Value Creation & Value Capture in Indian Garment Sector Bargaining

Russell Morton*

January 30, 2026

Click Here for Latest Version

Abstract

Low prices paid by large exporters to their many small suppliers raise concerns about unequal distribution of gains from exporting. However, when input or complementary markets (*e.g.*, insurance, credit) function poorly, buyers (exporters) adopt contractual and organizational strategies that shape prices beyond standard buyer power markdowns. This paper quantifies how prices reflect two distinct sources of buyer power arising from these firm responses to input market imperfections. First, contracts create surplus by mitigating frictions (*e.g.*, providing insurance or credit); then, low prices reflect buyer power over surplus created. Second, partial vertical integration enables surplus capture through threats to replace suppliers with in-house production. While buyer power both over surplus created and from surplus capture can lower prices, their distributional implications differ: only surplus creation aligns incentives for contracts to mitigate frictions. I estimate a structural bargaining model to decompose prices into surplus creation and capture using proprietary data on the universe of fabric purchases by a large Indian garment exporter. Estimates illustrate that buyer power over surplus creation through insurance-like contract terms, not from surplus capture, explains low prices; the buyer retains ∼60% of surplus created. Counterfactuals illustrate that introducing a missing profit insurance product raises prices paid to small risk-averse suppliers, but increasing buyer competition has limited effects.
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 Early-Career Workshop, 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), policymakers raise concerns that exporters pay low prices to their small suppliers, capturing most gains (Boudreau et al., 2023; Goldberg, 2023). Such concerns are especially compelling in LMIC supply chains, where typically a few large exporters purchase inputs from many suppliers (Zavala, 2022; Dominguez-Iino, 2023). However, when input markets—or complementary markets like insurance and credit—function poorly, large buyers (exporters) adopt contractual and organizational strategies that shape prices beyond standard buyer power markdowns.

This paper analyzes how vertical contract prices reflect two distinct forms of buyer power arising from contractual and organizational responses to poorly functioning markets in the context of Indian garment exporting. First, buyers and suppliers make agreements that mitigate frictions in input or complementary markets—a contractual strategy. For example, fabric suppliers face variable, lumpy orders, generating quantity and profit risk. Agreements can stabilize quantities, providing insurance.[1] When insurance markets are frictional, suppliers may accept lower prices in exchange for quantity stability; prices combine the payment for the input with an insurance premium. In other settings, agreements address different frictions, such as trade credit given financial frictions.[2] Then, prices bundle the input payment with interest. Therefore, discounts represent the buyer's share of surplus created by mitigating frictions inside the agreement; this surplus share reflects buyer power over *surplus creation*. Second, many buyers, including the garment exporter studied, reduce external sourcing through partial vertical integration—an organizational strategy.[3] Partial vertical integration may enable buyers to *capture surplus* through threats to replace suppliers with in-house production.[4] These threats reduce prices through buyer power from surplus capture.

Though buyer power both *i)* over surplus creation and *ii)* from surplus capture can lower prices, distinguishing between them is necessary to understand the distributional effects of vertical contracts. When low prices embody surplus created through vertical contract terms

[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) also stabilize quantities.

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

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

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

that mitigate frictions, discounts represent the buyer's share of that surplus. Conversely, replacement threats reduce prices purely through rent extraction, which I refer to as "*capture*." Crucially, only buyer power over surplus creation aligns incentives for agreement terms that mitigate frictions, such as insurance-like terms.

Buyer power in agreements has two components: bargaining weights and outside options. Bargaining weights govern the buyer's share of surplus created, while outside options govern surplus capture through replacement threats. Canonical oligopsony models (*e.g.,* Cournot or Bertrand) feature a single wedge between input prices and marginal revenue product. Therefore, such models cannot evaluate the distributional and policy concerns tied to whether buyer power reflects surplus creation or capture.

I develop a framework to decompose prices into buyer power both over surplus creation and from surplus capture that leverages proprietary data on the universe of fabric purchases by a large export-oriented Indian garment manufacturer. This decomposition overcomes two central identification challenges. First, the agreements studied are informal, with no official contract that clearly specifies terms. Therefore, I infer terms by comparing transactions across agreement and spot suppliers with identical price-relevant characteristics. Comparisons leverage rich data on ∼25,000 fabric transactions with highly granular observables (∼6,200 product categories). Two stylized facts emerge: *i)* agreements reduce, without eliminating, quantity variability; and *ii)* agreements feature nonlinear prices,[5] while spot suppliers receive flat, constant prices. Because both quantity assurance and nonlinear prices smooth profits, I interpret them as insurance-like agreement terms.

The second identification challenge arises because prices simultaneously reflect agreement surplus, outside options, and surplus sharing (the buyer bargaining parameter). I develop a new approach to separate the *level* of surplus created from the *sharing rule* that divides it.[6] My strategy leverages exogenous supplier-specific demand shocks over time. As price adjustments to demand shocks provide insurance, their magnitude identifies supplier preferences for the insurance-like terms that create agreement surplus. Conversely, average price levels identify buyer bargaining parameters.

I build a Nash bargaining model of the vertical contract that incorporates both surplus creation from insurance-like terms and surplus capture through threats to replace external

[5]The nonlinear prices are U-shaped: they are highest when profits would be lowest with constant prices, and vice versa.

[6]When the buyer does not threaten to replace the supplier with in-house production, both parties' outside option is the spot market because *i)* agreements with quantity *stability* cannot expand and *ii)* new agreements cannot emerge in the short run given the time and effort required to establish them. When the buyer threatens replacement with in-house supply, I estimate its outside option using cost data from integrated suppliers.

suppliers with in-house production. The model formalizes how price effects of insurance-like terms relate to supplier preferences for smooth profits. Such preferences can reflect various microfoundations, including borrowing and saving frictions that lead suppliers to fund frequent outflows (*e.g.*, payroll) from current profits.[7] For brevity, I refer to such preferences as "risk aversion." More risk-averse suppliers derive greater value from insurance-like terms, so agreements with such suppliers create more surplus. Because the buyer receives its surplus share through discounts, creation of larger surplus reduces prices. Therefore, surplus creation from insurance-like terms makes prices decrease in supplier risk aversion.

The model not only extends standard results about replacement threats from external suppliers (Ho and Lee, 2019; Ghili, 2022) to in-house production but also introduces a novel cutoff condition: threats reduce prices only when suppliers are sufficiently risk averse. When risk aversion falls below this cutoff, replacement threats break agreements, as the insurance-like terms generate insufficient surplus. Because quantities from agreement and integrated suppliers are inframarginal and cannot easily expand,[8] breakdown pushes sourcing to the expensive spot market rather than to in-house production.[9] Therefore, the buyer "as if" endogenously chooses its outside option to preserve agreement discounts. It uses in-house supply as a replacement threat with more risk-averse suppliers while maintaining spot supply as the outside option otherwise.[10] Overall, my model highlights how spot sourcing constrains replacement threats, in contrast to canonical models of settings with only bilateral negotiations (*i.e.,* without spot or integrated supply) where replacement threats have large effects.

I find that surplus creation alone, rather than surplus capture, generates all price reductions relative to spot prices, with the buyer retaining $\sim$60% of surplus created. This result emerges even though the setting appears favorable for buyer power from surplus capture given the large, visible integration episode ($\sim$ 20% of fabric needs). However, all supplier risk aversion estimates fall below the model-implied cutoff above which the buyer makes replacement threats. The median estimated risk aversion aligns with estimates for LMIC

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

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

[9]Integrated supplier cost convexity makes the buyer prefer increasing spot sourcing over expanding in-house production.

[10]I show that the buyer endogenously choosing its outside option is a reduced form of either of two games with both cooperative (*i.e.,* Nash bargaining) and noncooperative stages in the spirit of Miller and Watson (2013). In one game, the buyer signals its residual in-house capacity, which is private information that varies stochastically and determines the actual outside option. In the second, the buyer commits to residual in-house capacity, mirroring other agreement commitments (*e.g.,* agreement prices sometimes exceed spot).

coffee mills (Blouin and Macchiavello, 2019).[11] On average, agreements provide a 1.5% discount relative to spot prices, saving the buyer ∼0.9 million USD annually. This discount is meaningful relative to spot markups of 7% over costs.

To validate that replacement threats from in-house production do not affect prices, I use a difference-in-differences (DiD) design. Identification leverages both *i)* increased vertical integration—the buyer builds a new fabric mill—and *ii)* heterogeneity in suppliers' exposure to it.[12] This validation complements the structural analysis because the two approaches rely on distinct variation: exogenous demand shocks identify structural parameters. Consistent with model predictions, as all supplier risk aversion estimates fall below the cutoff for surplus capture through replacement threats, DiD estimates show no effects of integration on prices. Furthermore, DiD estimates are inconsistent with prices predicted by a standard replacement threat model (Ho and Lee, 2019) of a setting without a spot market. This difference highlights how sourcing marginal quantities from spot markets limits the effects of replacement threats.

Through the lens of my model, I quantify both *i)* the cutoff in risk aversion above which the buyer makes replacement threats with in-house production and *ii)* the price effects of such threats. The model-implied risk aversion cutoff exceeds both the estimates for agreement suppliers and those from other settings, although it is not implausibly higher than the latter.[13] In (hypothetical) agreements with suppliers sufficiently risk averse to face replacement threats, discounts reach 6.7% of spot prices. This magnitude is economically meaningful—it exceeds Indian value-added tax for apparel in my sample period (5%).[14] Moreover, such highly risk-averse suppliers are likely small firms with policy and empirical relevance: 95% of Indian industrial units have 10 employees or fewer (Rajagopalan and Shah, 2024).

Last, I use the model to evaluate two counterfactual policies to increase prices, focusing on the role of supplier risk aversion. First, I consider an increase in buyer competition (*e.g.*, from export promotion), modeled as reducing buyer power through a higher supplier bargaining weight.[15] For a supplier with median estimated risk aversion, increased competition reduces the agreement discount relative to spot prices by 75%. However, when supplier risk aversion exceeds the replacement threat cutoff, surplus from the agreement relative to

[11] These mills report average revenues of ∼3.5 million USD per year. Although large by LMIC standards, these mills are much smaller than the buyer, which buys ∼180 million USD of fabric annually.

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

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

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

[15] This change in the bargaining weight is interpreted as a surplus share change to prevent long-run switches to another buyer. As agreements require time and effort to establish, short-run switches are not possible.

in-house production is so small that even if the supplier received all surplus, prices barely shift. Therefore, ignoring replacement threats overstates the benefits of competition because, absent replacement threats, suppliers' outside option is the spot market. Then, with more surplus available to distribute, reallocation meaningfully affects prices.

Second, introducing the missing product for profit insurance against spot market quantity risk makes prices increase, especially for risk-averse suppliers. Insurance enables suppliers to smooth profits outside the agreement, strengthening their outside option and removing the buyer's ability to extract surplus through replacement threats. Intuitively, this counterfactual eliminates buyer power from being the sole provider of insurance. While moral hazard may limit the feasibility of profit insurance, policies could support profit smoothing by improving financial market access (facilitating self-insurance), supporting cooperatives to stabilize profits, and developing products to hedge risk (*e.g.*, derivatives). The effectiveness of insurance highlights that helping risk-averse firms to smooth profits delivers benefits not only by equalizing utility across states but also by raising prices. Therefore, insurance delivers not only well-known efficiency gains from smoothing, but also advances equity by counteracting market power from replacement threats.

Overall, this paper shows that concerns about low supplier prices should distinguish whether prices reflect surplus *creation* inside vertical agreements or surplus *capture* through rent extraction. Low prices remain a distributional concern regardless of their source, but when discounts compensate suppliers for insurance- or credit-like terms, they need not imply anticompetitive behavior in input markets. Accordingly, policies should address the missing complementary markets as their absence gives rise to buyer power over surplus creation. This result emphasizes the importance of this form of buyer power, which is understudied in LMIC supply chains (*e.g.*, Boehm and Oberfield (2020); Zavala (2022); Boudreau et al. (2023); Dominguez-Iino (2023); Bergquist et al. (2023); Startz (2025)).

This paper contributes to our understanding of buyer power in vertical contracts (*e.g.,* Grennan (2013); Rubens (2023); Treuren (2025); Demirer and Rubens (2025)) . I emphasize how poorly functioning spot markets lead vertical contract prices to reflect buyer power both *i)* over surplus created by mitigating frictions within the agreement and *ii)* from threats to replace suppliers with in-house production. Distinguishing between these sources matters because their distributional implications differ. To quantify how surplus creation affects prices, I incorporate risk aversion into my empirical model, as insurance creates surplus in my setting. I develop a new identification strategy to separate risk aversion from bargaining parameters that leverages how demand shocks affect both price levels and curvature. Al-

though insurance is central to vertical relationships (Lafontaine and Slade, 2007), empirical models often abstract from risk aversion and insurance (*e.g.*, Cuesta et al. (2019); Lee et al. (2021)). Such models may underestimate the more risk-averse party's bargaining parameter, generating biased price estimates. Additionally, my model shows that marginal sourcing from spot markets constrains replacement threats, compared to their large effects in standard models of settings without coexisting spot markets (Ho and Lee, 2019; Ghili, 2022).

Last, I advance the literature on LMIC vertical contracts, which shows that contracting frictions lead agreements to use relational enforcement from future surplus rather than formal enforcement from legal institutions (*e.g.*, Macchiavello and Morjaria (2015, 2020); Ghani and Reed (2022); Brugues (2025)). Instead of focusing on enforcement, this paper both: *i)* highlights how agreements create surplus by mitigating frictions other than contracting, such as missing insurance markets; and *ii)* quantifies sharing of agreement surplus. To do so, I estimate a bargaining model rather than using standard take-it-or-leave-it (TIOLI) relational models. While TIOLI models are well-suited to study enforcement, they cannot address distributional questions because they allocate all surplus to one party. Additionally, I examine how vertical agreements respond to changes in the *vertical* organization of production, while prior work generally considers effects of changes to market structure—the *horizontal* organization of production. As vertical integration—in contrast to market structure—is a strategic firm choice that can facilitate rent extraction via replacement threats, it introduces novel strategic considerations. This co-existence of integration and relational contracts is empirically relevant as both are responses to frictional markets (Morton, 2023).

The remainder of the paper is organized 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 provide insurance to suppliers. I then describe the increase in vertical integration from the buyer's construction of a new mill. Last,

I depict the buyer's overall fabric sourcing, illustrating that spot suppliers provide *marginal* volume.

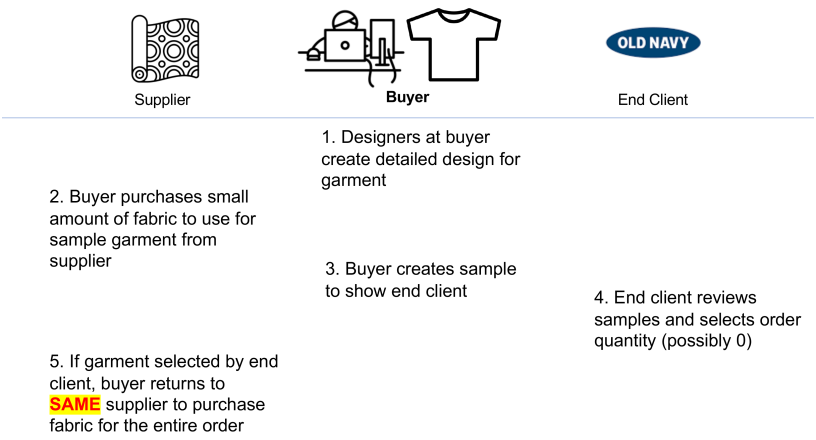
2.1 Data

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; therefore, I directly observe their markups.[16]

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

2.2 Context: Fabric Procurement and Garment Manufacturing

Figure 1: Garment Manufacturing Process



[16]Cost estimates correspond to a mill that is not capacity constrained and represent average variable costs.
[17]Data are a division-by-month panel from March 2019 to April 2020, where divisions refer to specific processes for fabric production without substitution across each other, such as weaving, dyeing, and knitting.

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, approves a subset, and specifies quantities. 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.[18] 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.[19]

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.[20] (Online Appendix OA1.4 discusses why prices for the full order are not set before sampling.)

2.3 Stylized Facts of Buyer-Supplier Agreements

These two features of garment manufacturing—quantity risk and high prices from hold-up—suggest that agreements providing suppliers insurance in exchange for lower prices could benefit both parties. Consistent with this intuition, I document two stylized facts: *i)* agreements feature quantity stability, and *ii)* their prices are U-shaped in capacity utilization, with a discount on average for the buyer.

Because agreements are informal and unwritten, I infer their terms empirically by comparing transactions across agreement and spot suppliers. As the data do not directly indicate supplier type, I classify suppliers by transaction frequency, as is standard in the literature.[21]

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

[19]In the rare cases when the buyer switches suppliers after sampling, production is typically brought in-house and for orders with unusually distant delivery. This strategy facilitates experimentation and control to match the fabric.

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

[21]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 reflects that agreements are relational contracts; therefore, agreements require repeated interactions to sustain the future surplus necessary for self-enforcement. I classify suppliers transacting every month as agreement suppliers.[22] 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

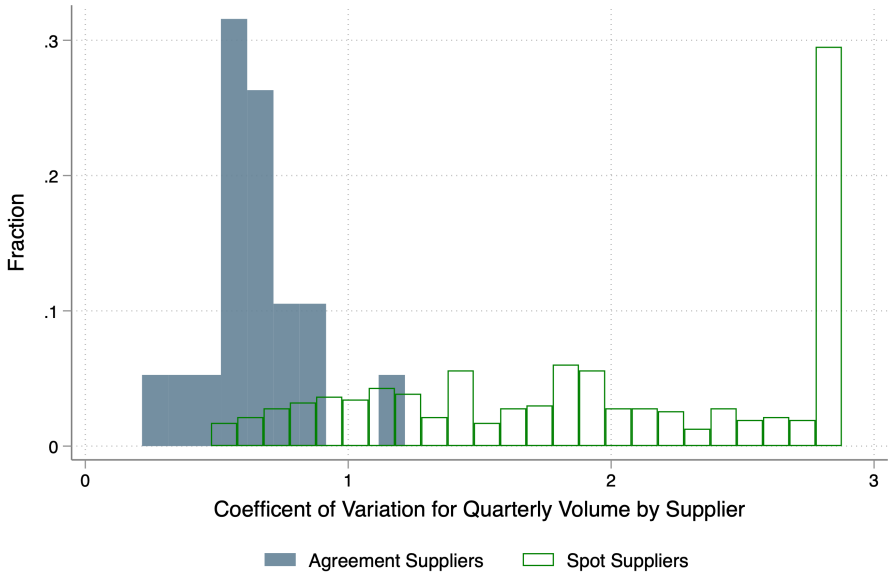
First, I document that agreements provide quantity assurance (*i.e.,* reduce quantity variability). Figure 2 documents that the coefficients of variation for quarterly volume per supplier are noticeably lower for agreement than for spot suppliers.[23] This result is not entirely mechanical: supplier classification depends on transaction frequency, not volume, so agreement suppliers could have volatile quantities. 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.[24] The buyer can provide quantity assurance through at least two channels: *i)* increasing the number of samples and *ii*) matching the supplier to end clients with different expected quantities.

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

[23]The difference in average coefficient of variation by supplier type is statistically significant ($p < 0.001$).

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

Figure 2: Coefficient of Variation and Supplier Type



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 ($N = 471$). The difference in average coefficient of variation by supplier type (agreement vs. spot) is statistically significant ($p < 0.001$).

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[25] and fabric input prices.[26] Quantities are expressed as quarterly volumes relative to supplier productive capacity—effectively, quarterly supplier-specific *capacity utilization*—to account for heterogeneity in supplier scale.

[25]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).

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

Figure 3: Pricing Function: Binned

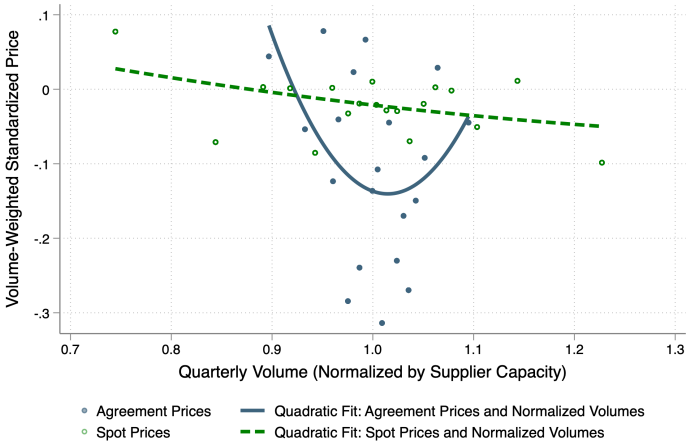
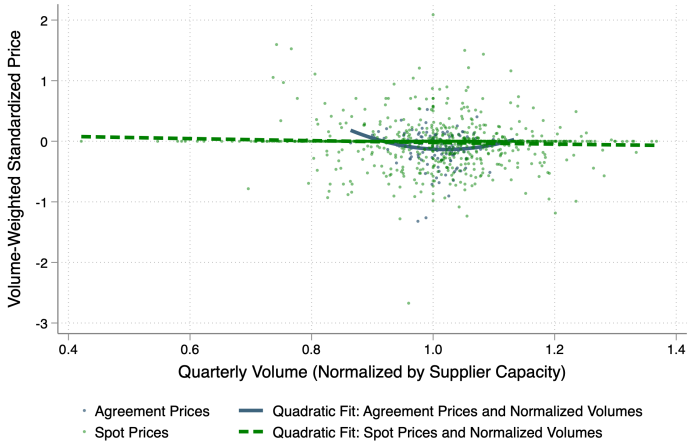


Figure 4: Pricing Function: Raw



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 4.1.2). 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. Appendix Section A.1.3 confirms that only agreement, and not spot, prices statistically exhibit a U-shape, and that this pattern is only stronger with supplier and quarter fixed effects. Furthermore, average agreement prices are below spot prices, highlighting that agreements include a discount (shown in Online Appendix OA1.1).

2.4 Interpretation as Insurance-Like Agreement Terms

I interpret both quantity assurance and the U-shaped pricing function as insurance-like agreement terms that smooth supplier profits. First, quantity variability generates profit variability.[27] Therefore, quantity assurance stabilizes profits (but agreements do not fully

[27]Quantity variability generates profit variability when quantity changes are not offset by markup changes.

eliminate variability because orders are lumpy and cannot be split across suppliers).[28]

Second, the U-shaped pricing function also reduces profit variability relative to constant prices. When quantities are low, profits with constant prices would be relatively low because few units are sold; higher prices for low quantities smooth profits. When quantities are high, high costs (*e.g.,* due to capacity constraints, and verified empirically using data from in-house suppliers in Figure 10),[29] erode margins under constant prices. Higher prices at high capacity utilization offset these costs, further stabilizing profits.[30] 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.[31] (Separately, I interpret constant spot prices as indicating that spot suppliers hold up the buyer. Therefore, spot prices are determined by the *buyer's* outside option, which is independent of a single supplier's capacity given its small size relative to the buyer.)

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 must simultaneously account for the facts that: *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, flat spot prices reject that the U-shape reflects a constant markup over U-shaped total cost curves.[32]

[28]Quantity assurance is limited because the buyer's primary instruments to stabilize quantities—frequency of sampling and the distribution of likely quantities for a given end client matched to the supplier—do not perfectly match a specific quantity.

[29]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.

[30]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.

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

[32]Even more complex stories also struggle to rationalize these facts. I continue with the example where U-shaped prices reflect constant per-unit markups over U-shaped cost curves. Then, to rationalize high, flat spot prices, agreement suppliers must have *differential* market power—less market power than spot suppliers, who hold up the buyer. Quantity assurance, in that case, could simply reflect the buyer trying to minimize prices by keeping suppliers at capacity levels with low costs. However, if the buyer has relatively more market power with agreement suppliers, what is the source of the buyer power? The most natural explanation is that the buyer provides some additional benefit to such suppliers. My analysis suggests that insurance is this extra benefit. Fabric differentiation seems unlikely to create market power: the median (mean) agreement supplier competes with spot suppliers for 76% (69%) of its volume.

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). More efficient suppliers might pass on cost savings to the buyer, reducing average prices. However, efficiency alone does not explain why agreements feature more stable, rather than larger, quantities nor U-shaped prices. Regardless, to test this hypothesis, I replicate the comparisons between agreement and spot suppliers, but restricted to only high-volume spot suppliers ("placebo suppliers").[33] Alternative plausible explanations of unobserved heterogeneity largely imply that supplier volume, rather than transaction frequency, should shape transactions. If supplier efficiency explains lower prices, then suppliers with large volumes—regardless of transaction frequency—should charge lower prices. Instead, Online Appendix OA1.2 shows that differences in quantities and prices persist, which implies that they arise from agreements rather than correlates of size.[34]

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.[35] 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 the hold-up problem through supplier competion on prices *before* sampling (possibly allowing quantity-contingent prices given cost convexity)? Second, why does the buyer not source only through agreements to avoid high spot prices?

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

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

[35] Specifically, they shift at the threshold between transacting every month and transacting nearly every month.

2.5 Context: Vertical Integration at a Large Buyer

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 average 7% over average variable costs.) Figure 5 shows the new mill. The organization of the mill emphasizes the buyer's sophistication (and scale), contrasting with the description of messy and disorganized Indian garment manufacturers in Bloom et al. (2012).

Figure 5: Integrated Supplier Constructed in June 2018

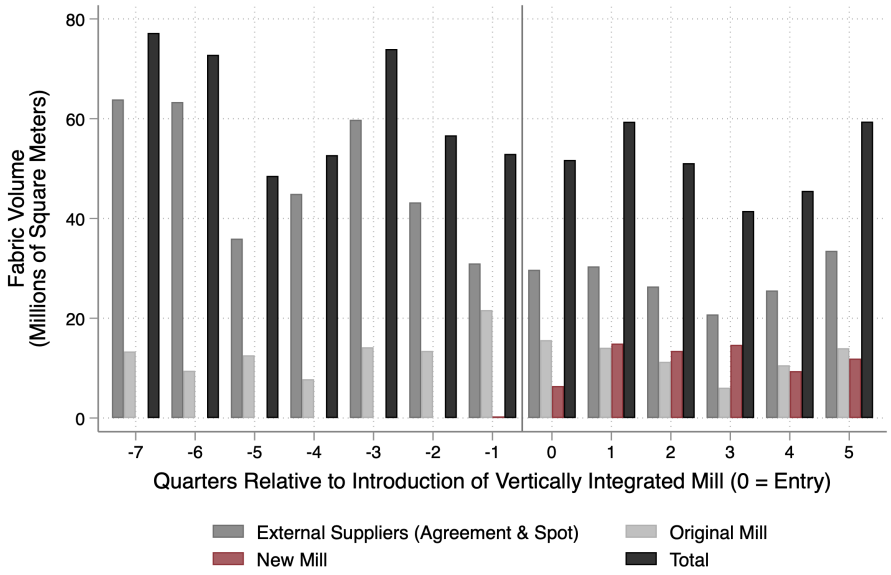


2.6 Context: Overall Sourcing

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.[36]

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

Figure 6: Fabric Sourcing Over Time

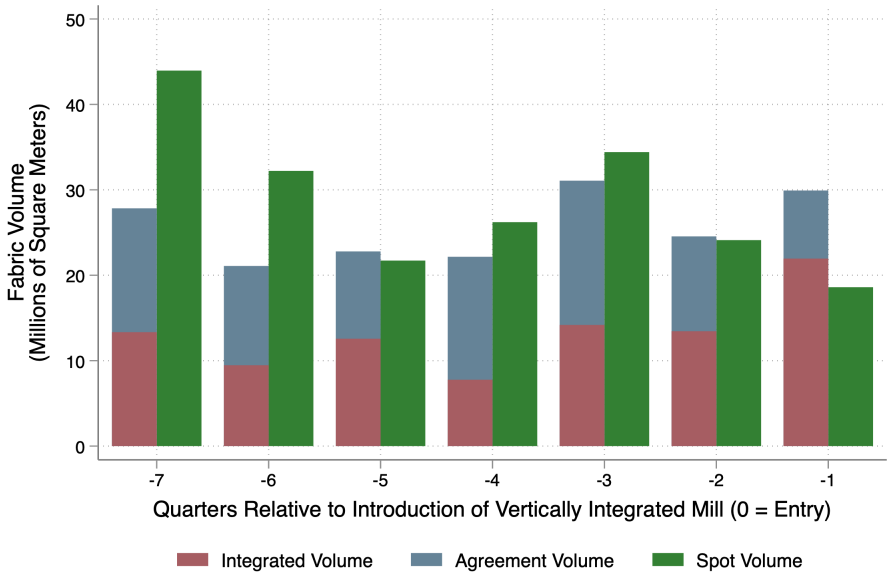


Note: Data from a panel constructed from the universe of fabric transactions by the buyer at the supplier-type-by-quarter level. Bars represent the quantity of fabric purchased during the quarter from each supplier type, summing across all suppliers of the type. Volume from external suppliers sums across 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 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 $\sim$20% of the buyer's fabric, while its largest external supplier provides only $\sim$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.[37]

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

Figure 7: 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 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 ($R^2 = 0.09$).[38] These results reflect the buyer using agreement and integrated suppliers for *inframarginal* sourcing, while spot suppliers fulfill *marginal* orders.[39]

Two central institutional features rationalize the use of integrated and agreement supply

[38]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%.

[39]Online Appendix 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.

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

This section develops a theoretical model that formalizes how prices reflect both *i)* surplus creation from insurance-like agreement terms and *ii)* surplus capture through threats to replace suppliers with in-house production. The model extends standard models of replacement threats by incorporating threats from in-house, rather than external, suppliers. It also reveals a new result: replacement threats reduce prices *only* when supplier risk aversion exceeds a cutoff. In agreements with suppliers whose risk aversion falls below this cutoff, the buyer as-if endogenously refrains from making replacement threats because such threats would break the agreement.

The model incorporates the empirically observed insurance-like agreement features—both quantity assurance and U-shaped prices—and buyer threats to replace suppliers with in-house production. Recall that quantity assurance in the agreement provides *stable*, but not necessarily *larger* quantities. Accordingly, I model both agreement and spot quantities as random draws from distributions with the same mean, but higher variance for spot quantities. U-shaped prices arise as the buyer and supplier bargain over a per-unit agreement pricing *schedule* $p^A(q^A)$ rather than a single scalar price, as standard in Nash bargaining. This pricing function $p^A(q^A)$ specifies, before quantities are known, how unit prices vary with the period's quantity realization. Last, the buyer's replacement threats emerge from its outside option: integration enables the buyer to use in-house production instead of the spot market.

While the model is static to transparently isolate surplus creation from surplus capture, Appendix A.3.1 verifies results are robust to incorporating dynamic incentive compatibility constraints for relational enforcement (as in Macchiavello and Morjaria (2023)). The static model can be interpreted as bargaining over long-run surplus in established agreements at their steady state, with relational contracts enforcing the cooperative agreement modeled.

3.1 Model Setup

Equation 1 specifies the bargaining problem.

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

Buyer surplus is the difference between the cost of purchasing quantity q^A at price $p^A(q^A)$ in the agreement and the outside option value (O_B). I model the buyer as risk neutral given its large size and sophistication.[40] Buyer surplus reflects only procurement costs, as revenue from end clients (each garment's price and quantity) is determined before fabric procurement.

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

Two insurance-like agreement terms generate surplus, thereby shaping agreement prices. First, the U-shaped pricing function smooths profits, increasing supplier utility by reducing profit variability. Second, quantity assurance generates additional surplus through both insurance and cost reduction; the latter means that agreements benefit even risk-neutral suppliers. With convex costs,[43] Jensen's inequality implies that reducing quantity variance lowers average costs: high costs at large quantities are not correspondingly offset by lower costs at small quantities. Cost convexity is documented in Figure 10.

Bargaining parameters determine surplus shares. The buyer bargaining parameter is α_B. Higher values indicate that the buyer receives a larger share of surplus created by the

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

[41] Alternatively, firms may inherit owners' preferences over profits because small suppliers' profits are likely their owners' primary income source. I do not take a stance on the specific source of these preferences.

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

[43] 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.

agreement. As the buyer and supplier split the surplus, their bargaining parameters sum to one (*i.e.,* the supplier bargaining parameter is $1 - \alpha_B$).

The outside option for both the buyer (pre-integration) and supplier is transacting in the spot market. Because agreements are relational contracts, which require time to establish, neither the buyer nor the supplier can switch to a new agreement partner in the short run. Furthermore, because agreements feature quantity assurance, which reduces variability without necessarily increasing volumes, average agreement quantities are fixed.

I model the spot market as perfectly elastic with constant price p^S, consistent with flat spot prices in Figure 4.[44] Flat spot prices reflect that spot suppliers extract all surplus because they can hold-up the buyer as the only viable suppliers once the end client selects the sample. When spot suppliers hold-up the buyer, spot prices make the buyer have no surplus; therefore, spot prices depend on the buyer's outside option, which does not vary with supplier capacity utilization. Furthermore, even though the buyer is large, it is small relative to the Indian fabric market: its market share is less than 0.5%.[45]

After integration, the buyer can use in-house supply as its outside option. I model the integrated supplier having the same production technology as external suppliers (*i.e.,* the same cost function), as the buyer reports that all relevant suppliers employ (near-)frontier technology. Morton (2023) empirically documents strong substitutability between integrated 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 heterogeneity between internal and external suppliers (see Appendix OA1.5.2).[46]

Quantities reflect stochastic end-client demand rather than endogenous buyer or supplier choices. Agreements specify the *ex-ante distribution* of quantities, with quantity assurance leading to lower variance (but not necessarily lower average volume). Formally, I model spot quantities as a mean-preserving spread of agreement quantities, with mean quantities for both distributions at 100% of capacity utilization—the supplier optimal capacity utilization.

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

[45]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).

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

Quantities can be interpreted relative to the capacity that the supplier reserves for the buyer. Therefore, the importance of this specific agreement can affect supplier risk aversion (*i.e.,* suppliers with diversified buyers might be less risk averse over profits from any one buyer).[47]

Formally, the outside options are:

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

The bargaining model flexibly includes both buyer and supplier market power. In bargaining models, market power is shaped by both surplus shares (*i.e.,* the bargaining parameters) and outside options. I interpret these two objects that encode market power differently. The buyer bargaining parameter represents the buyer's share of surplus created, thereby aligning incentives for insurance provision when positive. Conversely, replacement threats that improve the buyer's outside option affect prices purely by extracting rents, capturing surplus. Supplier market power appears in the model both in the agreement through the supplier bargaining parameter (one less the buyer bargaining parameter) and in the spot market through hold-up, which leads to high spot prices.

The model does not consider buyer marginal costs of insurance provision for two reasons. First, as agreement quantities are inframarginal, the buyer can fulfill agreement quantity targets by reallocating sourcing across suppliers. Therefore, providing quantity assurance need not induce any marginal costs to the buyer.[48] Similarly, U-shaped agreement prices that also act as insurance need not change the buyer's average sourcing costs. It follows that U-shaped prices do not generate marginal insurance provision costs to the buyer. Second, theoretically, (non-decreasing) marginal costs of insurance provision only magnify the importance of surplus creation. With such costs, the buyer provides insurance until its marginal benefit meets marginal cost. Then, buyer power over surplus creation actually *increases* insurance provision because the buyer captures more of the incremental surplus created by insurance. Therefore, buyer power over surplus creation has even larger price effects with

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

[48] The buyer's insurance provision costs are presumably fixed costs per agreement supplier to meet quantity guarantees (such as coordination costs). These fixed costs likely increase as the sum of agreement supplier quantity targets approaches the buyer's overall volume. Such fixed costs are sunk for incumbent agreement suppliers, and help explain incomplete sourcing from agreement suppliers (see Appendix OA1.5).

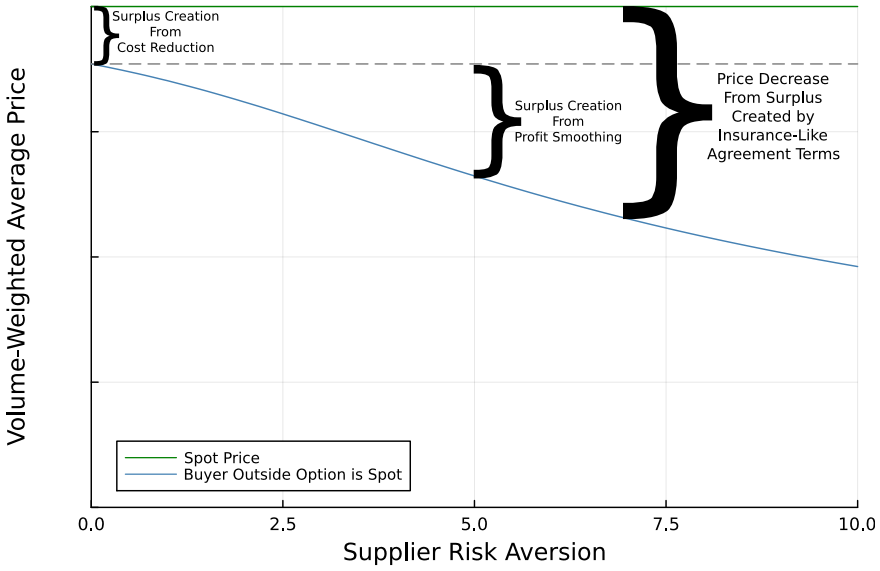
marginal insurance provision costs.

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 Surplus Creation & Capture

Figure 8 illustrates how supplier risk aversion affects the buyer's expected price (*i.e.,* the volume-weighted average fabric price), holding the buyer bargaining parameter constant.[49] Purchases from the spot market do not incorporate bargaining, so prices do not depend upon supplier risk aversion. When the buyer's outside option is the spot market, surplus creation in the agreement generates discounts relative to spot prices.[50] Then, some surplus created in the agreement accrues to the buyer as a discount. The agreement discount with a risk-neutral supplier (*i.e.,* when the CARA risk aversion coefficient is 0) reflects surplus generated only through cost reductions from quantity assurance. As supplier risk aversion increases, the agreement creates additional surplus by smoothing profits through both insurance-like terms: *i)* quantity assurance and *ii)* U-shaped prices. Consequently, the discount grows.

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



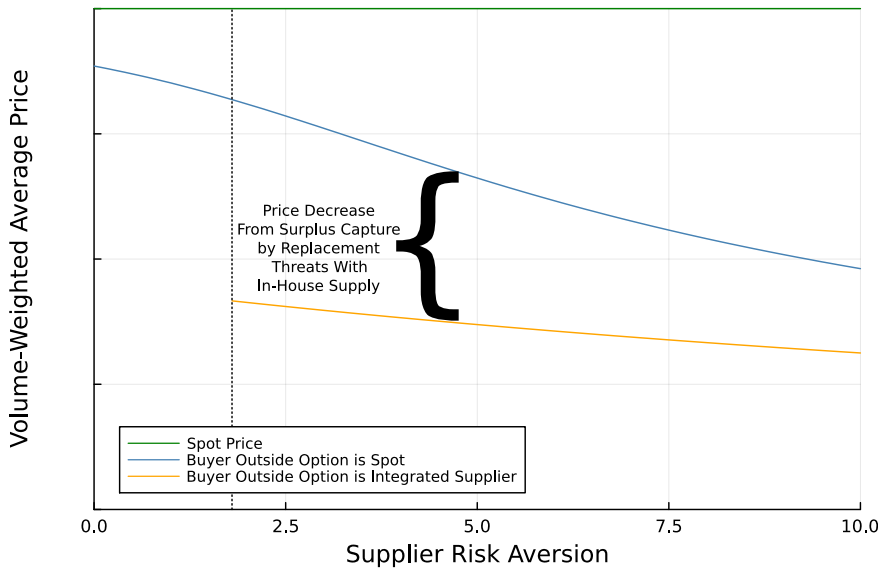
[49] Online Appendix OA1.7 shows similar results but varying the buyer bargaining parameter and holding risk aversion constant. Although both influence prices, only risk aversion affects surplus creation from insurance-like agreement terms and capture through strategic replacement threats.

[50] Discounts emerge for any positive buyer bargaining parameter—the sufficient condition for buyer power over surplus creation.

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.

Figure 9 shows the price effects of surplus capture through threats from the buyer to replace the supplier with in-house production—when the buyer's outside option switches to its integrated supplier. The contracting environment constrains the benefits of replacement threats: there is a cutoff level of risk aversion below which the agreement breaks in response to replacement threats. This breakdown occurs because suppliers with low risk aversion receive insufficient surplus from insurance-like agreement terms. Consequently, no pricing function generates positive surplus for both parties.

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).

When supplier risk aversion falls below this cutoff, the buyer would prefer to retain its pre-integration agreement because the agreement includes a discount. However, in Nash bargaining, outside options are not choices but primitives that represent the value of the disagreement outcome. But, in practice, bargaining parties may have incentives to strategically present specific outside options, especially when the credibility of the outside option is private information of one party. To formalize this intuition, I show that the buyer endogenously choosing its outside option can be interpreted as the reduced-form representation of two distinct extensions of the Nash bargaining game with different information structures and timing assumptions. Both extensions consider a *contractual equilibrium* that includes both the cooperative Nash bargain and a noncooperative phase (Miller and Watson, 2013).

The first extension introduces a noncooperative signaling game about residual capacity at the in-house supplier before cooperative bargaining occurs, consistent with signaling game models of incomplete information bargaining (*e.g.,* Backus et al. (2019)). Residual capacity determines whether the in-house supplier is the buyer's true outside option. It is also the buyer's private information, as only the buyer knows other orders at the in-house supplier, and varies stochastically with end-client demand. The signaling game incorporates beliefs about residual capacity in a principled way. (While Nash bargaining is axiomatic with perfect information, Lee et al. (2021) highlight that it has similarities with passive beliefs). I establish the existence of a weak perfect Bayesian Equilibrium where the buyer claims that its in-house supplier is not capacity constrained, but only when bargaining with suppliers that value nonprice agreement terms above a cutoff (shown in Appendix A.5). This cutoff valuation of nonprice agreement terms corresponds to the cutoff level of supplier risk aversion above which the buyer captures surplus through replacement threats, as risk aversion determines surplus from insurance-like terms that smooth, without necessarily increasing, profits.

The second extension, in Appendix A.6, considers a noncooperative game where the buyer commits to a level of residual capacity before cooperative bargaining. This commitment mirrors the buyer's commitment elsewhere in the agreement, such as paying agreement prices over spot prices for some capacity realizations. I similarly show the existence of an

equilibrium where the buyer makes replacement threats only to suppliers whose surplus from nonprice terms exceeds a cutoff.

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

4 Structural Model

This section builds and estimates a structural model closely tied to the theoretical model to decompose agreement discounts, relative to spot prices, into surplus creation from insurance-like terms and surplus capture through replacement threats with in-house production. The decomposition requires estimating the supplier risk aversion and buyer bargaining parameters that rationalize observed prices. These estimates facilitate the empirical decomposition. I then use the model to examine how the decomposition of discounts varies for different supplier types. Last, I evaluate policy counterfactuals to increase prices that small suppliers receive through the lens of the model, with a focus on the importance of supplier risk aversion.

4.1 Estimation Approach

4.1.1 Buyer and Supplier Primitives

The structural model features three parameter groups: supplier-specific risk aversion parameters; agreement-specific buyer bargaining parameters; and a common tuning parameter. Estimation proceeds in two steps. First, an inner optimization recovers the risk aversion and bargaining parameters that best fit prices, given the tuning parameter. The outer optimization selects a common tuning parameter using out-of-sample post-integration data. I use a tuning parameter to penalize deviations from risk neutrality for two reasons. First, firms should not deviate extremely from risk neutrality. Second, I guard against overfitting the data given the limited number of observations per supplier.[51]

The inner optimization has three components. First, for each agreement, I derive the model-implied agreement pricing function $\hat{p}^A(q^A; \mathbf{\Omega}_s)$ that maximizes the Nash product in (1), where $\mathbf{\Omega}_s$ represents the parameter vector for each relationship (*i.e.,* the supplier risk aversion and buyer bargaining parameters). The agreement pricing function depends on

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

these parameters because their values affect the Nash product and, consequently, the optimal pricing function.[52] Second, I estimate prices for observed agreement quantities with the model-implied pricing function. Last, I compute the objective function for each supplier—the root-mean-squared error (RMSE) between estimated and observed prices in pre-integration quarters. The objective function adds a LASSO-style penalty—the absolute value of the risk aversion parameter weighted by the tuning parameter—to discourage extreme deviations from risk neutrality (zero in CARA utility).

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

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

S denotes the set of suppliers, $\{q^A_{s,t}\}$ are the observed quantities for agreement supplier s,[53] 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.

The outer optimization finds the tuning parameter that maximizes out-of-sample fit in thte post-integration quarters, similar to a holdout sample used 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:

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

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

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

4.1.2 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
$\frac{\mathbb{V}[q^M]}{\mathbb{V}[q^{RC}]}$	Variance in Quantities for Spot Suppliers Compared to Variance in Quantities for Agreement Suppliers	Transaction Data
p^M	Ratio of Prices from Spot Suppliers to Estimates from Integrated Supplier	Transaction Data

I estimate the cost function using panel data on capacity and costs from integrated supply. As the buyer's procurement staff indicate that agreement and integrated suppliers use similar technology, their cost functions should be comparable. I estimate the total cost function as $c(q) = \phi_0 + \phi_1 q + \phi_2 q^2$, where ϕ_2 represents convexity. I do not impose convexity (*i.e.,* $\phi_2 > 0$), and estimates are equivalent to a second-order Taylor approximation of the true cost function. The panel data on cost and capacity are at the division-by-month level from April 2019 to March 2020, where division corresponds to a broad fabric group (*e.g.,* knits and striped fabrics).[54]

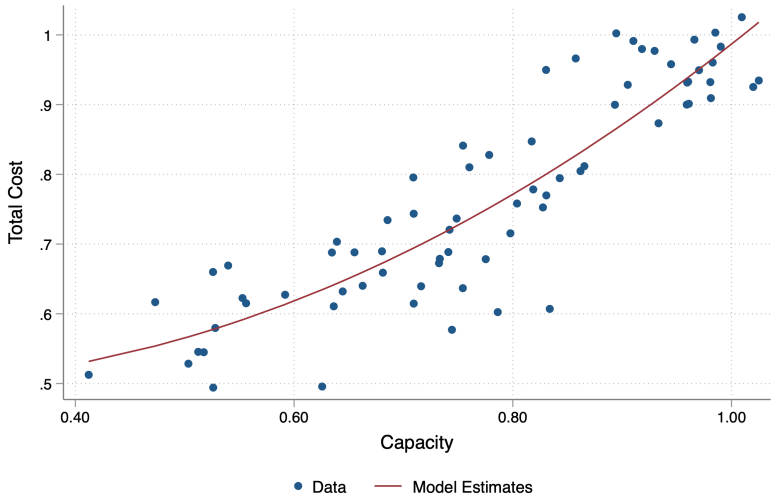
Identification of the cost function leverages exogenous capacity shocks that arise because sampling fixes the supplier before final order quantities are realized. Therefore, realized

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

quantities reflect stochastic end-client demand rather than buyer choice.[55] As deviations from the quantity target are plausibly exogenous, ordinary least squares (OLS) regressions provide causal estimates of the effect of capacity on costs.

Estimates demonstrate that the model fits the data well, as shown in Figure 10. The R^2 is .795 and the RMSE is .07, which is small relative to overall cost variation. Additionally, the estimated cost function is convex, with the coefficient on q^2 being .7815 (p-value = 0.02).[56]

Figure 10: Cost Convexity of the Integrated Supplier



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 the volume-weighted average variable cost for that division-month, expressed as a percent of the average volume-weighted cost for the division for comparability across divisions. Capacity utilization for that division-month is reported directly in the data as a percentage of capacity used.

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

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

[56]I transform costs into percentages relative to mean costs for each fabric group, as costs are reported by fabric group. This conversion facilitates simultaneous analysis of all fabrics. The strong fit suggests that heterogeneity by fabric type is unlikely to be a central feature of the cost function.

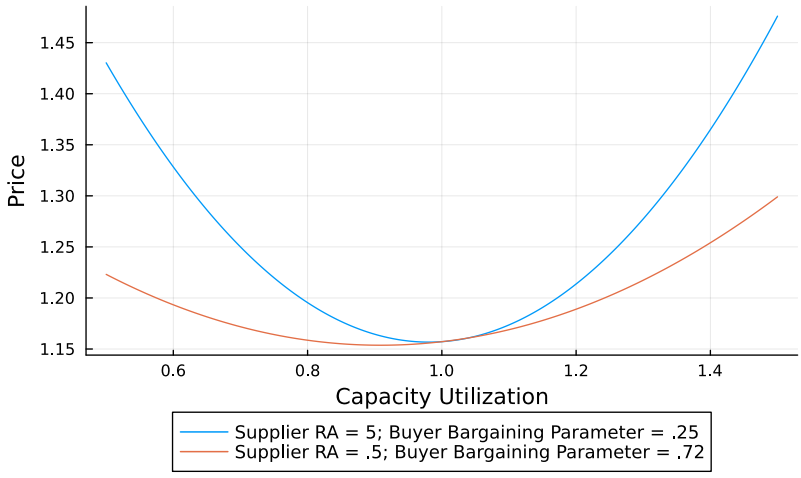
4.2 Identification of Risk Aversion & Bargaining Parameters

Average agreement prices alone cannot separately identify supplier risk aversion from the buyer bargaining parameter as increasing either parameter individually reduces prices. When the supplier is more risk averse, the agreement creates more surplus; since the buyer receives a share of that surplus, higher risk aversion lowers prices. Similarly, when the buyer bargaining parameter is higher, the buyer's share of surplus increases, reducing prices.

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

With only one price-capacity tuple, it is impossible to separate supplier risk aversion from the buyer bargaining parameter. For any 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 0.25 and a supplier risk aversion of 0.5 with a buyer bargaining parameter of 0.72 result in the same price when capacity utilization is 1.

Figure 11: Convex Pricing Function for Different Parameters



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 0.25. Then, I shift the risk aversion to .5 and find the buyer bargaining parameter that results in the same price when capacity utilization is 1.0, which is 0.72. 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, a *second* price-capacity tuple for a *different* capacity enables identification. Lower risk aversion yields a flatter, less U-shaped contract, while a higher bargaining parameter shifts prices downwards.[57] Figures 12 and 13 illustrate this identification: shifting supplier risk aversion primarily affects the shape of the pricing function, whereas shifting the buyer bargaining parameter mainly changes the average price level.

Figure 12: Shift Supplier Risk Aversion

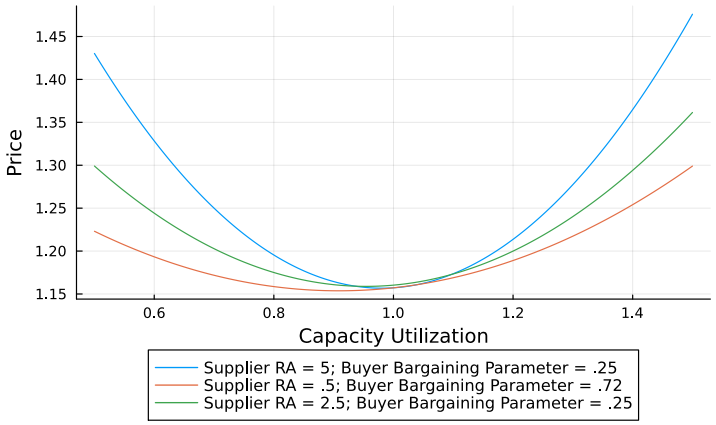
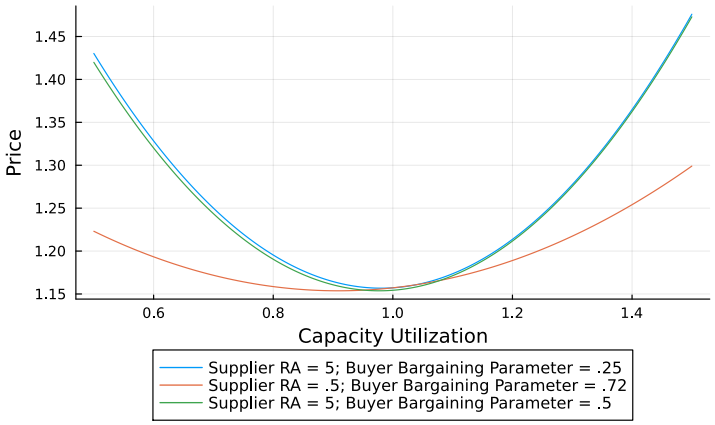


Figure 13: Shift Buyer Bargaining Parameter



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 shifts either 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. The nature of garment production—where the buyer and supplier commit to working together before knowing which garment(s) will be ordered, if any, nor what their quantities will be—limits the extent to which agreements can provide quantity assurance. Namely, agreements only reduce, without eliminating, quantity variance. I interpret

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

deviations from targeted agreement volumes as exogenous because neither the buyer nor the supplier chooses which garments are ordered (nor their quantities); instead, orders reflect demand faced by end clients. Furthermore, end clients choose garments after suppliers are chosen through sampling. Appendix A.8.1 addresses the main threat to exogeneity: that agreements include dynamic quantity assurance across quarters (*i.e.,* a quarter following one with an extreme capacity realization is closer to the target).

4.3 Risk Aversion and Bargaining Parameter Estimates

Parameter estimates in Table 2 (confidence intervals in Appendix Table A.10) show that risk aversion among agreement suppliers resembles that of the LMIC coffee mills in Blouin and Macchiavello (2019) (.0085 for the median agreement supplier and .0068 for coffee mills).[58] These coffee mills, described as "large firms by developing country standards," have ∼$3.5 million in annual sales and ∼$2 million in total assets. Given similarities to the suppliers I study in terms of (likely) size, sophistication, and setting—with respect to both market institutions (as firms in LMIC) and economic activity (contracts with buyers in an export-oriented setting)—the similarity of the risk aversion estimates seems reasonable. Although these results suggest that agreement suppliers are large firms, especially by LMIC standards, the buyer studied is much larger: it purchases ∼$180 million yearly on fabric, and its fabric expenses must be smaller than revenue.[59]

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

[58] Blouin and Macchiavello (2019) and other studies estimate risk aversion using different utility functions. For comparability, I convert reported risk aversion parameters to a CARA coefficient given their estimate and utility functions. Appendix Table A.3 shows (converted) CARA risk aversion parameters.

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

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, with higher values indicating that the buyer receives a higher 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 receives slightly more than half of the surplus created by the agreement.[60] This result aligns with the risk aversion estimates as both suggest that agreements involve relatively sophisticated, large suppliers (with the exception of one highly risk-averse supplier).[61] The large size of agreement suppliers is consistent with firms needing sophisticated capabilities to maintain complex contractual agreements (Macchiavello and Morjaria, 2023).

Because estimation includes outside moments, I prioritize model validation in Online Appendix OA1.8. First, the model matches an untargeted moment. Second, the model outperforms two alternatives in post-integration quarters; this comparison serves as an out-of-sample test because estimation primarily relies on pre-integration quarters.[62] (As the DiD directly estimates the price effects of surplus capture through buyer threats to replace suppliers with in-house production, DiD validation follows presentation of the analogous model estimate.)

4.4 Decomposing Prices into Surplus Creation and Capture

Table 3 decomposes agreement discounts into surplus creation from insurance-like agreement terms and capture through replacement threats with in-house production. To do so, I first determine whether each supplier's risk aversion surpasses the cutoff above which the buyer makes replacement threats. Formally, I test whether any pricing function yields weakly positive surplus for both parties when the buyer's outside option is the integrated supplier. The discount when the buyer's outside option is the spot market reflects the buyer's share of agreement surplus created, while any additional discount arises from surplus captured through replacement threats.

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

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

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

Estimates illustrate that agreement discounts are entirely attributable to buyer power over surplus creation, with no surplus capture through replacement threats. This result reflects agreement suppliers' low, albeit still present, risk aversion. Agreement prices, on average, are 1.5% below spot prices. This discount is a meaningful reduction as compared to spot markups near 7%. These savings decrease the buyer's annual fabric procurement costs by ∼$0.87 million. I further decompose agreement discounts from surplus creation into cost reductions from quantity assurance and profit smoothing (*i.e.,* insurance). To isolate the role of cost reductions alone, I compute discounts for a risk-neutral version of the supplier, holding all else constant. Profit smoothing then explains the remainder of the discount. Cost reductions and profit smoothing each account for approximately half of the total discount.

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 market, 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 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).

To assess robustness to the one outlier supplier with high risk aversion, I compute discounts with the *i)* median and *ii)* mean values of the supplier risk aversion and buyer bargaining parameters.[63] The main finding holds: all discounts arise entirely from surplus creation. Because all suppliers have risk aversion below the cutoff above which the buyer

[63]I consider two hypothetical suppliers each with the average capacity variance of agreement suppliers. The average discounts calculated instead for each agreement supplier but changing risk aversion and bargaining parameters to their median and mean yields almost identical results (less than .0001 p.p. different).

captures surplus through replacement threats, both the median and mean parameters also lie below that threshold. However, the level and composition of discounts differ. Discounts decrease overall because parameters affect prices nonlinearly: doubling both the supplier risk aversion and buyer bargaining parameter more than doubles the discount.[64] This change reflects that decreasing the influence of one highly risk-averse supplier reduces the relative contribution of profit smoothing (insurance) to the discount.[65] Discounts are slightly higher with the median rather than mean parameter values; the larger median than mean buyer bargaining parameter more than offsets the lower median than mean risk aversion.

4.5 Model Validation: Difference-in-Differences Estimates

I validate the structural model by comparing results with a DiD design that directly estimates the effects of surplus capture through threats to replace suppliers with in-house production.[66] 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 integrated supplier. I consider suppliers who produce fabric that is subsequently made by the new mill as treated because the buyer can credibly threaten to replace their output with in-house production; remaining suppliers are control. 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; in my setting, comparisons condition on suppliers' exposure to trends in demand at the fabric level that reflect evolving fashion tastes (detailed in Appendix $A.11$).[67] Because neither difference identifies structural parameters, comparisons of the model-implied replacement-threat effect with the DiD estimate leverage that the two empirical approaches rely on distinct sources of identifying variation.

The theoretical model generates sharp predictions of the effects of vertical integration, given supplier risk aversion estimates. As even the most risk-averse agreement supplier has risk aversion below the cutoff above which the buyer makes replacement threats (see Appendix Figure A.5), the model does not predict any price change for agreement suppliers.

[64] 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 the surplus size and the buyer's share.

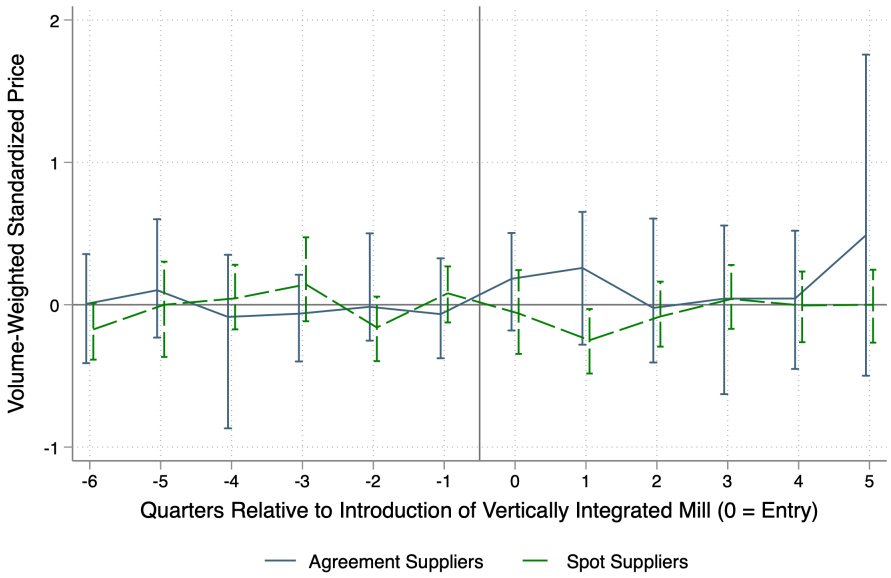
[65] Given that identification leverages the shape of the agreement pricing function, I confirm that the U-shape is robust to removing the most risk-averse supplier (Appendix Figure OA1.14).

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

[67] I confirm that the adjustment attains pre-treatment balance. Furthermore, as outcomes are *not* matched pre-integration, pretrends remain informative.

Figure 14 confirms that prices of agreement suppliers do not respond to integration. It also illustrates the absence of pretrends.

Figure 14: 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 a semiparametric difference-in-differences design (Abadie, 2005) detailed in Appendix *A*.11 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 buyer's sourcing of *marginal* quantities from spot suppliers suggests that vertical integration should decrease spot rather than agreement volumes.[68] Appendix Figure A.8 confirms that agreement quantities do not change, but spot quantities do. Furthermore, the significant decline in spot quantities suggests that the DiD design has sufficient statistical power to detect economically meaningful effects.

Appendix Table A.6 shows average treatment effects pooled across periods.[69] Pooled results confirm that agreement suppliers do not experience any statistically significant effect on quantities or prices. (The point estimates are near zero, although not always precise.)

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

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

However, vertical integration causes a statistically significant and economically meaningful decrease in quantities for spot suppliers. The quantity decline averages .4 standard deviations (∼120,000 square meters of fabric per quarter on average).[70] Online Appendix OA1.10.2 shows robustness to the synthetic DiD approach from Arkhangelsky et al. (2021).

Online Appendix OA1.10.3 provides an additional validation exercise that compares DiD event study coefficients using actual prices with equivalent estimates using model-implied prices, computed for observed quantities. Event study estimates using model-implied prices in Figure OA1.19 are all within the 95% confidence intervals of event study coefficients using actual prices.

4.6 DiD Estimates Are Inconsistent with Model without Spot Market Sourcing

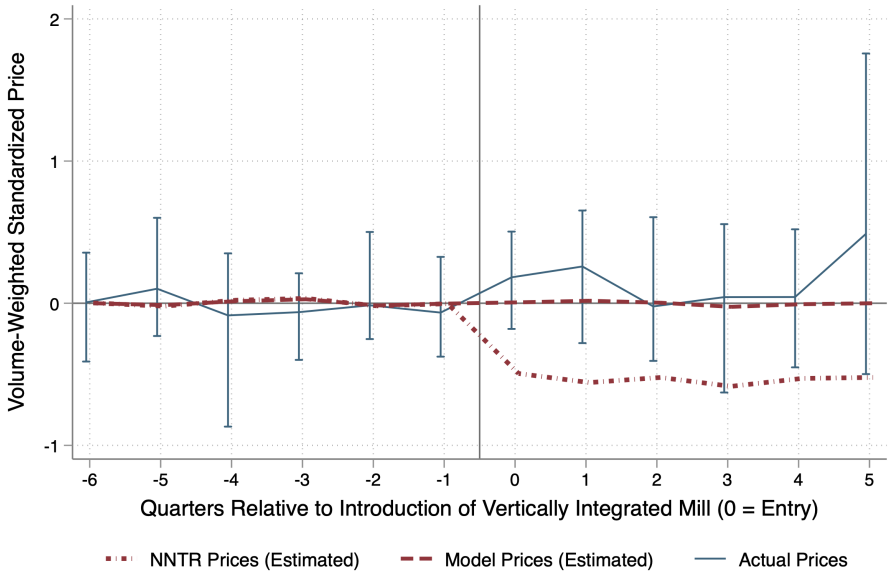
To illustrate how spot markets constrain replacement threats, I compare my DiD results with estimates from prices implied by a standard model of replacement threats that omits spot markets: the Nash-in-Nash with Threat of Replacement (NNTR) model from Ho and Lee (2019). Ho and Lee (2019) study a bargaining environment with neither a spot market nor integrated supply—agreements between employers and health insurers in California. Therefore, NNTR does not incorporate marginal sourcing from spot suppliers, constraints on the formation of new agreements (reflecting relational enforcement), or replacement threats through in-house production. The differences between the predictions of my model and NNTR do not imply a general critique of the latter but rather emphasize the importance of the contracting environment for the external validity of the effects of replacement threats.

Figure 15 illustrates that price decreases predicted by the NNTR model are inconsistent with the DiD estimates of price changes from replacement threats.[71]

[70] Additionally, I confirm robustness to an alternative measure of volume with desirable properties, including scale-invariant t-statistics: the quartic root (Thakral and Toh, 2023).

[71] My NNTR estimation approach adopts the economic environment in Ho and Lee (2019) while retaining features of my model for comparability. As there is no spot market in NNTR, I model the buyer's outside option of not trading with any party as the loss in profit from no trade. The calculation of no trade profit relies on estimated margins of Bangladeshi garment manufacturers in Cajal-Grossi et al. (2023).[72] I model the integrated supplier as one alternative supplier available to the buyer, where the integrated supplier has a risk aversion parameter of zero and the buyer captures all surplus (*i.e.,* the buyer bargaining parameter is one). Additional details on my application of the NNTR model are in Appendix A.12.

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*.11. NNTR prices estimate the model in Ho and Lee (2019) (described in Appendix A.12). 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).

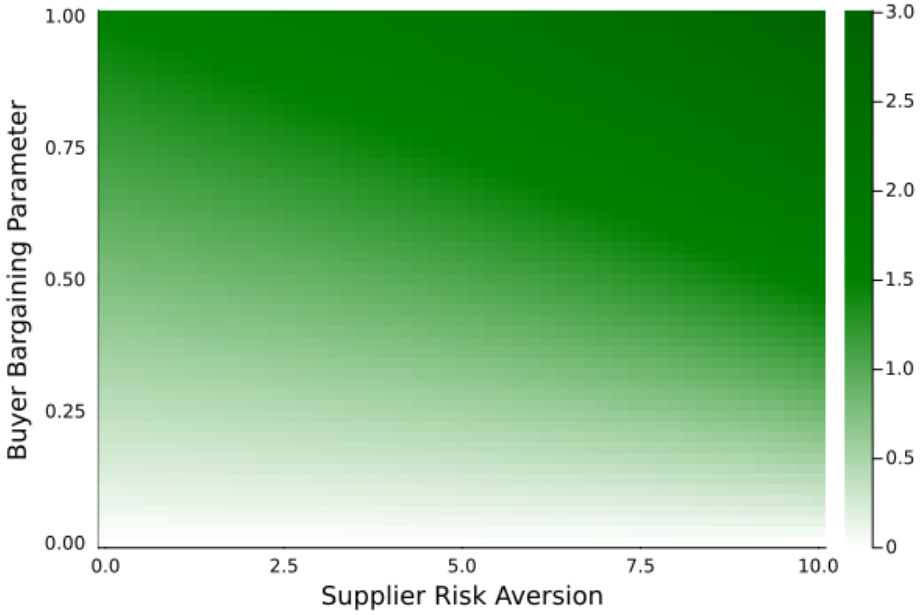
4.7 Model-Implied Prices for Small Suppliers

To study how surplus creation and capture affect prices for policy-relevant small suppliers (*i.e.,* suppliers smaller than observed agreement suppliers), I simulate agreement prices across a wide range of supplier risk aversion and buyer bargaining parameters using the model.

Figure 16 reports how agreement discounts, as a percent of spot prices, vary with supplier risk aversion and the buyer bargaining parameter. While discounts increase in both, the price change from higher risk aversion is modest. At the median estimated buyer bargaining parameter, the discount for a risk-neutral supplier is less than one percentage point smaller than the discount from a supplier at the median estimated risk aversion of observed agreement suppliers. For all risk aversion parameters shown, discounts reflect surplus creation, rather than surplus capture; the buyer always uses the spot market as its outside option when supplier risk aversion falls below the cutoff. Because discounts do not increase much in risk

aversion, cost reductions—rather than profit smoothing—drive observed discounts.

Figure 16: Agreement Discounts (% of Spot Price)



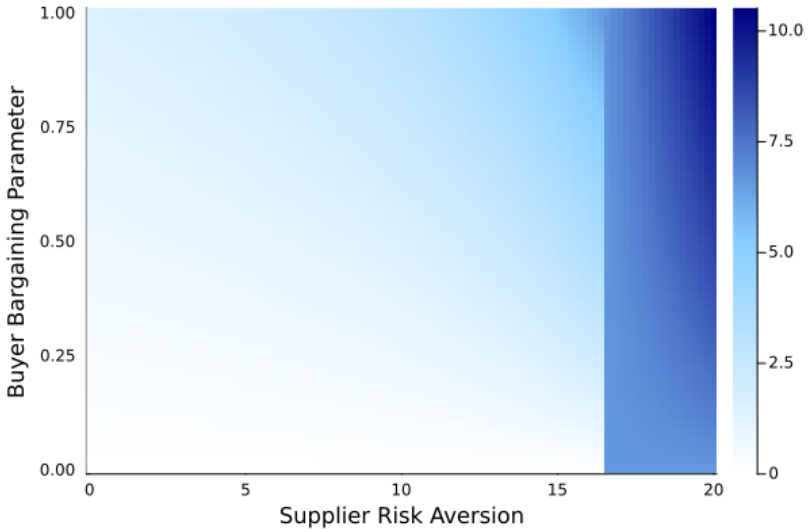
Note: The heatmap shows the model-implied discount in the agreement relative to the spot price, expressed as a percentage of the spot price, for different values of the supplier risk aversion and buyer bargaining parameters. The heatmap legend is to the right of the heatmap and units are percentages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

Figure 17 shows discounts for suppliers that are sufficiently risk averse for the buyer to capture surplus through replacement threats with in-house production. The cutoff level of risk aversion above which the buyer makes replacement threats 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. As the risk aversion threshold is not itself a structural parameter but depends upon them, 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).[73] LMIC smallholder farmers may satisfy these conditions and rely on bargained agreements (Macchiavello and Morjaria, 2020); they

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

could be highly risk averse and have convex production costs as a key input is fixed (land).

Figure 17: 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. Compared to Figure 16, this heatmap includes more risk-averse suppliers.

The discount increases meaningfully once supplier risk aversion surpasses the cutoff where the buyer makes replacement threats. For example, when the buyer bargaining parameter implies agreement surplus is shared equally between the parties, replacement threats increase the discount by 76% from 3.8% to 6.7% of spot prices. This difference is economically meaningful, as agreement discounts average 1.5% of spot prices. The 6.7% discount exceeds Indian VAT for apparel in the sample period (5%)[74] and is large relative to markups in the spot market from hold-up (7.36%).

Online Appendix OA1.11 incorporates relational enforcement for a conservative calibration of the discount factor, documenting negligible differences in estimated prices.[75] This result suggests that suppliers are at least as patient as the calibrated discount factor, which is based on microfinance interest rates. A high discount factor for agreement suppliers is con-

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

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

sistent with low estimated risk aversion—both suggest agreement suppliers are large firms.[76]

5 Counterfactual Policy Evaluation

I analyze counterfactual policies aimed at increasing prices paid to small suppliers through the lens of the model. I focus on small suppliers because of their empirical importance as the majority of LMIC firms and employers (Hsieh and Olken, 2014), as well as distributional motivations given policymaker concerns that gains from exporting do not reach up the supply chain. Consistent with this focus on small firms, ubiquitous policy initiatives target small and medium enterprises (SMEs)[77] Furthermore, discounts from surplus creation and capture both increase in risk aversion, and small suppliers are likely the most risk averse.

I evaluate two policies to increase prices paid to small, risk-averse suppliers. First, motivated by concerns about market power of the few, large buyers (exporters), I increase competition between buyers over supplier output. Second, given the importance of missing markets in LMIC settings, I create the missing insurance product so risk-averse firms can smooth profit without relying on the agreement.

5.1 Increase Downstream Buyer Competition

I analyze how increased downstream buyer competition affects prices. Many policies can encourage development of large exporting firms including reducing red tape and other administrative costs and facilitating connections to international end clients ("export promotion").

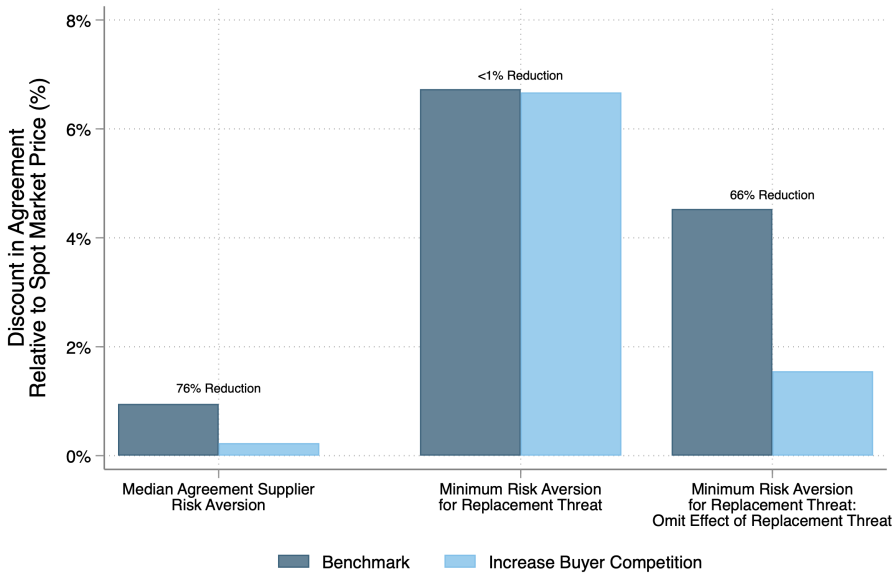
I model increased downstream buyer competition by halving the buyer bargaining parameter. While competition from additional buyers should improve suppliers' outside options in the long run, competition does not change suppliers' short-run outside option (the spot market) because suppliers cannot quickly establish a new agreement with relational enforcement. (Relational enforcement requires time and effort to develop and maintain (Cajal-Grossi et al.,

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

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

2023; Brugues, 2025).) Therefore, I model the buyer responding to competition by offering suppliers a larger share of surplus to guarantee loyalty. This interpretation mirrors efficiency wages, as both incentivize loyalty through a higher surplus share (Katz, 1986). Empirical patterns support modeling buyer competition as a reduction in the buyer bargaining parameter: more valuable suppliers receive larger surplus shares (*i.e.,* have lower estimated buyer bargaining parameters).[78]

Figure 18: Downstream Buyer Competition and Agreement Prices



[78]Compared to other agreement suppliers, those agreement suppliers producing fabrics with higher concentration are likely more valuable to the buyer because supplier-fabric concentration affects the ease of providing quantity assurance. Consider an extreme case in which all agreement suppliers sell the same fabric. Then, when aggregate quantity for that fabric is low, the buyer cannot provide quantity assurance to all suppliers. Therefore, it is easier to provide quantity assurance to suppliers in relatively concentrated fabric markets. Appendix Section A.13 describes the empirical approach to establish that supplier concentration in fabric markets is significantly negatively correlated with estimated buyer bargaining parameters. Additionally, this proxy for agreement supplier value is not significantly correlated with estimated risk aversion parameters. This result mitigates concerns that the correlation between supplier value and buyer bargaining parameters is spurious.

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 threshold above which the buyer captures value through strategic replacement threats. 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 18 illustrates that increased downstream buyer competition meaningfully reduces the discount from a supplier as risk averse as the median agreement supplier: discounts decrease by 76%, from 0.95% to 0.23%. However, for suppliers with risk aversion above the cutoff at which the buyer uses in-house supply as a replacement threat, the discount falls by less than 1%: from 6.73% to 6.67%. Omitting surplus capture through replacement threats *overestimates* the benefits of increased competition; without it, increased competition reduces the discount by 66%—from 4.5% to 1.5%.

Intuitively, increasing downstream buyer competition proves ineffective as a policy because, once the buyer makes replacement threats with in-house production, agreement surplus is small relative to the buyer's participation constraint. When surplus is small, changing the sharing rule has minimal effects on prices. This intuition clarifies not only why abstracting from surplus capture generates incorrect policy conclusions for small risk-averse firms, but also why this policy effectively increases prices paid to suppliers with low risk aversion. In agreements with suppliers with low risk aversion, there is enough agreement surplus that changes to surplus sharing meaningfully affect prices.[79]

5.2 Create Missing Insurance Product

The importance of missing markets in LMIC settings, and the limited impact of increased buyer competition, motivate analyzing the effects of adding the missing profit insurance product. I evaluate how access to actuarially fair insurance over spot market profits affects agreement discounts. Insurance mitigates how supplier risk aversion increases discounts by eliminating profit risk in the spot market.[80] To do so, I leverage that with actuarially fair

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

[80]Blouin and Macchiavello (2019) also explore how missing insurance markets interact with contracting behaviors, highlighting how suppliers' incentives for strategic default lead trading partners to use index-priced contracts to remove such temptations. While Blouin and Macchiavello (2019) focus on the efficiency implications of the missing insurance markets, I focus on distributional implications.

insurance in the spot markets, suppliers purchase insurance such that utility is constant across states. Therefore, each supplier's outside option is $O_S = U(\mathbb{E}_{q^S}[p^S q^S - C(q^S)]$.

Although moral hazard concerns make governments unlikely to provide such insurance, policies can support smoothing.[81] Enhancing access to financial markets can enable small firms to hedge risk. Access to credit allows firms to borrow to cover expenses when revenue is low instead of accepting low agreement prices. Policies that enable small firms to receive credit at favorable rates are common, even in high-income countries. Beyond credit, improving access to saving can help suppliers accumulate reserves over time, effectively allowing them to self-insure. Furthermore, policies can facilitate access to derivatives and other financial products that provide direct risk hedging, though financial literacy and transaction costs may limit take-up. Additionally, policies can support trade associations and cooperatives to provide insurance. These organizations may be able to mitigate moral hazard, for example, through indexing the quantity of insurance offered to historical capacity or revenue.

Figure 19: Agreement Discounts Without Insurance

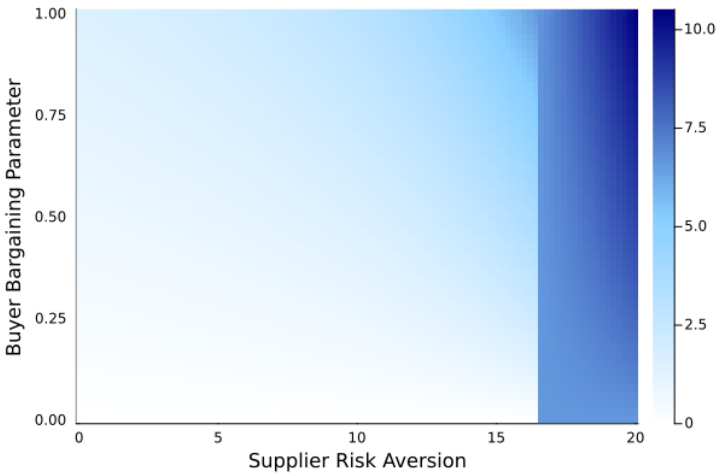
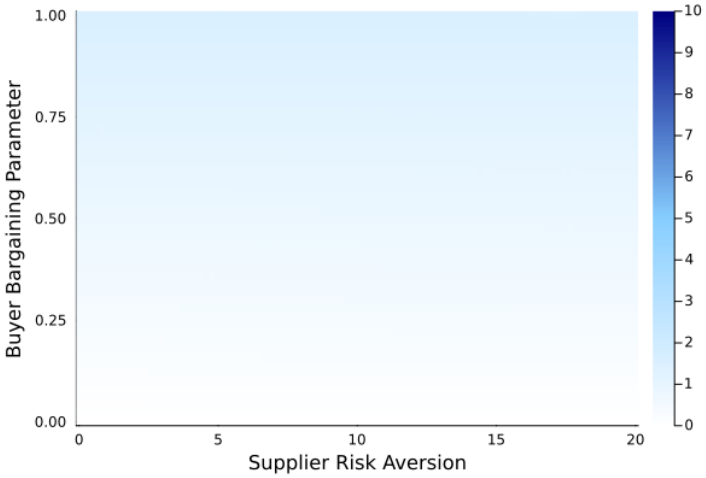


Figure 20: Agreement Discounts With Insurance



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 19 reproduces Figure 17, while Figure 20 shows discounts when suppliers have access to actuarially fair insurance over profits when selling to the spot market.

When suppliers buy insurance, discounts decline because the buyer cannot capture sur-

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

plus through replacement threats. This result reflects that the surplus capture derives from the buyer serving as the de facto monopolist of profit smoothing because insurance is only available inside agreements; adding the missing insurance market breaks the monopoly.

6 Conclusion

This paper highlights how low vertical contract prices can arise from two distinct forces with different welfare implications: surplus creation inside the agreement from mitigating market frictions and surplus capture through replacement threats with in-house supply. Studying the Indian garment supply chain, I show that agreements include two insurance-like terms: quantity assurance and nonlinear prices. A theoretical model formalizes the price effects of surplus creation and capture; both make (average) prices decrease in supplier risk aversion, with surplus capture only affecting prices when risk aversion exceeds a cutoff. Estimates from a structural model, identified by plausibly exogenous demand shocks, indicate that observed agreement suppliers all have risk aversion below this cutoff, so the entire agreement discount ($\sim$1.5% of spot prices) reflects surplus creation. Through its buyer power over surplus creation, the buyer receives $\sim$60% of the gains from surplus created by insurance-like terms. More broadly, the decomposition emphasizes that low prices reflect the buyer's share of gains from addressing market frictions within the vertical relationship rather than pure rent extraction. Correspondingly, policies to increase the prices that small suppliers receive should consider targeting the frictions that lead small firms to rely on agreements rather than exclusively addressing competition in input markets.

More broadly, this paper provides additional evidence that challenges to smoothing, especially relevant in LMIC settings, not only decrease welfare by generating different marginal utility across states, but also shape prices. For example, LMIC agricultural producers buy and sell at the same time as other nearby producers because they face challenges storing output. Simultaneous sales generate negative price effects—producers "sell low and buy high" (Burke et al., 2018). Similarly, challenges in smoothing over productivity shocks lead agricultural wage workers to increase labor supply simultaneously, magnifying the decrease in their wages (Jayachandran, 2006). In my model, small risk-averse firms accept lower average prices to smooth profits because of missing risk-smoothing institutions. Together, these results highlight that policies that address exposure to risk can benefit both *efficiency*—by equating marginal utility (and revenue, *etc.*) across states—and *equity*—by improving the outside option for those least able to bear risk.

References

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

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

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

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.

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

Demirer, Mert and Michael Rubens, "Welfare Effects of Buyer and Seller Power," *Working Paper*, 2025.

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).

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

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

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

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

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

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

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

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

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

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

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

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

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

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

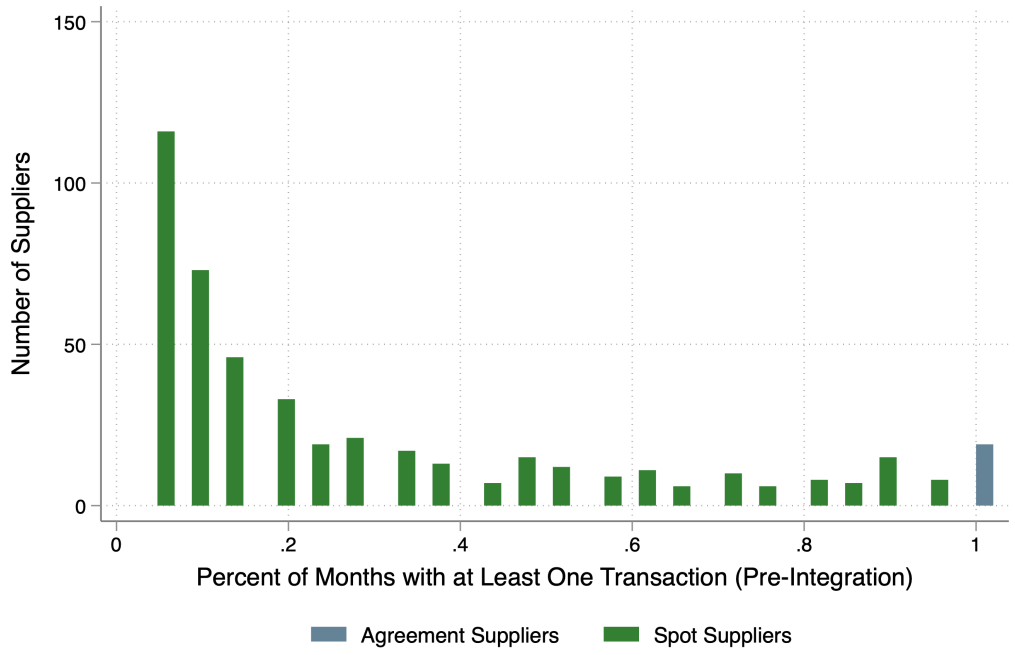
Zavala, Lucas, "Unfair Trade? Monopsony Power in Agricultural Value Chains," *Working Paper*, 2022.

A Appendix

A.1 Fabric Procurement Agreements

A.1.1 Supplier Classification: Agreement versus Spot

Figure A.1: Distribution of Percentage of Months with Transactions (Pre-Integration)



Note: Data from 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.

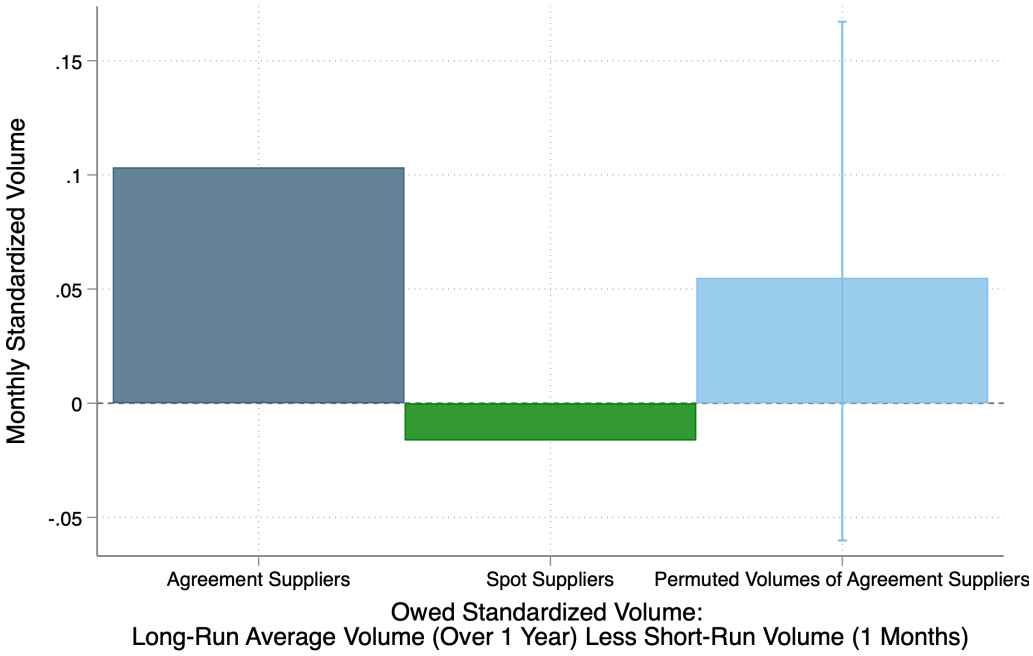
A.1.2 Quantity Assurance: Owed Quantity Guarantees

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.[82] 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 standardized volume in the prior month less the average of monthly standardized volume in the prior year, not including 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-

[82]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.[83] 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.

[83] 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.

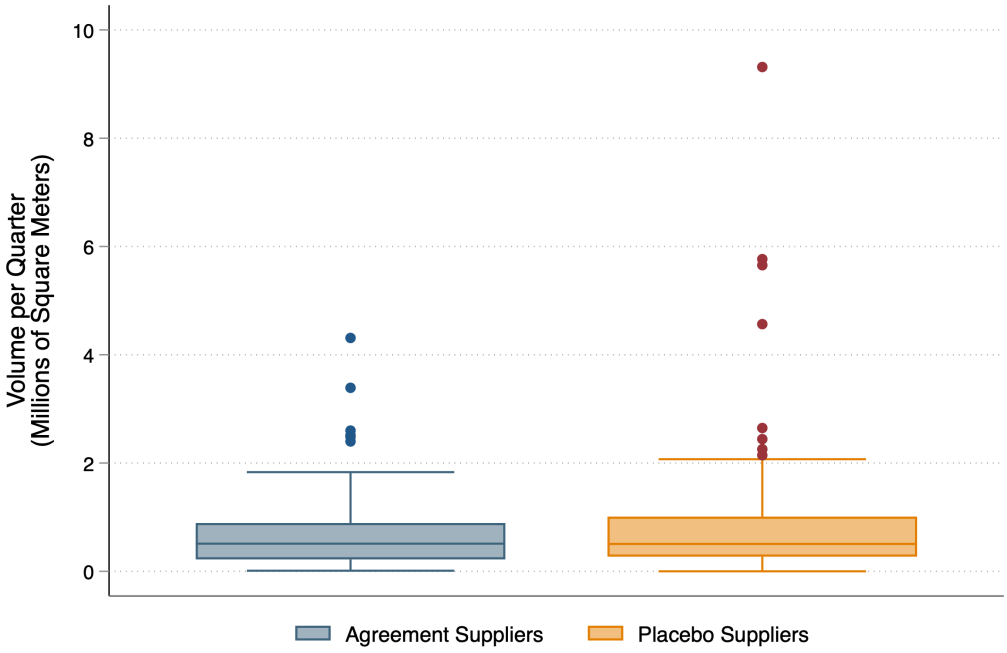
Table A.1: Pricing Function Convexity by Supplier Type

	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
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 4.1.2. 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).[84]

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:

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

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

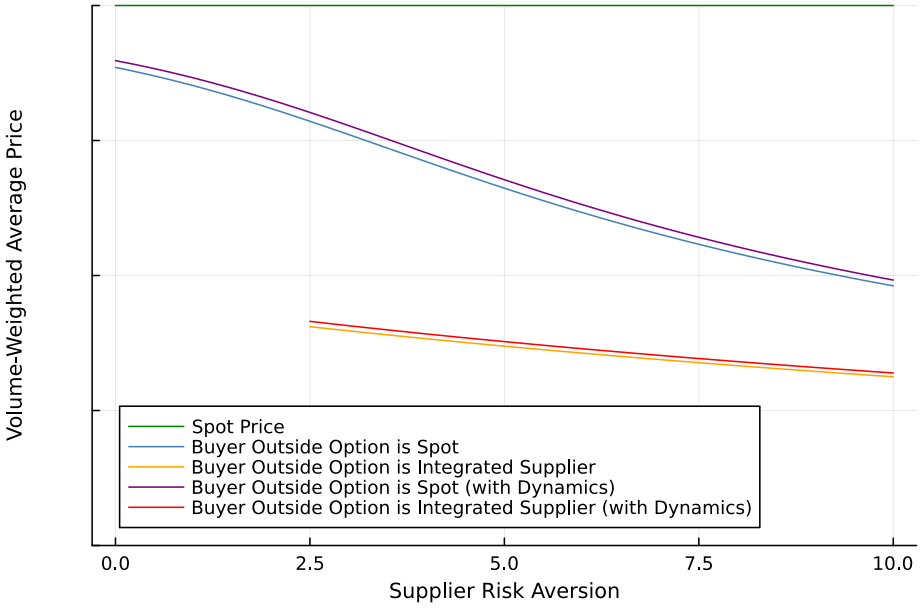
Note that for δ^S, δ^B sufficiently large, dynamic incentives are irrelevant as only long-run payoffs shape behavior.

The analysis below shows that results are robust to evaluating a dynamic relational contract model of the agreement.[85]

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

[85] 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

Consider a signaling game that proceeds as follows:

1. Nature draws a state of the world, which represents the residual capacity remaining at the buyer's in-house supplier. Residual capacity at the buyer's in-house supplier is both *i)* the buyer's private information and *ii)* stochastic given the institutional details of garment manufacturing (*i.e.,* the buyer does not know which garments will be ordered). Furthermore, capacity at the in-house supplier is what determines the buyer's true outside option. I binarize the state of residual capacity to highlight the key cutoff: whether it would be lower cost for the buyer to produce the relevant volume internally or to source it through the spot market. Therefore, $\tilde{q} = 0$ means that the in-house supplier is capacity unconstrained, such that production is lower cost in-house, while $\tilde{q} = 1$ means that the in-house supplier is capacity constrained, such that production is more costly in-house than sourcing from the spot market. I express the buyer's cost function as $C_B(\tilde{q}, 1)$ to denote the cost of producing one unit when residual capacity is $\tilde{q}$. It follows that $C_B(1, 1) > p^S > C_B(0, 1)$.

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

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

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

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

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

7. Additionally, if the supplier rejects the buyer's claim ($\hat{q}$), then costly verification (*e.g.,* a factory tour) reveals the true state of the world $\tilde{q}$. In this case, a transfer γ is made within the relationship that punishes either the buyer for claiming an inaccurate $\hat{q}$ or the supplier for rejecting an accurate $\hat{q}$. In the model, γ broadly captures (in a reduced-form way) any incentives that prevent parties from trying to misrepresent the state of the world. Note that no model results depend upon the γ being a transfer; although not shown, results are qualitatively identical if γ is exclusively a penalty.[86]

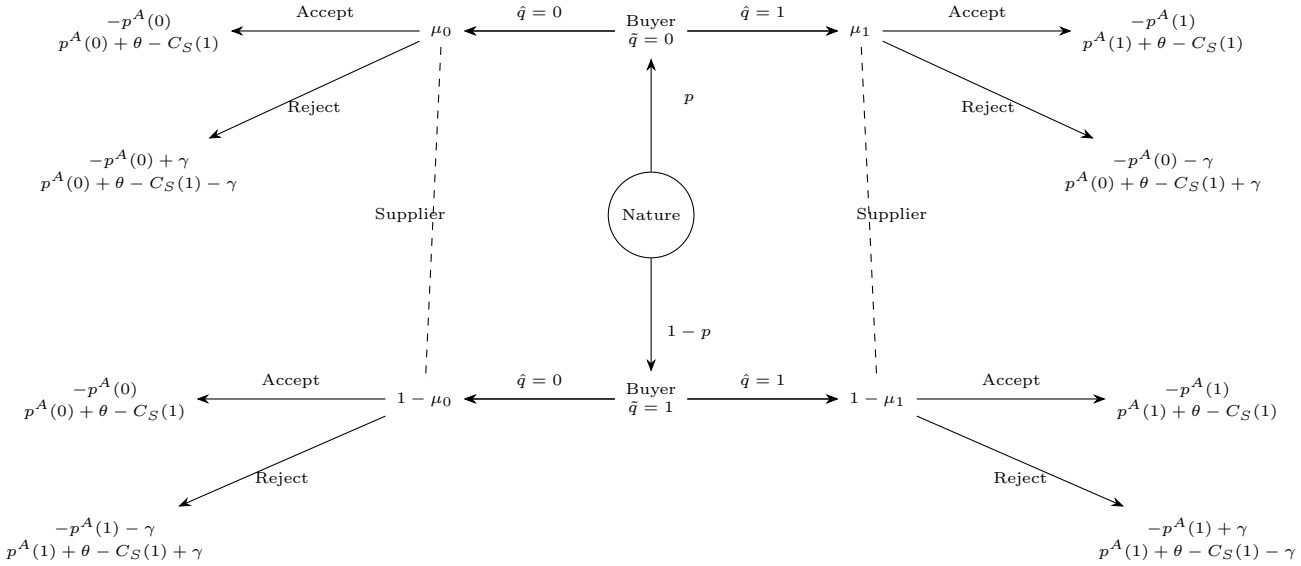
8. Nash bargaining prices are computed assuming that the buyer and supplier equally split the surplus (in other words, bargaining parameters are .5). I assume that the transfer γ does not occur if the agreement dissolves.

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

As I use weak Perfect Bayesian Equilibrium as the solution concept, denote the supplier's belief that they are at the node where $\tilde{q} = 0$ as μ_0 after the buyer plays $\hat{q} = 0$ and as μ_1 after the buyer plays $\hat{q} = 1$. Note that replacing cooperative Nash bargaining with an infinitely repeated non-cooperative game yields similar qualitative results.

A.5.1 Replacement Threat: Pooling Equilibrium at $\hat{q} = 0$

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



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

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

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

Now, solve the Nash bargaining problem:

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

The F.O.C. then implies:

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

Similarly,

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

The F.O.C. then implies:

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

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

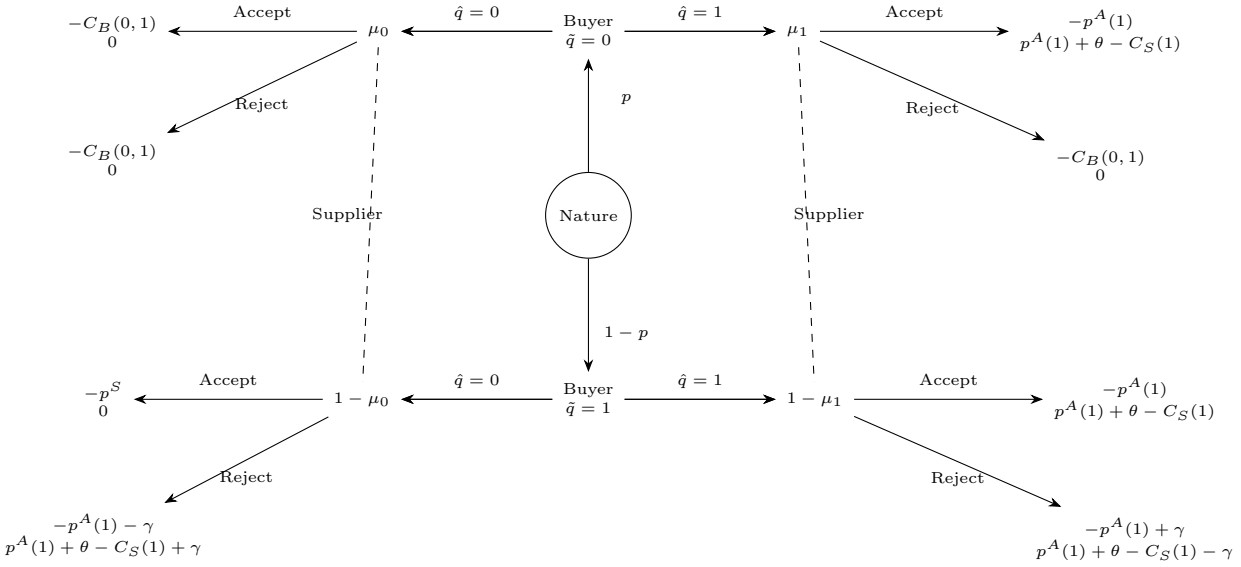
Formally, the equilibrium is that the buyer plays $\hat{q} = 0$ for both types, the supplier

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

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

A.5.2 No Replacement Threat: Pooling Equilibrium at $\hat{q} = 1$

Now, consider the case where θ is small enough, representing a supplier with low risk aversion, such that there is no surplus relative to $\tilde{q} = 0$. However, θ is sufficiently large that there are gains to trade possible when $\tilde{q} = 1$ (note that $\theta < 0$ could still be sufficiently large). The payoffs then change, with the agreement breaking down when the relevant buyer's outside option is that it is unconstrained, either because the supplier accepts $\hat{q} = 0$ or rejects $\hat{q} = 1$. This breakdown reflects that the supplier's participation constraint is $p^A(0) : p^A(0) + \theta - C(1) \geq 0$. As θ decreases, no surplus remains, and the agreement can break. Then, the game can be expressed as:



Now, I show that the pooling equilibrium at $\hat{q} = 0$ can break for this lower θ using a

proof by contradiction. First, note that the supplier's best response is to accept after the buyer plays $\hat{q} = 0$ as long as:

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

As $\gamma > 0$,

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

But, this contradicts the definition of θ as large enough that there are gains relative to $\tilde{q} = 1$. Therefore, there is no equilibrium where the supplier accepts after the buyer plays $\hat{q} = 0$.

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

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

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

A.5.3 Discussion

This analysis shows a game where the buyer is able to threaten replacement with in-house production—by signaling that the in-house supplier is *not* capacity constrained—to improve its bargaining position, but this equilibrium only exists when the buyer bargains with a sufficiently risk-averse supplier. Although not shown, the pooling equilibrium at $\hat{q} = 1$ also exists when θ is high. If the buyer can select the equilibrium, consistent with its agenda-setting role and ability to coordinate expectations as the larger firm that invites suppliers for meetings, the buyer would always prefer the pooling equilibrium at $\hat{q} = 0$ because $p^A(0) < p^A(1)$.

A.6 Replacement Threats in a Game with Commitment

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

1. The buyer commits to a capacity level at its integrated supplier of $q^C \in \{0, 1\}$. 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 (as in Appendix A.5), $p^A(0) = \frac{C_S(1) - \theta + C_B(0,1)}{2}$. 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, as in Appendix A.5), 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.6.4 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 (at their equilibrium values) relative to their total demand from end customers shapes residual supply available at the in-house supplier. Alternatively, the game with commitment in Appendix A.6 directly incorporates the spot market for additional sourcing at the end of the game. Extending either model from a bilateral game to a sequential multilateral game that repeats the bilateral game would yield the same equilibrium as the bilateral case in both models regardless of the bargaining order with agreement suppliers.[87]

This equivalence reflects that, on path, the buyer never actually displaces an agreement supplier with in-house production. Therefore, prior agreements with in-house suppliers do not change the credibility of in-house production as a replacement threat because residual

[87] 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).

in-house capacity remains unchanged.[88] Similarly, a simultaneous (rather than sequential) multilateral game, such as Nash-in-Nash, would also produce the same multilateral equilibrium. The equivalence of the simultaneous game reflects that neither the signaling game nor the commitment game need create any contracting externalities across negotiations, as none of the bilateral equilibria actually use any in-house capacity.[89]

A.7 Outside Estimation Details

As the cost function is measured in capacity units, transaction quantities must be converted into capacity units to estimate the structural model. Two assumptions underlie my conversion 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 increase linearly with quantity (see Online Appendix Table OA1.3). Expressing quantities in logs captures this pattern in a reduced-form manner without introducing additional parameters to govern how the distribution of quantities over time affects capacity.

The second assumption is that the supplier's capacity, at least for agreement suppliers, is the mean capacity in the pre-period. This assumption reflects that the agreement provides consistent capacities around an agreed-upon capacity target. I model quantity realizations in the agreement as 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.[90] 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

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

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

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

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 that 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.[91] Placebo suppliers are the closest counterfactual for agreement suppliers overall given their similar volumes.[92] Encouragingly, the estimate of counterfactual variance using placebo suppliers is reasonable: 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.[93] The spot price, which incorporates hold-up, has an estimated margin of 7 percentage over costs. This estimate also seems quite plausible: it is a large enough markup to create incentives for strategic sourcing through both vertical integration and agreements but not so large as to be implausible.

A.8 Structural Bargaining Model: Evidence Supporting Model Assumptions

A.8.1 Quantity Exogeneity

Although capacity utilization for any supplier is certainly not random, as the buyer provides quantity assurance in the agreement, 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

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

[92] Even using placebo suppliers, data limitations prevent using the obvious empirical analogue: capacity variance at the supplier level for placebo suppliers. Given possible errors in this moment, I take model fit seriously in section OA1.8 to mitigate concerns that errors in outside moments lead to poor model performance.

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

observed quantity assurance in the agreement.

However, buyer behavior that reduces *dynamic* capacity variation would bias estimates, as I would underestimate how much quantity assurance the agreement provides. It follows that the estimation approach would overestimate at least one of risk aversion and the supplier bargaining parameter (which is one less the buyer bargaining parameter).[94]

To mitigate such concerns, I directly test the most likely examples of such behavior in Table A.2. Specifically, if the buyer provides additional dynamic capacity variance reduction, it would presumably take the form of ensuring that extreme capacity shocks are unlikely to occur repeatedly. I directly test this hypothesis by examining the autocorrelation 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.

Table A.2: 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 4.1.2) 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 OA1.8.1 should not find that model prices are unbiased.

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

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

Table A.3: Comparison with Other Risk Aversion Estimates

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

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

Table A.4 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.

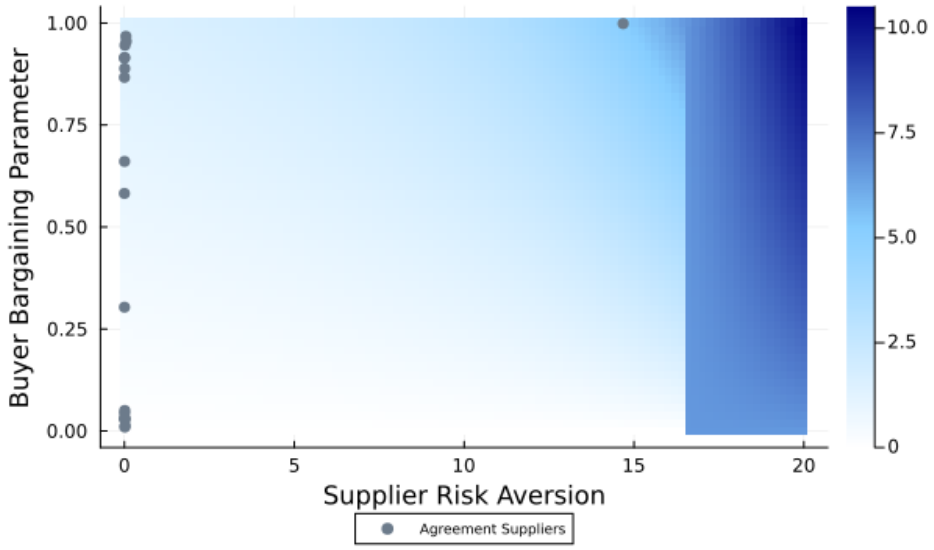
Table A.4: Buyer Bargaining Parameter and Supplier Risk Aversion: 95% CI

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, with higher values indicating the buyer receives a higher surplus share. Confidence intervals are estimated from 100 bootstrap replications clustered by supplier, where each replication re-samples normalized quantity and prices for each supplier (with replacement). 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.

A.10.1 Model-Implied Surplus Capture from Agreement Suppliers

Figure A.5: Surplus Capture from Agreement Suppliers



Note: This figure plots agreement supplier types over figure 17. Therefore, this figure shows that all agreement suppliers have risk aversion below the cutoff for the buyer to make replacement threats.

A.11 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)).[95] 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

[95]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 makes 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).

exposure to vertical integration via heterogeneity in suppliers' fabric portfolios. Therefore, I adopt a DiD design with propensity-score adjustment that accommodates non-random treatment (Abadie, 2005), detailed below in Appendix A.11.1. 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.[96] Furthermore, as outcomes are *not* matched pre-integration, pre-trends support the validity of parallel trends.[97]

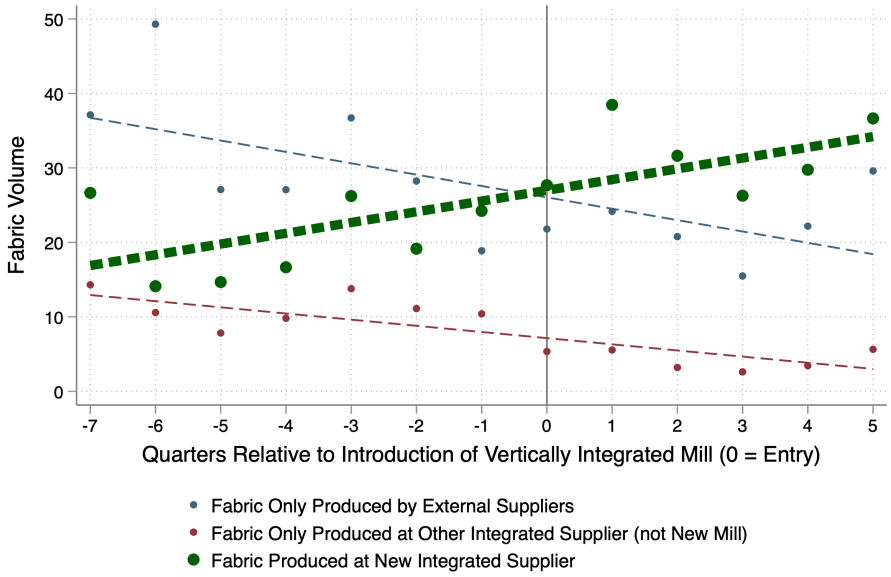
A.11.1 Semiparametric Difference-in-Differences Design and Identification

Figure A.6 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 fantastic 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).

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

[97]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.6: 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.[98] 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.[99] Formally, this assumption can be expressed as:

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

[98]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.

[99]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 $\boldsymbol{X}_i$,[100] 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.5 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,[100] it is not mechanical that this propensity score approach would achieve balance. I also document in Figure A.7 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.[101]

Table A.5: Balance on Untargeted Outcomes

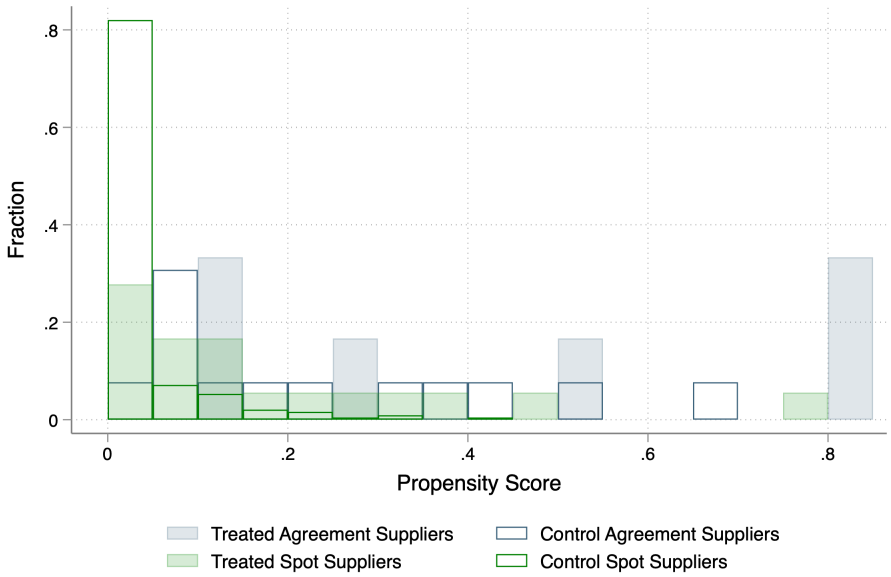
	Coefficient on Treatment Indicator:	
Outcome Variable	Without IPW Adjustment	With IPW Adjustment
Pre-Integration Volume	3033768	1085459
	[1189072, 5489585]	[−1004042, 3383491]
Pre-Integration Transaction Count	72.93	−3.35
	[44.49, 106.45]	[−46.3, 15.1]
Pre-Integration Fabric Count	26.71	−.98
	[17.08,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

[100]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.

[101]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.

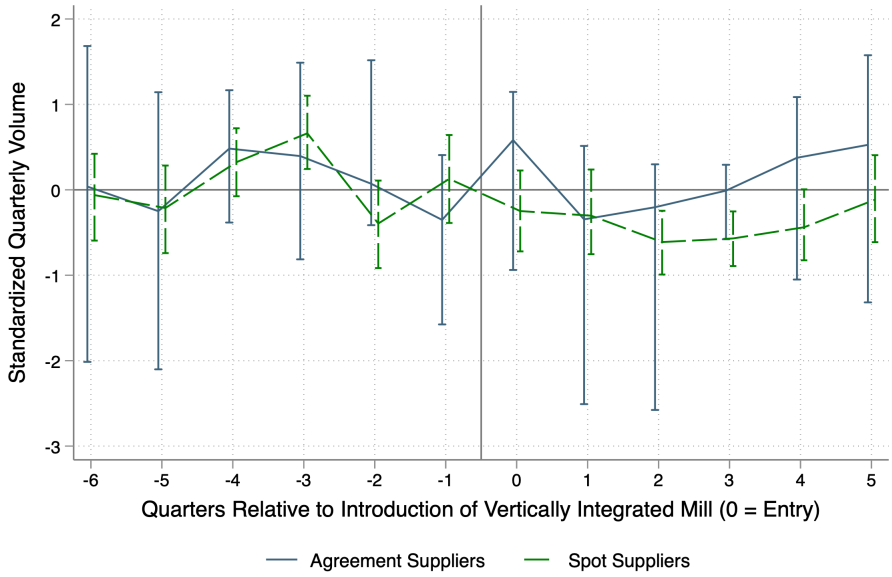
Figure A.7: Distribution of Propensity Scores by Supplier Type



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.11.2 DiD: Effects on Quarterly Volume

Figure A.8: 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 a semiparametric difference-in-differences design (Abadie, 2005) detailed in Appendix *A*.11 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).

A.11.3 Pooled DiD Estimates

Table A.6: 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*.11 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. 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. 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).

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

A.12 NNTR Implementation

My NNTR estimation approach aims to capture the strategic incentives for threats of replacement in Ho and Lee (2019). 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 the buyer model in three key ways. First, I model the buyer's outside option of no agreement as losing the profit from the order entirely. This assumption reflects that, absent the spot market, the buyer cannot source fabric to fulfill the order. To do so, I calibrate profit margins for a garment manufacturing using estimates from Cajal-Grossi et al. (2023).[102] Second, I model the integrated supplier as an additional external supplier, as the bargaining environment in Ho and Lee (2019) does not include any integrated production. I assume that the integrated supplier has zero risk aversion, as it is part of a large firm, and the buyer can extract all surplus from it. Then, I implement the bargaining model as in NNTR, finding the average price that would make the buyer indifferent between keeping a supplier and switching to the excluded supplier for each excluded supplier. Inside agreements, suppliers charge the minimum of the bargained price and the agreement price. I find that the buyer actually uses the integrated supplier as the sole excluded supplier, as doing so brings down prices from included suppliers.

A.13 Buyer Bargaining Parameter and Supplier Market Power

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

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

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,[103] and the trade-off of using volume from one order to fulfill quantity assurance across suppliers exclusively applies to agreement suppliers.

Table A.7: 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.7 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.8: 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

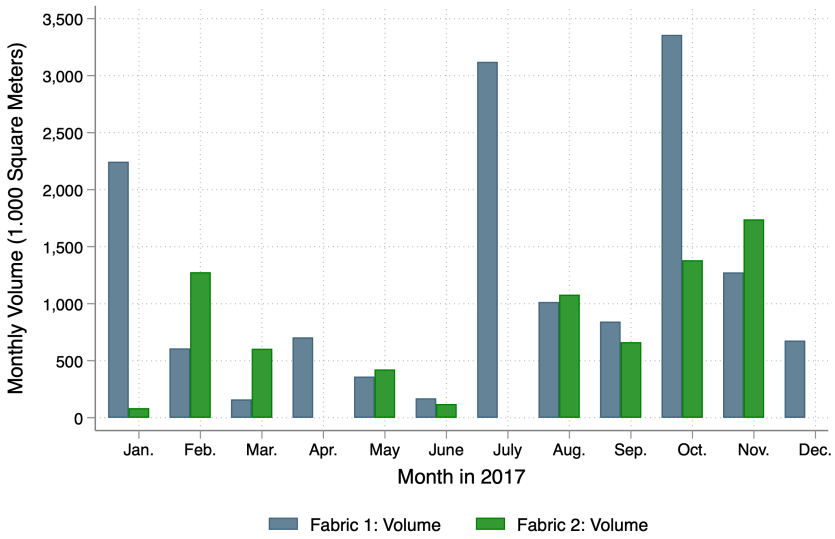
Note: Data from universe of fabric transactions by the buyer and model estimates. The coefficients shown are from a regression of the risk aversion parameters estimated from the structural model on supplier-fabric concentration, where supplier-fabric concentration is the supplier-fabric HHI across all 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.

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

OA1 Online Appendix

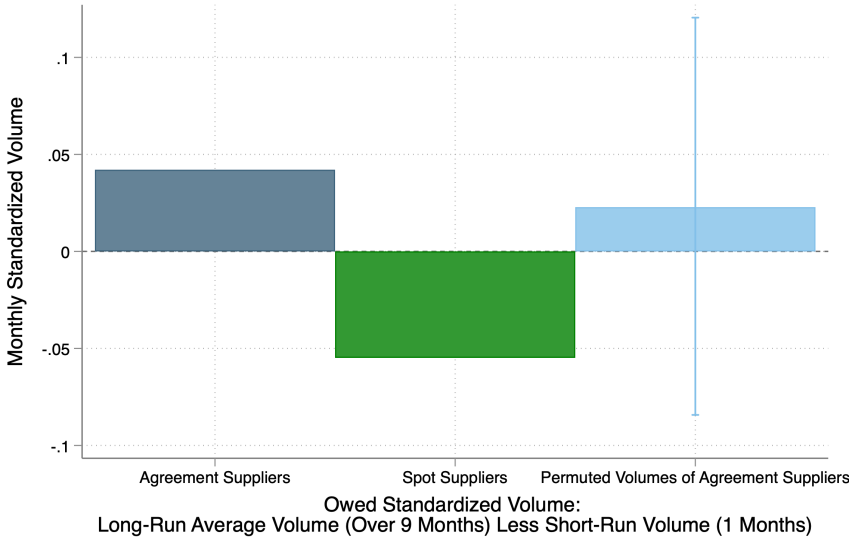
OA1.1 Additional Descriptives About Buyer-Supplier Agreements

Figure OA1.1: Variation in Monthly Quantity by Fabric



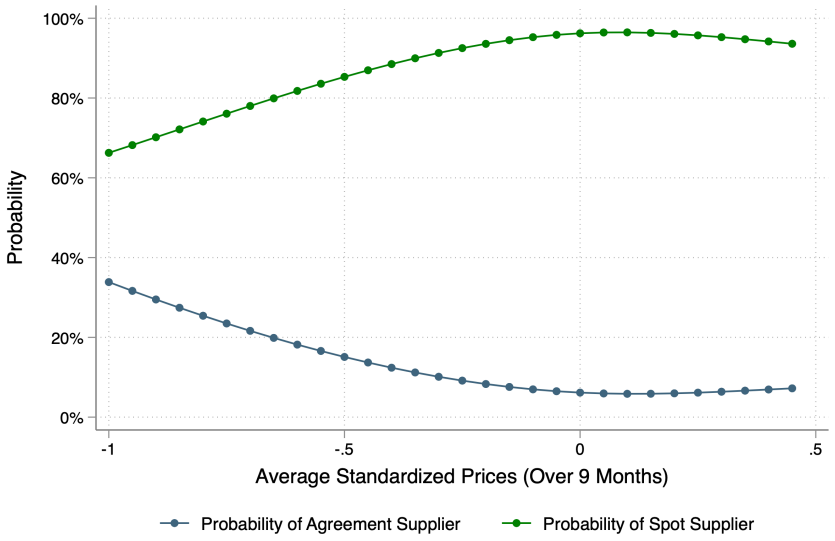
Note: Data from universe of fabric transactions by the buyer. Bars represent the quantity of fabric purchased during the month in 2017 for two selected fabrics. Fabrics were selected based on data availability to illustrate variation.

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 standardized volume in the prior month less the average of monthly standardized volume in the prior three quarters, not including 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.

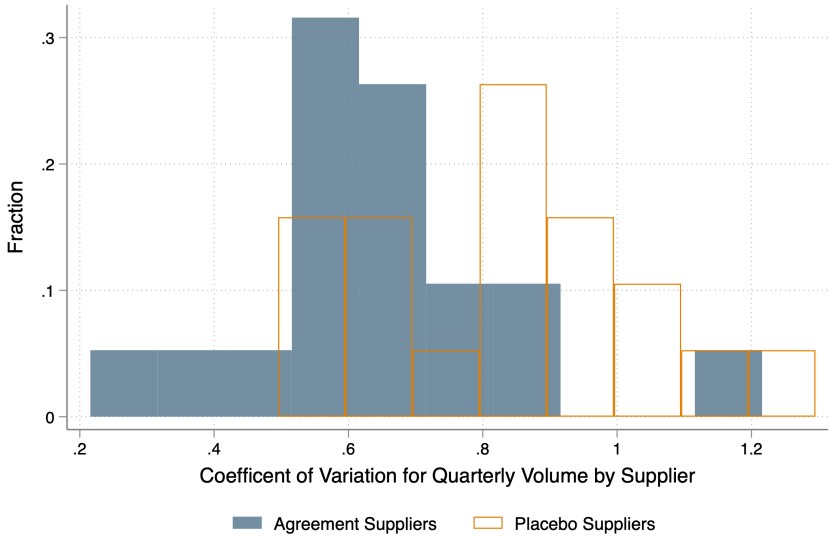
Figure OA1.3: Supplier Types and Prices



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.

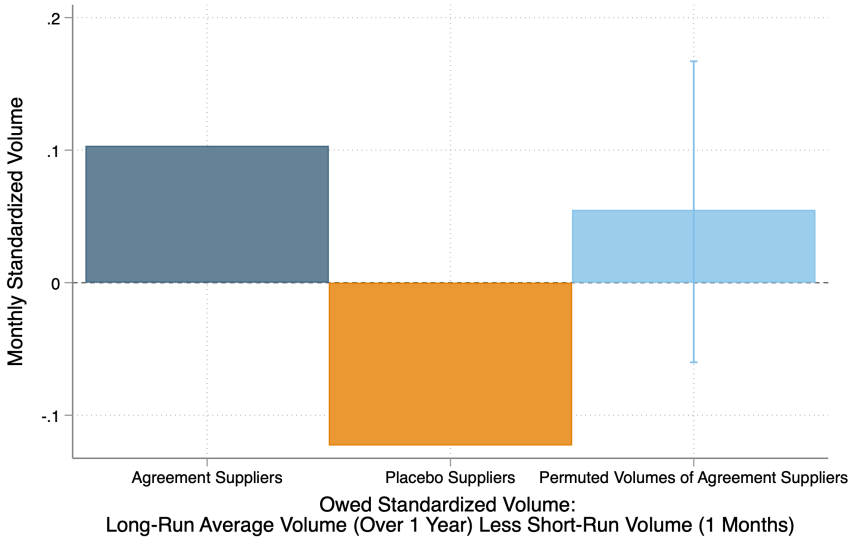
OA1.2 Placebo Supplier Analysis

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 standardized volume in the prior month less the average of monthly standardized volume in the prior year, not including 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

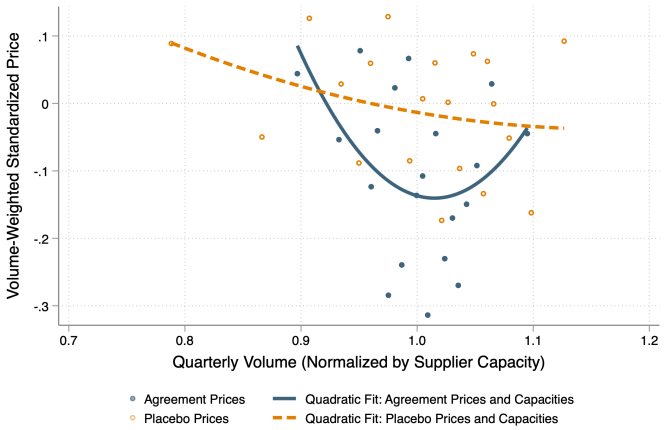
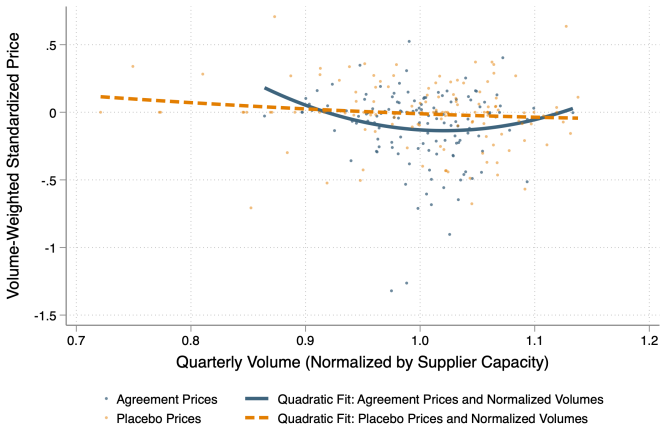


Figure OA1.7: Pricing Function (Placebo): Raw



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 4.1.2). 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
Implied Quantity for Minimum Price	1.021	1.295	1.014	1.014	1.008	1.005
95% CI	$[.981, 1.06]$	$[-1.365, 3.956]$	$[.994, 1.033]$	$[.88, 1.148]$	$[.984, 1.033]$	$[.894, 1.116]$
Slope at Minimum Capacity	-4.055	$-.596$	-5.415	-1.864	-4.662	-2.079
95% CI	$[-6.828, -1.282]$	$[-2.354, 1.162]$	$[-9.541, -1.29]$	$[-3.958, .23]$	$[-8.865, -.459]$	$[-4.236, .079]$
Slope at Maximum Capacity	2.914	$-.164$	4.333	.788	4.04	.973
95% CI	$[-.349, 6.178]$	$[-1.756, 1.429]$	$[.778, 7.887]$	$[-1.14, 2.715]$	$[.431, 7.648]$	$[-1.024, 2.97]$
R^2	.034	.015	.462	.25	.497	.281
N	133	130	133	130	133	130

Note: Data from 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 4.1.2. 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?

There are (at least) four factors that can prevent explicit contracts from resolving the hold-up problem in this setting. First, if the explicit contract is not enforceable because of weak contracting institutions, then the explicit contract is effectively irrelevant. Research in India suggests that contracting parties may be concerned that courts are unable to effectively enforce 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

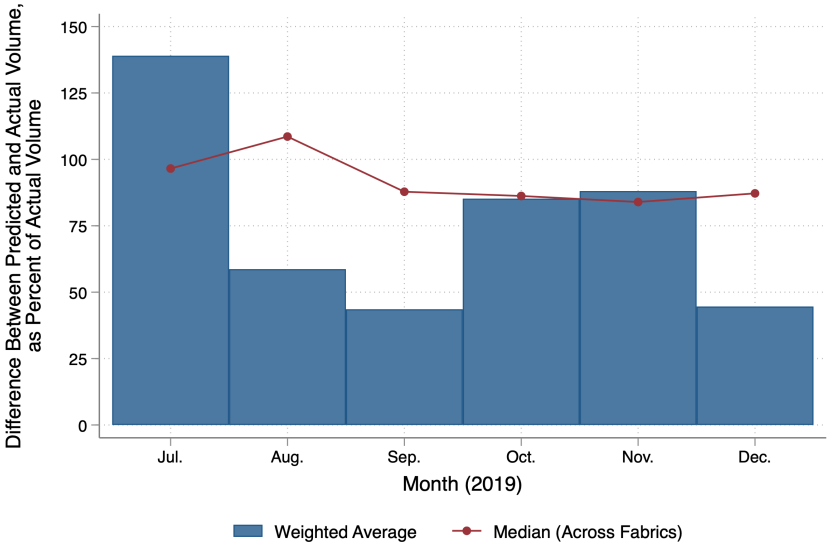
many different states, even abstracting from enforcement concerns. The just-in-time nature of the procurement process only exacerbates this problem as the complex contingent contract would need to be written in a short time period. Third, even if a contract could specify prices as a function of the supplier's 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. Fourth, capacity utilization is likely not contractible, as it is not observable by the buyer.

OA1.5 Why Not Exclusively Agreement Suppliers?

Given the evidence that agreement suppliers offer lower prices than 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 7.

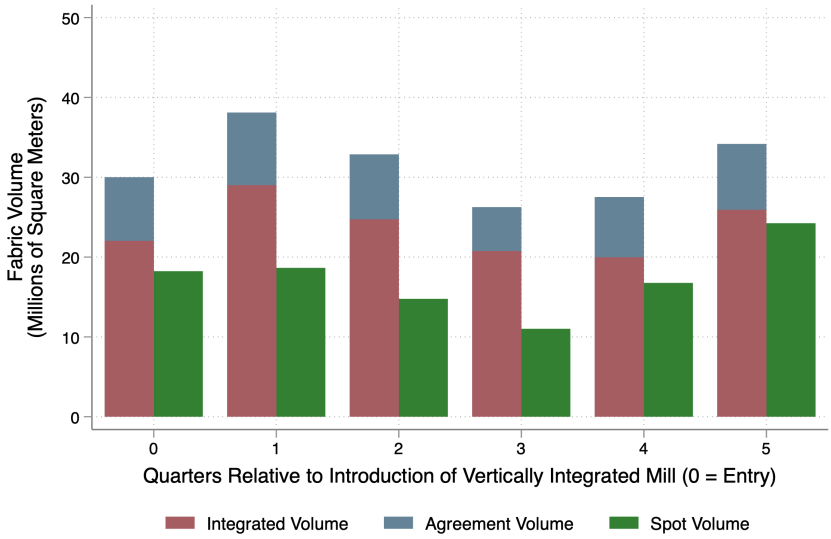
Figure OA1.8: Forecast Error in Volume Predictions



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

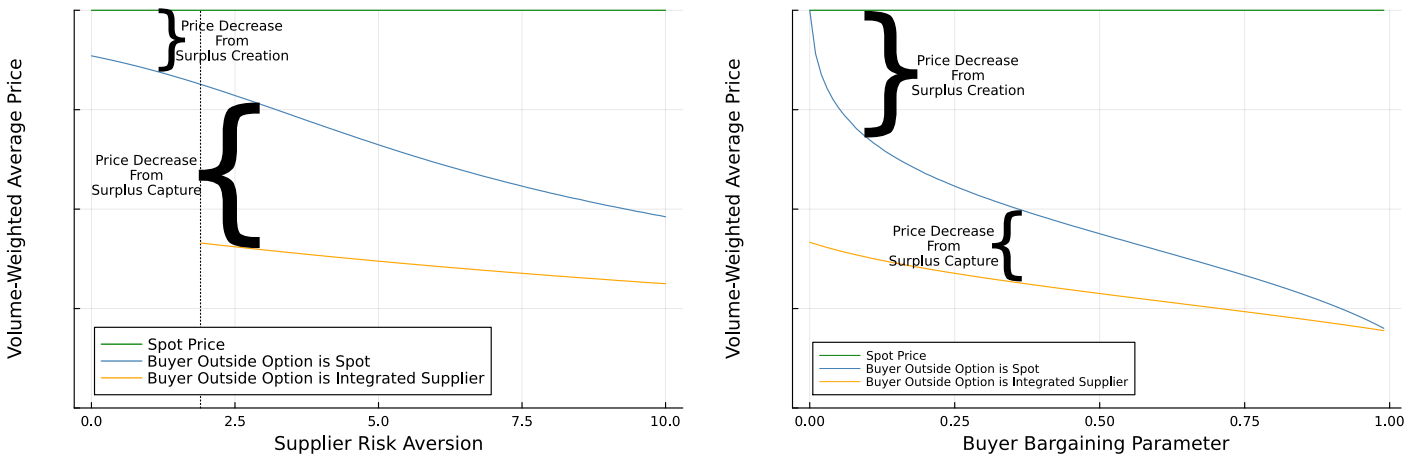


Note: Data from a panel constructed from the universe of fabric transactions by the buyer at the supplier group-by-quarter level. Bars represent the quantity of fabric purchased during the quarter for *i)* integrated and agreement suppliers and *ii)* spot suppliers, summing across all suppliers in the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

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

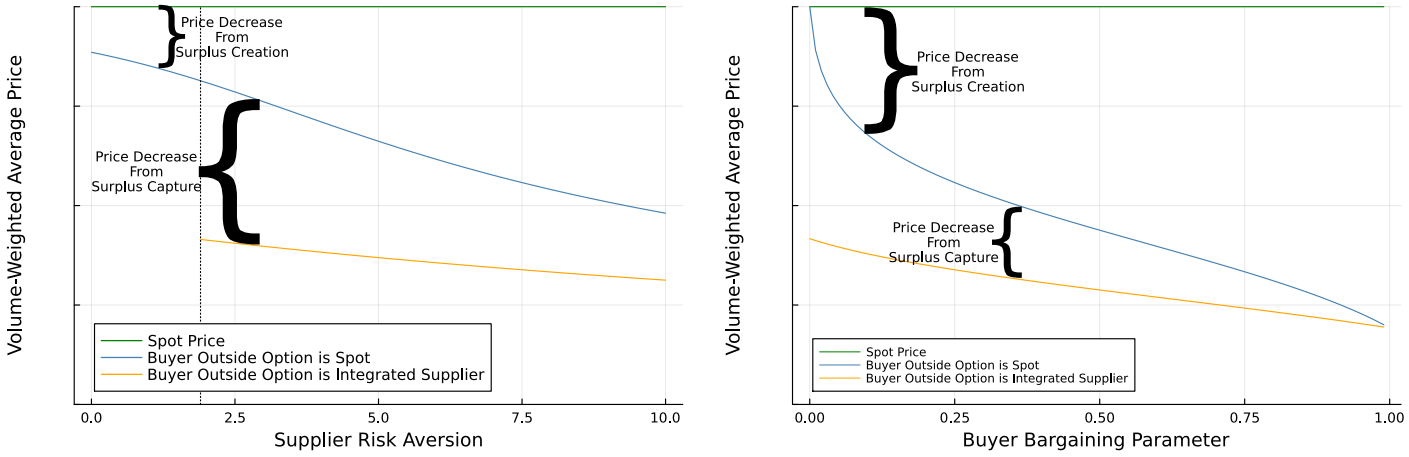
Results are robust to the external supplier having lower or higher costs than the integrated supplier, shifting costs by 10% in either direction. To ensure that optimal quantity is still $q = 1$ without changing market prices, I only shift the fixed part of costs by 10%, keeping convexity the same. Results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

OA1.6 Signaling Game: A Numerical Example

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

Consider when the buyer bargains with a supplier of type $\theta = 1$. This value of θ is sufficiently large, in the economic environment considered, to generate agreement surplus regardless of the buyer's residual capacity state. The Nash bargaining prices in such a setting are:

$$\begin{aligned} p^A(0) &= .5 \\ p^A(1) &= 1 \end{aligned}$$

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

$$\begin{aligned} .5 + 1 - 1^2 &\geq \mu_0(.5 + 1 - 1^2 - \gamma) + (1 - \mu_0)(1 + 1 - 1^2 + \gamma) \\ .5 &\geq \mu_0(.5 - \gamma) + (1 - \mu_0)(1 + \gamma) \\ (2\mu_0 - 1)\gamma &\geq .5 - .5\mu_0 \\ \frac{2\mu_0 - 1}{1 - \mu_0}\gamma &\geq .5 \end{aligned}$$

Belief consistency for a weak Perfect Bayesian Equilibrium requires $\mu_0 = p$ (recalling that p represents the probability that the in-house supplier is *not* capacity constrained). Therefore,

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

Although I do not observe p directly from the data, note that $\frac{2p-1}{1-p}$ increases in $p \forall p \in [0, 1)$; therefore, for p, γ sufficiently high, this inequality holds. Formally, as long as the inequality holds (meaning p, γ are sufficiently high), then the equilibrium features the buyer playing $\hat{q} = 0$ and the supplier accepts after $\hat{q} = 0$.

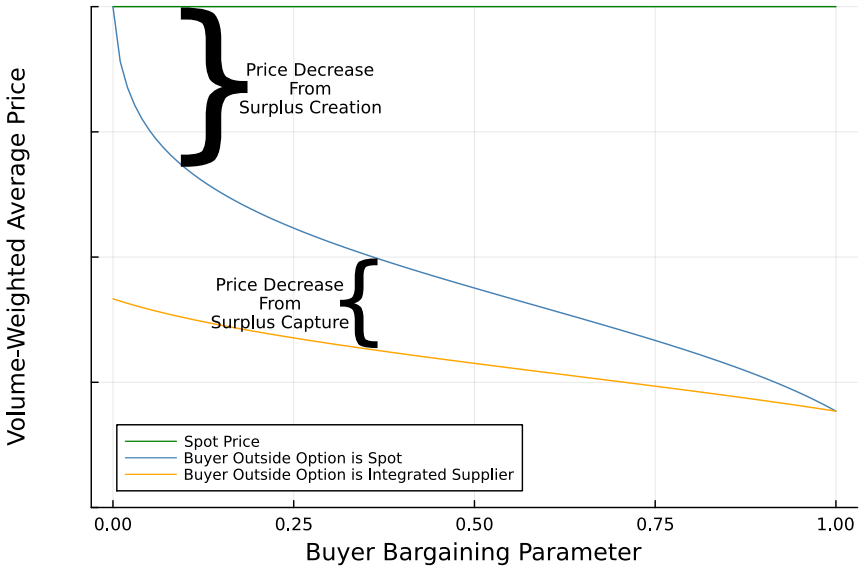
Alternatively, consider when the buyer bargains with a supplier with $\theta = -.5$, as occurs for a supplier with low risk aversion. (Recall that θ is normalized such that it represents the value of the nonprice agreement components **less profits from the spot market**, so negative values are reasonable). Then, there is surplus only when the buyer is capacity constrained (*i.e.,* $\tilde{q} = 1$): when $\tilde{q} = 0$, surplus is $-p^A - (-C_B(0, 1)) + p^A + \theta - C_S(1) = \theta < 0$. Therefore, the agreement dissolves when $\theta = -.5$ and the buyer's outside option in bargaining is in-house supply with slack capacity. However, surplus generated by the agreement, when the buyer is capacity constrained in-house and would source from the spot market, is $\theta - C_S(1) + p^S = -.5 - 1 + 2 = .5$. Then, the bargained price is $p^A(1) = 1.75$. However, when $\tilde{q} = 0$, the surplus from cooperation is $\theta - C_S(1) + C_B(1) = -.5 - 1 + 1 < 0$.

In that case, the agreement breaks down.

Now, I show that the supplier does not accept when the buyer plays $\hat{q} = 0$. The supplier's payoff from accepting $\hat{q} = 0$ is 0 given that the agreement breaks down. However, the payoff from rejecting is $p(0) + (1-p)(p^A(1) + \theta - C_S(1) + \gamma) = (1-p)(.25 + \gamma) > 0 \forall p \in [0, 1)$. Therefore, rejecting dominates accepting. It follows that the supplier rejects after $\hat{q} = 0$. And, as accepting after $\hat{q} = 1$ is a dominant strategy, the supplier will do so. Therefore, the buyer chooses between playing $\hat{q} = 0$, where the supplier rejects, and playing $\hat{q} = 1$, which the supplier accepts. The buyer prefers then to play $\hat{q} = 1$ for a cost savings of .25 relative to playing $\hat{q} = 0$.

OA1.7 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.

OA1.7.1 Outside Estimation: Log Quantities

Table OA1.3: Log Quantity and Capacity

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

Note: Data from 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.

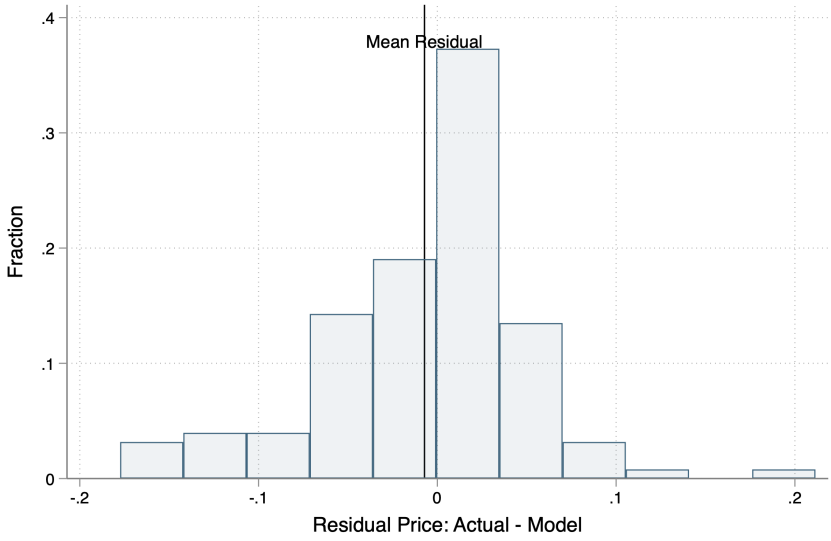
OA1.8 Model Validation and Fit

OA1.8.1 Model Validation: Untargeted Moment

As the estimation routine minimizes root-mean-squared-error (RMSE) between model-implied and actual prices, estimated prices could be biased. The possibility of bias emerges because minimizing RMSE targets a combination of 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.[104] 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.

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).

OA1.8.2 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 the analogous distribution from two alternative estimators, as shown in Table OA1.4.

First, I compare model estimates to the mean price for the supplier in pre-integration quarters ("pre-integration mean"). Note that the pre-integration mean approximates price

[104]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.

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.

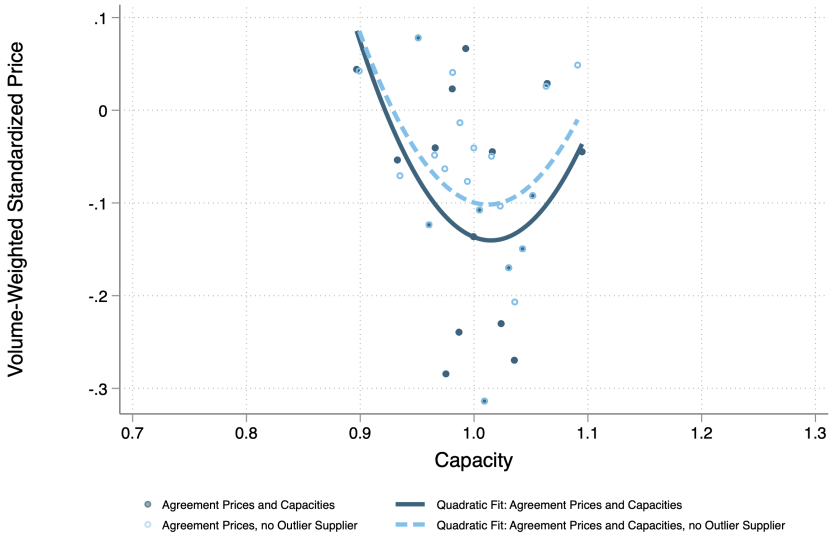
Table OA1.5: Out-of-Sample Fit Comparison (OLS with Tuning)

	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).

OA1.9 Pricing Function Robustness to Outlier Supplier

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

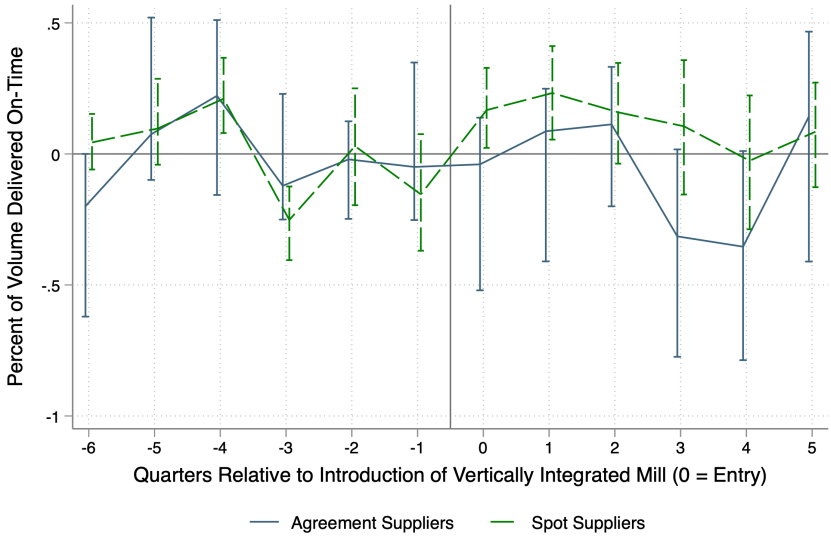


Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Figures show both the quadratic fit and a scatter plot. The quadratic fit plots fitted values from a regression of volume-weighted standardized price on both normalized quarterly volume and 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 4.1.2). The scatter plot in Figure OA1.14 shows binned prices and volumes using 20 quantiles.

OA1.10 DiD Supplemental Analysis and Robustness Checks

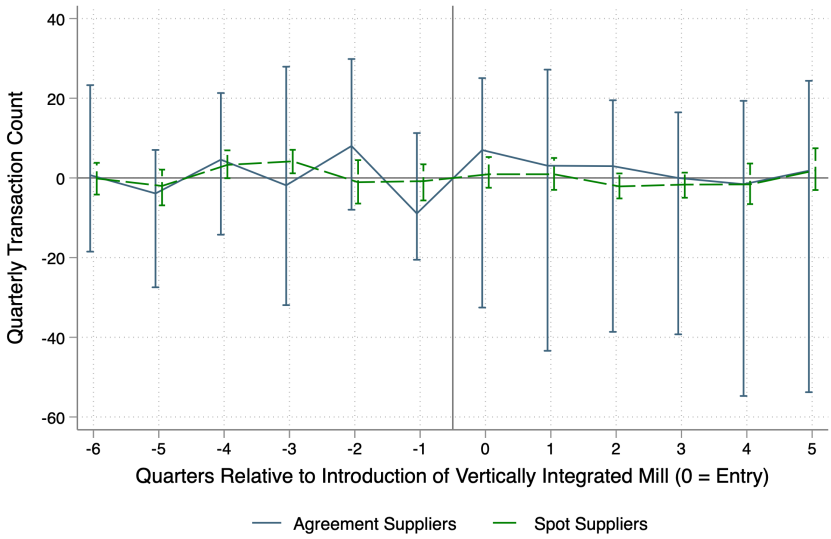
OA1.10.1 Additional Event Studies

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*.11 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.

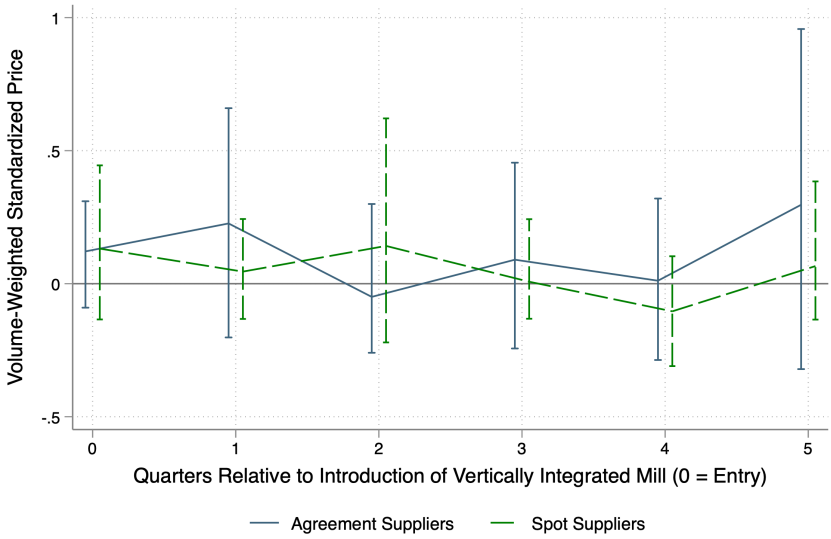
Figure OA1.16: DiD Estimates: Effects on Transaction Count



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*.11 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.

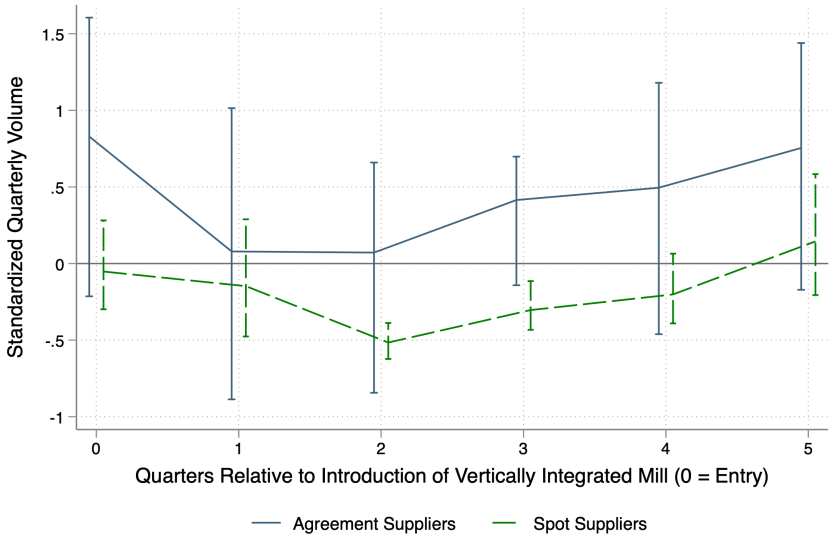
OA1.10.2 Synthetic DiD

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



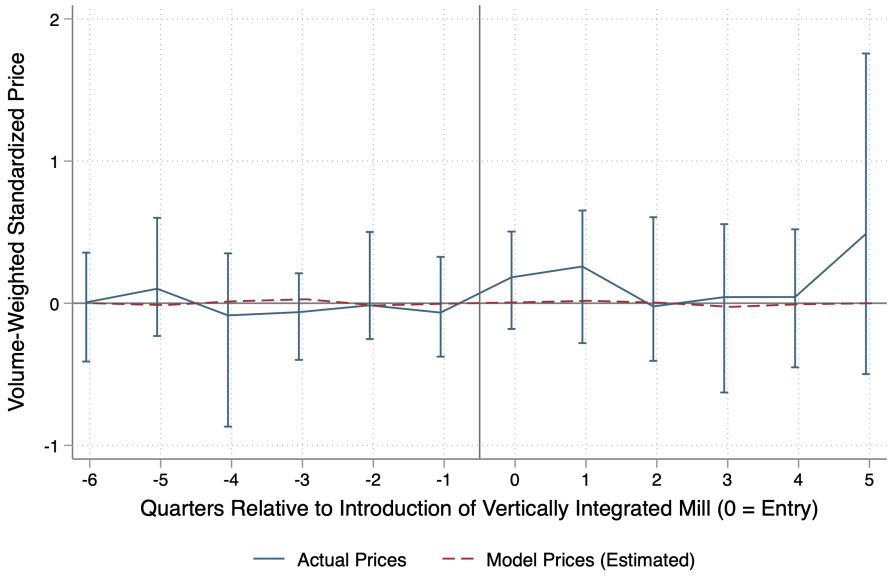
Note: Data from a supplier-by-quarter panel constructed from the universe of fabric transactions by the buyer. Event study estimates use the synthetic difference-in-differences design in Arkhangelsky et al. (2021) and are estimated separately for agreement and spot suppliers. Suppliers are considered 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).

OA1.10.3 DiD Validation: Event Study Estimates as Untargeted Moments

As structural model estimation does not use the variation leveraged in the DiD design,[105] 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 95% confidence intervals of event study coefficients using actual prices.

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

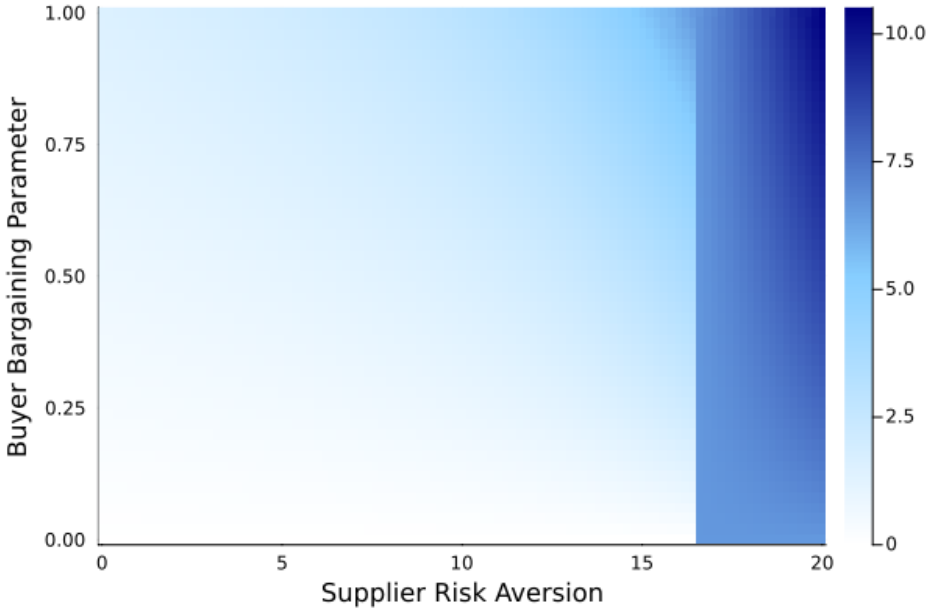
Figure OA1.19: DiD Coefficients as Untargeted Moments



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*.11 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.

OA1.11 Agreement Discounts Incorporating DICC

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 microfinance institution insurance rates in India; 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).

OA1.12 Horizontal Merger of Upstream Suppliers

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

merger due to capacity variation reduction, but the effect is small—less than 1%.[106]

I also consider the effects of a horizontal acquisition where the acquired firm is a small constrained supplier with the minimum risk aversion such that the buyer captures value through replacement threats. In this scenario, the horizontal acquisition increases prices by 24.6%. Decomposing the change as above, I find that reducing the variance increases price by 23.6% and reducing risk aversion similarly increases prices by 22.5%. These results have similar magnitudes because both of these effects of the merger are sufficient to stop the buyer from using replacement threats. This result highlights the role of scale economies in effectively reducing the value of quantity assurance, as unfavorable capacity states are less likely to be realized after integration due to the increase in quantities.

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

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