 - \frac{\theta}{2}$. When θ is high, $p_N = 1.0, p_T = 0.5$; when θ is low, $p_N = 1.55$ and there is no p_T because the Nash bargain dissolves.

First, recall from Appendix A.5 that the supplier always accepts the signal N (consistent with positive gains from trade).

OA1.8.1 Pooling on T When θ is High

The condition to check is that the supplier does not want to deviate; it is clear that regardless of the buyer type, p_T is lower than all other procurement costs so the buyer would not deviate. The supplier's on-path payoff, for the high θ supplier, is $p_T(\theta)+\theta-c^S = .5+1-1 = .5$. If the supplier deviates, it rejects the threat. Then, sequential rationality dictates that the expected chance the buyer has slack capacity in-house is $.5\times.2+.5\times.3 = .25$. The disagreement payoff for the supplier is then $.25\times p_T+(1-.25)\times p_N-\kappa = .25\times.5+.75\times1-.4 = .875-.475 = .4$. The supplier does not deviate as $.5 > 0.475$.

OA1.8.2 Pooling on T When θ is Low

When θ is low, the supplier's on-path payoff is 0. As above, if the supplier deviates, then the supplier rejects the threat. When it is revealed that the buyer has slack capacity, Nash bargaining dissolves and the supplier payoff is 0 (even without the required compensation). When it is revealed that the buyer is capacity constrained, the supplier's surplus from $p_N(\theta)$ before the compensation would be $1.55 + \theta - c^S = 1.55 - .1 - 1 = .45$. Incorporating compensation, the supplier's payoff is $.45 - \kappa = -.025$. Therefore, bargaining breaks down. The supplier has the same payoff from deviation and acceptance, but both involve no bargaining and the buyer does not receive a discount. As noted, this is an equilibrium that is far from buyer optimal: the buyer sources from the spot market every period.

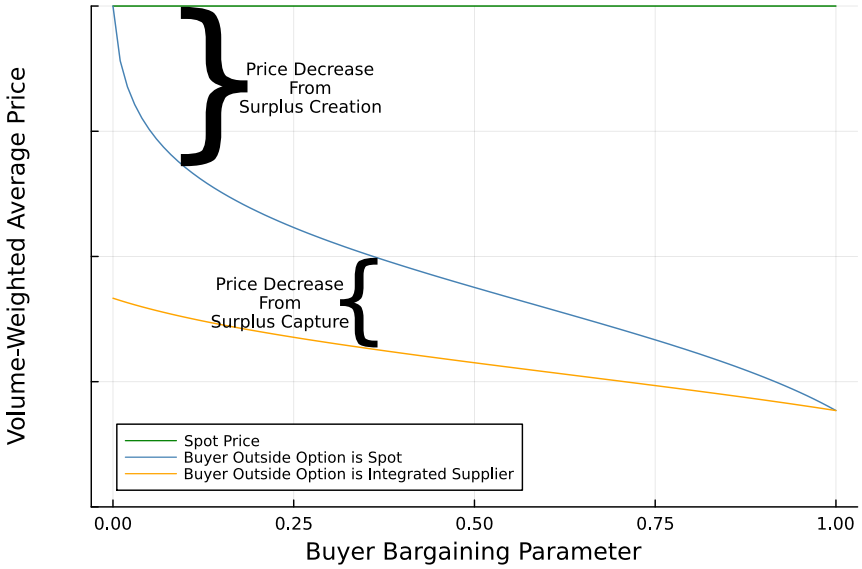
OA1.8.3 Pooling on N When θ is Low

In this case, rather than the supplier having a temptation to deviate, the buyer who receives a forecast that they are relatively likely to have slack capacity has the strongest temptation to deviate. The buyer's on-path procurement cost is $p_N(\theta) = 1.55$. If the buyer deviates

and threatens replacement, then the supplier rejects as θ is low. And, as shown above, bargaining breaks down off-path once θ is low. Therefore, the buyer's deviation payoff is $.3c^B + (1 - .3)p^S = .3 + 1.4 = 1.7$. It follows that the procurement cost is higher under deviation—1.7—than on-path—1.55. It follows that the buyer pools.

OA1.9 Surplus Creation & Capture: Buyer Bargaining Parameter

Figure OA1.12: Bargained Prices and the Buyer Bargaining Parameter



Note: Costs represent volume-weighted average procurement costs that reflect prices from solving the model as parameterized in Appendix A.4. Spot prices are flat because there is no bargaining on the spot market; therefore, supplier risk aversion does not affect spot prices.

Figure OA1.12 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 receive all surplus when the buyer bargaining parameter is zero; therefore, the bargained price is identical to the spot price. As the buyer bargaining parameter increases, the buyer receives a share of the surplus created by the nonprice features of the agreement in the form of a discount. 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 surplus capture is visible in the graph as the difference between the bargained price when the outside option is the spot market (the blue line) and the price when the outside option is the integrated supplier (the orange line). When the buyer bargaining parameter is near one, changing the buyer's outside option has a negligible effect on bargained prices since prices reflect the supplier's outside option (*i.e.,* the supplier's individual rationality constraint), which does not change when the buyer adds integrated capacity. Equivalently, surplus 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. Therefore, improvements in the buyer's bargaining position do not affect bargained outcomes.

OA1.10 Alternative Replacement Threat Model: TIOLI Offers

Alternatively, one might posit that vertical integration enables the buyer to make a take-it-or-leave-it offer backed by the threat of replacing the supplier with in-house production to avoid breaking the agreement. However, this model is both internally inconsistent and rejected by the empirical results. Theoretically, if the buyer can use an outside option to discipline bargaining, that outside option must be a credible replacement option (as in Ho and Lee (2019)). It seems implausible that the buyer could convince suppliers to increase the discount by threatening replacement with in-house production but, simultaneously, the buyer is willing to accept a smaller discount than in-house production costs imply because it does not actually want to use that replacement option. This logic points to a more natural alternative model in which the spot market, as the source of marginal volumes, remains the only credible outside option even after partial integration. That model would imply an even stronger result than the proposed model: spot markets *always* constrain replacement threats when they provide marginal volumes.

Second, the TIOLI model has a stark empirical implication rejected by the data. Buyer bargaining parameter estimates are for the pre-integration period and are below one (see Table 2). Therefore, pre-integration, the buyer does not extract all surplus through TIOLI offers. This alternative model accordingly implies large price responses to partial vertical integration as there is meaningful scope for the buyer to extract more surplus. However, the DiD results (see Section 4.6) show no price effects of partial vertical integration.

OA1.11 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 the 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 using transaction-level observations. Columns from left to right add fixed effects: supplier, fabric group, and month. Robust standard errors are shown in parentheses and $*, **, ***$ denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively.

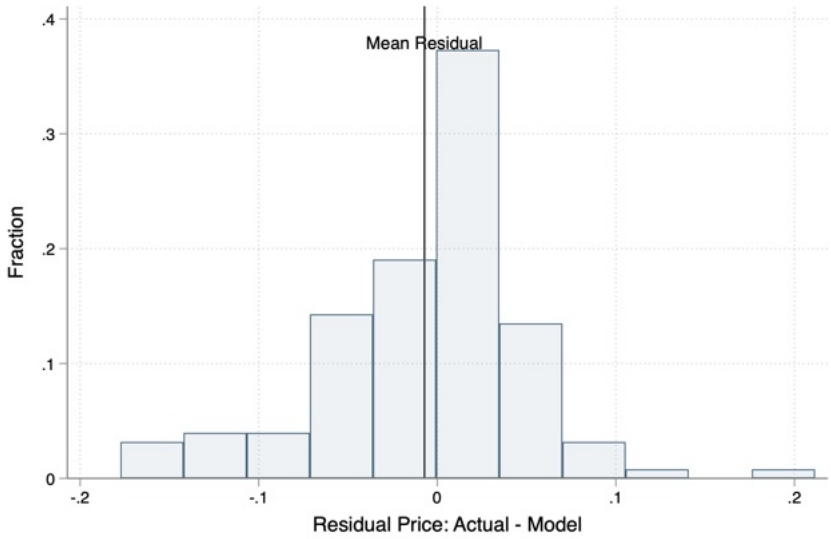
OA1.12 Model Validation and Fit

OA1.12.1 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 variance and bias. For example, an outlier high price can lead the model to pick parameters that increase all estimated prices 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.[62] Figure OA1.13 demonstrates that the distribution of residuals is symmetric with the mean residual near 0 at $-.007$, demonstrating that the model closely fits this untargeted moment. Moreover, the median residual, which is also untargeted, is even closer to 0 at .002.

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

Figure OA1.13: Untargeted Moments: Bias



Note: The histogram shows the distribution of residuals from comparing the model-implied prices to the actual prices. Each observation is the difference in price for one supplier-quarter, with the height of the bar representing the fraction of supplier-quarters with a residual in the bin. The mean residual is near zero at $-.007$ (even though the model targets RMSE, which does not necessarily deliver an unbiased estimator).

OA1.12.2 Model Validation: Out-of-Sample Fit

I analyze out-of-sample fit by comparing the distribution of supplier RMSE in the post-integration period from the model with two alternative estimators in Table OA1.4. I refer to this exercise as out-of-sample because estimation of supplier-specific risk aversion and buyer bargaining parameters exclusively use pre-integration data, conditional on the regularization tuning parameter (λ). However, because estimation of λ estimation evaluates fit in the post-integration data, post-integration prices still indirectly affect final parameter estimates even though they are not the primary variation that identifies them.

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 and only has a buyer bargaining parameter. This approximation reflects that the bargaining parameter is largely determined by the mean price. Therefore, a model with just a bargaining parameter would 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 explaining prices. 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 (1)	Supplier Mean (2)	OLS (3)
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: The table shows summary statistics for supplier RMSE in the post-integration period for different estimation routines. The supplier RMSE is the sum of the RMSE in prices for that supplier across all post-integration quarters; all model estimates primarily consider pre-integration data. Column (1) shows results from the structural bargaining model. Column (2) shows results using the pre-integration supplier-specific mean price. Column (3) shows results from a regression of prices on quantities for each supplier.

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. On average, OLS with a tuning parameter outperforms the other alternative estimators, but still has mean (median) RMSE 7% (10%) over my model.

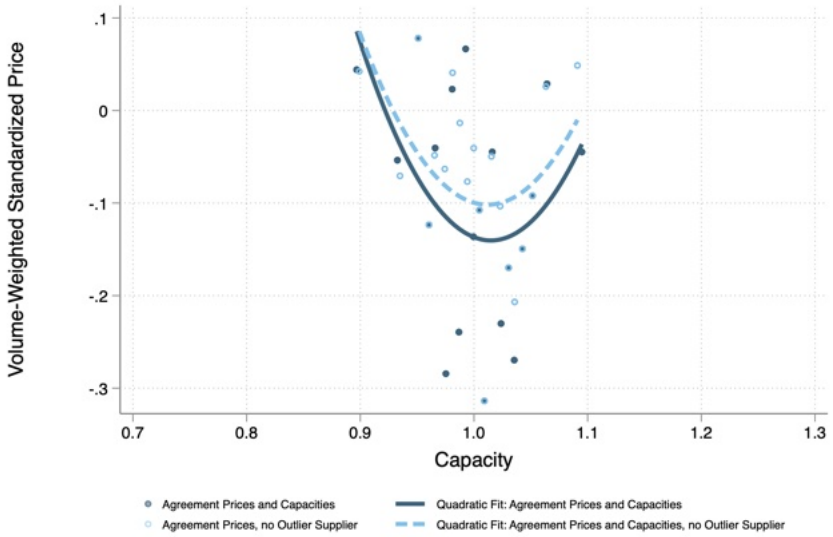
Table OA1.5: Out-of-Sample Fit Comparison (OLS with Tuning)

	Structural Model (1)	Supplier Mean (2)	OLS with Tuning (3)
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: The table shows summary statistics about the distribution of supplier RMSE in the post-integration period for different estimation routines. The supplier RMSE is the sum of the RMSE in prices for that supplier across all post-integration quarters; all model estimates primarily consider pre-integration data. Column (1) shows results from the structural bargaining model. Column (2) shows results using the pre-integration supplier-specific mean price. Column (3) shows results from a regression of prices on quantities for each supplier, incorporating a tuning parameter on the slope to minimize overfitting (mirroring the tuning parameter in the structural model).

OA1.13 Pricing Function Robustness to Outlier Supplier

Figure OA1.14: Pricing Function & Outlier Supplier: Binned

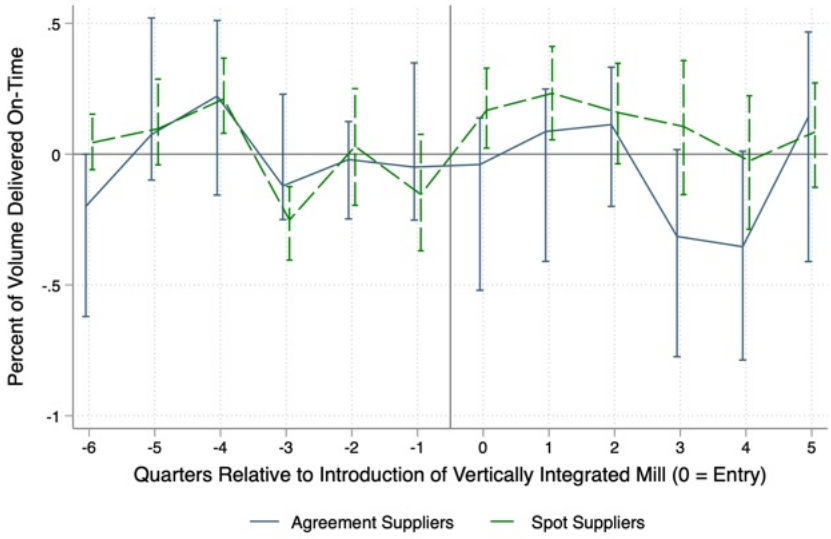


Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Figures show both the quadratic fit and a scatter plot. The quadratic fit plots fitted values from a regression of volume-weighted standardized price on both normalized quarterly volume and normalized quarterly volume squared. Volume-weighted average prices are standardized by fabric and month. Normalized quarterly volume represents capacity utilization relative to supplier-specific capacity, where 1.0 = 100% capacity (details on measurement in 4.2). The scatter plot in Figure OA1.14 shows binned prices and volumes using 20 quantiles.

OA1.14 DiD Supplemental Analysis and Robustness Checks

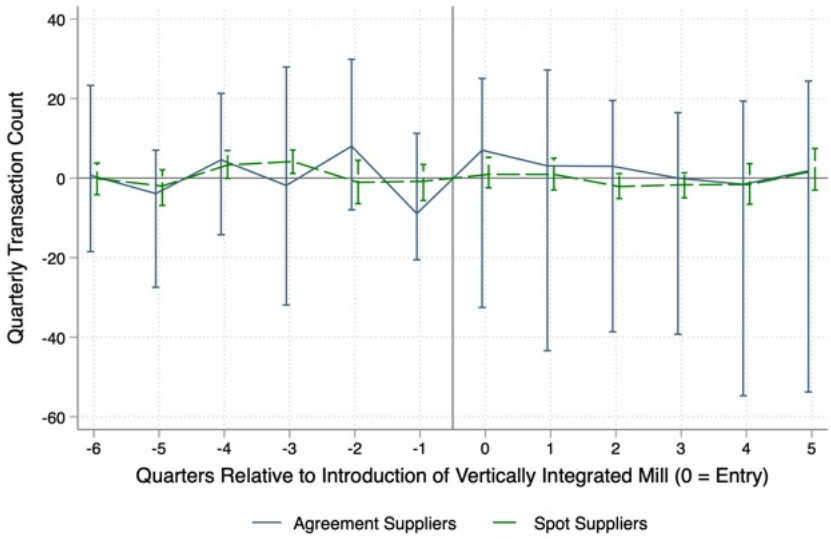
OA1.14.1 Additional Event Studies

Figure OA1.15: DiD Estimates: Effects on Reliability



Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Event study estimates use a semiparametric difference-in-differences design (Abadie, 2005) detailed in Appendix *A*.12 and are estimated separately for agreement and spot suppliers. Suppliers are considered treated if they produce a fabric that is also made at the newly built integrated fabric mill. Confidence intervals are bootstrapped to account for the two-stage estimation and clustered by supplier. Reliability is measured as the percentage of volume delivered on-time in the supplier-quarter.

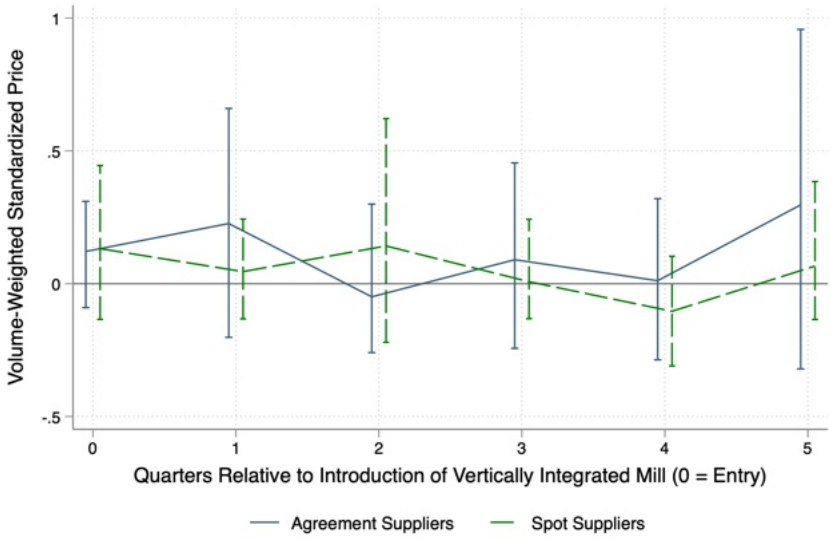
Figure OA1.16: DiD Estimates: Effects on Transaction Count



Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Event study estimates use a semiparametric difference-in-differences design (Abadie, 2005) detailed in Appendix *A*.12 and are estimated separately for agreement and spot suppliers. Suppliers are considered treated if they produce a fabric that is also made at the newly built integrated fabric mill. Confidence intervals are bootstrapped to account for the two-stage estimation and clustered by supplier.

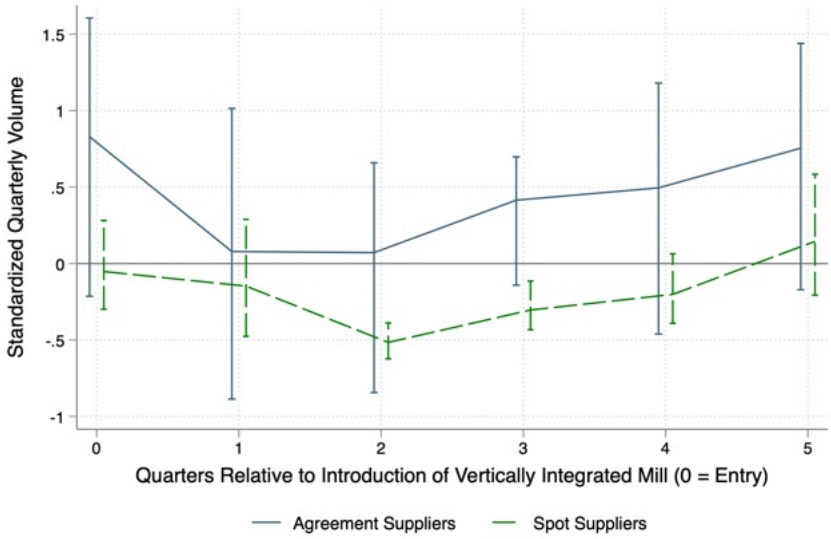
OA1.14.2 Synthetic DiD

Figure OA1.17: Synthetic DiD Estimates: Effects on Prices



Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Event study estimates use the synthetic difference-in-differences design in Arkhangelsky et al. (2021) and are estimated separately for agreement and spot suppliers. Suppliers are considered treated if they produce a fabric that is also made at the newly built integrated fabric mill. Confidence intervals are bootstrapped to account for the two-stage estimation and clustered by supplier. Volume-weighted average prices are standardized by fabric and month (as described in Section 2.3.2).

Figure OA1.18: Synthetic DiD Estimates: Effects on Quarterly Volumes



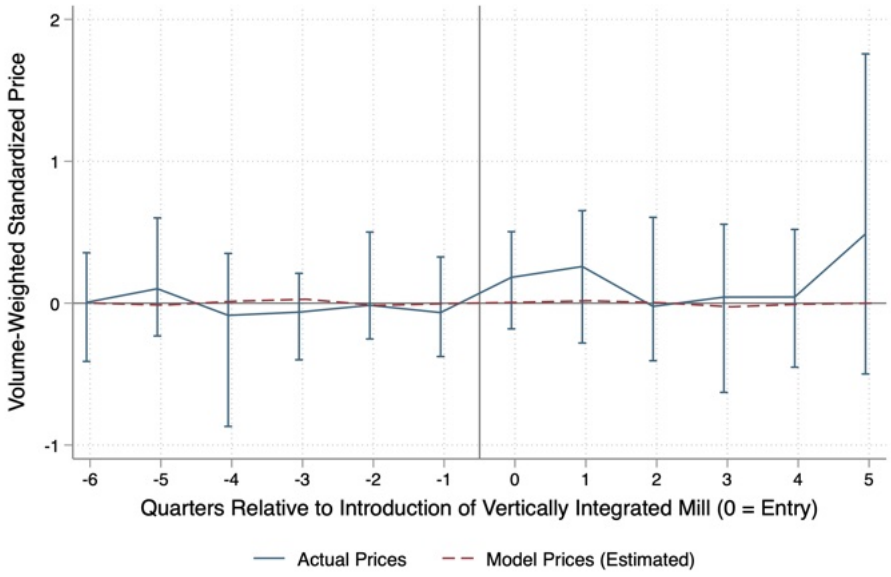
Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Event study estimates use the synthetic difference-in-differences design in Arkhangelsky et al. (2021) and are estimated separately for agreement and spot suppliers. Suppliers are considered treated if they produce a fabric that is also made at the newly built integrated fabric mill. Confidence intervals are bootstrapped to account for the two-stage estimation and clustered by supplier. Standardized quarterly volume is standardized at the supplier level (*i.e.,* each observation is the z-score of volume for that supplier, where the mean and standard deviation are calculated using pre-integration data only).

OA1.14.3 DiD Validation: Event Study Estimates as Untargeted Moments

As structural model estimation does not use the variation leveraged in the DiD design,[63] 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 OA1.19 are all within the 90% confidence intervals of event study coefficients using actual prices.

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

Figure OA1.19: DiD Coefficients as Untargeted Moments

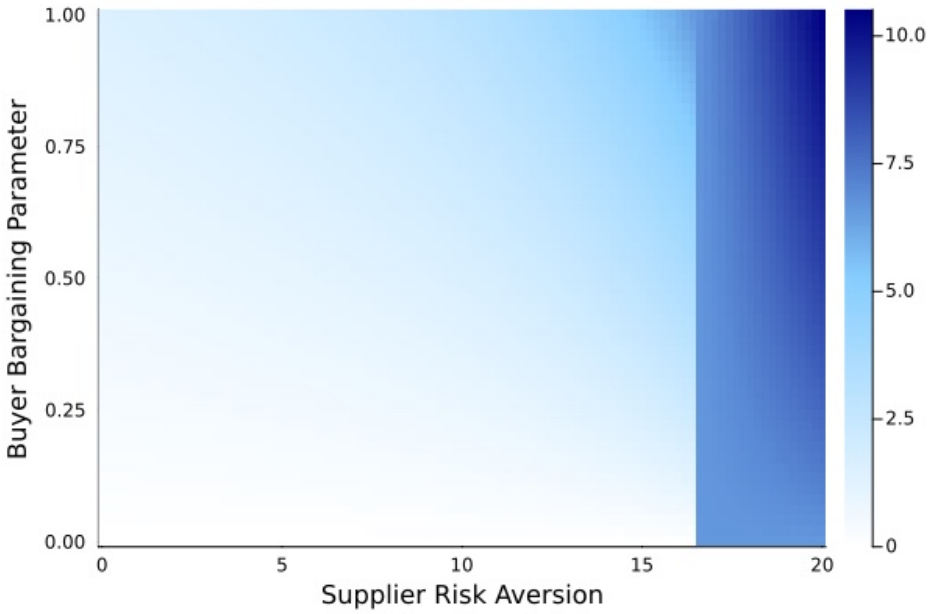


Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Event study estimates use a semiparametric difference-in-differences design (Abadie, 2005) detailed in Appendix *A*.12 and are estimated separately for agreement and spot suppliers. Suppliers are considered treated if they produce a fabric that is also made at the newly built integrated fabric mill. Confidence intervals are bootstrapped to account for the two-stage estimation and clustered by supplier. Volume-weighted standardized prices are standardized by fabric and month (as described in Section 2.3.2). Model prices are the volume-weighted average prices in standardized units implied by the structural model.

OA1.15 Agreement Discounts Incorporating DICC

The model features a cutoff level of the discount factor above which dynamic incentive compatibility constraints are slack. 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.

Figure OA1.20: Agreement Discounts (% of spot price)



Note: The heatmap shows the model-implied discount in the agreement relative to the spot price, expressed as a percent of the spot price, for different supplier risk aversion and buyer bargaining parameters, incorporating the DICC in Section A.3.1. The heatmap legend is to the right of the heatmap and units are percentages of spot prices. Model estimates consider the average capacity and quantity assurance across agreement suppliers. The discount factor for both parties is calibrated based on interest rates of Indian microfinance institutions; to be conservative, I use a low discount factor associated with 150% of market interest rates in Indian microfinance such that the quarterly discount factor is .88 (such that the yearly discount factor is .60).