Essays on Industrial Organization in Low-Income Countries

by

Russell P. Morton

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2024

Doctoral Committee:

Professor Ying Fan, Chair
Professor Lauren Falcao Bergquist, Yale University
Professor Zach Brown
Professor Dean Yang

Russell P. Morton
rpmorton@umich.edu
ORCID iD: 0000-0003-0668-5816

© Russell P. Morton 2024

ACKNOWLEDGEMENTS

I am immensely thankful to have benefited from the mentorship of a fantastic team of advisors who have been encouraging, thoughtful, patient, and wise beyond expectation. Working with Lauren on coffee has truly been a central part of my experience, reminding me what I love about economics research and helping build confidence and trust in my training and skills, especially at times when I found it hardest to do so. Through both advising and working on coffee, Lauren has improved how I approach all facets of research–from designing and running field experiments to connecting assumptions and mechanisms in models to the empirical setting and understanding what sources of variation identify key parameters. Although I initially only took the IO field class because theory was not offered, Zach's class undoubtedly changed my path forever. He highlighted how powerful the IO toolkit is for quantifying economic mechanisms and evaluating policies–and I never looked back. As an advisor, Zach has been a wonderful resource, from helping me slow-down and focus on the big picture to technical advice for building and estimating structural models that will stay with me for the rest of my career. Zach has also been a much-need calm voice of encouragement, reassurance, and support. Ying has guided me throughout this process, providing invaluable advice on all stages of research–from going through the core of my model and estimation approach to listening to countless versions of my introduction. Ying has also shown how much she cares about her advisees as people, reassuring me about my work and sharing insights on life as a researcher. I am also immensely thankful to Dean, whose approach to economics and encouragement has been instrumental in my experience–I often hear his voice in my head (what I affectionately call 'Dean'-isms) when people talk to me about research. Furthermore, Dean has been the linchpin of the development community at Michigan, where he has created a collaborative environment that has made me feel at home even when my research is not the most typical reduced-form micro-development. Last, Jie Bai has effectively been an unofficial fifth member of my committee. She has supported me through both fantastic constructive feedback on research, listening to my job market paper presentation many times, as well as believing in me even when I did not. Beyond my advisors, many other faculty provided helpful feedback on my research and encouragement; while I will not name them all, Hoyt Bleakley, Francine Lafontaine, David Miller, and Ameet Morjaria were especially helpful.

I also am extremely grateful for fellow Ph.D. students that have improved my experience both within and outside Lorch Hall. It is hard to imagine that I would have been able to complete the Ph.D. program without Max Huppertz, who has been a better colleague and friend than I could have ever hoped. From working together on problem sets at Jefferson Market during first year to being a sounding board for both technical questions and high-level framing and presentation issues, Max is always my first stop when I do not know where to go. He unceasingly provides thoughtful and constructive feedback (including even coding simulations), even when he is incredibly busy. He is also a great friend, from office coffee breaks to movie nights, tennis, and raclette, and my time in Ann Arbor outside the office would not have been nearly as enjoyable without him. I am also very fortunate that the coffee project connected me with Yulu Tang. Not only has Yulu has listened to me present my job market paper multiple times–and provided useful, constructive feedback–but she has also commiserated with me throughout the job market, and helped me feel positive even when it was hard to do so. Many other fellow Ph.D. students have had important positive effects on my experience–both inside and outside Lorch Hall–including James Allen, Luis Baldomero-Quintana, Steph Karol, Emir Murathanoglu, and Shwetha Raghuraman.

Outside of academia, I am incredibly fortunate to have had such a great support system from family and friends. My parents have always taught me to believe that I could accomplish anything I set my mind to (as long as I work hard enough) and that I should find a profession that feels like an avocation rather than just a vocation. My sister has helped remind me that life does not end when I close STATA for the evening, indulging my complaints about academia when I felt low and then jumping to buy a Rochester t-shirt the next day. And words cannot express my thanks to Emilie–for her steadfast trust in and support for me, her willingness to pack-up her life for my career, her patience listening to me practice over and over again, and her kindness to (mostly) tolerate my collection of coffee mugs strewn about the house in busy times.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS ii

LIST OF FIGURES vii

LIST OF TABLES ix

LIST OF APPENDICES xi

ABSTRACT xii

CHAPTER

1 Vertical Integration and Relational Contracts: The Threat Point Effect . 1

1.1 Introduction 1
1.2 Data and Context 7
1.2.1 Data 7
1.2.2 Context: Vertical Integration at a Large Buyer 8
1.2.3 Context: Fabric Procurement and Garment Manufacturing 10
1.2.4 Context: Relational Contracts in Fabric Procurement 12
1.2.5 Alternative Possible Explanations for Observed Patterns: Supplier Heterogeneity and Analysis of Placebo Suppliers 17
1.3 Conceptual Model of the Threat Point Effect 19
1.3.1 Model Description and Set-Up 19
1.3.2 Threat Point Effect and Buyer Bargaining Parameter 23
1.3.3 Threat Point Effect and Supplier Risk Aversion 24
1.4 Structural Model 26
1.4.1 Why a Structural Model 26
1.4.2 Estimation Approach 27
1.4.3 Separate Identification of Buyer Bargaining Parameter and Supplier Risk Aversion 33
1.4.4 Buyer Bargaining Parameter and Supplier Risk Aversion Estimates 36
1.4.5 Model Validation and Fit 37
1.4.6 Quantifying the Threat Point Effect by Buyer and Supplier Characteristics 40
1.4.7 Model Validation using Difference-in-Difference Estimates 43
1.5 Policy Implications: Reallocating Surplus to Small Suppliers 49

1.5.1 Increase Downstream Buyer Competition . . . 49
1.5.2 Create Missing Market for Insurance . . . 52
1.6 Conclusion . . . 53

2 Vertical Integration and Relational Contracts: Empirical Evidence on Substitutability from Indian Garment Manufacturing . . . 56

2.1 Introduction . . . 56
2.2 Relational Contracts and Vertical Integration as Substitutes: Which Fabrics are Brought In-House? . . . 60
2.2.1 Estimation and Identification . . . 60
2.3 Ruling Out Threats to Causal Interpretation: Omitted Variable Bias . . . 64
2.3.1 Supplier Market Power . . . 65
2.3.2 Production Cost Heterogeneity . . . 67
2.3.3 Fabric Specific Comparative Advantage . . . 68
2.3.4 Reliability . . . 70
2.4 Ruling Out Threats to Causal Interpretation: Placebo Supplier Analysis . . . 71
2.5 Model . . . 73
2.5.1 Model Set-Up . . . 73
2.5.2 Hold-up Prices . . . 74
2.5.3 Vertical Integration Resolves Hold-Up . . . 75
2.5.4 Relational Contracting Resolves Hold-Up . . . 75
2.6 Policy Implications and Conclusion . . . 77

3 Quality Incentives and Upgrading in Uganda's Coffee Supply Chain . . . 79

3.1 Introduction . . . 79
3.2 Data and Context . . . 84
3.2.1 Data: Network Mapping . . . 84
3.2.2 Data: High-Frequency Sampling Surveys . . . 84
3.2.3 Data: Producer and Intermediary Surveys . . . 85
3.2.4 Context: Coffee Quality and Production in the Rwenzori Region of Uganda . . . 86
3.3 Intermediation, Quality Incentives, and Quality Upgrading . . . 88
3.3.1 Intermediation . . . 88
3.3.2 Quality Incentives: The Quality Premium . . . 88
3.3.3 Quality Investments . . . 92
3.4 Barriers to Entry, Productive Substitutability, and the Diminished Quality Premium Upstream . . . 92
3.4.1 Barriers to Entry to High-Quality Intermediation . . . 93
3.4.2 Productive Substitutability . . . 94
3.4.3 Interaction of Barriers to Entry and Productive Substitutability . . . 94
3.4.4 Alternative Explanations for a Diminished Quality Premium Upstream 95
3.5 Supply Chain Model of Quality Investments . . . 95
3.5.1 Upstream Profit Maximization . . . 96
3.5.2 Downstream Profit Maximization . . . 97

3.5.3 Market Clearing 99
3.5.4 Markowns with Value-Added Along the Chain 100
3.6 Calibrated Model Results 101
3.7 Conclusion 106

APPENDICES 107

BIBLIOGRAPHY 161

LIST OF FIGURES

FIGURE

1.1 Integrated Mill Constructed in June 2018 8
1.2 Fabric Production Over Time by Supplier Type 9
1.3 Fabric Production Over Time for Relational Contract and Market Suppliers . . 10
1.4 Garment Manufacturing Process 11
1.5 Discounts and Supplier Type 15
1.6 Coefficient of Variation and Supplier Type 16
1.7 Volume Comparison: Placebo and Relational Contract Suppliers 18
1.8 Relational Contract Prices and the Buyer Bargaining Parameter 23
1.9 Relational Contract Prices and Supplier Risk Aversion 24
1.10 Cost Convexity at the Integrated Supplier 31
1.11 Relational Contract Prices for Different Parameters 35
1.12 Shift Supplier Risk Aversion 36
1.13 Shift Buyer Bargaining Parameter 36
1.14 Untargeted Moments: Bias 39
1.15 Threat Point Effect Magnitude (% of pre-integration relational price) 41
1.16 Buyer Gains from Relational Contracts (% of market price) 42
1.17 Difference-in-Difference Estimates for Effect of Vertical Integration on Standardized Prices 45
1.18 Difference-in-Differences Estimates for Effect of Vertical Integration on Quarterly Volume 46
1.19 Untargeted Moments: Difference-in-Differences 48
1.20 Downstream Buyer Competition and Relational Contract Prices 51
1.21 RC Discount Without Insurance 53
1.22 RC Discount With Insurance 53

3.1 Transaction Data Example 85
3.2 Downstream Activity Choices and Cost Function Realization 98
3.3 Markdowns with Productive Investments by Downstream 100

A.1 Fabric Sourcing by Supplier Type 107
A.2 Distribution of Percentage of Months with Transactions (Pre-Integration) . . . 108
A.3 Variation in Monthly Quantity by Fabric 109
A.4 Discounts and Relational Contract Suppliers 110
A.5 Discounts and Supplier Type 111
A.6 Discounts and Supplier Type 112
A.7 Coefficient on Owed Quantity by Supplier Type 113

A.8 Volume by Supplier Type (Pre-Integration) . . . 114
A.9 Coefficient on Owed Quantity and Supplier Type . . . 115
A.10 Forecast Error in Volume Predictions . . . 118
A.11 Discounts and Supplier Type: Placebo Suppliers . . . 119
A.12 Discounts and Supplier Type: Placebo Suppliers . . . 120
A.13 Discounts and Supplier Type: Placebo Suppliers . . . 121
A.14 Discounts and Supplier Type: Placebo Suppliers . . . 122
A.15 Coefficient of Variation by Supplier Type: Placebo Suppliers . . . 123
A.16 Coefficient on Owed Quantity by Supplier Type: Placebo Suppliers . . . 123
A.17 Distribution of Owed Quantity by Supplier Type: Placebo Suppliers . . . 125

B.1 Relational Contract Prices and the Buyer Bargaining Parameter with Dynamics 127
B.2 Relational Contract Prices and Supplier Risk Aversion with Dynamics . . . 128
B.3 Relational Contract Prices, the Buyer Bargaining Parameter, and Supplier Risk Aversion when External Supplier has Lower Costs . . . 129
B.4 Relational Contract Prices, the Buyer Bargaining Parameter, and Supplier Risk Aversion when External Supplier has Higher Costs . . . 129
B.5 Distribution of Supplier Squared Residuals by Estimation Method . . . 139
B.6 Distribution of Supplier RMSE by Estimation Method . . . 140
B.7 Threat Point Effect for In-Sample Suppliers . . . 141
B.8 Trends in Fabric Demand and Vertical Integration . . . 142
B.9 Distribution of Propensity Scores by Supplier Type . . . 144
B.10 Difference-in-Difference Estimates for Effect of Vertical Integration on Reliability 145
B.11 Difference-in-Difference Estimates for Effect of Vertical Integration on Transaction Count . . . 146
B.12 Synthetic Difference-in-Differences Estimates for Effect of Vertical Integration on Quarterly Volume . . . 147
B.13 Synthetic Difference-in-Differences Estimates for Effect of Vertical Integration on Standardized Prices . . . 148

D.1 Distribution of Chain Lengths . . . 154
D.2 Quality Observability . . . 157

LIST OF TABLES

TABLE

1.1 Additional Moments for Estimating the Structural Model 30
1.2 Parameter Estimates from Structural Model 36
1.3 Out-of-Sample Fit Comparison 40
1.4 Pooled Pre- and Post-Treatment Effects 47

2.1 Substitutability of VI & RC 63
2.2 Substitutability of VI & RC: Supplier Market Power Robustness 66
2.3 Substitutability of VI & RC with Fabric Group FE 68
2.4 Substitutability of VI & RC: Buyer Fabric Comparative Advantage Robustness 69
2.5 Substitutability of VI & RC: Reliability Robustness 71
2.6 Substitutability of VI & RC: Placebo Suppliers Robustness 72

3.1 Summary Statistics: Producers 87
3.2 Summary Statistics: Intermediaries 87
3.3 Diminished Quality Incentives Upstream 89
3.4 Diminished Quality Incentives Upstream: Within Respondent Variation 91
3.5 Transactions by Layer: Quality Investments Along the Chain 92
3.6 Effects of Barriers to Entry & Productive Substitutability on Prices, Quality Incentives, and Markdowns 102
3.7 Effects of Barriers to Entry & Productive Substitutability on Quality Investments and Distrubtion of Surplus 104
3.8 Effects of Barriers to Entry & Productive Substitutability on Task Allocation . 105

B.1 Log Quantity and Capacity 134
B.2 Log Quantity and Prices 135
B.3 Capacity Patterns 136
B.4 Capacity Stability 137
B.5 Comparison with Other Risk Aversion Estimates 137
B.6 Buyer Bargaining Parameter and Supplier Risk Aversion: 95% CI 138
B.7 Out-of-Sample Fit Comparison (OLS with Tuning) 139
B.8 Balance on Untargeted Outcomes 144
B.9 Buyer Bargaining Parameter and Supplier Concentration 150
B.10 Supplier Risk Aversion and Supplier Concentration 150

C.1 Substitutability of VI & RC: OLS 151
C.2 Substitutability of VI & RC: Placebo and Supplier Market Power Robustness . 152

D.1 Big Exporters More Than Double the Quality Premium 155
D.2 Downstream Have Better Access to Big Exporters, Especially for High Quality . 156
D.3 Quality Premium Not Associated with Possible Relational Contracts 158
D.4 Upstream Sells to Many Distinct Buyers 159

LIST OF APPENDICES

A Relational Contracts in Indian Garment Manufacturing 107

Input Sourcing by Supplier Type 107
Relational Contract Discounts 107
Owed Quantity Guarantees 114
Why Not Exclusively Relational Contracts? 117
Placebo Analysis: Results 118

B Vertical Integration and Relational Contracts: The Threat Point Effect 126

Robustness Checks for Conceptual Model 126
Credibility of Displacing Relational Contract Supply with Integrated Supply . . . 129
Multilateral Bargaining 132
Parameterization and Solution of Conceptual Model 133
Structural Model of Relational Contract: Evidence Supporting Modelling Assumptions 135
Structural Model of Relational Contract: Comparison of Risk Aversion Estimates 137
Confidence Interval for Buyer Bargaining Parameter and Supplier Risk Aversion . 138
Model Validation: Out-Of-Sample Fit RMSE Density 138
Model Validation: Out-Of-Sample Fit Comparison with OLS with Tuning Parameter 139
Validation using Difference-in-Differences Estimates 140
Difference-in-Differences Supplemental Analysis and Robustness Checks 144
Horizontal Merger of Upstream Suppliers 148
Buyer Bargaining Parameter and Supplier Market Power 149

C Vertical Integration and Relational Contracts: Empirical Evidence on Substitutability from Indian Garment Manufacturing 151

Additional Analysis of Substitutability of Vertical Integration and Relational Contracts 151
Relational Contract Robustness to Buyer Costs to Providing Demand Assurance . 152

D Quality Incentives and Upgrading in Uganda's Coffee Supply Chain . . 154

Intermediation: Supply Chain Lengths 154
Market Access 155
Alternative Explanations for Reduced Quality Premium Upstream 156
Search Costs 158
Model Calibration 159

ABSTRACT

This dissertation consists of three chapters studying supply chains in low-income countries.

The first chapter studies firm responses to contracting frictions, which lead firms to rely on vertical integration or relational contracts rather than spot markets. Empirical evidence highlights that vertical integration and relational contracts co-exist both within industries and individual firms. What are the economic and policy implications of this co-existence? I identify and quantify a new mechanism called the *threat point effect*, which is the contract change from firms improving their bargaining position through partial vertical integration. I build and estimate a structural model to quantify the threat point effect in the context of a large Indian garment manufacturer that adds integrated capacity bargaining with relational fabric suppliers. Model estimation uses the universe of the manufacturer's fabric purchase transaction data. The threat point effect reduces input prices by 6.7% for small constrained suppliers that highly value the reduced exposure to demand shocks in the relational contract. Policy analysis using the model shows that supporting small firms' ability to smooth is more effective at reducing relational discounts that increasing downstream buyer competition.

The second chapter studies whether firms consider vertical integration and relational contracts to be alternatives, as empirical evidence shows these governance regimes result in similar transactions. Therefore, vertical integration and relational contracts are likely to be substitutes–an exogenous increase in the price of procuring inputs through one form increases the quantity procured through the other. However, vertical integration and relational contracts might not be substitutes. Different transaction governance forms could be specifically matched to different frictions or transaction characteristics, increased competition from the integrated supplier could break the relational contract, or the integrated firm may not be able to simultaneously meet commitments under the relational contract and achieve targeted internal capacity utilization. I test for substitutability using novel data on the universe of fabric transactions from a large Indian garment manufacturer. I find strong evidence of substitutability, leveraging exogenous variation at the product level to mitigate concerns about omitted variable bias. I conclude by discussing the policy implications of substitutability, highlighting that substitutability raises concerns that relational contracts can result in anticompetitive behavior.

The third chapter studies quality upgrading as many low and middle-income countries (LMICs) include upgrading to attain export premia in their development strategies. In LMIC supply chains, producers rarely sell directly to world markets; instead, products travel through intermediaries. We study how intermediation affects incentives for quality production and quality investments in the context of a Ugandan coffee supply chain. First, we map the supply chain, documenting that coffee changes hands, on average, twice before reaching exporters. Next, leveraging high-frequency transaction-level data, including lab quality assessments, we show that quality premium diminishes up the supply chain. To rationalize the diminished quality premium, we build a model with two key economic mechanisms. First, barriers to entry for intermediation in the high-quality segment result in less demand and more buyer power. Second, both producers and intermediaries engage in quality investments, with the substitutability of these investments reducing intermediary demand for high quality upstream. Both mechanisms can explain the reduced quality premium but have different implications for efficiency and distribution of surplus. We use a loosely calibrated version of the model to examine the efficiency implications and distributional consequences of barriers to entry and productive substitutability.

CHAPTER 1

Vertical Integration and Relational Contracts: The Threat Point Effect

1.1 Introduction

In many low-income countries, institutions to support interfirm trade are limited. For example, firms might not trust courts to enforce contracts effectively (Banerjee and Duflo, 2000). As a response to these contracting and other frictions,[1] firms often rely on either vertical integration or relational contracts to circumvent input markets. When firms vertically integrate, they bypass input markets by sourcing in-house. Alternatively, firms can use relational contracts to facilitate cooperative, rather than opportunistic, trading that creates surplus for both parties. These contracts prevent opportunistic behavior by punishing it through ending the relationship, blocking access to future surplus (Baker et al., 2002). Empirical evidence highlights that vertical integration and relational contracts often co-exist within an industry, and even within individual firms.[2]

What are the economic and policy implications of the co-existence of vertical integration and relational contracts within a firm? I propose that co-existence improves the integrated firm's bargaining position within its relational contracts. Consider a buyer bargaining with external suppliers, although similar effects apply when suppliers integrate. When the buyer integrates, its threat point (outside option) changes from other external suppliers to sourcing internally. This shift in the buyer's threat point alters its bargaining position with external suppliers, empowering the buyer to appropriate more of the relational contract surplus.[3]

[1] Other frictions include, for example, input market power (Hortaçsu and Syverson, 2007), and quality non-contractability (Hansman et al., 2020a).

[2] Many firms in low-income countries are partially integrated (*i.e.*, they source inputs both internally and externally). For example, 62% of integrated firms in Karnataka source some inputs externally (Garg et al., 2023) and Breza and Liberman (2017) study a partially integrated Chilean retailer with relational suppliers. As integration suggests the presence of frictions, many external suppliers for partially integrated firms are likely relational suppliers.

[3] When a supplier partially integrates, its threat point changes to selling to the integrated buyer rather

This contract change caused by vertical integration altering bargaining positions has policy implications for policy makers with distributional preferences. Bargaining positions shape how buyers and suppliers share gains from trade, but we know little about how this surplus is shared in relational contracts, despite their ubiquity in low-income countries (Macchiavello, 2022; Macchiavello and Morjaria, 2023). Furthermore, understanding the distribution of relational contract surplus speaks to concerns that global value chains (GVC), which often rely on relational contracts, are inequitable with low prices for low-income country producers (Boudreau et al., 2023)—the World Bank notes that "gains from GVC participation are not distributed equally."

This paper studies how vertical integration affects bargaining in the context of relational contracts between a large Indian garment manufacturer (hereafter, buyer) and its external fabric suppliers. In their relational contract, the buyer receives a discount relative to market prices in exchange for reducing suppliers' quantity variability (*i.e.*, demand assurance); many relational contracts similarly include quantity assurance and a wedge relative to market prices. Suppliers benefit as reduced quantity variability helps them control costs and stabilize profits, which can be especially valuable for risk-averse firms in low-income countries.[4] I leverage a combination of rich transaction-level data on the universe of the garment manufacturer's fabric purchases and an increase in vertical integration, as the buyer partially vertically integrates by constructing a new in-house fabric mill, to analyze how integration affects contracts with 471 suppliers, both relational and market.

I define the threat point effect as the change in contracts with external firms caused by the change in outside option (threat point) due to vertical integration. As the buyer benefits from the relational contract through discounts, the threat point effect in my setting is the additional decrease in relational prices after the buyer adds integrated capacity (*i.e.*, the extra discount the buyer receives in the relational contract after partial integration).

Then, I build a structural model to quantify the threat point effect. Motivated by both reduced-form evidence of the relational contract and institutional details of the empirical setting, the model features a buyer and a supplier bargaining cooperatively over a relational contract pricing function that determines surplus. Pre-integration, trading with the market is the threat point (outside option) for both the buyer and supplier; after integration, the buyer's threat point becomes its integrated supplier. Vertical integration allows the buyer to credibly threaten to use its integrated supplier to displace the external supplier.[5] The

than market buyer, improving contracts with external buyers.

[4]Examples of risk aversion shaping firm behavior in low- and middle-income countries include coffee mills (Blouin and Macchiavello, 2019) and Ghanaian agriculture producers (Karlan et al., 2014).

[5]I support this assumption in a companion paper (Morton, 2023a) that empirically documents substitutability between vertical integration and relational contracts leveraging exogenous variation in relational

model microfounds the importance of demand assurance in the relational contract through both supplier risk aversion and convexity of the production technology; I model the buyer, a large sophisticated firm purchasing nearly 230 million USD of fabric yearly, as risk neutral. A buyer bargaining parameter that governs surplus allocation closes the model.

The model illustrates that the threat point effect is largest for small constrained (*i.e.,* risk-averse) firms. Because these firms highly value the relational contract's demand assurance, the buyer can leverage its threat point change to reduce relational contract prices. However, the model also shows that the threat point effect can be small, even zero, for some cases due to the relational contracting environment. When pre-integration relational contract surplus is small, as occurs for suppliers with low-risk aversion, increased competition from the integrated supplier can break the relational contract, resulting in no threat point effect.

I estimate the structural model to quantify the threat point using unusually rich transaction data on the universe of fabric purchases by the buyer, where I observe the price, quantity, detailed product information, date, and supplier identity (including purchases from integrated suppliers). The model enables quantifying the threat point effect for varied buyer and supplier types, including out-of-sample types, which is economically important given the heterogeneity of the threat point effect highlighted in the model for different buyer and supplier characteristics (*e.g.,* supplier risk aversion). This heterogeneity is also empirically relevant as a wide range of firms use relational contracts—from major U.S. airlines (Gil et al., 2022) to small-holder Rwandan coffee farmers (Macchiavello and Morjaria, 2020). I provide a novel approach to separately identify the buyer bargaining parameter from supplier risk aversion based on the shape of the relational contract pricing function. Risk-averse suppliers prefer contracts that equalize profits; therefore, relational contracts with more risk-averse suppliers provide more profit stability against demand shocks by adjusting prices. Conversely, the buyer bargaining parameter shifts prices regardless of demand shock.

I find that the threat point effect (*i.e.,* the contract change caused by vertical integration) can decrease relational contract prices by as much as 6.7% for risk-averse suppliers. At 6.7%, the threat point effect is economically meaningful–larger than Indian VAT for apparel at the time.[6] I validate the model by showing it produces both a close match to untargeted moments and superior out-of-sample fit than two alternative estimators. I also validate the model using the buyer's construction of a new mill as a natural experiment. I compare contracts with suppliers before and after the new mill is active, leveraging heterogeneity in supplier exposure to integration.[7]

contracting costs at the product level.

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

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

This finding has important distributional implications, as vertical integration decreases relational contract prices most for small constrained suppliers that highly value demand assurance in the relational contract. Small firms (suppliers in this setting) in low-income countries are also highly policy relevant as they comprise the majority of firms in low- and middle-income countries (Hsieh and Olken, 2014) and receive specific attention from policy makers, including the IMF and World Bank. Furthermore, relational contracts are likely especially important for small firms as they may face additional frictions and have limited access to contracting and other institutions.

I evaluate two counterfactual policies: *i)* increasing downstream buyer competition, inspired by the focus on market power in industrial organization; and *ii)* creating the missing market for suppliers to insure profit against demand shocks given evidence of the deleterious effects of missing markets from development economics. Increasing downstream buyer competition has a trivial effect on prices paid to small risk-averse suppliers. The buyer responds to enhanced buyer competition by increasing the share of surplus that suppliers receive—raising prices—to prevent them from switching to another relational buyer in the long-run (as relational contracts require time and effort to establish, suppliers cannot switch in the short-run). However, once the buyer's threat point is its integrated supplier, the surplus in the relational contract is so small that even allocating all of it to the supplier does not meaningfully shift prices. It follows that omitting the threat point effect results in large overestimates of the benefit of this policy—when the buyer's threat point is market supply, there is sufficient surplus that reallocation changes prices.

Second, creating the missing market for profit insurance against demand shocks increases relational prices and effectively eliminates the threat point effect. Suppliers value insurance because it reduces their reliance on the relational contract to smooth profits, thereby mitigating the buyer's ability to use its threat point change to reduce relational prices. This result suggests policies to improve small firms' ability to smooth profit, such as increasing access to savings accounts to facilitate self-insurance or financial products to hedge exposure (*e.g.*, derivatives), can generate large benefits through both facilitating intertemporal smoothing and shifting bargaining positions and, thereby, profits.

I contribute to a rich literature studying how relational contracts react to shocks to economic conditions, such as supply or demand, or market structure and competition–specifically, whether the relational contracts survive or breakdown as a result (Macchiavello and Morjaria, 2015, 2020, 2023; Ghani and Reed, 2022; Gil et al., 2022). I focus on surplus sharing as the primary outcome of interest and analyze how buyer and supplier characteristics shape the level of and sharing rule for relational contract surplus, as well as how firms can behave strategically to increase their share of the surplus by integrating. Ad-

ditionally, I study the effects of a different shock–*vertical* organization of production–rather than shocks to market structure, which is *horizontal* organization of production.

The contracting literature (*e.g.,* Harris and Nguyen (2023)) also studies how transaction governance forms (*i.e.*, whether transaction terms are determined by spot markets, relational contracts, or an integrated firm) interact. Specifically, prior work shows that effective spot markets can, on the one hand, make relational contracts less desirable but, on the other hand, support relational contracts (*e.g.*, by providing transparent information about the value of the outside option). This paper also studies interactions between transaction governance forms, but analyzes the interaction between vertical integration and relational contracts and considers interaction within a firm rather than across firms,[8] illustrating that vertical integration can make relational contracts more desirable by improving contract terms. Additionaly, this paper contributes to the empirical bargaining literature studying how firms use threats of replacement to improve their bargaining position (*e.g.,* Ho and Lee (2019); Ghili (2022))–I show that integrated firms can also be used as threats of replacement (rather than just alternative external firms) and analyze how relational contracts can limit firms' ability to use replacement threats out of concerns of breaking the relational contract.

This research also complements prior work on the motivations for and effects of vertical integration. Both theoretical and empirical studies illustrate that frictions which limit the effectiveness of market transactions result in vertical integration (Williamson, 1979; Lafontaine and Slade, 2007; Acemoglu et al., 2009; Forbes and Lederman, 2009; Macchiavello, 2012; Breza and Liberman, 2017; Hansman et al., 2020a; Garg et al., 2023).[9] The literature on the effects of vertical integration focus on quantities, such as increasing input purchases due to reducing double marginalization (Hortaçsu and Syverson, 2007), or on other downstream buyers, such as foreclosure to or increased input prices for downstream rivals to weaken downstream competition (Asker, 2016). I contribute to these literatures by proposing a novel effect of vertical integration–changing bargaining position and contracts with upstream firms–that creates an additional rationale for integration. This mechanism can affect either upstream or downstream firms, depending on which firm integrates, although it applies only to partially integrated firms.

Understanding the motivations for and effects of vertical integration influences efficiency through competition policy, which can have large welfare effects in low-income countries

[8]My companion paper (Morton, 2023a) discusses interaction effects across firms, highlighting possible anticompetitive effects of relational contracts.

[9]Other studies provide evidence that firms integrate to transfer and spread intangible inputs, such as management practices and organizational capabilities (Atalay et al., 2014), across the firm. These two motivations are similar as, arguably, using vertical integration for intangible assets reflects the difficulties in contracting for intangible inputs given that they are typically firm-specific and hard to measure.

where many markets are fragmented and uncompetitive (Bergquist and Dinerstein, 2020; Banerjee et al., 2022; Bergquist et al., 2024). Because the threat point effect can shift input market prices, it can have either pro- or anti-competitive effects downstream depending on the direction of the input market price change–the threat point effect decreases input prices when a buyer partially integrates and vice-versa when a supplier partially integrates. However, in settings where relational contract quantities are fixed and are inframarginal, as occurs when spot market input transactions persist after integration, the threat point effect has negligible direct effects on competition because it does not shift marginal input costs–only inframarginal input prices change.

Because the trade-off between incentives and insurance (hold-up and demand assurance in my setting) is central to many vertical relationships (Lafontaine and Slade, 2007), my empirical model incorporates insurance and takes seriously both estimating risk aversion and separately identifying it from bargaining parameters. Standard models of vertical relationships (*e.g.*, Cuesta et al. (2019); Lee et al. (2021)) omit risk aversion, resulting in underestimates of the bargaining parameter for the more risk-averse party. And, even with correct parameter estimates, models generate incorrect prices if one firm provides insurance (*e.g.*, quantity assurance) but it is outside the model. In my empirical setting, omitting the demand assurance the supplier receives in the relational contract (a vertical relationship) would lead to underestimated supplier surplus and overestimated prices.[10] Additionally, I develop a novel strategy to separately identify bargaining parameters from risk aversion by leveraging variation in both price levels and the convexity of prices with respect to capacity.

Last, I contribute to the literature estimating structural models of relational contracts by developing an approach that does not require observing relational contracts breaking down (Galenianos and Gavazza, 2017; Startz, 2021; Harris and Nguyen, 2023). I directly estimate the value of the outside option leveraging data on costs and commonly available spot market prices. As relational contracts should generally not break down in equilibrium (and breakdowns may suggest the presence of changes in the economic environment), this approach can have broad applicability and generalize to other settings where either no relational contract breakdowns occur or breakdowns are not clearly observable.

The remainder of the paper is organized as follows. Section 2 describes the data and context. Section 3 presents the conceptual model of the relational contract pre- and post-vertical integration. In section 4, I estimate and validate the conceptual model. Then, I use it to quantify the threat point effect for different buyer and supplier types. Section 5

[10]Fortunately, the biases from omitting risk aversion and from omitting insurance offset each other. However, they are unlikely to counteract each other perfectly and could be different orders of magnitude, biasing price estimates.

discusses the policy implications of the threat point effect and analyzes policies to reallocate surplus to small firms. Section 6 concludes and considers broader implications.

1.2 Data and Context

My empirical setting focuses on an Indian garment manufacturer purchasing fabric from internal and external suppliers, where external suppliers include both relational contract suppliers and market suppliers not in a long-term relationship with the buyer. I first describe the data before providing context on the firm and the increase in integrated capacity at the buyer. Last, I describe the hold-up problem in this setting that motivates the relational contract, including documenting empirical patterns in the data consistent with my characterization of the relational contract.

1.2.1 Data

I use two data sources: transaction data and cost data. The rich transaction data provide information about the universe of fabric purchases completed by the buyer between September 1, 2016 and December 31, 2019, including the transaction price, the date, the supplier which provided the fabric (including for internal suppliers), the quantity, and detailed information about the specific fabric. The 34,681 transactions represent over one billion square meters of fabric purchased from 622 distinct suppliers (471 pre-integration). There are two integrated suppliers, with one built during the timeframe of the study and the other joining the firm many years prior to the start of my data. Additionally, the transaction data include estimated production costs for fabrics which are produced at the internal mills; these estimates are for a mill that is not capacity constrained and roughly represent average variable costs.

The cost data from the internal mills provide data on capacity utilization and costs at the two internal mills operated by the buyer. These data are at a monthly level from March 2019 to April 2020 and include a broad support of levels of capacity utilization from 41% to 102%.

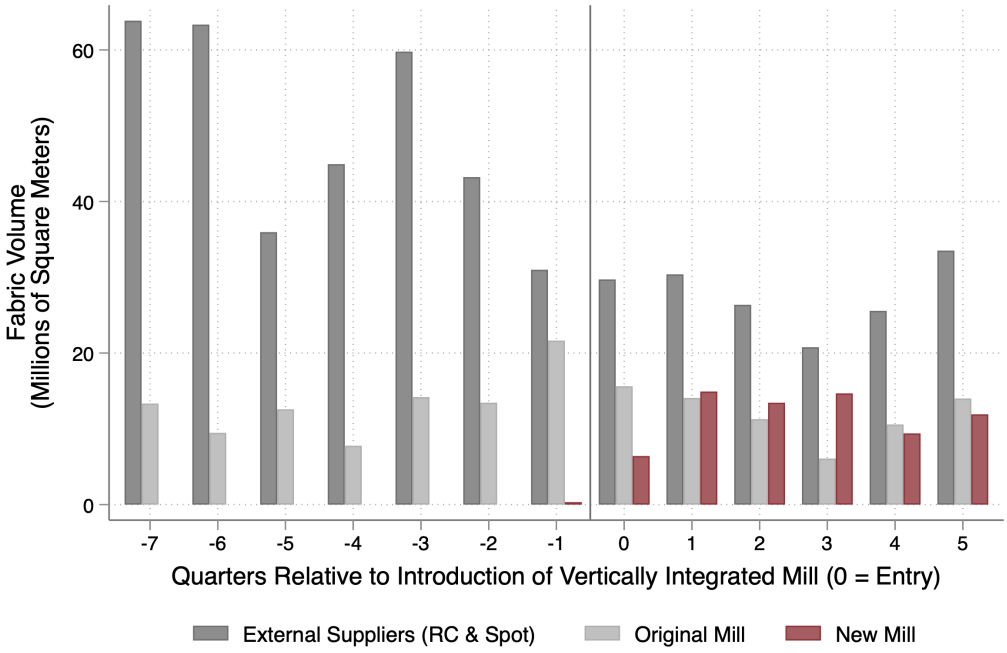
1.2.2 Context: Vertical Integration at a Large Buyer

Figure 1.1: Integrated Mill Constructed in June 2018



Prior to building its integrated mill, the buyer sources $\sim$ 250 million square meters of fabric per year (worth $\sim$ 228 million USD over 10,000 distinct orders) both at a pre-existing internal mill ($\sim$ 21% on average) and from external suppliers, both relational contract suppliers and market suppliers. Figure 1.2 illustrates three key facts about this integration that make this case study well-suited to analyze the threat point effect. First, the integration is large enough for the buyer to credibly threaten to bring production in-house. After integration, the percent produced internally increased greatly from $\sim$ 21% to $\sim$ 44%, with roughly half produced at the newly built internal mill. For comparison, the largest external supplier produced only 7%. Second, sourcing from external suppliers remains important after integration, as external suppliers still produce the majority of fabric. It follows that prices from external suppliers are still economically meaningful for the buyer, such that the buyer still cares about bargaining with them. Last, downstream demand from end clients (*e.g.,* garment retailers like Old Navy) does not change much (decreasing slightly to $\sim$ 215 million square meters of fabric per year), such that the buyer must make strategic decisions to reduce some combination of the number of external suppliers and the average volume per external supplier.

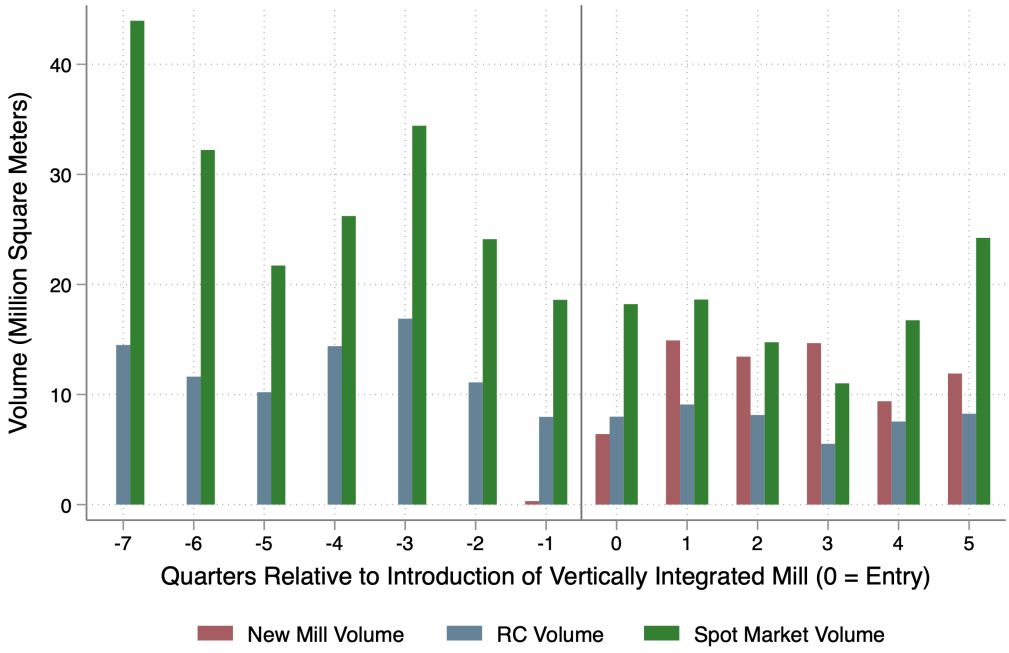
Figure 1.2: Fabric Production Over Time by Supplier Type



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

However, the integration is not so large that the buyer must reduce relational quantities– the buyer could reduce volumes from only market suppliers. Figure 1.3 illustrates this point, as the quantity produced internally at the new mill is much less than the pre-integration quantities sourced from market suppliers. It follows that the buyer can (and, as I later show, does) displace market rather than relational suppliers with integrated supply.

Figure 1.3: Fabric Production Over Time for Relational Contract and Market Suppliers

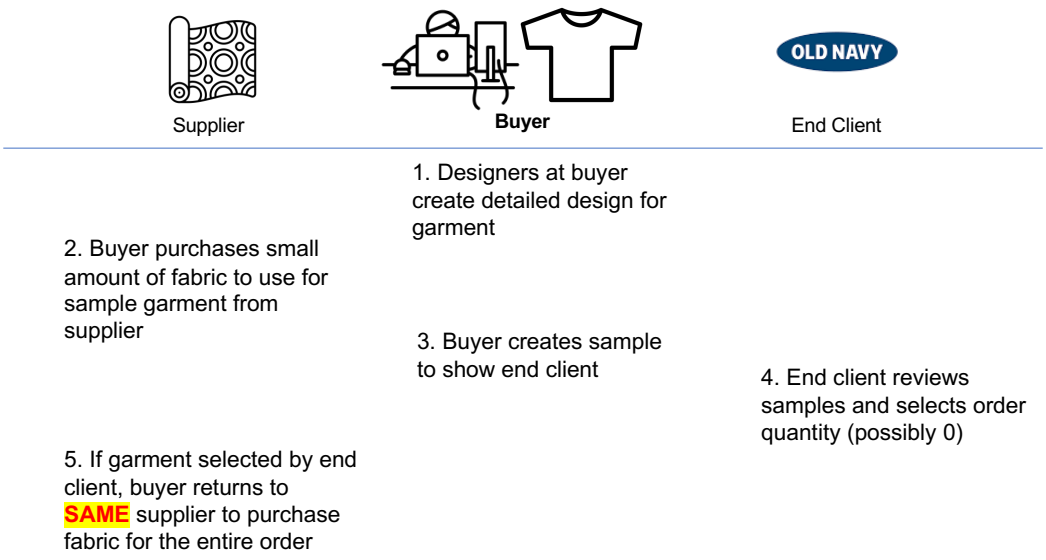


Note: Data from universe of fabric transactions by buyer. Bars represent the quantity of fabric purchased during the quarter from each supplier type.

1.2.3 Context: Fabric Procurement and Garment Manufacturing

Garment manufacturing, in this context, begins with designers who work for the garment manufacturer (the buyer) conceptualizing and planning the garment, sketching out the shape and color, and choosing the broad type of fabric (shown in Figure 1.4). The buyer then selects a supplier (a mill) to provide fabric for a small number of examples of the garment, called "samples"; although the buyer does pay for this fabric, this cost is largely negligible given the small quantity of fabric ordered and does not constitute a commitment from the supplier to the same price for future orders even for the same fabric. Fabric procurement costs for sample garments are so trivial that the buyer does not even record them in a centralized database. The samples are subsequently shown to an end client (*e.g.*, Old Navy), who then decides which samples to order and the quantities for any order it places. Only once the order is placed by the end client does the buyer return to the initial supplier to purchase the fabric necessary for the garment. Finalized designs are confirmed just-in-time (*i.e*, as close to shipment and sell date as possible) to ensure that they reflect recent trends. In general, there are only roughly 30 days between the initial fabric purchase for the sample and sourcing fabric for approved orders, in between which the sample is shown to the end

Figure 1.4: Garment Manufacturing Process



client and the quantity is confirmed.

A key feature in this process is that the supplier which provided fabric for the sample almost always receives the order for the fabric because the buyer does not want to change the fabric after the end client has seen it.[11] This supplier rigidity reflects that many subtle features of fabric are difficult to specify and copy, leading the buyer to believe that the end client might not accept a garment with a similar, but not identical, fabric made by another supplier. Additionally, timelines are short, preventing switching to a cheaper supplier and engaging in trial-and-error over time to match the fabric. The buyer's emphasis on ensuring that the garment produced for the end client is the same as the sample emphasizes the importance of the end client to the buyer. The importance of the end client is consistent with other research about the same buyer demonstrating that the buyer's production process prioritizes serving end clients over maximizing productive efficiency (Adhvaryu et al., 2019). Therefore, when the buyer selects a supplier to provide a sample, it is selecting the supplier for the entire order despite the fact that neither the buyer nor the supplier know if the

[11]Based on discussions with fabric procurement managers at the buyer, the probability of a supplier receiving a fabric order given that it provided the sample and is a large enough mill is at least 80%.

garment will be selected nor what the quantity would be (although both the buyer and the supplier likely have some beliefs over purchase probabilities and volumes from the history of transactions with the end client). Unsurprisingly, in the rare cases where the supplier is changed, production is typically brought in-house where quality can presumably be controlled intensely and for orders with long time before delivery to ensure that there is sufficient time to match the fabric.

The rigidity in supplier selection once the supplier has provided fabric for the sample creates a classic hold-up problem, allowing the supplier to charge a price that extracts at least the static surplus (*i.e.,* profit) the buyer receives from the order. The supplier could likely even extract some of the dynamic surplus from the relationship between the buyer and the end client given the importance of the relationship with the end client to the buyer.[12]

1.2.4 Context: Relational Contracts in Fabric Procurement

The hypothesized relational contract between the buyer (the garment manufacturer) and its suppliers (the external mills) centers around the buyer receiving a discount from the supplier, relative to the high price in the input market where the supplier holds up the buyer, in exchange for demand assurance in the form of reduced capacity variance as compared to transactions in the spot market. This is a relational contract as the cooperative equilibrium occurs, rather than the static equilibrium where the buyer is held up, and it is enforced by the continuation value of future transactions rather than an explicit contract that governs any individual transaction.

Even risk-neutral suppliers value demand assurance due to the connection between capacity and costs. Empirical evidence of the importance of capacity for costs includes that the buyer purchases a non-trivial amount of fabric at prices below estimated costs for a mill operating at reasonable capacity (19% of transactions and 26% of volume), using average variable cost estimates from the buyer's own internal mills. This result is consistent with suppliers charging lower prices when capacity utilization is low, and are not driven by new suppliers trying to signal low prices before increasing prices afterwards–the frequency of

[12]It is unlikely in this setting for buyer to be able to hold up the seller given the timing of the garment production process. Specifically, the buyer and supplier negotiate the transaction price once the end client has approved the garment and specified the quantity. At this point, while the specific fabric is very valuable to the buyer, the supplier could still sell the relevant capacity to other buyers, such as garment manufacturing firms targeting domestic markets and even as stuffing for pillows, etc. The buyer's main opportunity to hold up the supplier would be after the supplier has paid the sunk cost to actually produce the fabric, at which point the buyer could try to renegotiate prices. However, the buyer is unlikely to do so as the explicit contract for the specific transaction is likely sufficiently simple and straightforward to be easily enforceable in court (and, even if not, the buyer may have concerns about developing a reputation for not upholding prior price agreements which could make future contracting difficult; as a large firm, the buyer has strong incentives to maintain its reputation).

fabric purchased below cost for suppliers who have long histories of transactions with the manufacturer is similar to the overall frequency (19% of transactions and 23% of volume).

The importance of capacity for supplier costs likely reflects the environment having demand heterogeneity over time, due to the seasonality of the fashion industry, and input adjustment costs associated with regulations and financial constraints.[13] It follows that suppliers which can predict future quantities accurately can use that information to accept or reject various sampling (and subsequent trading) opportunities with other buyers to manage capacity, and thereby costs. Discussions about fabric procurement with the buyer are consistent with this story. For example, a fabric procurement manager at the buyer who works with H&M states that he even provides quarterly estimates of future transaction volume to the most important suppliers and that past volume "is the most important factor whenever we address and negotiate the prices and lead time."[14]

The central features of the relational contract between the buyer and the supplier, specifically the discount provided to the buyer in exchange for a quantity guarantee to the supplier, have testable empirical implications. First, the discount in the relational contract suggests that suppliers which charge lower average markups should be more likely to be in a relational contract compared to other suppliers. Second, the quantity guarantee in the relational contract suggests that volumes for relational contract suppliers should be more stable than volumes for other suppliers, as long as the quantity guarantee does not change meaningfully over the course of my sample.

As suppliers' relational contract status, long-run markups, and implicit quantity guarantees are not directly observable in the data, the econometrician must define these objects, or proxies for them, to bring these hypotheses to the data. Therefore, I define relational contract suppliers as suppliers which transact every month.[15,16] Economically, focusing on

[13]For example, the Factories Act of 1948 specifices that employers must pay double the standard wage to overtime workers.

[14]Similarly, the fabric procurement manager for Target highlights the relevance of capacity, noting that he tracks "in a spreadsheet style wise/Mill wise/ season wise how much business a mill can handle considering the mill capacity" based on previous transactions with the supplier.

[15]Defining relational contract status based on transaction frequency is common in the empirical contracting literature. For example, Macchiavello and Miquel-Florensa (2017) similarly defines relational contracts based on consistent trading every period ("at least three consecutive seasons") and Macchiavello and Morjaria (2015) also use transaction frequency. Cajal-Grossi et al. (2023) assign a buyer to a contracting strategy, relational versus spot, using the ratio of shipments to suppliers as a measure of supplier concentration. This approach is quite similar to using transaction frequency in practice as there should be a strong correlation between transaction number and transaction frequency, especially as it seems unlikely that buyers in their study would concentrate purchases in time given that end clients purchase garments throughout the year and their demand is based on somewhat unpredictable trends, except for the case where the buyer is purchasing an uncommon fabric that is "on-trend", which is not likely to be part of a relational contract. Additionally, both approaches reflect a similar economic logic: relational contracts require future transactions for enforcement and should, therefore, have many transactions.

[16]I only consider months prior to the construction of the new internal supplier for defining relational

transaction frequency reflects that relational contracts enforce themselves through future transactions; therefore, parties in relational contracts must transact frequently. I classify 19 of 471 external suppliers as relational contract suppliers. These 19 suppliers account for 25% of fabric volume prior to integration.

Because the relational contract includes a discount, suppliers who provide a discount should be much more likely to be in relational contracts. There is little reason for other suppliers to offer a discount in equilibrium, although some suppliers may do so when costs are low due to idiosyncratic capacity shocks and the supplier underestimates industry-wide capacity utilization and, therefore, market prices.[17] In this setting, as production technology is anecdotally similar across suppliers (at least for the same fabric and among the set of suppliers that have passed the buyer's inspections and review process to be in the set of suppliers considered), there is little scope for suppliers to have persistently different costs, *except* through targeting capacity.

To test this hypothesis empirically, I estimate the relationship between supplier prices and relational contracting status. I measure supplier prices as the average standardized price per supplier. I standardize prices both at the fabric level, as costs differ across fabrics due to varying production times and input prices, and at the monthly level, as industry-wide capacity changes over time (importantly, while some seasons generally have higher capacity utilization, such as the early fall for December holiday shopping, individual fabrics experience large unpredictable heterogeneity, as shown in Appendix Figure A.3).

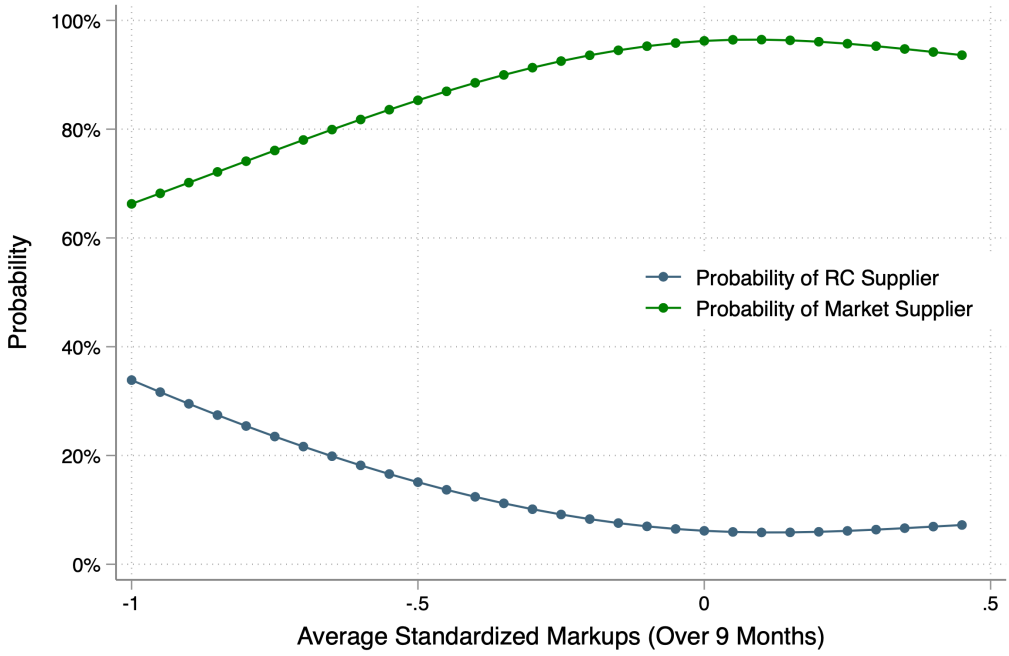
This measurement approach aims to capture the extent to which suppliers hold up the buyer. For example, consider two fabric purchases from different suppliers that occur at the same time for different fabrics with different production costs due to different material costs. If the two suppliers charge the same price, the supplier of the high cost fabric is not holding up the buyer to the same extent. It follows that using unadjusted prices as a measure of supplier prices would not account for cost heterogeneity across fabrics and over time. An alternative approach would be to measure markups directly using the buyer's cost data for fabrics they produce. However, this approach would both limit the sample, recalling that the buyer's cost data are only available for the subset of fabrics they produce, and fail to adjust for industry-wide seasonality in capacity usage.

contract status. This time restriction ensures that subsequent analyses of the effects of increased vertical integration on relational contracts do not allow treatment (*i.e.,* the creation of the internal mill) to also impact supplier classification as a relational contract supplier. Appendix Figure A.2 shows the distribution of the percentage of months with transactions across suppliers.

[17] It is possible that suppliers trying to become relational contract suppliers might offer discounts to try to signal their interest in forming a relational contract. This concern is unlikely in this setting given that the buyer is a large established firm that presumably cannot add more relational contracts, otherwise it would. I discuss this point in more depth in A.4.

Figure 1.5 illustrates the connection between long-run markups (relative to other suppliers, rather than marginal cost), proxied as the volume-weighted average standardized markup over the same nine month window,[18] and relational contracting. It shows a spline estimate of the association between long-run markups and whether a supplier is a relational contract or market supplier. Consistent with the connection between relational contract suppliers and discounts, the analysis illustrates that suppliers that do not provide discounts relative to other suppliers are very likely to be market suppliers (at least 90%). Additionally, the probability of being a relational contract supplier decreases in long-run markups, while the probability of being a market supplier increases in long-run markups. Similar patterns emerge when using one year rather than nine months to calculate long-run markups (see Appendix Figure A.4) and using a bin-scatter rather than a spline (see Appendix Figures A.5, A.6).

Figure 1.5: Discounts and Supplier Type

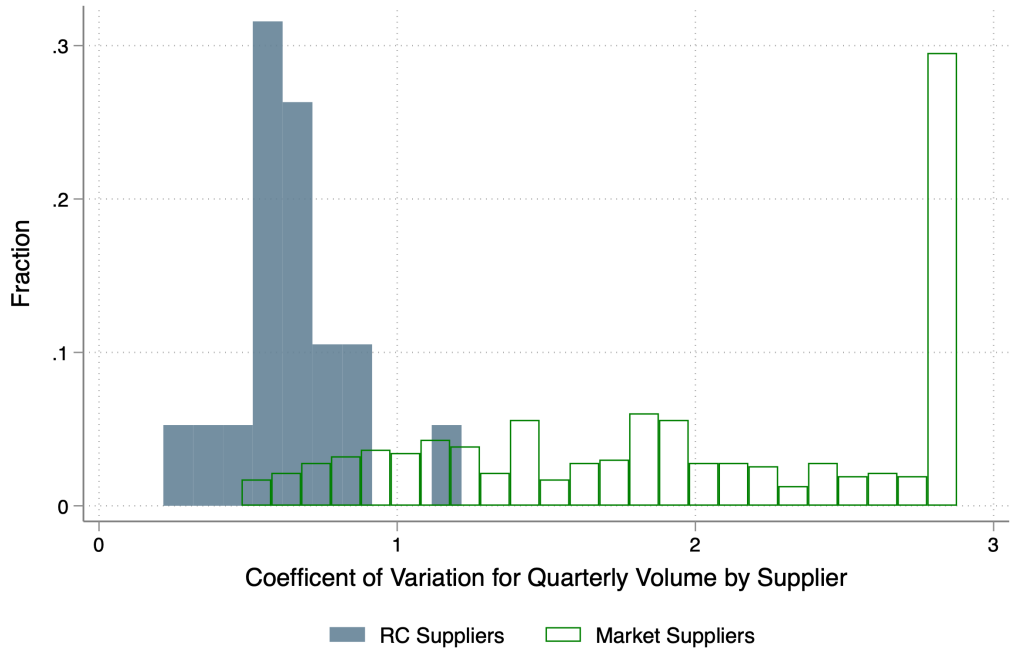


Note: Data from universe of fabric transactions by buyer. Probability estimates represent marginal effects from a non-parameteric series estimator calculated using transactions from the three quarters prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

[18]Specifically, the data are from the nine months prior to the introduction of the new vertically integrated mill.

Suppliers benefit from the relational contract through the demand assurance that facilitates suppliers both controlling capacity and, therefore, costs, and smoothing profits. Demand assurance means relational contract suppliers receive relatively stable volumes over time, conditional on the quantity level promised in the relational contract not changing in an economically meaningful way over the course of the sample. To bring this hypothesis to the data, I first compare the coefficient of variation for relational contract suppliers versus market suppliers. As the quantity guarantee is designed to stabilize capacity, relational contract suppliers with quantity guarantees should experience more stable capacity utilization than other suppliers. Figure 1.6 illustrates that suppliers in relational contracts have less volatile demand than market suppliers. The difference in average coefficient of variation by supplier type is statistically significant ($p < .001$). Importantly, this result is not entirely mechanical, as the criterion to define relational contracting status only incorporates transaction frequency and not volume. It follows that there could be suppliers identified as relational contract suppliers with extremely volatile demand, as long as those suppliers receive at least one order per month. In Appendix A.3, I show that quantity patterns for relational contract suppliers are consistent with the buyer directing volumes to fulfill quantity guarantees.

Figure 1.6: Coefficient of Variation and Supplier Type



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

An explicit contract at the time of sampling that specifies a price for the fabric should

the end client order the garment seems, intuitively, like an attractive alternative approach to resolve the hold-up problem. In Appendix A.3.1, I discuss why the buyer and seller rely on a relational rather than explicit contract. As the buyer prefers sourcing fabrics from relational contract suppliers to market suppliers, in principle, the buyer should only use relational contract suppliers. Appendix A.4 clarifies why the buyer is unable to source exclusively from relational contract suppliers in this setting.

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

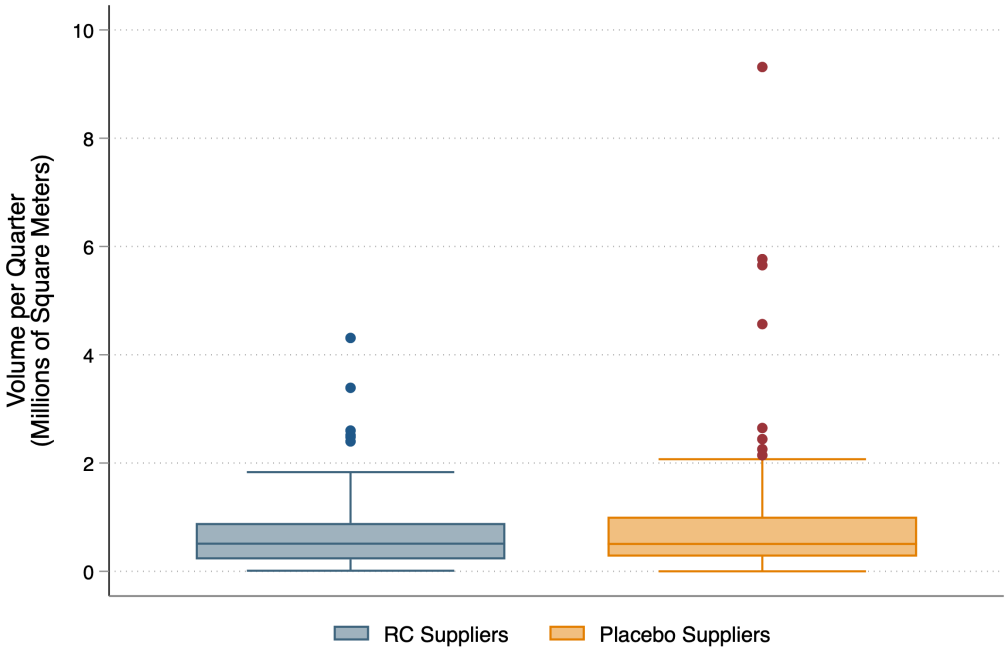
Although the evidence is consistent with the hypothesized relational contract, I conduct additional analysis to mitigate concerns that unobserved supplier heterogeneity explains the patterns in the data rather than relational contracts. For example, some suppliers could be more efficient, perhaps due to better management given that technology is fairly standardized (at least among the suppliers that produce similar fabrics and meet the buyer's screening criteria). More efficient suppliers might pass through some of their cost savings to the buyer, leading to lower prices, and might receive more (although not necessarily more stable) volume. That said, this scenario seems unlikely as there are no incentives for efficient suppliers to pass on any cost savings due to efficiency outside a relationship–statically, even very efficient firms would optimally charge the high hold-up price.

To mitigate concerns that unobserved supplier heterogeneity, rather than relational contracting, explains observed patterns in the data, I reproduce the analysis but compare relational contract suppliers to a subset of market suppliers that receive large orders but do not transact as frequently with the buyer ("placebo suppliers"). As the relational contract both depends on transaction frequency for the value of the future and provides demand assurance, placebo suppliers are plausibly not in a relational contract. Importantly, many stories of unobserved heterogeneity suggest that supplier volume, rather than transaction frequency, should differentiate contracting, such as the example of heterogeneity in supplier efficiency. In other words, plausible unobserved supplier heterogeneity is likely highly correlated with supplier volume. It follows that finding different patterns in the data for placebo and relational contract suppliers suggests that unobserved heterogeneity is not a valid alternative explanation for the empirical patterns. Note that this analysis alone does not suggest that relational contract suppliers are more or less efficient than other suppliers, but rather that the discounts they offer are due to the relational contract rather than any possible efficiency advantages they might have.

I construct a set of placebo suppliers by finding the suppliers with the largest volumes pre-

integration that are not relational contract suppliers; to keep the set of suppliers as directly comparable to the set of relational contract suppliers, I select the same number of placebo suppliers as relational contract suppliers. I then verify that volumes are indeed similar for relational contract and placebo suppliers, as this approach to defining placebo suppliers does not mechanically ensure that placebo suppliers receive similar volumes to relational contract suppliers.

Figure 1.7: Volume Comparison: Placebo and Relational Contract Suppliers



Note: Data from universe of fabric transactions by buyer. Boxplots show volume per supplier per quarter using all quarter prior to integration. Placebo suppliers are defined as the suppliers with the largest average quarterly volume that are not relational contract suppliers regardless of transaction frequency chosen such that the number of placebo suppliers matches the number of relational contract suppliers.

Figure 1.7 illustrates that relational contract and placebo suppliers have similar quarterly volumes and that the distributions of quarterly volumes have fairly similar support. If anything, the placebo suppliers have slightly *larger* volumes than the relational contract suppliers, with higher 25th and 75th percentiles, higher mean, and almost identical medians ($\approx$ 1% different).

Appendix A.5 documents findings that placebo suppliers both charge higher prices and have less stable volumes than relational contract suppliers.

1.3 Conceptual Model of the Threat Point Effect

This section presents a conceptual model of the relational contract between a buyer and supplier to illustrate how buyer and supplier characteristics shape the magnitude of the threat point effect, which is the difference between the relational contracting equilibrium outcomes (quantity and prices) before and after integration. I first describe the model, which is motivated by the features of the empirical setting and microfounds relational contract surplus through supplier characteristics (risk aversion), demand assurance, and the production technology. I then study threat point effect heterogeneity through the lens of the model with respect to buyer and supplier characteristics (*i.e.,* relationship type).

When the vertical integration is sufficiently large to more than crowd out all market suppliers, quantities in the relational contract must change (holding downstream demand faced by the buyer fixed). In this setting, the prices in the relational contract will also likely change as the new equilibrium relational contract quantities decrease, which would presumably increase relational prices by reducing the value of the demand assurance provided. However, my empirical setting allows me to isolate the change in relational contract *prices* due to the change in threat point as the buyer still purchased over half of its fabric inputs from external suppliers after integration, which is more than pre-integration relational contract quantities ($\sim 25\%$). The effect of the change in threat point on prices applies broadly regardless of the level of integration and should impact relational contract prices even in settings where the buyer would need to reduce relational contract quantities.[19]

1.3.1 Model Description and Set-Up

I model the prices in the relational contract as solving the Nash bargaining problem between the buyer and supplier, where the bargain is over a relational contract pricing function that maps quantities to prices. For simplicity and to focus on the key economics of this problem, the main version of the model presented does not include any dynamic incentives to govern the relational contract and make it self-enforcing. Extending the model in Appendix B.1.1 to incorporate standard relational contract dynamics (specifically, the dynamic incentive compatibility constraint as described in Macchiavello and Morjaria (2023)) does not change any key results. Rather, relational contract dynamics serve as the unmodeled mechanism that facilitates the buyer and supplier reaching a cooperative equilibrium; without dynamics, the supplier's optimal strategy would be to hold up the buyer in every period. Furthermore, the buyer and supplier bargain over a function that maps quantities to prices. As relational

[19]Excepting complete integration, which renders external suppliers irrelevant.

contracts require future trade, such an agreement is especially important in the context of a relational contract.

$$\underset{\substack{p^{RC}(q^{RC}) \\ \in \Gamma:\mathbb{R}^+ \to \mathbb{R}^+}}{\operatorname{argmax}} \left(\underbrace{-\mathbb{E}_{q^{RC}}[p^{RC}(q^{RC})q^{RC}] - O_B}_{\text{Buyer Surplus}} \right)^{\alpha_B} \left(\underbrace{\mathbb{E}_{q^{RC}}[U(p^{RC}(q^{RC})q^{RC} - C(q^{RC}))] - O_S}_{\text{Supplier Surplus}} \right)^{1-\alpha_B} \tag{1.1}$$

The relational contract solves the Nash bargaining problem in (1.1), specifying the price schedule in the relational contract for each observed quantity (capacity). The buyer bargaining parameter is α_B; higher values of α_B result in relational contract pricing functions that allocate more surplus to the buyer as its surplus factors more heavily in the objective function. As the buyer and supplier split the surplus, their bargaining parameters must add to one. Therefore, the supplier's bargaining parameter is $1 - \alpha_B$.

The buyer's surplus is the difference in profit when buying fabric in the relational contract–purchasing quantity q^{RC} at price $p^{RC}(q^{RC})$–versus buying fabric from the outside option–with value O_B. As revenue from the buyer's sales to the end client do not depend upon fabric prices since the negotiation over production prices and quantities precedes bargaining over fabric purchasing, buyer revenue is the same in both the relational contract and with the outside option. Therefore, I can express the buyer surplus from the relational contract as the difference in costs from purchasing the fabric in the relational contract as compared to the outside option. I assume the buyer is risk neutral as the buyer is a large firm. Additionally, the buyer requires suppliers to purchase some insurance for the main risk of production delays, and each supplier's quantity is small relative to the buyer's overall quantity, such that the buyer is reasonably risk neutral, especially for any individual relational contract.

The supplier's surplus is the utility of profit in the relational contract–with CARA utility function $U(\cdot)$ and production costs given capacity $C(q)$–less the value of the outside option O_S. I assume that suppliers are risk averse, as in Blouin and Macchiavello (2019) and Karlan et al. (2014). Supplier risk aversion could derive from, among other possibilities, owners' preferences over profits mirroring individuals' preferences over consumption as small supplier profits are likely the sole income source for owners. Alternatively, firms may need to make at least some payments every period (*e.g.,* payroll) and financial market frictions make it challenging or expensive to borrow or save to pay such expenses rather than using contemporaneous cash flows. I assume the supplier cost function is increasing and has increasing marginal costs (*i.e.,* is convex). This assumption reflects that marginal costs can increase due to machine inputs being fixed in the short-run; therefore, increasing production requires

more congestion at machines. Additionally, under Indian labor laws, labor costs could increase dramatically as quantity increases–the Factories Act of 1948 requires employers to pay double wages for overtime work.[20] Using data from the integrated supplier, I also empirically validate cost function convexity (see Figure 1.10).

I now describe the outside option for both the buyer and the supplier, starting with unit prices and costs and then describing quantities in the model. Before vertical integration, the outside option for both the buyer and the supplier is to transact with the market. Even though the buyer is a large firm, the buyer (and suppliers) are small relative to Indian fabric markets (let alone global fabric markets).[21] Therefore, I model the market as perfectly elastic with a constant price p^M, which includes hold-up, faced by both the buyer and the supplier.[22]

After integration, the outside option for the buyer changes to use the same cost function as the supplier, reflecting that the buyer and its relational contract suppliers presumably use similar technology–anecdotally, these firms are all users of (near-)frontier technology rather than firms with large research and development departments trying to innovate and compete by advancing fabric production technology. Allowing for some cost heterogeneity between the buyer and supplier does not change results meaningfully (see Appendix B.1.2); intuitively, the lack of change reflects that shifting costs impacts both the inside and outside options in a fairly similar way such that the overall change is minimal.[23] This change in the outside option assumes that the buyer treats their integrated supplier and their relational contract suppliers as alternatives (*i.e.,* substitutes). Morton (2023a) shows strong evidence of substitutability in this setting, leveraging exogenous variation at the fabric level. It follows that the integrated supplier is likely a credible threat to displace external relational contract suppliers, which corresponds with the investment in integrated capacity as a meaningful and

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

[21]Recall the buyer uses fewer than 300 million square meters of fabric per year, which is less than 0.5% of the 71 billion square meters of fabric produced in India in 2019, and India is the sixth largest producer by country globally (India Brand Equity Foundation).

[22]Note that incorporating market power for the buyer when purchasing inputs in the spot market should not affect the threat point effect in the context of the bargain between the buyer and a single relational contract supplier. The buyer cannot both threaten to use the integrated supplier to displace the relational contract supplier and use it to decrease the quantity purchased in the spot market simultaneously; therefore, changing the outside option does not affect market prices as it does not change market quantities. Overall, aggregating across suppliers, if the buyer has market power in the spot market, then prices in the spot market would decrease as the buyer replaces market supply with integrated supply, but it would not impact the threat point effect directly as only market prices would change. That said, incorporating indirect effects through the outside option changing in the Nash bargain would lead to underestimates of the threat point effect, as the outside option improves for the buyer and worsens for the supplier as lower overall demand reduces market prices.

[23]Furthermore, large heterogeneity in costs functions seems implausible–if the buyer were much more efficient, it should completely integrated and if the supplier were much more efficient the buyer should never integrate.

visible sunk cost. Appendix B.2 provides additional details on the contracting environment and supplier beliefs about capacity and demand that are relevant for this assumption and Appendix B.3 discusses what structure needs to be imposed on the multilateral bargaining processes for multiltateral negotiations to result in the same equilibrium as the individual case.

Quantities in the model reflect the stochastic nature of demand in the industry, with quantities determined by the end client before bargaining for fabric prices rather than by the buyer or supplier as part of the relational contract. As the relational contract is designed to provide demand assurance rather than large demand volume, I model market quantities as a mean-preserving spread of demand quantities. I set mean quantities (for both the relational contract and market) to 100% capacity (and ensure that, given market prices and productive technology, the supplier's optimal capacity utilization is 100%). This assumption does not imply that the supplier only works with the buyer. Instead, it could be interpreted as the capacity utilization relative to the capacity saved by the supplier for the buyer, such as effectively reserving floor space or time for the buyer. In this sense, how important the relational contract is to the supplier relative to their contracts with other buyers is incorporated through the risk aversion parameter (*i.e.,* if the supplier has other income streams, then they might be less risk averse over production for this one buyer in the model.)[24] Formally, the outside options are:

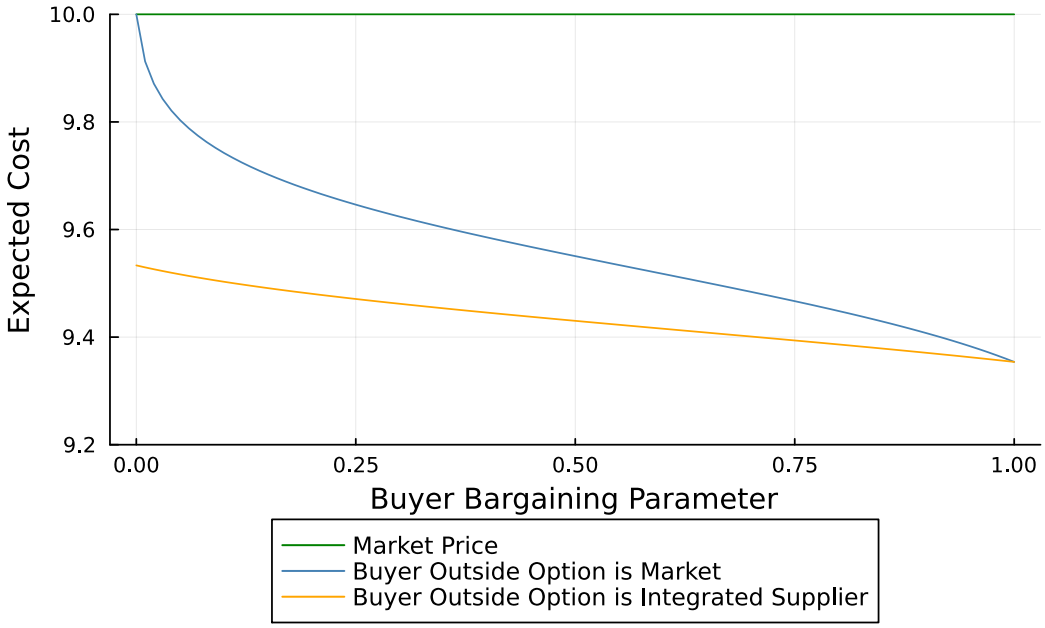
$$\begin{aligned}
\text{Buyer before integration: } & O_B = -p^M \mathbb{E}_{q^{RC}}[q^{RC}] \\
\text{Buyer after integration: } & O_B = -\mathbb{E}_{q^{RC}}[C(q^{RC})] \\
\text{Supplier before and after integration: } & O_S = \mathbb{E}_{q^M}[U(p^M q^M - C(q^M))]
\end{aligned}$$

For the exact parameterization of the model and a discussion of how equilibrium relational contract prices are computed, please see Appendix B.4; parameters and functional forms are chosen to illustrate the economics of the threat point effect rather than calibrated to match data or parameter estimates from other studies.

[24] In theory, the supplier working with other buyers could also affect the cost convexity for the supplier. For example, if the supplier has a diversified buyer pool, then the supplier might be able to balance out demand shocks from different buyers; although, seasonality in industry demand reduces any individual suppliers' ability to smooth demand by contracting with many buyers (*i.e.*, demand shocks across buyers are correlated). Recall, however, that changes in costs do not affect results for the magnitude of the threat point effect, as shown in Appendix B.1.2.

1.3.2 Threat Point Effect and Buyer Bargaining Parameter

Figure 1.8: Relational Contract Prices and the Buyer Bargaining Parameter



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

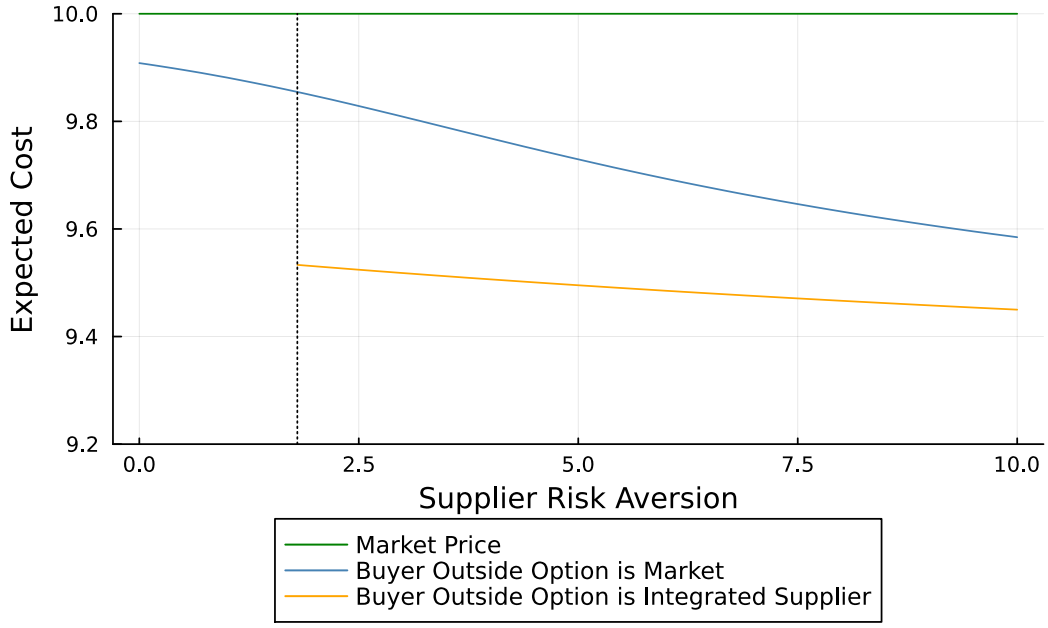
Figure 1.8 shows how the buyer's expected fabric procurement cost in the relational contract varies with respect to both the buyer bargaining parameter and the buyer's outside option, holding supplier risk aversion constant. When the buyer bargaining parameter approaches zero and the buyer's outside option is the market, the expected cost to the buyer approaches the market price. The economic intuition for this result is that the buyer's outside option determines the relational contract price when the buyer bargaining parameter is low–the supplier extracts all surplus so relational contract prices increase until expected buyer surplus is zero. As the buyer bargaining parameter increases, the buyer's share of the surplus increases through lower relational contract prices, decreasing expected cost. When the buyer's outside option is the integrated supplier, as the buyer's threat point improves (*i.e.,* is lower cost), relational contract prices decrease. As the buyer bargaining parameter increases, the supplier's outside option plays an increasingly important role in determining relational contract prices. Therefore, when the buyer bargaining parameter is one, the supplier's outside option determines relational contract prices and the buyer's outside option

is effectively irrelevant, resulting in the same relational contract prices (and expected cost) regardless of buyer outside option.

The threat point effect is visible in the graph as the difference between the relational contract price when the outside option is the market (the blue line) and the relational contract price when the outside option is the integrated supplier (the orange line). The threat point effect can be quantitatively important (over 5% of market prices) when the buyer's outside option influences relational contract prices. However, when the buyer has a high bargaining parameter, changing the buyer's outside option has a negligible effect on relational contract prices since prices are effectively determined by the supplier's outside option (*i.e.,* the supplier's individual rationality/participation constraint), which does not change when the buyer adds integrated capacity. Equivalently, the threat point effect is small when the buyer bargaining parameter is large because the buyer already extracts all the surplus pre-integration–there is no scope for the buyer to extract more surplus such that improvements in the buyer's bargaining position do not affect relational contract outcomes.

1.3.3 Threat Point Effect and Supplier Risk Aversion

Figure 1.9: Relational Contract Prices and Supplier Risk Aversion



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

Figure 1.9 illustrates how the expected cost to the buyer varies with respect to both supplier risk aversion and the buyer's outside option, holding the buyer bargaining parameter constant. First, the expected cost in the relational contract when the buyer's outside option is the market and supplier risk aversion is zero highlights that risk aversion is not the only feature of the model that makes demand assurance valuable for suppliers–the combination of cost convexity and lower quantity variance in the relational contract compared to the market also creates surplus. As supplier risk aversion increases, the value of the outside option, with high quantity, cost, and profit variance, decreases relative to the inside option, creating higher surplus in the relational contract (holding prices fixed). As the buyer bargaining parameter is not zero in Figure 1.9, some of this surplus is shared with the buyer through lower relational contract prices. When the buyer outside option is the integrated supplier, surplus in the relational contract decreases dramatically because the buyer's outside option improves.

Given the lower quantity variance in the relational contract and cost convexity, there is a cutoff level of risk-aversion below which the relational contract does not create any surplus and breaks down, recalling that the supplier does have some surplus when selling to the market as market prices are greater than costs due to high hold-up prices. In this scenario, assuming the buyer has the ability to threaten the outside option of its choice after vertical integration (either market or integrated supplier),[25] the buyer would prefer to convince the supplier that the relevant outside option is the market to keep the relational contract. In the empirical context studied, for example, the buyer could argue that the integrated supplier is irrelevant for bargaining by clarifying that the integrated supplier displaces volume from other suppliers or even claiming that the integrated supplier produces different fabrics. If the supplier is sufficiently risk averse for a relational contract to exist when the buyer outside option is the integrated supplier, then the buyer prefers to use the integrated supplier as the outside option to reduce prices.

As in Figure 1.8, the threat point effect is the difference between the relational contract price when the outside option is the market (the blue line) and the relational contract price when the outside option is the integrated supplier (the orange line). First, there is no threat point effect for low levels of supplier risk aversion, as the buyer strategically keeps the market as the outside option and does not try to leverage an improved bargaining position due to vertical integration because doing so would break, rather than improve, the relational contract. Second, when supplier risk aversion is above the cutoff for the relational contract to

[25]A model where the buyer selects which outside option to use in bargaining, even if buyer can only threaten to use the integrated supplier as an outside option if there is capacity at the integrated supplier, would not change this result. Appendices B.2 and B.3 provides additional details on assumptions regarding outside options and multilateral bargaining underpinning the proposed bilateral negotation equilibrium.

exist with the integrated buyer as the outside option, the threat point effect decreases slightly as risk aversion increases. This result reflects that total surplus differentially increases more in risk aversion when the buyer's threat point is the market due to the algebra of the Nash product.[26] As some of this surplus accrues to the buyer in the form of lower prices, relational contract prices decrease more quickly with respect to risk aversion when the buyer threat point is the market.

1.4 Structural Model

1.4.1 Why a Structural Model

I estimate the conceptual model structurally, recovering the price schedule that solves the Nash bargaining problem specified in equation 1.1. This structural modelling approach is useful for two key reasons: *i)* quantifying the threat point effect for heterogeneous buyer and supplier types, reflecting the importance of heterogeneity in the conceptual model, and *ii)* evaluating policies to reallocate surplus to small constrained suppliers.

Although reduced-form methods can measure the threat point effect in a case study (and I later validate my structural estimates using a difference-in-difference approach that directly estimates the threat point effect in this setting), they do not facilitate quantifying the threat point effect for heterogeneous buyer and supplier types. Specifically, conducting a rich analysis that incorporates heterogeneity by relationship type with only reduced-form methods similar to my approach requires data and a setting with not only enough relationships to estimate the effect for relationships of that type but also a variety of different relationship types (and knowledge of what type each relationship is). Finding data and a setting for such analysis would be challenging, if not practically impossible–it would require finding case studies where vertical integration occurred (and without an increase in downstream demand for the buyer), with buyers sufficiently large that it purchases externally from both market and relational contract suppliers after integration, with the desired heterogeneity in buyer bargaining parameters and supplier risk aversion, and with meaningful heterogeneity in supplier exposure to integration. As large buyers presumably intentionally select their relational contract suppliers based on specific characteristics, their relational contract suppliers likely resemble each other, making it largely implausible to find sufficient supplier heterogeneity from any reduced-form analysis.

[26]Specifically, because total surplus is a *product* rather than a *sum*, total surplus is much higher when the buyer's threat point is the market. Formally, a marginal increase in supplier surplus increases total surplus more when the level of buyer surplus is higher.

My structural approach enables quantifying the threat point effect for out-of-sample relationships and conducting counterfactual policy analysis through, although requires imposing additional assumptions. As other settings demonstrate that even small firms, such as smallholder Rwandan farmers in Macchiavello and Morjaria (2020) and small fishing enterprises in Sierra Leone (Ghani and Reed, 2022), use relational contracts, it is important to consider the quantitative importance of the threat point effect for these presumably out-of-sample firms. Specifically, as buyer and supplier characteristics determine the threat point effect, threat point effects for these firms could be very different than what is observed for in-sample fabric suppliers.

Not only does the structural approach facilitate quantifying the threat point effect for small constrained suppliers, it also provides a framework for evaluating policies to reallocate surplus to them. Specifically, I consider two policies to improve the relational contract for small firms: *i)* increasing downstream buyer competition; and *ii)* creating the missing market that generates demand for relational contracts from the small suppliers by allowing them to purchase insurance. I validate the model three ways to ensure that model captures key patterns in the data and produces reasonable price estimates.

I begin by discussing estimation and identification of the model before describing the parameter estimates and model validation against untargeted moments and out-of-sample fit. I then quantify the threat point effect for heterogeneous buyer and supplier (*i.e.*, relationship) types. Last, I validate the model using a difference-in-differences approach that compares the model-implied threat point effect, which depends upon parameter estimates for in-sample suppliers, with difference-in-difference estimates of the threat point effect. Additionally, I show that event study coefficients using model-implied prices match untargeted event study coefficients using actual prices.

1.4.2 Estimation Approach

1.4.2.1 Buyer and Supplier Primitives

I use the structural model to estimate the buyer bargaining parameter and supplier risk aversion for each supplier relationship governed by the relational contract. Additionally, I estimate a tuning parameter to penalize deviations from risk neutrality, motivated both by economic intuition that firms should not have extreme deviations from risk neutrality and econometric concerns of overfitting given that there are relatively few observations per supplier. Overall, my approach reflects that quantifying the threat point effect for heterogeneous types is largely an out-of-sample exercise; therefore, I develop an estimation strategy that naturally allows for testing out-of-sample fit. Specifically, I primarily estimate parameters

using pre-integration data and test out-of-sample fit using post-integration data. Although this approach is admittedly out-of-sample in a different sense, namely quarters for in-sample suppliers rather than for out-of-sample relationship types, it provides at least supportive evidence that the model performs reasonably well. I first describe the "inner" optimization used to estimate the buyer bargaining parameter and supplier risk aversion parameter given the tuning parameter and then describe the "outer" optimization to estimate the tuning parameter given the previously estimated buyer bargaining parameter and supplier risk aversion.

The "inner" optimization has three steps. First, I derive the model-implied relational contract pricing function $p^{RC}(q^{RC}; \mathbf{\Omega}_s)$ for each relationship by finding the function that optimizes the Nash product as in (1.1), where $\mathbf{\Omega_s}$ represents the parameter vector for each relationship: the buyer bargaining parameter and supplier risk aversion parameter. The relational contract pricing function depends upon the buyer bargaining parameter and the supplier risk aversion parameter through the Nash product, as shifting these parameters shifts the value of the Nash product and, therefore, the optimal relational contract prices.[27] Second, I calculate the price estimates for the quantities observed in the relational contract for each supplier implied by the model-derived relational contracting pricing function. Last, I compute the objective function per supplier, which is the root mean squared error of the price estimates from the relational contracting pricing function using *only the pre-integration data*.

I penalize extreme estimates of the risk aversion parameter by adding the absolute value of the risk aversion parameter, weighted by the tuning parameter, to the objective function. As I use CARA utility, a risk aversion parameter of zero is risk neutral, meaning that deviations from risk neutrality in either direction are equally penalized. This approach is similar to the set-up of a LASSO regression by constraining extreme parameter values. Incorporating the tuning parameter to push parameter estimates towards risk neutrality reflects both economic motives, as firms are unlikely to have extreme deviations from risk neutrality, as well as econometric concerns around overfitting given the relatively few observations per supplier.

[27] I use a second-order Taylor approximation, meaning $p^{RC}(q^{RC}; \mathbf{\Omega}_s) = \beta_{0,s} + \beta_{1,s} q^{RC} + \beta_{2,s}(q^{RC})^2$ to find the $\beta_{0,s}, \beta_{1,s}, \beta_{2,s}$ that maximize the Nash product. As suppliers typically operate near full capacity, I estimate the second-order Taylor approximation around a capacity of one. Identification of $\beta_{0,s}, \beta_{1,s}, \beta_{2,s}$ follows the logic of identification of buyer bargaining parameter and supplier risk aversion in 1.4.3. This logic is clearest for $\beta_{0,s}$, which shifts the level of relational contract prices, and $\beta_{2,s}$ which shapes the convexity of relational contract prices. Including $\beta_{1,s}$ incorporates additional flexibility for the relational contract to respond to increasing marginal costs. For example, average cost per unit is larger for capacity utilization of 1.2 than for capacity utilization of .8, but the relational contract price would be the same without including $\beta_{1,s}$. It follows that supplier risk aversion, along with convexity in costs, together lead to larger $\beta_{1,s}, \beta_{2,s}$ as these parameters help smooth profit, with $\beta_{1,s}$ and $\beta_{2,s}$ makes the relational contract stabilize supplier profits across different capacity realizations.

Formally, the "inner" optimization finds the buyer bargaining parameter and supplier risk aversion parameter, given the tuning parameter, to minimize the sum of the per supplier objective functions:

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

In this optimization problem, S represents the set of suppliers, $\{q_{s,t}^{RC}\}$ represents the set of observed quantities for relational contract supplier s across pre-integration periods,[28] T is the number of periods, $p_s^{RC}(q_{s,t}^{RC}; \mathbf{\Omega}_s)$ is the price estimated by the relational contract pricing equilibrium for supplier s for quantity $q_{s,t}^{RC}$, $p_s^{RC}(q_{s,t}^{RC})$ is that actual relational contract price for supplier s for quantity $q_{s,t}^{RC}$, λ is the tuning parameter, and θ_s is the supplier risk aversion parameter.

The "outer" optimization finds the tuning parameter than maximizes fit in post-integration quarters. Specifically, the optimization routine finds the tuning parameter that leads to the relational contract pricing function that best fits the post-integration data, recalling that the relational contract pricing function is determined using only the pre-integration data. The tuning parameter impacts post-integration fit by influencing the estimated relational contract pricing function: a larger tuning parameter pushes estimated supplier risk aversion parameters towards risk neutrality. Therefore, the "outer" optimization solves:

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

[28] While it would be appealing to use monthly rather than quarterly data for this analysis to increase the sample size, the relational contracts appear to operate at the quarterly level rather than the monthly level. Not only are the reduced form patterns consistent with the relational contract clearest at the quarterly level, but anecdotally staff procuring fabric at the buyer emphasize that meetings with key suppliers about capacity projects are once per quarter (and not once per month).

1.4.2.2 Outside Estimation

To estimate the model, it is necessary to find additional moments from the data. Table 1.1 describes the additional moments used:

Table 1.1: Additional Moments for Estimating the Structural Model

Parameter(s) in Model	Moment(s) to Match	Data
Cost Function	Actual Costs and Capacity	Cost and Capacity Data from Integrated Supplier
Mean Capacity	Average Log Quantity	Transaction Data
$\frac{\mathbb{V}[q^M]}{\mathbb{V}[q^{RC}]}$	Variance in Quantities for non-RC Suppliers Compared to Variance in Quantities for Relational Contract Suppliers	Transaction Data
p^M	Ratio of Prices from Market Suppliers to Estimates from Integrated Supplier	Transaction Data

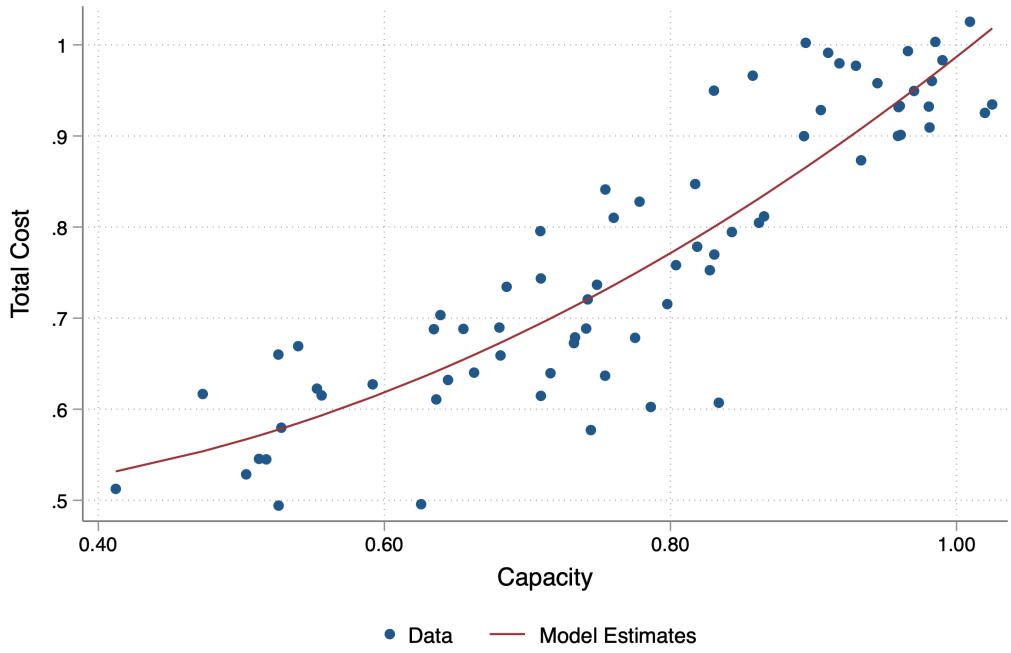
I use data on capacity and costs at the integrated supplier to estimate the convexity of the cost function. As, anecdotally, relational contract suppliers use similar technology to the integrated suppliers, their cost functions should be similar to the cost function at the integrated supplier (nevertheless, I include a robustness check that incorporates supplier cost heterogeneity). Therefore, I estimate the total cost function for suppliers as $c(q) = \phi_0 + \phi_1 q + \phi_2 q^2$. I include q^2 given the importance of convexity, although noting that I impose no structure to force the model to yield $\phi_2 > 0$, which would be consistent with convexity. Estimates are equivalent to a second-order Taylor approximation of the true cost function.

The data used are the cost data from the integrated supplier, which are a panel at the division (*i.e.*, fabric type, such as knits and striped fabrics) and month level from April 2019 to March 2020.[29] Additionally, note that, especially at the month level, it is reasonable to treat these estimates as recovering the true causal effect of capacity on costs because quantities are, at least locally, exogenous. Quantity exogeneity derives from the institutional features of how quantities are determined in garment manufacturing. Recall that the fabric supplier is locked in at the time of making the sample garment, which is *before* the end quantity for the order is determined. It follows that the exact quantity used at the integrated supplier is

[29]For convexity, it is important to use the most disaggregated data temporally rather than combine. To fix ideas, consider a supplier that operates for two months at 100% capacity and another supplier than operates first at 80% capacity and another month at 120%. For both suppliers, the average capacity is 100%. Therefore, a regression of costs on capacity at an aggregated level would show different costs for the same capacity level that do not correspond to actual capacity differences apparent in the disaggregated data.

determined by stochastic downstream demand rather than choices by the buyer.[30] Therefore, as quantities are (at least locally) exogenous, OLS regressions yield causal estimates of the effect of capacity utilization on costs.

Figure 1.10: Cost Convexity at the Integrated Supplier



Note: Data from integrated suppliers at division by month level. Costs used are the volume-weighted average cost, expressed as a percent of average volume-weighted cost for the divsion. Capacity is reported directly in the data.

The estimates show that the model fits the data well. The R^2 is .795 and the RMSE is .07, which is small relative to the variance in costs overall in the data. Additionally, the estimate finds strong convexity, with the coefficient on q^2 at .7815 and a p-value of .02. I transform costs so units are in percentage terms relative to the mean costs for the fabric group, as costs are reported by fabric group, facilitating analysis combining data from all fabrics. The strong fit of the model suggests that heterogeneity by fabric type is not likely a central feature of the cost function.

As the cost function is measured in capacity units, transaction quantities need to be converted into capacity units to estimate the structural model. Two assumptions underlie

[30]Based on previous transactions with end clients, the buyer presumably has some sense of how common certain volumes are from various end clients. The buyer likely uses this knowledge when sampling to try to ensure that capacity utilization at the integrated supplier is not too extreme. However, even with this knowledge, for short time periods, it is impossible to achieve exactly targeted capacity usage–consistent with the heterogeneity in capacity usage estimated.

my conversion: first, I assume that log quantities are relevant for capacity usage. This assumption reflects that orders with larger quantities typically also have longer time between order date and delivery, such that capacity usage does not increase linearly with quantity (see Appendix Table B.1). Therefore, increases in quantity result in smaller increases in capacity usage. Expressing quantities in logs captures this pattern in a reduced-form manner without requiring estimating any additional parameters that govern how quantities are spread over time to map to capacity. Additionally, the log specification is also reasonable for revenue, as larger orders are associated with lower prices (see Appendix Table B.2), which may reflect lower costs from learning-by-doing for the specific fabric on the supplier side, consistent with evidence in the industry in Adhvaryu et al. (2023).

The second assumption is that the supplier's total capacity for the relational contract is the mean capacity in the pre-period. This assumption reflects that the relational contract is conceptualized as providing consistent capacity usage around an agreed upon capacity. To estimate the pricing function in the relational contract, I assume that relational contract quantities are drawn from the gamma distribution based on the method of moments estimator for the observed relational contract quantities. I chose a gamma distribution because it is a simple two parameter distribution with positive support and fits the data well.[31] Furthermore, using a gamma distribution nests many possible distributions, including other commonly used distributions (*e.g.,* chi-squared and exponential distributions are special cases of a gamma distribution). Overall, this approach incorporates scale economies, as suppliers with larger capacity experience the same deviation in actual quantity as a smaller change in capacity. It follows that large suppliers who receive many orders can likely reduce their capacity variance as they receive more stochastic end client demand shocks that, on average, tend to offset each other over many orders.

Additionally, it is necessary to estimate the counterfactual capacity variance that relational contract suppliers would experience if they were not in the relational contract. As the relational contract is designed to reduce capacity variance, I assume that relational contract suppliers' counterfactual market capacity utilization distribution would have the same mean as the relational contract but higher variance, as in the conceptual model. Because I do not have data from suppliers about total capacity utilization by buyer type (*i.e.,* relational contract and market), I assume the ratio of capacity variances by buyer type matches the empirical ratio of average within-supplier capacity variance for non-relational contract

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

suppliers relative to relational contract suppliers.

Using all non-relational contract suppliers would likely lead to overestimates of the variance that relational contract suppliers would face outside the relational contract given that non-relational contract suppliers work with other buyers and may even be in relational contracts with other buyers. Therefore, I restrict the set of non-relational contract suppliers for this calculation to the placebo suppliers.[32] Focusing on placebo suppliers to define the relevant set reflects that placebo suppliers are arguably the closest counterfactual for relational contract suppliers.[33]The estimates support this interpretation, as the counterfactual variance when working with the market is estimated to be 8.8 times larger due to demand assurance in the relational contract reducing average standard deviation of within-supplier capacity utilization to only .047. For comparison, using all market suppliers to estimate the counterfactual variance would result in an unreasonably large estimate of the variance that the relational contract suppliers would face when selling to the market–971 times the variance that relational contract suppliers experience in the relational contract.

Last, I estimate the market price relative to costs using transaction data. Specifically, for each transaction, I compute the margin charged by the supplier as the ratio of the transaction price over the estimated production costs for the desired quantity for a mill at 100% capacity utilization, where estimates use cost data from the integrated supplier.[34] The market price, with incorporates hold-up, has an estimated margin of 7 percentage points over costs, which seems intuitively reasonable: it is a large enough to create incentives for strategic sourcing through vertical integration and relational contracts but not so large as to be implausible.

1.4.3 Separate Identification of Buyer Bargaining Parameter and Supplier Risk Aversion

It is not immediately apparent that separate identification of the buyer bargaining parameter and supplier risk aversion is possible as increasing either buyer bargaining parameter alone or increasing supplier risk aversion alone both reduce relational contract prices. With a high buyer bargaining parameter, the buyer receives a larger share of the surplus generated by the relational contract through lower prices. Similarly, when supplier risk aversion increases, holding the buyer bargaining parameter fixed, relational contract surplus increases. As some of this surplus is shared with the buyer (as long as the buyer bargaining parameter is not

[32]As $log(0)$ is undefined, I interpret periods with no orders as having zero capacity utilization.

[33]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 1.4.5 to mitigate concerns that errors in outside moments lead to poor model performance.

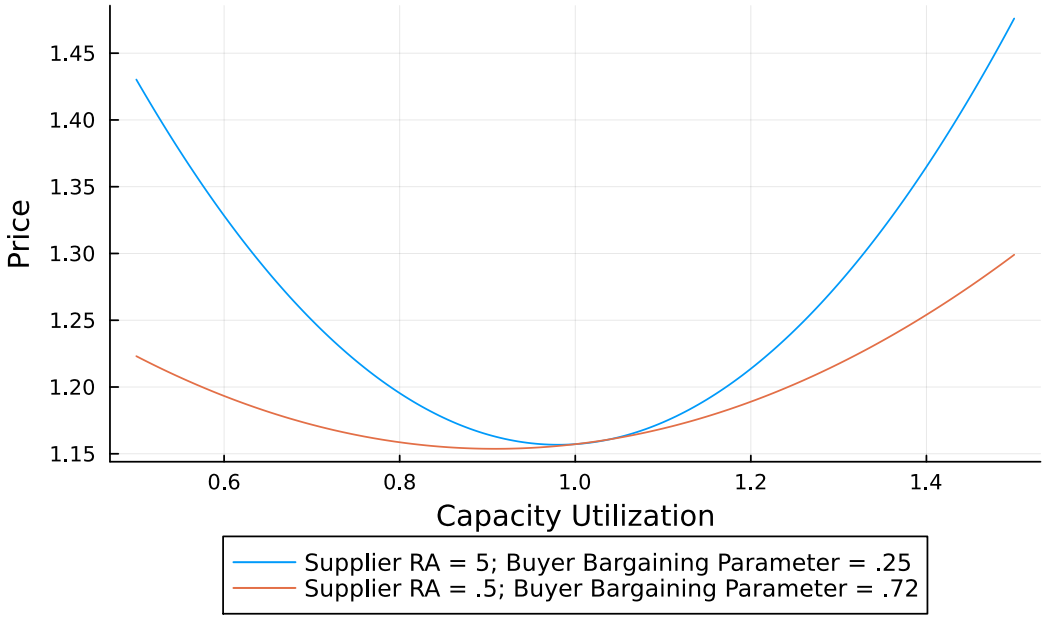
[34]Using the internal cost estimates means that market prices are constructed using the fabrics that are produced internally as the buyer does not estimate costs for fabrics not produced internally.

zero), higher risk aversion also decreases prices.

My identification strategy leverages two key sets of variation in the data. First, there is variation in how prices change with respect to capacity. The covariance of prices and capacity helps identify supplier risk aversion, as shifts to supplier risk aversion alter the *convexity* of the prices in the relational contract. More convex contracts are associated with more risk averse suppliers because convex prices equalize profit across fluctuations in capacity. This logic reflects that profit is highest when capacity is near one. Specifically, when capacity is low, even though cost per unit is low, small quantities require high prices to equalize profits. When capacity is high, cost per unit is high due to increasing marginal costs; therefore, prices need to be high to equalize profits. Second, there is variation in the level of relational contract supplier prices relative to market prices. The variation in the level of relational contract prices is key for identifying the buyer bargaining parameter, as shifting this parameter reduces prices for all capacity realizations.

If only one price capacity tuple is observed, it is impossible to separately identify buyer bargaining parameter and risk aversion. Specifically, for any price capacity tuple, given a buyer bargaining parameter, shifting risk aversion changes prices (with higher risk aversion leading to a lower price and vice-versa). Analogously, given any risk aversion, changing the buyer bargaining parameter can shift the price for a given quantity to reach the same price. As shown in Figure 1.11, both a buyer bargaining parameter of .25 with a supplier risk aversion of 5 and a buyer bargaining parameter of .72 with a supplier risk aversion of .5 result in the same price when capacity is one.

Figure 1.11: Relational Contract Prices for Different Parameters



However, once a *second* price capacity tuple is observed for a different capacity, it becomes possible to separately identify buyer bargaining parameter and supplier risk aversion. Figure 1.11 highlights how lower supplier risk aversion results in a less convex contract. Therefore, as the second price capacity tuple pins down the convexity of the contract, both parameters are separately identified.

To further visualize this identification strategy, figures 1.12 and 1.13 show the effects of individually shifting supplier risk aversion and buyer bargaining parameter, respectively. They highlight that shifting the risk aversion parameter mostly shifts convexity while shifting the buyer bargaining parameter mostly shifts levels.

Figure 1.12: Shift Supplier Risk Aversion

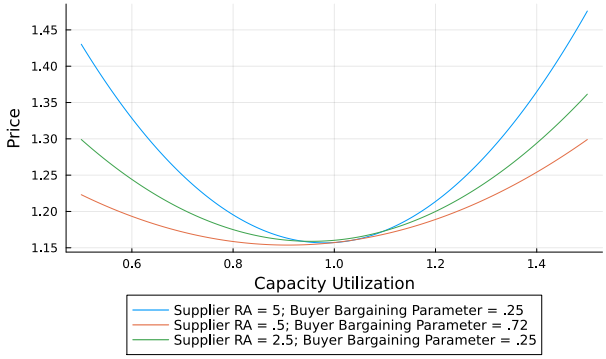
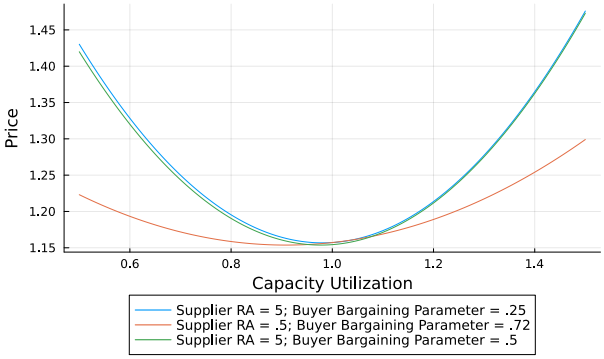


Figure 1.13: Shift Buyer Bargaining Parameter



To bring this identification argument to the data, I leverage supplier-specific capacity shocks that reflect stochastic end client demand. Specifically, although the relational contract has a target capacity, institutional features of the garment production process–specifically, that the buyer knows neither which samples will be ordered by the client nor what quantities would be conditional on an order being placed–mean that the relational contract reduces capacity variance but is unable to hit the capacity target every period. In Appendix B.5.1, I discuss this variation in more detail and provide additional evidence that these capacity shocks serve as exogenous variation.

1.4.4 Buyer Bargaining Parameter and Supplier Risk Aversion Estimates

Table 1.2: Parameter Estimates from Structural Model

	Buyer Bargaining Parameter	Supplier Risk Aversion
Minimum	.0094	0
25th Percentile	.0301	.0026
Median	.6214	.0085
75th Percentile	.9155	.0223
Max	.9985	14.6811
Mean	.5113	.8282
Standard Deviation	.4286	3.4573
Correlation Between Supplier Risk Aversion & Buyer Bargaining Parameter	.2847	

Note: Estimates from model.

Parameter estimates in Table 1.2 (confidence intervals in Appendix Table B.7) illustrate that the buyer typically receives slightly more than half of the relational contract surplus,[35] It follows that the selected relational contract suppliers are also likely reasonably large firms, consistent with the low estimates of supplier risk aversion (with the exception of one outlier supplier with high risk aversion and a high buyer bargaining parameter). Together, these results suggest that relational contract suppliers are at least fairly large sophisticated firms, consistent with positive selection into relational contracting. This positive selection matches the intuition that creating and maintaining relational contracts may require organizational capabilities within the firm, at least for more complex relational contracts as in this setting (Macchiavello and Morjaria, 2023). The estimate of the tuning parameter λ is small at .002, suggesting that a large tuning parameter is not the reason why estimated risk aversion is low.

I compare estimated risk aversion with estimates from other studies in Appendix Table B.5.[36] Estimated risk aversion for the relational contract fabric suppliers is close to that of exporting coffee mills studied in Blouin and Macchiavello (2019) (.0085 for fabric suppliers and .0068 for coffee mills), who describe the mills as "large firms by developing country standards" which average over $3.5 million a year in sales and about $2 million in total assets. As these firms are likely the most similar in size, sophistication, and setting both with respect to market institutions (as firms in low- and middle-income countries) and stakes of risk (contracts with buyers) to the fabric suppliers studied in this paper, it is promising that risk aversion estimates are quite similar. Table B.5 also highlights that risk aversion estimates in the literature vary widely based on the methodology, the setting, the economic agent, and stakes of the risky activity. Specifically, as the stakes increase, estimated risk aversion increases widely, reaching all the way up to 12.90.

1.4.5 Model Validation and Fit

Given the importance of quantifying the threat point effect for heterogeneous buyer and supplier types, including out-of-sample types, I prioritize model validation. First, I illustrate that the model matches an untargeted moment. Second, I show that the model has better performance in the post-integration period than two alternative estimators, recalling that post-integration observations are largely out-of-sample as model estimates primarily use pre-integration data.[37] Validation using difference-in-difference estimates follows the discussion

[35] I explore one possible microfoundation for heterogeneity in buyer bargaining parameter in B.13.

[36] As many studies do not report CARA coefficients, I convert the reported risk aversion to a CARA coefficient to be comparable given the estimate and utility function used.

[37] Specifically, supplier risk aversion and buyer bargaining parameter estimates use only pre-integration data, conditional on the tuning parameter.

of the threat point effect magnitude as a function of buyer and supplier characteristics, as it is necessary to quantify the threat point effect for the estimated in-sample buyer and supplier types before comparing it to the difference-in-difference estimates.

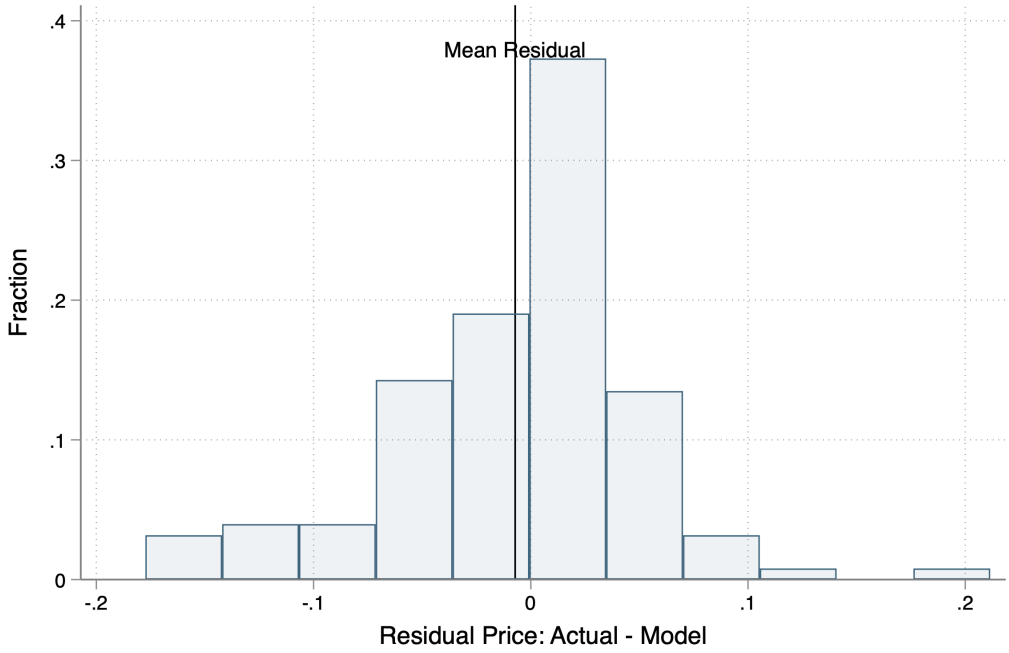
1.4.5.1 Model Validation: Untargeted Moment

As the model is non-linear and estimated parameters minimize root mean squared error of model-implied prices, estimated prices could be biased. Specifically, minimizing root mean squared error minimizes a combination of variance and bias. For example, an outlier high price can lead the model to increase all price estimates because root mean squared error penalizes large deviations heavily. In that case, the mean (and median) residual price would be biased, as the model overestimates other prices to reduce the residual for the outlier high price, and the residuals would not be symmetrically distributed around zero.[38] Therefore, I study the distribution of residuals to validate model fit. Figure 1.14 highlights that the distribution of residuals is reasonably symmetric and the mean residual is nearly 0 at $-.007$. I conclude that the model performs well for an untargeted moment: the mean residual.[39]

[38]Simulations of this exact scenario verify that indeed the model does result in biased average prices. The bias can be positive or negative depending on what quantities have outlier prices, as outlier prices also influence the convexity of the relational contract. If outlier high prices are in the middle of the quantity distribution, then the model reduces risk aversion to reduce convexity and, concordantly, reduces the buyer bargaining parameter to prevent price levels from decreasing too much. As price levels shift up on average due to the buyer bargaining parameter change, the mean bias is positive. On the other hand, when the outliers occur for extreme capacities, then risk aversion tends to increase as convexity increases, leading to lower buyer bargaining parameters. As price levels then tend to decrease, mean bias is negative.

[39]The median residual, which is also untargeted, is even closer to 0 at .002.

Figure 1.14: Untargeted Moments: Bias



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

1.4.5.2 Out-of-Sample Fit

I also analyze out-of-sample fit, comparing the distribution of supplier root mean squared error in the post period from the model with the distribution of supplier root mean squared error in the post period using two alternative estimators in Table 1.3. First, I compare model estimates to the mean price for the supplier from the pre-integration quarters. My relational contract bargaining model performs better than the mean pre-integration price broadly throughout the distribution of supplier root mean squared errors, including at the minimum, the max, and all quartiles. Specifically, mean and median supplier root mean squared error are seven percent higher using the mean pre-integration price than using the model.

I also estimate out-of-sample prices using an OLS regression of prices on capacity and a constant in the pre-period. This estimation routine has two parameters per supplier, namely the coefficient on capacity and the constant, meaning that it uses the same number of parameters per supplier as the structural model (buyer bargaining parameter and supplier risk aversion). It follows that gains relative to these OLS estimates emphasize the benefit of the structural model. Median and mean root squared error using the OLS estimates are,

respectively, 23 and 12 percent higher than using the model. Appendix Figure B.5 shows the distribution of root mean squared error across suppliers for each estimator.

Table 1.3: Out-of-Sample Fit Comparison

	Structural Model	Supplier Mean	OLS
Minimum	.0146	.0182	.0183
25th Percentile	.0482	.0551	.0683
Median	.0741	.0790	.0913
75th Percentile	.1003	.1004	.1091
Max	.1818	.1900	.1701
Mean	.0787	.0844	.0883
Standard Deviation	.0441	.0417	.0366

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

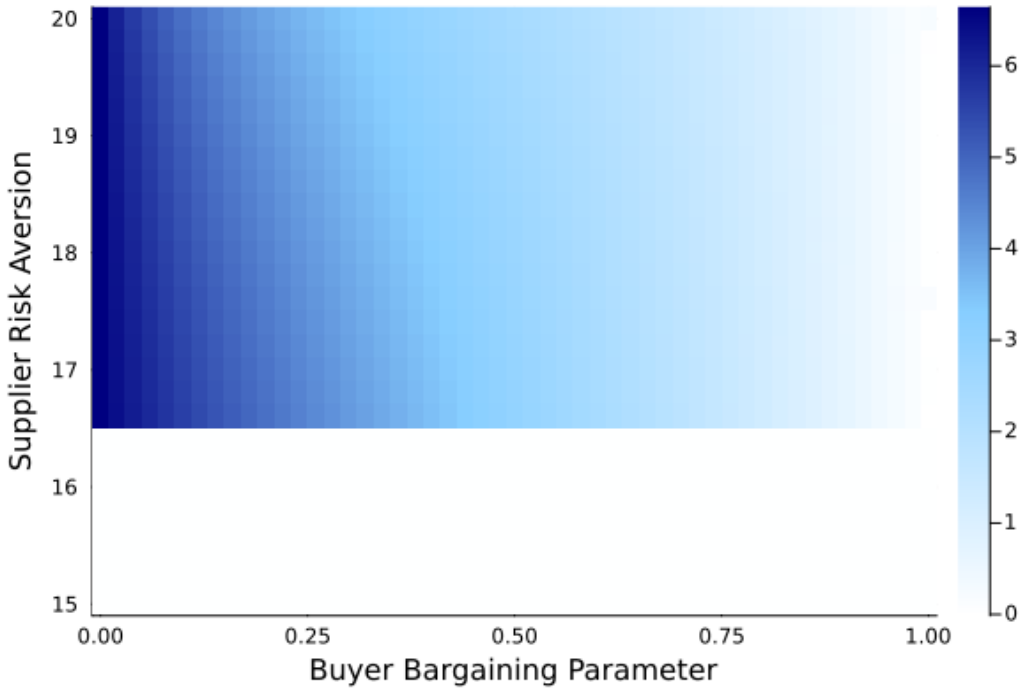
As the OLS model does not have a natural economic motivation for the tuning parameter, whereas economic intuition suggests extreme values for firm risk aversion are unlikely, OLS results do not incorporate a tuning parameter. However, as there is an econometric motive for including the tuning parameter, namely to reduce overfitting, Appendix B.9 compares model fit to OLS estimates with a tuning parameter. The fit of the OLS estimates with a tuning parameter is between the pre-integration supplier mean and standard OLS estimates without the tuning parameter.

1.4.6 Quantifying the Threat Point Effect by Buyer and Supplier Characteristics

I then study threat point effect heterogeneity with respect to buyer and supplier characteristics for different buyer and supplier types, both in-sample and out-of-sample. Figure 1.15 shows the estimated threat point effect, expressed as a percent of average pre-integration relational contract price. Estimates show that the minimum risk aversion such that the threat point effect influences relational contract prices in this setting is 16.6, which is slightly larger than the upper bound of estimates in the literature (12.9 in Gandelman and Hernandez-Murillo (2014)) but the same order of magnitude. However, this threshold level would decrease in settings with more cost convexity or where the relational contract provides more risk reduction relative to transacting in the market, as both of these factors decrease the cut-off level of risk aversion below which there is no relational contract. Relational contracts can provide more risk reduction in environments where they mitigate risk on both the demand and supply side of the market, such as for agricultural workers facing productivity risk as in Jayachandran (2006). This example is relevant as agricultural settings in low-

and middle-income countries can feature relational contracts even for smallholder farmers as in Macchiavello and Morjaria (2020); these farmers are also likely to have highly convex production costs given that a key input is mostly fixed (land) as well as higher levels of risk aversion as low-income individuals in settings without strong risk-coping institutions. Additionally, CARA risk aversion parameters also tend to increase for larger stakes. Estimates of supplier risk aversion in this study are for one quarter's worth of orders presumably for a large supplier; the same volume would presumably be a larger stake for a smaller firm with lower volumes, magnifying the potential for small firms to be more risk averse in general.

Figure 1.15: Threat Point Effect Magnitude (% of pre-integration relational price)



Note: Model estimates for the change in relational contract price as a percent of pre-integration relational contract price. Heatmap legend is to the right of the heatmap. Model estimates use the average capacity and demand assurance across in-sample relational contract suppliers.

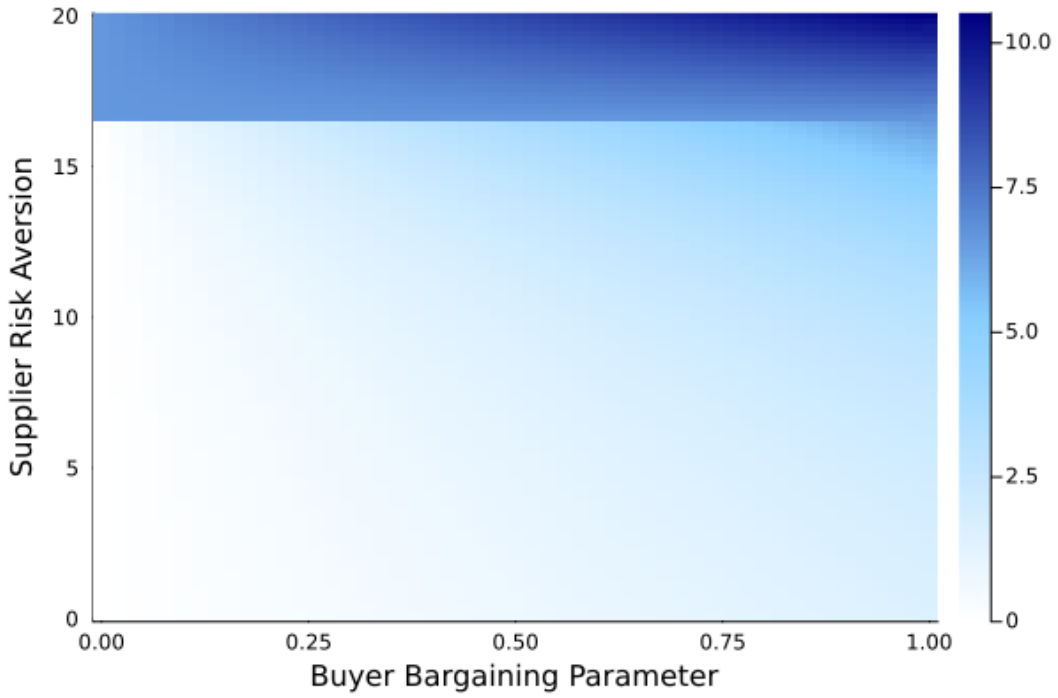
As illustrated in the conceptual model, the threat point effect is largest when the buyer bargaining parameter is small, as then the supplier's outside option, which is unaffected by vertical integration, determines prices. The threat point effect also decreases slightly as risk aversion increases due to the algebra of the Nash product.

The results highlight that the threat point effect can be quantitatively important, reducing prices by up to 6.7% of pre-integration relational contract when the market is the buyer's threat point. For comparison, 6.7% is larger than the VAT in India during the study time

period for apparel, which was only 5%.[40] As another comparison, when suppliers operate at the profit maximizing capacity, market prices are 7.36% higher than unit costs at 100% capacity, reflecting suppliers' ability to hold-up the buyer. Therefore, the threat point effect is also large relative to markups charged by suppliers holding up the buyer.

We can also analyze the threat point effect magnitude by examining the gains from relational contracts, from the buyer's perspective, relative to market prices in figure 1.16, which shows the discount relative to market prices expressed as a percent of market prices. Overall, the discount increases in both the buyer bargaining parameter, as the buyer's larger share of the surplus comes through lower prices, and in supplier risk aversion, as the surplus from the relational contract is larger for the same prices when supplier risk aversion is higher, some of which is distributed to the buyer for any positive bargaining parameter in the form of lower prices.

Figure 1.16: Buyer Gains from Relational Contracts (% of market price)



Note: Model estimates for the discount in the relational contract relative to the market as a percent of the market price. Heatmap legend is to the right of the heatmap. Model estimates use the average capacity and demand assurance across in-sample relational contract suppliers.

The heatmap emphasizes the quantitative importance of the threat point effect, with a large discontinuity visible when the threat point effect reduces relational contract prices.

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

This heatmap explicitly highlights how relational contracts, even when effective, magnify incentives to vertically integrate. For example, when surplus is split equally between the buyer and the supplier, the gain from the relational contract increases by 76% from 3.8% to 6.7% relative to market prices when the threat point effect starts to reduce prices. This difference is economically meaningful–relational contract prices for in-sample relational contract suppliers pre-integration are on average only 2.2% lower than market prices and this relatively small gain is still sufficient for the buyer to find it worthwhile to establish and maintain relational contracts.

These results highlight how that the threat point effect can be quantitatively important, creating incentives for buyers to vertically integrate even when relational contracts resolve the hold-up problem.

1.4.7 Model Validation using Difference-in-Difference Estimates

In addition to the validation and fit exercises in 1.4.5, I also validate the structural model using a difference-in-difference design. The difference-in-difference design directly estimates the threat point effect by comparing suppliers exposed to vertical integration based on the fabrics they produce with suppliers unexposed to vertical integration before and after the additional integrated capacity is added. Specifically, suppliers are considered exposed to vertical integration (*i.e.,* "treatment" suppliers) if the fabrics they produce pre-integration can also be made at the integrated supplier, such that the buyer switching to the integrated supplier is a credible threat to displace these suppliers. Alternatively, suppliers whose pre-integration fabrics do not overlap with the integrated supplier are not exposed to integration (*i.e.,* "control" suppliers). Importantly, all comparisons are within supplier type, meaning that relational contract suppliers are compared to other relational contract suppliers.

I then validate the model by showing that the direct estimate of the threat point effect using the difference-in-ddifference model coincides with the model-implied threat point effect for in-sample suppliers. In addition to validating the magnitude of the threat point effect, the difference-in-difference event study coefficients also serve as additional untargeted moments to validate the model as the structural model only targets the difference between model-implied and actual prices without incorporating any comparison across suppliers based on exposure to integration. As the difference-in-difference estimates are for both the pre-integration and post-integration quarters while the structural estimates use only pre-integration data, validating the model with these untargeted moments includes both in-sample and out-of-sample validation.

1.4.7.1 Difference-in-Difference Estimation Approach

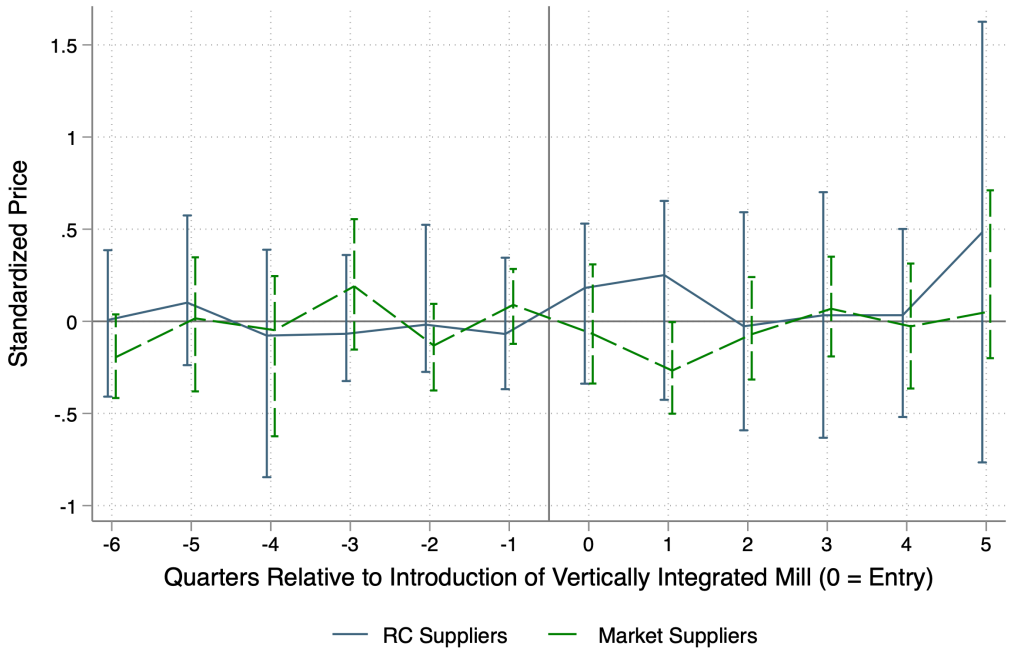
In a standard difference-in-differences design in this setting, the identifying assumption would be that counterfactual potential outcomes evolve in a parallel pattern for treatment and control suppliers. This assumption is *not* that assignment of suppliers to treatment status is random, which could be problematic as the analysis in Morton (2023a) demonstrates that fabrics were not randomly chosen to be integrated: fabrics with more relational contract suppliers are less likely to be brought in-house because vertical integration and relational contracts are substitutes. However, as suppliers are not selected into relational contracts based on individual fabrics (as discussed in Morton (2023a)), suppliers are unlikely to be intentionally selected to be exposed to integration.

Furthermore, the long history included in the data facilitate comparing trends in outcomes for many pre-treatment periods to support the assumption of parallel counterfactual outcomes. However, the data suggest that fabrics were brought in-house based on trends in fabric demand over time, creating non-random assignment in exposure to vertical integration. Therefore, I use a doubly-robust difference-in-difference design, described in detail in Appendix B.10.2. This design weakens the parallel assumptions trend to conditional parallel trends; in my empirical setting, comparisons are conditional on supplier exposure to trends in fabric demand, where trends are for fabric demand across all suppliers.

1.4.7.2 Validation: Threat Point Effect Magnitude Comparison

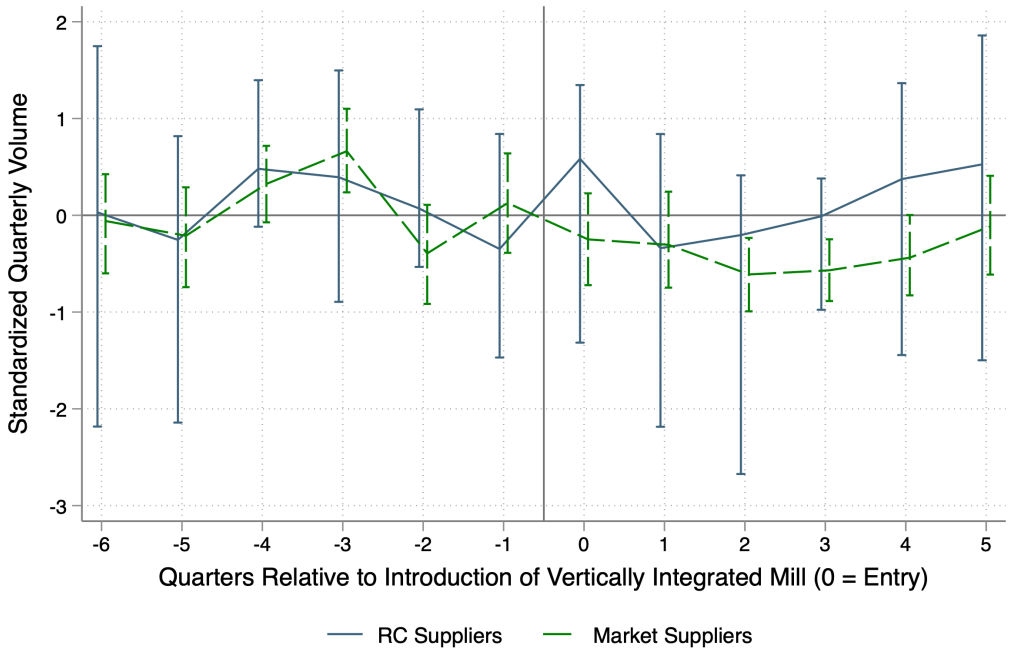
The model makes sharp predictions for the effects of vertical integration on both prices and quantities. As the largest estimated supplier risk aversion is below the threshold for the threat point effect to exist (see Appendix Figure B.7), difference-in-difference estimates should not show any meaningful change in prices for relational contract suppliers. Additionally, the model predicts that there should be no change in quantities for relational contract suppliers, with volume reduced from market suppliers rather than relational contract suppliers given that downstream demand from end clients does not increase. I find reduced-form estimates of the threat point effect consistent with the model.

Figure 1.17: Difference-in-Difference Estimates for Effect of Vertical Integration on Standardized Prices



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

Figure 1.18: Difference-in-Differences Estimates for Effect of Vertical Integration on Quarterly Volume



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

The difference-in-differences results in Figure 1.17 provide evidence that vertical integration did not affect prices for either supplier type, matching model predictions given estimated buyer bargaining parameters and supplier risk aversion parameters. Additionally, Figure 1.18 highlights that quantities (standardized by supplier) for relational contract suppliers did not decrease while quantities for non-relational contract suppliers did decrease. For both outcome variables and both supplier types, there are no obviously apparent pre-trends– the most visible pre-trend is that volumes for relational contract suppliers were *decreasing* pre-integration. Although this pre-trend is fairly weak, and there is no clear reason why it would occur, the reversal in sign after integration at least suggests that relational contract volumes did not decrease due to vertical integration.

As the few relational contract suppliers in the data limits statistical power, I also analyze the average treatment effect across all pre and post periods to increase power in Table 1.4.[41] The pooled results illustrate that non-relational contract suppliers experience a statistically significant, and economically meaningful at .4 standard deviations (on average $\sim$ 120,000

[41]Bootstrapped confidence intervals account for clustering at the supplier level, as highlighted in Bertrand et al. (2004), as well as uncertainty in the propensity score and regression outcome models.

square meters of fabric per quarter), decrease in quantities. However, relational contract suppliers do not have any statistically significant treatment effect for quantities or prices; although confidence intervals are not always small, point estimates are economically close to zero. Additionally, I show that results are robust to using an alternative measure of volume that has many desirable properties, including scale-invariant t-statistics: the quartic root (Thakral and Toh, 2023).[42]

I also verify that vertical integration does not have unanticipated effects on other outcome variables by analyzing effects on reliability, measured as the percent of on-time deliveries. I also analyze effects on transaction counts, as the decrease in market supplier quantity can reflect a combination of fewer transactions or reallocating sampling opportunities with smaller volumes to market suppliers. I do not find statistically significant results for any pooled coefficient for these outcome variables. I include the full event study plots for these outcome variables in Appendix B.11.1, which also document no treatment effect on these outcomes.

Table 1.4: Pooled Pre- and Post-Treatment Effects

	Standardized Volume	Volume (Quartic Root)	Volume-Weighted Standardized Price	Percent of Volume Delivered On-Time	Transaction Count
RC Suppliers Pre	.07	.85	−.02	−.02	−.28
	[−.28, .23]	[−.83, 1.57]	[−.12, .12]	[−.09, .13]	[−3.54, 7.45]
RC Suppliers Post	.15	.91	.16	−.06	2.21
	[−1.38, .79]	[−9.98, 6.59]	[−.42, .59]	[−.41, .17]	[−49.59, 20.66]
Market Suppliers Pre	.09	1.16	−.02	−.01	.69
	[−.05, .22]	[.01, 2.18]	[−.16, .07]	[−.24, .03]	[−.55, 1.83]
Market Suppliers Post	−.38	−4.88	−.05	.14	−.28
	[−.75, −.01]	[−7.77, −1.78]	[−.21, .21]	[−.17, .28]	[−3.48, 3.04]

Note: Data from universe of fabric transactions by buyer. Estimates represent the sum of doubly-robust difference-in-differences pre- and post-integration event study coefficients with bootstrapped 90% confidence intervals that account for the two-stage estimation in the design.

In Appendix B.11.2, I document robustness to using the synthetic difference-in-differences approach proposed in Arkhangelsky et al. (2021). This methodology reproduces the results

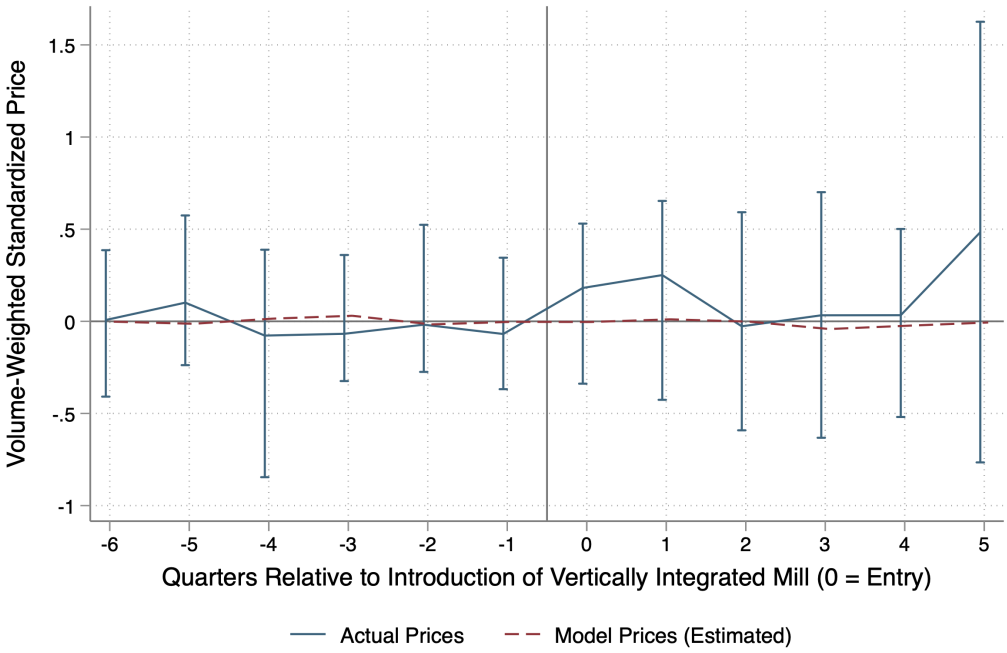
[42]The positive and significant effect for market suppliers pre-integration for the quartic root is consistent with a small anticipation effect. This anticipation effect corresponds with the integrated supplier starting to produce fabrics a the end of the last pre-integration quarter before it started 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.

that quantities decrease only for market suppliers and prices do not change for either relational contract or market suppliers.

1.4.7.3 Validation: Difference-in-Difference as Untargeted Moments

As the model does not use the difference-in-differences moments,[43] specifically the estimation does not leverage any variation in supplier exposure to integration based on fabrics produced, I compare difference-in-differences event study coefficients using model prices with untargeted analogous estimates using actual prices. The model predicted prices in Figure 1.19 are all within the 95% confidence interval.

Figure 1.19: Untargeted Moments: Difference-in-Differences



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

[43]Recall that these estimates are based on exposure to vertical integration based on fabrics before and after integration.

1.5 Policy Implications: Reallocating Surplus to Small Suppliers

The threat point effect decreases relational contract prices most for small constrained (*i.e.*, risk-averse) suppliers, motivating policy remedies given that policy makers likely prioritize small firms due to their large presence in the firm size distribution and role as the majority of employers in low- and middle-income countries (Hsieh and Olken, 2014). Prioritization of small firms could also reflect preferences to redistribute to reduce inequality, as reallocating profit to small firms could increase aggregate welfare if low levels of profit (and employee income) result in high marginal utility from additional profit. Last, preferences favoring small firms could reflect dynamic efficiency concerns in environments where productive small firms need to finance growth from contemporaneous profits given credit market frictions. Policy focus on small firms is evident empirically as initiatives targeting small and medium enterprises (SMEs) are omnipresent, such as the World Bank's policy advising on finance tools for SMEs[44] and the IMF's 2019 conference on SME financing.[45]

I consider two avenues for policy makers interested in redistributing surplus to small, risk averse suppliers. First, motivated by concerns common in industrial organization about market power, I consider increasing downstream buyer competition. Second, given the focus on missing markets in development economics, I create the missing market for insurance as insurance allows small firms to smooth profit without relying on the relational contract.

1.5.1 Increase Downstream Buyer Competition

I first study the effect of increasing downstream buyer competition, estimating relational contract prices in a different market structure where there are additional potential relational contract buyers that the supplier could trade with instead of the buyer studied. Policy makers have many tools to support the development of large exporting firms, including reducing red tape to export, supporting network formation between large firms and international end clients, changing tariffs, and even possibly subsidizing large firms to begin exporting. Moreover, policy makers could target barriers that constrain small firm growth, such as helping them access capital and other inputs necessary for growth.

To interpret the policy of increased downstream buyer competition through the lens of the model, I conceptualize increased downstream buyer competition as decreasing the buyer bargaining power parameter.[46] More relational buyers should improve the supplier's long-

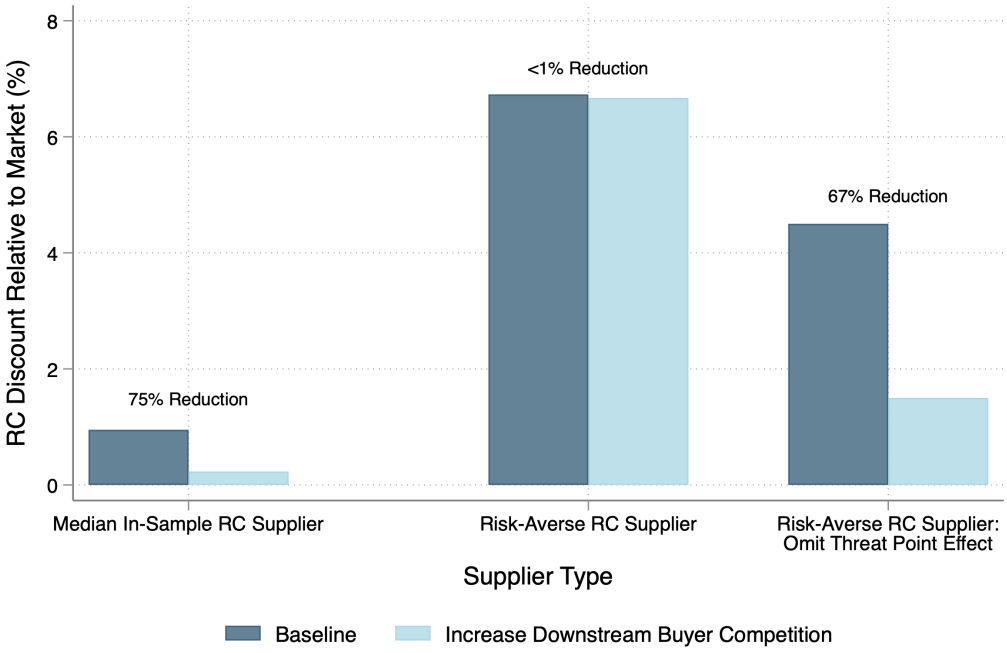
[44] World Bank SME Finance
[45] IMF SME Financing
[46] I halve the buyer bargaining parameter.

run outside option, but not the short-run outside option of selling on the market because it is not a credible threat for the supplier to add a relational contract buyer in the short-run. As long as the increase in buyer competition is known by the buyer, I propose that the buyer responds to competition by increasing the share of surplus suppliers receive to ensure supplier loyalty. This argument reflects similar economics as efficiency wages because both incentivize loyalty by offering a higher share of the surplus (Katz, 1986).

Suggestive empirical patterns based on the correlation between supplier attributes and estimated buyer bargaining parameters support this interpretation. Suppliers who provide fabrics with higher concentration (compared to other relational contract suppliers) are likely most valuable to the buyer. For example, if the buyer needs to procure a fabric for a specific garment at a discount relative to the market, then it needs to have a relational contract for that fabric. Additionally, suppliers who provide fabrics with higher concentration presumably make it easier for the buyer to fulfill the quantity guarantees in the relational contracts. To illustrate this idea, consider an extreme case where all suppliers sell the same fabric such that concentration is low. In this setting, if the aggregate quantity for that fabric is low, then the buyer will not be able to fulfill its quantity guarantees to all suppliers because assigning an order to one supplier to meet its quantity guarantee inherently means not assigning that order to another supplier. Therefore, suppliers who produce more concentrated fabrics, calculating concentration only among relational contract suppliers, need not compete with other suppliers for the same volumes.

I find that suppliers who sell products with low concentration are in relationships with lower buyer bargaining parameters, consistent with the buyer ensuring that these relational suppliers remain loyal as they are especially valuable. Appendix Section B.13 provides more information on the empirical approach used to show that that supplier concentration in fabric markets, relative to other relational contract suppliers, is negatively correlated with the estimated buyer bargaining parameters. Additionally, this proxy for relational contract supplier value is not significantly correlated with estimated risk aversion parameters, suggesting that the correlation between it and buyer bargaining parameters is not simply spurious.

Figure 1.20: Downstream Buyer Competition and Relational Contract Prices



Note: Model estimates. Figure shows change in relational contract discount from halving the buyer bargaining parameter from the estimated median for relational contract suppliers. Risk aversion for the median in-sample relational contract supplier is the median of estimated risk aversion parameters. Risk aversion for the risk-averse supplier is the minimum risk aversion such that there is a threat point effect. Model estimates use the capacity and demand assurance for a median in-sample relational contract supplier.

Figure 1.20 illustrates the change in discount in the relational contract relative to the market, expressed as a percent of the market price, induced by increased downstream buyer competition. The effect of this policy is small, especially for risk-averse suppliers. Specifically, while halving the median estimated buyer bargaining parameter results in a large 75% reduction in discount for a supplier at the median risk aversion of in-sample suppliers–from .95% to 0.23%–the same reduction in buyer bargaining parameter has a negligible effect for suppliers with the minimum risk aversion for the threat point effect to matter, reducing the discount by less than 1% from 6.73% to 6.67%. Importantly, ignoring the threat point effect *overestimates the benefits* of reducing the buyer bargaining parameter, as without the threat point effect the discount would decrease by 66% from 4.5% to 1.5%.

Intuitively, increasing downstream buyer competition is not an effective policy because, once the buyer uses the integrated supplier as its threat point, prices are so close to the buyer's participation constraint that there is little room to shift relational contract prices. It follows that changing the surplus sharing rule has minimal effects on prices. This intuition

clarifies why omitting the threat point effect results in incorrect policy conclusions for small risk-averse firms and why this policy is effective for large suppliers–when the buyer's threat point is the market, there is enough surplus that changing surplus sharing effectively reduces relational contract prices.

Appendix B.12 considers the effects of horizontal mergers of suppliers. Although horizontal consolidation does reduce the relational contract discount, consolidation effectively assumes away the problem of the existence of small constrained firms in the first place.

1.5.2 Create Missing Market for Insurance

The importance of missing markets in low- and middle-income country settings, as well as the limited effects of increasing downstream buyer competition, motivate analyzing the effects of adding the missing market for insurance. Specifically, I evaluate the change in the relational contract discount that would occur if suppliers had access to actuarially fair insurance over profits when trading with the market. This insurance would effectively blunt the effect of risk aversion on discounts the buyer receives because suppliers can smooth profit now without relying on the relational contract.[47] As all uncertainty is on the capacity and cost side in this setting, insurance over prices (or, equivalently, revenue) would function equally well, although such a contract would need to be non-linear as profits are not linear in capacity due to cost convexity.

While the government is presumably unlikely to offer such insurance (although price floor policies for agricultural goods are common and effectively insure agricultural producers against some adverse prices and could also improve producers' threat points), policies could support trade associations and cooperatives that provide price or revenue insurance (or at least support), and potentially even implement policies to mitigate moral hazard, such as pegging the volume of insurance offered to a running average of historic capacity or revenue.[48] Alternatively, the government could support policies to improve access to financial markets for suppliers to reduce their exposure to profit fluctuations. For example, improving access to credit and savings markets might help suppliers self-insure by building savings over time, increasing credit access for suppliers could allow them to borrow rather than need relational contracts, and supporting the development of markets for derivatives or other financial prod-

[47] Blouin and Macchiavello (2019) also shows how missing insurance markets interact with contracting behaviors, focusing on how supplier incentives to strategic default lead to lower levels of insurance provided by using index-priced contracts. While Blouin and Macchiavello (2019) focus on efficiency effects of the missing insurance markets, I focus on its implications for the distribution of surplus.

[48] As all uncertainty in this model for suppliers is from revenue, and not costs, price insurance would be isomorphic to profit insurance, although the insurance contract would need to non-linear in capacity to smooth profit given convex costs.

ucts could allow suppliers to hedge their risk more effectively (although, admittedly, financial literacy and transaction costs could limit take-up). And, policies supporting small firms to access financing at more favorable rates are common in many settings, including high-income countries.

Figure 1.21: RC Discount Without Insurance

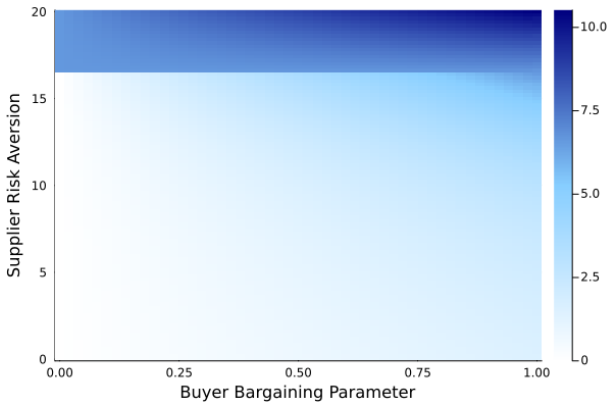
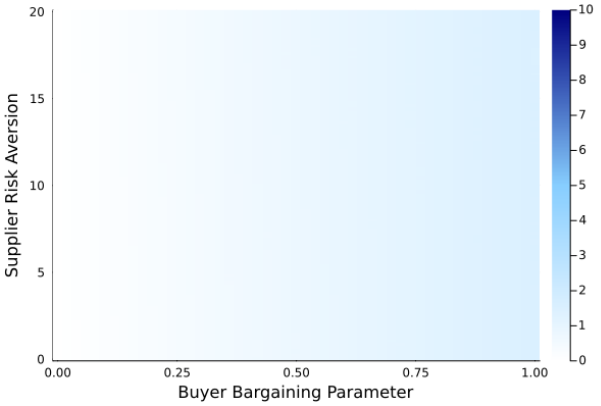


Figure 1.22: RC Discount With Insurance



Note: Model estimates for the discount in the relational contract relative to the market as a percent of the market price. Heatmap legend is to the right of the heatmap. Model estimates use the average capacity and demand assurance across in-sample relational contract suppliers.

Creating the missing market for insurance reduces the discount offered by risk-averse suppliers in the relational contract. Furthermore, it effectively eliminates the threat point effect, with risk-averse suppliers no longer receiving significantly lower prices. Adding insurance is effective because it prevents the buyer from using the integrated supplier as its threat point to prevent the relational contract from breaking, as the supplier's threat point improves. In other words, the most effective approach to counteract the buyer's improved bargaining position due to a change in the buyer's threat point is symmetric: improve the supplier's bargaining position by improving its threat point.

1.6 Conclusion

This paper highlights that firms can vertically integrate to improve their bargaining position with external firms. The intuition is that vertical integration changes a firm's threat point, which improves its bargaining position vis-à-vis external firms. This improvement in bargaining position can improve contracting with external firms, allowing the integrated firm to extract more surplus–I call this change in contracts caused by the change in threat point due to vertical integration *the threat point effect*. I build a conceptual model studying

how vertical integration affects relational contracts. The details of the model are motivated by the empirical setting of a large Indian garment manufacturer which added an integrated fabric supplier bargaining with external relational contract suppliers. The conceptual model highlights that the magnitude of the threat point effect depends upon the sharing rule for and level of surplus in the relational contract. The threat point effect can reach over 5% of market prices when suppliers are risk averse, as supplier risk aversion means supplier surplus is large even when prices are low. Therefore, the buyer can leverage its threat point change to further increase the relational contract discount. However, if the integrating party extracts all surplus pre-integration, then there is no threat point effect as there is no additional surplus available for the buyer to extract. If the level of surplus in the relational contract is so small pre-integration such the relational contract would break with the integrated firm as the threat point, then there also is no threat point effect. In this case, the integrating firm would prefer to keep the pre-integration relational contract, so the relational contract equilibrium would not change.

I estimate the model, finding that the threat point effect magnitude is quantitatively important for small risk-averse firms, reaching up to 6.7% of pre-integration relational contract prices. Given that small firms are both economically and policy relevant as the majority of firms and employers in low- and middle-income countries and experience the largest threat point effect, I evaluate two policies to reallocate surplus to them. First, motivated by concerns about competition and market power common in industrial organization, I analyze the effects of increasing downstream buyer competition. Although this policy helps large firms, it has a trivial effect for small firms. Importantly, not accounting for the threat point effect results in both *underestimates* of the level of the discount offered by small risk-averse firms and *overestimates* of the benefits of this policy. Second, motivated by the focus in development economics on missing markets, I create the missing market for insurance to allow small risk-averse firms to smooth profit outside the relational contract. This policy improves small firm surplus and eliminates the threat point effect.

More broadly, my results contribute to the growing body of evidence that challenges to smoothing profits in low-income country settings not only decrease welfare through preventing equalizing marginal utility across states, but also have important price effects. For example, agricultural producers in low-income country settings tend to buy and sell at the same time as other nearby producers due to limitations in their ability to store agricultural output, resulting in negative price effects as they "sell low and buy high" (Burke et al., 2018). And, inability to smooth over productivity shocks leads wage workers in low-income agricultural settings to increase their hours, magnifying the decrease in their wages as productivity shocks are correlated over space and time (Jayachandran, 2006). In my setting,

small risk-averse firms accept lower average prices in order to smooth profits. Therefore, policies reducing exposure to profit (or income) fluctuations can have important welfare benefits not only through equalizing marginal utility but also by improving prices. Future research should seek to identify which policies are most effective at assisting small firms and producers in low-income settings to cope with risk, measuring policy effectiveness with respect to both smoothing across states and prices.

Last, my results also emphasize that increasing competition in low- and middle-income country markets with relational contracts may not always achieve desired policy goals of improving market performance and enhancing welfare. For example, increased competition reduces output and small producer (*i.e.,* farmer) welfare in Rwandan coffee markets (Macchiavello and Morjaria, 2020). Similarly, Brugues (2020) argues that mitigating market power alone would lead to welfare losses given contracting frictions in the context of relational contracts in the Ecuadorian manufacturing supply chain. My counterfactual analysis suggests that policies directly targeting the underlying missing markets that create demand for the relational contract may be more effective at meeting policy goals than increasing competition. For example, in the same empirical setting in Macchiavello and Morjaria (2020), improving farmers' access to financial markets to procure inputs could reduce their need for trade credit and allow them to benefit from downstream price competition for their coffee. The standard intuition of the benefits of competition would likely apply once firms rely less on relational contracts. Keeping in mind concerns that changes in second-best institutional settings can have unintended adverse effects and that demand for relational contracts in some settings emerges due to transaction rather than market features, future research should identify and evaluate policies creating missing markets (and, more broadly, mitigating market failures) that generate demand for relational contracts from small firms.

CHAPTER 2

Vertical Integration and Relational Contracts: Empirical Evidence on Substitutability from Indian Garment Manufacturing

2.1 Introduction

Firm respond to input market frictions, such as hold-up problems where relationship-specific investments lock-in one party and thereby enable the other party to extract surplus, by relying on transaction-specific governance to organize behavior and prices rather than transacting at arm's length on the market (Williamson, 1979). There are two forms of transaction-specific governance: vertical integration ("unified governance") and relational contracts ("bilateral governance"). Vertical integration resolves the hold-up problem by eliminating incentives for opportunistic behavior as the firm maximizes total profits and allocates inputs by command rather than prices. Relational contracts are informal agreements that facilitate cooperative action, rather than opportunistic behavior, such as holding up a party after a relationship-specific investment is sunk, through the value of the future (Baker et al., 2002). Relational contracts achieve contractual performance through the presence of surplus in the relationship for both parties relative to transacting with the market and punishing opportunistic behavior by preventing future access to that surplus through ending the relationship. As relational contracts rely on the value of the future rather than courts or other third-party enforcement, they are likely especially important in low- and middle-income countries (Macchiavello, 2022).

Because both vertical integration and relational contracts can resolve input market frictions, it seems natural to think that firms would treat them as substitutes.[1] Formally,

[1]Williamson (1979) suggests that anticipated need to adapt the contract pushes the theoretically predicted governance structure towards vertical integration, as an integrated firm can adapt whenever needed. Economies of scale (combined with imperfect input specificity) push towards relational contracting, as one supplier can produce specialized inputs for many firms and all parties benefit from the scale economies.

substitutes means that an exogenous decrease in the price of procuring an input through a relational contract should decrease the quantity of inputs procured through integrated suppliers (and vice-versa). This hypothesis is consistent with prior evidence that transactions governed by vertical integration and relational contracts resemble each other under typical circumstances (Macchiavello and Miquel-Florensa, 2017).

However, there are also many reasons why vertical integration and relational contracts might not be substitutes. For example, the different contract governance tools could be so strongly matched to transaction characteristics or market institutions that firms cannot switch between forms. Alternatively, the increase in competition within the relational contract from the new integrated firm could cause the relational contracts to break down because the integrated firm changes outside options, leading the buyer to replace relational contract volume with integrated volume rather than replacing market volume with integrated volume. Or, the integrated firm may no longer be able to meet its commitments in the relational contract due to directing volume to the integrated firm, which is especially relevant for the many relational contracts that include quantity assurance. Then, the relational contract dissolves once the integrated firm stops providing quantity assurance.

This paper empirically tests the substitutability of vertical integration and relational contract in the context of the Indian garment sector. Identification of substitutability leverages exogenous variation in contracting costs using a novel dataset of the universe of fabric transactions between a large Indian garment manufacturer (hereafter, buyer) and its market, relational contract, and integrated suppliers.

Substitutability of vertical integration and relational contracts is a policy relevant because vertical integration can have anticompetitive effects, such as incentivizing foreclosure and raising rivals' costs to decrease downstream competition (Asker, 2016). If relational contracts function similarly to vertical integration, competition authorities should also evaluate possible anticompetitive effects of relational contracts. Such competition concerns can be especially important in low- and middle-income countries, where many markets are fragmented and uncompetitive (Bergquist and Dinerstein, 2020; Banerjee et al., 2022; Bergquist et al., 2024).

This setting is well-suited to empirically test substitutability for three key reasons. First, the large size of the buyer and the detailed product level data facilitate analysis at the product level within a single buyer, mitigating concerns of omitted variable bias due to heterogeneity across buyers in contracting costs, strategies, and institutions. Second, the buyer builds an integrated mill, which serves a shock to vertical integration that facilitates analyzing which products are brought in-house relative to pre-integration behavior. I interpret these integration decisions as a revealed preference measure of which fabrics the buyer prefers to

integrate. Third, the size of the integration is economically meaningful both in absolute terms, at roughly 20% of fabric sourced, and relative terms, doubling quantity produced internally; for comparison, the largest external supplier accounted for only 7% of fabric purchases pre-integration.

I find strong empirical evidence that vertical integration and relational contracts serve as substitutes, with an additional relational contract supplier causing a 30 to 50 percentage point reduction in integrated production. I then rationalize this result through a simple model motivated by the empirical setting that illustrates how vertical integration and relational contracts can both resolve the relevant hold-up problem.

Substitutability between vertical integration and relational contracts suggests refining Williamson (1979), supporting the conclusion that vertical integration and relational contracting are similar but showing that firms may use both because they are close substitutes rather than the existence of a single and unique optimal governance structure for the transaction. In principle, either increasing marginal costs within a governance structure due to diseconomies of scale in production within a governance type (*e.g.*, due to needing to access production with higher marginal costs by using more marginal suppliers, workers, or capital) or preferences for having many distinct suppliers for protection against unlikely but extreme shut-down or other issues could incentivize using multiple transaction governance forms simultaneously for the same input. Morton (2023b) explores another motivation for firms to use multiple transaction governance forms simultaneously, namely that partial integration can improve a firms' bargaining position with external firms.

I then discuss the policy implications of substitutability between vertical integration and relational contracts. Substitutability of vertical integration and relational contracts raises potential concerns that relational contracts have similar anticompetitive effects to vertical integration, specifically in regard to foreclosure and raising rivals' prices; I conclude that some of these concerns are still present, but that the degree to which they raise anticompetitive concerns depends on features of the relational contract, market structure, and production technology. I also argue that relational contracts should reduce concerns about double marginalization, similar to vertical integration.

Substitutability helps resolve the puzzle of low-correlation between levels of vertical integration and measures of contracting environment quality found in the literature, as theory predicts firms should vertically integrate in response to low-quality contracting institutions (Acemoglu et al., 2009; Macchiavello, 2012). Specifically, relational contracts are both likely more prevalent in settings where there is more vertical integration, as they are substitutes, and are presumably more prevalent in settings with weaker contracting institutions (Macchiavello, 2022). Therefore, analyzing the effects of contracting environments on levels of

vertical integration should either control for levels of relational contracting or consider the combination of these forms of transaction-specific governance together rather than vertical integration alone.

I also contribute to this literature by improving upon identification and measurement. By focusing on one firm at one point at time and leveraging exogenous variation at the fabric level around one event, my research design mitigates many of the identification challenges that limit causal interpretation of the cross-country correlations in prior work. For example, property rights could be correlated with both contracting costs, as strong property rights could facilitate contracting, and vertical integration, as strong property rights may increase incentives for entrepreneurs through reduced risk of expropriation after profitable integration. These concerns lead Acemoglu et al. (2009) to note that "results reported in the paper cannot be interpreted as causal relationships."

My unusually rich transaction data, including the universe of fabric purchases with detailed information on prices, quantities, product, supplier identity and date, allow me to observe integrated flows, unlike most prior studies. Without similar data, Acemoglu et al. (2009) proxy for vertical integration for firms across all countries by combining data on the firms' activities and applying U.S. input-output data while Macchiavello (2012) proxies for vertical integration as the ratio of value added over output. If U.S. input-output tables are not globally representative, or firms' activity codes do not map well to levels of vertical integration, then the approach in Acemoglu et al. (2009) would have measurement error. Similarly, if market power or quality differentiation influence observed value added beyond just what percent of production occurs in-house, then the approach in Macchiavello (2012) would have measurement error.

Breza and Liberman (2017) are closest to this paper, showing that an exogenous increase in the costs of using a relational contract to procure inputs results in more integration in the context of a large Chilean retailer. However, the policy variation used for identification not only generally increases the cost of procuring an input through relational contracts on the intensive margin but also both changes costs within the relational contract and legally makes the terms of pre-existing relational contracts invalid. It follows that such variation is far from ideal for cleanly testing substitutability as multiple facets of the contracting and cost structure change together to simultaneously influence observed equilibrium outcomes.

This paper also adds to the growing empirical literature showing that firm behavior and market organization in low- and middle-income country settings that arise from weak institutions can induce less distortion and welfare loss than anticipated. For example, Bassi et al. (2022) illustrate that small firm sizes do not necessarily prevent firms from using technology associated with scale economies because of an active rental market for large machines.

Brugues (2020) even argues that relational contracts achieve efficient outcomes once established. I add to this literature by suggesting that firms in low- and middle-income countries have multiple strategies available to mitigate market frictions, relying on both vertical integration and relational contracting to improve input trade. Therefore, lower levels of vertical integration than theoretically predicted may not inherently suggest that firms are unable to resolve input market frictions–firms use relational contract as a substitute.

The remainder of the paper is organized as follows. Data and context, including institutional details regarding the vertical integration in the study, fabric procurement, and a description of the relational contract with supporting reduced-form evidence are described in detail in 1.2. Section 3 presents the evidence of substitutability, discussing estimation and identification. In section 4, I analyze threats to exogeneity, providing evidence that such concerns are unlikely to result in overestimates of substitutability. Section 5 details a robustness check by documenting lack of substitutability between vertical integration and large suppliers not in relational contracts, and section 6 develops a simple conceptual model motivated by the empirical setting to illustrate how firms can use either vertical integration or relational contracts to resolve hold-up problems. Section 7 concludes and discusses policy implications.

2.2 Relational Contracts and Vertical Integration as Substitutes: Which Fabrics are Brought In-House?

As both vertical integration and relational contracts can resolve contracting frictions, I hypothesize that buyers should consider them as alternatives for mitigating the hold-up problem. In other words, relational contracts and vertical integration should be substitutes, meaning that an exogenous decrease in the cost of procuring an input through a relational contract (vertically integrated supplier) should decrease the quantity of inputs sources through vertically integrated suppliers (relational contract suppliers). I test this hypothesis using exogenous variation in contracting costs at the product level within a single buyer, mitigating concerns of endogeneity related to heterogeneous buyer types or preferences.

2.2.1 Estimation and Identification

I estimate the association between the percentage of the fabric that is brought in-house and the number of relational contract suppliers who produced the fabric pre-integration, which is a cost-shifter for purchasing the input through a relational contract, as in equation 2.1. As more than half of all outcome observations are zero, meaning no fabric is brought in-

house, I estimate all regressions using a tobit model.[2] Although point estimates from OLS are smaller,[3] they retain the sign and statistical significance of the estimates from the tobit model.

$$\begin{aligned}\%\text{ Fabric Brought In-House}_i = \alpha &+ \beta\Big(\#\text{ RC Suppliers Pre-Integration}_i\Big) \\ &+ \boldsymbol{\delta}\Big(\text{Pre-Integration Demand \& Cost Shifters}_i\Big) \\ &+ \epsilon_i \end{aligned} \tag{2.1}$$

I estimate specifications both with and without controls for demand and cost shifters that can influence how much fabric production is brought in-house. The control variables are: *i)* the percent of the fabric that is produced in-house at the other integrated supplier prior to integration, which presumably decreases demand for bringing fabric in-house but also may decrease the cost of integrating the fabric through learning-by-doing; *ii)* the fabric's percent of total volume prior to integration, which presumably increases demand for bringing fabric in-house as it captures demand for the fabric; and *iii)* the average quarterly growth rate for the fabric in the year preceding integration to control for trends in fabric demand influencing integration decisions. Last, I control for a supply shifter, specifically the total number of suppliers for the fabric, across both relational contract and market suppliers. If few suppliers can produce the fabric, then market power could lead to higher prices for the fabric (rather than the hold-up problem associated with contracting frictions in this setting), creating additional but unrelated incentives to bring production in-house. I only use pre-integration data for all controls to ensure that vertical integration is not endogenously affecting them.

The ideal variation to estimate this model is exogenous shocks to the relative costs of relational contracts and vertical integration–if decreases in the costs of procuring an input through relational contract suppliers (vertically integrated supply) result in decreases in quantities of inputs procured through vertically integrated supply (relational contract suppliers), then the two are substitutes. I leverage the number of pre-integration relational contract suppliers who can produce the product as plausibly exogenous variation in the costs of procuring an input through a relational contract.[4]

[2]The tobit model includes the natural bounds for a dependent variable that is a percentage: 0 (0%) and 1 (100%).

[3]See Appendix Table C.1.

[4]The data do not include information on the full set of fabrics that a supplier could produce; instead, I proxy for the set of fabrics that a supplier could produce by using the set of fabrics which they do produce. As this approach introduces some measurement error, I weight fabrics in subsequent analyses by their share of pre-integration volume, as fabrics with larger volume are likely to see closer to the full set of relevant

I posit that the number of pre-integration relational contract suppliers serves as an exogenous shifter of relative costs of relational contracts due to fixed costs of establishing a relational contract (as in Cajal-Grossi et al. (2023)), although this fixed cost argument is sufficient but not necessary for identification. Specifically, a fixed cost to establishing a new relational contract makes purchasing fabrics through a relational contract more expensive for fabrics that do not have pre-existing relational contract suppliers–it becomes necessary to create a new relational contract to source that fabric from a relational contract supplier. A fixed cost for creating a relational contract seems very plausible in this setting, with relational contracts presumably requiring an investment in communication and screening supplier quality and capabilities.

However, fixed costs are not the only economic rationale possible for interpreting the number of pre-integration relational contract suppliers as a cost shifter for procuring a fabric through a relational contract. For example, any cost convexity with respect to number of relational contract suppliers in procuring an input through a relational contract is sufficient. Cost convexity could be due to marginal effort to track quantities and prices in each relational contract increasing in the number of relational contract suppliers.

The identifying assumption is that fabrics are "as-if" randomly assigned to relational contract suppliers.[5] This assumption can hold even if relational contract suppliers are *endogenously* selected for a variety of reasons, such as cultural distance (Cevallos Fujiy et al., 2023), location, reliability, and working style, as long as this selection is not for a specific fabric. As suppliers typically produce many fabrics, selling at least 7 (15) different fabrics in the 3 (6) months prior to integration, with the median relational contract supplier in terms of number of fabrics produced selling 20 (42) distinct fabrics in the same timeframe, it seems unlikely that relational contract suppliers are matched to each fabric intentionally.

Additionally, economically, the buyer should not select relational contract suppliers due to performance for only a few fabrics, as relational contracts are long-lived but trends in fabric vary dramatically over time. It follows that supplier choice should prioritize broad performance capabilities. Furthermore, at the time of sampling, when suppliers are "locked-in", it is not known which orders will be selected by the end client, further exacerbating the challenge of intentionally targeting fabrics to relational contract suppliers to meet relational contract quantity guarantees.

Note that this identifying assumption is much weaker than previous work studying how contracting institutions influence vertical integration; for example, any cross-country analysis

suppliers produce the fabric and, therefore, have less measurement error.

[5]Equivalently, the identifying assumption is that there are no contracting frictions (or other factors) which affect both the pre-integration count of relational contract suppliers and levels of vertical integration.

likely has many possible omitted variables, such as differences in legal or other institutional factors, that are not a concern for an analysis focusing on one buyer in one country. Furthermore, this identification strategy uses the same intuition as Ciliberto and Williams (2014), which also leverages variation at the product level (airplane routes) induced by a long-term firm choice with a large fixed cost (airport gate rental).

It is reasonable to interpret this empirical strategy as effectively showing the reduced-form of an instrumental variables approach, as the pre-integration count of relational contract suppliers shifts the difference in marginal costs between purchasing an input from a relational contract supplier and purchasing it from an integrated supplier. However, because this cost difference is not observed in the data, a standard instrumental variables approach is not feasible.

Table 2.1: Substitutability of VI & RC

	% Vertically Integrated (Fabric)			
	(1)	(2)	(3)	(4)
Number of RC Suppliers	−.5369***	−.5231***	−.3109***	−.4773***
	(.154)	(.1704)	(.0776)	(.0979)
% Volume from Other		−.0627		.9708***
Integrated Supplier (Pre-Integration)		(.5203)		(.345)
Fabric % of Total Volume		41.0198***		18.0216*
		(9.0516)		(10.408)
Trend in Fabric Demand		0***		0
(Quarterly Growth)		(0)		(0)
Number of All Suppliers		.0653		.0562
		(.0592)		(.0357)
Months Pre and Post Integration	3	3	6	6
N	162	162	307	307
R^2	.1321	.5623	.0902	.2608
Mean Dependent Variable	.0737	.0737	.0852	.0852

Note: Data from universe of fabric transactions by buyer. Regressions shown consider, alternatively, transactions from three or six months pre- and post-integration. Robust standard errors in parentheses and weights by fabric percent of total volume. ***,** ,* indicate significance at the 1, 5, and 10 percent levels, respectively.

The statistically significant negative correlation between the number of relational contract suppliers who can produce a fabric pre-integration and the percent of fabric brought in-house after integration in Table 2.1 provides strong evidence of substitutability between relational contracts and vertical integration. The estimates suggest that for each additional relational contract supplier who produces the fabric, the percentage of fabric production

brought in-house decreases by approximately 30 to 50 percentage points, depending upon the specification. The magnitude is economically meaningful: the mean percentage of fabric brought in-house is roughly 8%, indicating that much of the effect is at the extensive margin of bringing fabric production in-house. This negative correlation is robust to adding controls and changing the length of the window before and after integration used to measure both volumes and the number of suppliers for the fabric pre- and post-integration.

Next, I provide supporting evidence that the estimated association between the quantity of inputs brought in-house and the number of pre-integration relational contract suppliers is causal. First, I discuss possible omitted variable bias concerns and conduct additional analysis to illustrate that they are unlikely to drive results. Second, I replicate the presented evidence but using placebo suppliers instead of relational contract suppliers to demonstrate that the observed findings are unlikely to reflect supplier heterogeneity associated with order volume, such as unobserved heterogeneity in supplier efficiency, rather than a relational contract.

2.3 Ruling Out Threats to Causal Interpretation: Omitted Variable Bias

I consider possible omitted variable bias that could influence the interpretation of the estimated parameter as indicative of substitutability of relational contracting and vertical integration. Two potential sources of omitted variable bias would lead the current approach to *underestimate* the degree of substitutability, meaning that an increase in the number of relational contract suppliers should be associated with even less fabric brought in-house.

First, the importance and definition of quality may vary at the fabric level; for example, print fabrics must have consistent patterns that are not a quality issue for solid fabrics. In this case, as both relational contracts and vertical integration are used to ensure high quality (Macchiavello and Morjaria, 2020; Hansman et al., 2020a), quality importance should be positively correlated with both. It follows that the estimate would be biased upwards, meaning that actual substitutability is even stronger (more negative) than estimated. Buyer market power may also vary by fabric; when the buyer has more market power there is less need to use either vertical integration or relational contracts to obtain low prices. As with quality, if buyer market power varies by fabric and is correlated with both vertical integration and relational contracting, it would bias the estimate upwards, also leading to underestimates (*i.e.*, less negative) of substitutability. Because these omitted variable stories, even if true, mean that there is even stronger substitutability between vertical integration and relational

contracts than shown in my estimates, I do not evaluate them further.

Other possible omitted variables could lead to bias that *overestimates* substitutability *overestimated*, meaning that actual substitutability is less than estimated. I consider supplier market power, production cost heterogeneity, fabric specific comparative advantage, and reliability.

2.3.1 Supplier Market Power

Supplier market power could influence both vertical integration and relational contracts, creating incentives to bring production in-house to avoid paying high prices associated with supplier market power and making relational contracting challenging as suppliers have a strong outside option due to relatively high profits from market sales. However, it is not necessarily the case that relational contracting is negatively correlated with supplier market power–there could still be a relational contract to reduce hold-up, potentially with higher prices. Regardless, given the potential for supplier market power to serve as an omitted variable as it could possibly be positively correlated with vertical integration and negatively correlated with relational contracting, I study supplier - fabric HHI and vertical integration. If suppliers have relatively high market power for a specific fabric, then there should be few sellers of the fabric, resulting in a high HHI for the fabric. Therefore, if supplier market power drives vertical integration, which is necessary for omitted variable bias, then there should be a positive correlation between supplier - fabric HHI and vertical integration.

To test this hypothesis, I estimate the following two regression models, with the first estimating the correlation between supplier - fabric HHI and the second assessing the stability of the relationship between the number of pre-integration relational contract suppliers for the fabric and the percentage of the fabric brought in-house, controlling for supplier - fabric HHI.

$$\begin{aligned}\%\text{ Fabric Brought In-House}_i &= \alpha + \gamma\Big(\text{ Supplier-Fabric HHI}_i\Big) \\ &\quad + \boldsymbol{\delta}\Big(\text{Pre-Integration Demand \& Cost Shifters}_i\Big) \\ &\quad + \epsilon_i \\ \%\text{ Fabric Brought In-House}_i &= \alpha + \gamma\Big(\text{ Supplier-Fabric HHI}_i\Big) \\ &\quad + \beta\Big(\#\text{ RC Suppliers Pre-Integration}_i\Big) \\ &\quad + \boldsymbol{\delta}\Big(\text{Pre-Integration Demand \& Cost Shifters}_i\Big) \\ &\quad + \epsilon_i\end{aligned}$$

Table 2.2: Substitutability of VI & RC: Supplier Market Power Robustness

	% Vertically Integrated (Fabric)					
	(1)	(2)	(3)	(4)	(5)	(6)
Supplier - Fabric HHI	.0266	−.5587	−.3089	−.3962	−.9331**	.0924
	(1.0757)	(.9668)	(.3201)	(.5225)	(.43)	(.4963)
Number of RC Suppliers		−.571***	−.5215***		−.3907***	−.4783***
		(.1794)	(.1663)		(.0775)	(.0982)
% Volume from Other			−.0314			.9794***
Integrated Supplier (Pre-Integration)			(.471)			(.3423)
Fabric % of Total Volume			42.9412***			17.9862*
			(8.9542)			(10.1126)
Trend in Fabric Demand			0***			0
(Quarterly Growth)			(0)			(0)
Number of All Suppliers			.0408			.0627
			(.0581)			(.0459)
Months Pre and Post Integration	3	3	3	6	6	6
N	162	162	162	307	307	307
R^2	0	.1437	.5699	.0095	.1468	.2612
Mean Dependent Variable	.0737	.0737	.0737	.0852	.0852	.0852

Note: Data from universe of fabric transactions by buyer. Regressions shown consider, alternatively, transactions from three or six months pre- and post-integration. Robust standard errors in parentheses and weights by fabric percent of total volume. ***,** ,* indicate significance at the 1, 5, and 10 percent levels, respectively.

Estimates in Table 2.2 illustrate that supplier-fabric HHI is only weakly associated with the

percentage of production brought in-house. It follows that supplier market power does not appear to shape vertical integration decisions. Additionally, controlling directly for supplier - fabric HHI does not meaningfully change the estimated relationship between the number of relational contract suppliers for the fabric and the percentage of production brought in-house, with an additional relational contract supplier associated with a 39 to 57 percentage point decrease in the production brought in-house.

2.3.2 Production Cost Heterogeneity

Production costs could also vary by fabric and this variation could potentially be associated with both relational contracting and vertical integration (recall that analysis is at the fabric level, not the supplier level). If there is heterogeneity in the fixed cost of producing various fabrics, then it would create differential supplier market power by fabric, as discussed in 2.3.1. Heterogeneity in variable costs by fabric may also influence decisions of which fabrics to integrate, although the motivation would have to be for some reason other than transaction prices because, in principle, the hold-up problem that motivates vertical integration is independent of variable costs. For example, consider a situation where the buyer is choosing to integrate either a fabric with a low production cost and relatively low hold-up margin or a fabric with a high production cost and a relatively high hold-up margin. As integration resolves the hold-up problem, the buyer should integrate the fabric with high production cost to avoid the hold-up margin; the buyer chooses the fabric with higher production cost because production costs increase prices regardless of supplier type. It follows that contracting frictions alone do not motivate a correlation between variable cost heterogeneity and vertical integration.

However, there could be other motivations to integrate lower cost fabrics. For example, the buyer tends to receive some trade credit from suppliers; if the trade credit is a percent of transaction value, then disproportionately integrating high-cost fabrics increases the average value of trade credit received, assuming quantities are equal. That said, it is alnot evidence that such heterogeneity impacts relational contracts, as relational contracts in this setting mitigate the hold-up problem which does not, in principle, relate to variable cost heterogeneity for the same reason as described above for vertical integration.

Nevertheless, to be thorough, I estimate a regression that includes fabric group fixed effects for the three main fabric groups (jeans, greige, and cotton fabrics), as variable costs are likely similar within fabric type, with standard errors clustered at the fabric group level.

Table 2.3: Substitutability of VI & RC with Fabric Group FE

	% Vertically Integrated (Fabric)			
	(1)	(2)	(3)	(4)
Number of RC Suppliers	−.122*	−.2747	−.1923**	−.2008***
	(.064)	(.1865)	(.0829)	(.0464)
% Volume from Other		.4861***		1.1555***
Integrated Supplier (Pre-Integration)		(.0758)		(.3499)
Fabric % of Total Volume		26.0795***		2.551
		(3.6844)		(15.2312)
Trend in Fabric Demand		0***		0***
(Quarterly Growth)		(0)		(0)
Number of All Suppliers		.0001		−.0111
		(.0547)		(.0069)
Fabric Group FE	Y	Y	Y	Y
Months Pre and Post Integration	3	3	6	6
N	162	162	307	307
R^2	.4484	.6415	.2587	.3488
Mean Dependent Variable	.0737	.0737	.0852	.0852

Note: Data from universe of fabric transactions by buyer. Regressions shown consider, alternatively, transactions from three or six months pre- and post-integration. Robust standard errors in parentheses and weights by fabric percent of total volume. ***,** ,* indicate significance at the 1, 5, and 10 percent levels, respectively.

Although coefficients in Table 2.3 decrease in magnitude, point estimates remain economically meaningful, with the smallest in magnitude at 12.2 percentage points, which is larger than the mean percentage of fabric brought in-house (7%). Additionally, the estimates remain statistically significant, with the exception of one three-month coefficient which is marginally insignificant ($p = .14$).

2.3.3 Fabric Specific Comparative Advantage

Buyer fabric-specific comparative advantage in production could also influence vertical integration, with the buyer integrating fabrics where it has a comparative advantage. However, it is unclear that relational contracts would be associated with this comparative advantage–if the buyer anticipated future comparative advantage when planning what fabrics to integrate, it is plausible that the buyer would prefer relational contract suppliers who work in other fabrics. Even though this scenario seems unlikely, given that relational contracts are long-standing and relational contract suppliers provide a wide array of fabrics, previous estimates control for production at the other integrated mill. This control is useful for mitigating

concerns about this omitted variable bias, as if fabric specific comparative advantage shapes choices about which fabrics to integrate, then it presumably also shaped what fabrics were brought in-house at the prior mill and is unlikely to change dramatically over time.

A "learning-by-doing" mechanism also makes this control useful for reducing concerns about fabric specific comparative advantage, as the buyer would presumably be more likely to have comparative advantage in fabrics it already produces. Additionally, if comparative advantage is similar for fabrics in the same fabric group, then the analysis in Table 2.3 also addresses concerns of fabric specific comparative advantage.

Nevertheless, I consider one additional test, using the logic above that persistence of comparative advantage over time or "learning-by-doing" suggests that volume at the older integrated supplier shifts comparative advantage, should it exist. Therefore, I study the correlation between relational contract volumes and volumes at the older integrated supplier by regressing the percent of fabric volumes procured from relational contract suppliers on the percent of fabric volumes procured through the older integrated supplier. A positive coefficient on the percent of volume from the older integrated supplier indicates that the buyer is more likely to use relational contracts even for fabrics where it likely has comparative advantage in production. Therefore, a positive coefficient would be inconsistent with substitutability results reflecting omitted variable bias related to buyer comparative advantage for specific fabrics.

Table 2.4: Substitutability of VI & RC: Buyer Fabric Comparative Advantage Robustness

	Relational Contract Supplier Makes Fabric			
	(1)	(2)	(3)	(4)
% Volume from Other	.8664**	1.1637**	.4886***	.6034***
Integrated Supplier (Pre-Integration)	(.3692)	(.5327)	(.1886)	(.2275)
Fabric % of Total Volume		−3.8407		.5321
		(8.0173)		(6.8823)
Trend in Fabric Demand		0		0
(Quarterly Growth)		(0)		(0)
Number of All Suppliers		−.0305		−.0231
		(.0324)		(.0258)
Months Pre Integration	3	3	6	6
N	541	541	1054	1054
R^2	.0118	.0118	.0057	.0057
Mean Dependent Variable	.0737	.0737	.0852	.0852

Note: Data from universe of fabric transactions by buyer. Regressions shown consider, alternatively, transactions from three or six months pre- and post-integration. Robust standard errors in parentheses and weights by fabric percent of total volume. ***,** ,* indicate significance at the 1, 5, and 10 percent levels, respectively.

I find a positive correlation in Table 2.4, suggesting that fabrics which are more likely to have buyer comparative advantage (should it exist) do not have less volume procured through relational contracts.

2.3.4 Reliability

Reliability could also drive both vertical integration and relational contracts. At the fabric level, some fabrics could be more difficult to make consistently or have more input supply frictions leading to less reliable production. The buyer could use vertical integration to control reliability directly or relational contracts to increase incentives for reliability by conditioning future volume on it (as in Hansman et al. (2020a)). However, reliability is not a clear omitted variable, as relational contracts could have a direct effect on reliability. It follows that a positive correlation between relational contracts and reliability could represent the causal effect of relational contracting on reliability rather than reliability driving transaction governance choice. Nevertheless, reliability can be observed in the data, measured as the percentage of on-time deliveries. Controlling directly for reliability in Table 2.5 does not meaningfully change the association between vertical integration and relational contracts, although the results do suggest that fabrics with more reliable delivery are indeed less likely to be brought in-house.

Table 2.5: Substitutability of VI & RC: Reliability Robustness

	% Vertically Integrated (Fabric)					
	(1)	(2)	(3)	(4)	(5)	(6)
Reliability (% On-Time Delivery)	−1.1999*	−1.3759**	−.8375**	−.3804	−.5716**	−.2973
	(.6936)	(.6533)	(.4169)	(.3098)	(.2408)	(.2486)
Number of RC Suppliers		−.6115**	−.5892***		−.3317***	−.3851***
		(.2507)	(.1555)		(.0819)	(.1033)
% Volume from Other			.8619			1.156***
Integrated Supplier (Pre-Integration)			(.6061)			(.3646)
Fabric % of Total Volume			41.181***			2.0099
			(12.4573)			(11.1446)
Trend in Fabric Demand			0**			0
(Quarterly Growth)			(0)			(0)
Number of All Suppliers			.0337			.0531
			(.0661)			(.0368)
Months Pre and Post Integration	3	3	3	6	6	6
N	161	161	161	305	305	305
R^2	.0417	.199	.5051	.0149	.1211	.2722
Mean Dependent Variable	.0697	.0697	.0697	.0834	.0834	.0834

Note: Data from universe of fabric transactions by buyer. Regressions shown consider, alternatively, transactions from three or six months pre- and post-integration. Robust standard errors in parentheses and weights by fabric percent of total volume. ***,** ,* indicate significance at the 1, 5, and 10 percent levels, respectively.

2.4 Ruling Out Threats to Causal Interpretation: Placebo Supplier Analysis

To further support the causal interpretation of this correlation, I reproduce prior analysis but replacing the number of relational contract suppliers with the number of placebo suppliers by estimating:

$$\begin{aligned}\%\text{ Fabric Brought In-House}_i = \alpha &+ \beta\left(\#\text{ Placebo Suppliers Pre-Integration}_i\right) \\ &+ \boldsymbol{\delta}\left(\text{Pre-Integration Demand \& Cost Shifters}_i\right) \\ &+ \epsilon_i\end{aligned}$$

Table 2.6: Substitutability of VI & RC: Placebo Suppliers Robustness

	% Vertically Integrated (Fabric)			
	(1)	(2)	(3)	(4)
Number of Placebo Suppliers	.1341	.3683*	.163**	.2884***
	(.1565)	(.2091)	(.0797)	(.0981)
% Volume from Other		.4332		1.7033***
Integrated Supplier (Pre-Integration)		(.7384)		(.4885)
Fabric % of Total Volume		47.9857***		.0124
		(10.8551)		(15.1192)
Trend in Fabric Demand		0**		0
(Quarterly Growth)		(0)		(0)
Number of All Suppliers		−.2137***		−.1091**
		(.0721)		(.0544)
Months Pre and Post Integration	3	3	6	6
N	162	162	307	307
R^2	.0091	.4466	.0364	.1662
Mean Dependent Variable	.0737	.0737	.0852	.0852

Note: Data from universe of fabric transactions by buyer. Regressions shown consider, alternatively, transactions from three or six months pre- and post-integration. Robust standard errors in parentheses and weights by fabric percent of total volume. ***,** ,* indicate significance at the 1, 5, and 10 percent levels, respectively.

The results in Table 2.6 show a positive and significant correlation between the number of placebo suppliers who can produce a fabric pre-integration and the percent of fabric brought in-house after integration. Importantly, this result differs from the estimates for the relational contract suppliers, where all estimated associations are negative and significant. The difference between the estimates for relational contract and placebo suppliers emphasizes that the observed correlation does not reflect the effect of large suppliers on vertical integration, thereby emphasizing the importance of relational contracts.

Additionally, the difference in results further restricts what omitted variables stories could plausibly drive observed substitutability results. For example, communication issues, such as different languages, could create contracting frictions that influence both relational contracts and vertical integration, if the relational contract improves communication. (That said, it is not clear why there would be communication issues at the fabric level, let alone that such issues would determine the total number of relational contract suppliers for a fabric, especially as the buyer has offices throughout India with staff who speak both English and local languages.) Nevertheless, the placebo analysis further provides evidence that such omitted variables are not likely to drive results, as there is presumably meaningful communication for large orders that have more deliveries. Therefore, if communication issues played an

important role in decisions to vertically integrate, then we would expect to see the same negative result for the number of placebo suppliers as observed for the number of relational contract suppliers. This logic generalizes to many other stories about selection on supplier unobservables, including the idea that relational contract supplier are really just better suppliers. Specifically, it seems implausible that the buyer would purchase such large quantities from placebo suppliers, especially given their higher prices than relational contract suppliers, if they were bad suppliers on other unobservable dimensions.

A final robustness check estimates a model with the number of relational contract suppliers, the number of placebo suppliers, and the supplier-fabric HHI all included (see Appendix Table C.2). Results from this model provide further evidence that relational contracts are indeed a substitute for vertical integration, with an additional relational contract supplier associated with a 35 to 55 percentage point decrease in the quantity of production brought in-house. While the point estimate for the association between the number of placebo suppliers and the percentage of production brought in-house remains positive, once controls are included, the magnitude is small and is not statistically significant, which is consistent with the buyer not differentiating meaningfully between placebo and other market suppliers (*i.e.,* treating all market suppliers similarly, at least among the market suppliers who pass the buyers' screening criteria).

2.5 Model

In this section, I write down a simple model of the hold-up problem motivated by the empirical setting studied. Then, in the context of the model, I illustrate the conditions under which the firm prefers vertical integration or relational contracts for transaction governance rather than the market. Illustrating that both vertical integration and relational contracting can resolve the hold-up problem in this setting provides a microfoundation for the substitutability between the two: both transaction governance regimes allow the buyer to avoid paying the hold-up price making them substitutes from the perspective of the buyer.

2.5.1 Model Set-Up

I describe the payoffs for the buyer and seller below, with each firm seeking to maximize profit. In the model, p^E represents the price paid by the end client (*e.g.,* Old Navy), q represents the order quantity, $c^B(q)$ represents the buyer's cost of garment production, p^F represents the fabric price paid by the buyer to the seller, and $c^S(q, \nu)$ represents the seller's cost of fabric production for a buyer where ν reflects the variability of capacity from the

relevant buyer.

$$\underbrace{\Pi^B}_{\text{Buyer's Profit}} = \underbrace{p^E q}_{\text{Buyer Revenue}} - \underbrace{(c^B(q) + p^F q)}_{\text{Buyer Costs: Production and Fabric}} \tag{2.2}$$

$$\underbrace{\Pi^S}_{\text{Seller's Profit}} = \underbrace{p^F q}_{\text{Seller Revenue}} - \underbrace{c^S(q,\nu)}_{\text{Seller Costs: Production}} \tag{2.3}$$

I make the following assumptions:

$$p^E q - c^B(q) - c^S(q,\nu) > 0 \text{ (positive surplus from orders)} \tag{2.4}$$

$$\frac{dc^B(q)}{dq} > 0 \text{ (marginal buyer production costs increase in quantity)} \tag{2.5}$$

$$\frac{\partial c^S(q,\nu)}{\partial q} > 0 \text{ (marginal supplier production costs increase in quantity)} \tag{2.6}$$

$$\frac{\partial c^S(q,\nu)}{\partial \nu} > 0 \tag{2.7}$$

(marginal supplier production costs increase in quantity variance) (2.8)

While assumptions about surplus and the relationship between costs and quantities (2.4, 2.5, and 2.6) are fairly standard, the assumption of the effect of quantity variance on costs (2.8) merits further discussion. Costs increasing in variance, in this setting, reflect the convexity of the production technology[6]. Specifically, with convex production costs, the expected cost from a mean-preserving spread of quantities increases in the magnitude of the spread.[7] In reference to the relational contract, the ν term microfounds the benefit of the demand assurance in the relational contract for the supplier. Figure 1.10 provides empirical evidence of cost convexity in this setting based on analysis of cost and capacity data from the integrated supplier.

2.5.2 Hold-up Prices

Because the buyer commits to the supplier at the time of sampling, the buyer has no more choice in the model and the supplier can hold-up the buyer. Through the lens of the model, the supplier picks the price so the buyer does not decide to cancel the order and have a payoff of zero. Therefore, the supplier sets p^F to meet the buyer's participation constraint.

[6]Convexity in production cost in this industry is described and estimated in Morton (2023b)

[7]By definition of convexity, $\mathbb{E}[c(q)] > c(\mathbb{E}[q])$.

It follows that:

$$p^E q - c^B(q) - p^F q = 0$$
$$p^F = p^E - \frac{c^B(q)}{q}$$

The fabric price in this setting extracts all surplus from the order, with the buyer not making any profit on the order. If the buyer also cares about their reputation with the end client, then the buyer's surplus also includes the change in their reputation from filling the order, in which case the supplier may be able to set an even higher price. Note that capacity variance plays no role in this transaction; capacity variance is not contractible as it is difficult to measure and changes with sector-wide trends.

2.5.3 Vertical Integration Resolves Hold-Up

Below, I illustrate the conditions under which the buyer prefers to vertically integrate, assuming that adding integrated capacity entails fixed cost F and that the buyer can build an integrated supplier with the same technology (*i.e.,* production costs) as the market supplier. Therefore, the buyer prefers to integrate as long as:

$$p^E q - c^B(q) - c^S(q,\nu) - F \geq p^E q - c^B(q) - p^F q$$
$$p^E q - c^S(q,\nu) - c^B(q) \geq F$$

This result highlights that for sufficiently low F the buyer will prefer to vertically integrate than to purchase from the market. The intuition for this result is that the left-hand side represents the difference in fabric procurement cost when integrated–if this cost is lower than the fixed cost, the buyer prefers to integrate.

In a dynamic model where the buyer only has to pay the fixed cost once in perpetuity, the condition becomes weaker with the inequality becoming $\frac{p^E q - c^S(q,\nu) - c^B(q)}{1-\delta} \geq F$ where δ represents the buyer's discount factor.

2.5.4 Relational Contracting Resolves Hold-Up

I model the relational contract as specifying $\{p^{RC}, \nu^{RC}\}$, as fabric quantities are determined by the end client and, therefore, are not affected by the relational contract. I assume that the punishment strategy to prevent the hold-up equilibrium in the relational contract is a grim

trigger punishment strategy. The buyer's incentive compatibility constraint for the relational contract is satisfied whenever the price in the relational contract, p^{RC}, is lower than the price charged by market suppliers who hold up the buyer, p^F; I assume that there is no additional cost to the buyer of including demand assurance ν^{RC} (although later relax this assumption). As the buyer has no profitable static deviations, the buyer prefers the relational contract in each period. On the other hand, the seller has a profitable static deviation: hold up the buyer to make it pay p^F while having the low costs associated with ν^{RC} for the current period. Therefore, the seller with discount factor δ prefers the relational contract as long as:

$$\begin{aligned}(1-\delta)\Big(p^{RC}-c^S(q,\nu^{RC})\Big)+\delta\Big(p^{RC}-c^S(q,\nu^{RC})\Big) &\geq (1-\delta)\Big(p^F-c^S(q,\nu^{RC})\Big)\\ &\quad+\delta\Big(p^F-c^S(q,\nu)\Big)\\ \frac{\delta}{1-\delta}\Big(c^S(q,\nu)-c^S(q,\nu^{RC})+p^{RC}-p^F\Big) &\geq p^F-p^{RC}\end{aligned}$$

For the buyer to prefer the relational contract, it must be that $p^{RC} \leq p^F$. The surplus gains from the relational contract in terms of cost reduction are $c^S(q,\nu)-c^S(q,\nu^{RC}) \equiv \Delta c^S$. Therefore, incentive compatibility for the buyer requires relational contract prices to be $p^{RC} = p^F - \phi\Delta c^S$ for some $\phi \geq 0$. Substituting for p^{RC}, the equation becomes:

$$\frac{\delta}{1-\delta}\Big(\Delta c^S + p^F - \phi\Delta c^s - p^F\Big) \geq p^F - (p^F - \phi\Delta c^S)$$

This equation simplifies to:

$$\delta \geq \phi$$

This result highlights that there exists a relational contract equilibrium that resolves the hold-up problem. The discount provided in the relational contract relative to the hold-up price depends upon the supplier's discount factor and other factors that shape ϕ, which could include bargaining parameters over surplus. As supplier impatience increases, the more tempting it is for the supplier to deviate, restricting the discount in the relational contract. The discount also increases in the cost reduction from demand assurance in the relational contract, as for the same ϕ the discount is larger when Δc^S is larger. In Appendix C.2, I consider an extension where there are costs to the buyer of providing $\nu^{RC} < \nu$; with this extension, there is a lower bound on the discount factor for both the buyer and the

seller, rather than only for the seller.

2.6 Policy Implications and Conclusion

The substitutability between vertical integration and relational contracts naturally raises questions about whether relational contracts have policy implications similar to vertical integration. I consider the implications of relational contracting for two common policy issues associated with vertical integration: foreclosure (and raising rivals' costs more broadly) and double marginalization.

Relational contracts, in principle, have the potential to create economic conditions leading to foreclosure and raising rivals' costs, although concerns, in general, are reduced compared to integrated firms. First, relational contracts have the potential to be anticompetitive explicitly by design. For example, the relational contract could specify that the supplier either not work with other buyers or charge them higher prices; as long as such behavior creates additional surplus in the relational contract (and is public information within the relationship), the relational contract could specify such behavior.

Even without anticompetitive behavior embedded in the relational contract, relational contracts may indirectly cause foreclosure or raise rivals' costs through two separate channels. First, relational contracts may distort quantities relative to the market equilibrium, as prices in relational contracts are often different from market prices. Then,if supplier marginal costs increase in quantities, relational contracts can result in higher input prices for other buyers as the supplier passes on some of the cost increase to other buyers. Second, the relational contract could reallocate production across suppliers in two ways that could lead to higher costs for rivals even without market quantities changing: *i)* the relational contract supplier could be positively selected with respect to costs and have capacity constraints such that the cost distribution of suppliers available for other buyers to trade with is higher; and *ii)* the residual capacity of the one relational contract supplier for other buyers in the market could be lower, effectively making the input market less competitive. Intuitively, this concern of foreclosure closely matches the intuition that relational contracts can be anticompetitive by effectively locking parties into their relationship and preventing new buyer-supplier links from forming as in Board (2011).

Importantly, both of these concerns depend on the extent to which investments are relationship specific. When investments are perfectly relationship specific, then the supplier effectively is creating a separate set-up for the buyer and the general impact of the relational contract on other buyers in the input market should be minimal. In fact, other buyers could even benefit from the relational contract's existence in settings with economies of scope be-

cause the production of the relationship specific input has spillovers that reduce costs for the non-specific input.[8] However, there are many settings, including the empirical setting of fabric markets in this paper, where relational contracts trade quantity assurance for pricing and the good traded (or at least the capacity and physical and human capital for producing those goods) is sufficiently general that concerns of foreclosure are relevant. For example, there are relational contracts for ice (Ghani and Reed, 2022), coffee cherries (Macchiavello and Morjaria, 2020), and roses Macchiavello and Morjaria (2015), all of which seem sufficiently general for concerns about foreclosure and raising rivals' costs to apply. In the empirical setting in this paper, such concerns due to relational contracts are small as, even though the capacity part of the relational contract is fairly general (*i.e.*, non-specific), the fabric market is sufficiently competitive that it seems implausible for a decrease in a few suppliers to affect market prices.

Like vertical integration, relational contracts are also likely to eliminate double marginalization. Relational contracts can maximize joint surplus between the buyer and seller (as in Brugues (2020)), which occurs at the same quantity regardless of input price: once the relational contract governs behavior, the input price is just the transfer price for determining how surplus is shared rather than influencing what quantity the buyer sells in the final goods market. It is not possible to test this hypothesis in the context of this paper's empirical setting because the end client sets quantities before prices are determined rather than following the structure of common models of double marginalization.

Overall, this paper documents that vertical integration and relational contracts are substitutes by leveraging exogenous product-level variation in costs by transaction governance form using novel data that captures input flows. I rationalize these results through a simple model that shows how both vertical integration and relational contracts can resolve the hold-up problem motivated by the empirical setting, providing a microfoundation for the substitutability of vertical integration and relational contracts. The policy concerns raised by this substitutability remain an important avenue for future work, which can seek to identify settings with exogenous changes in the level of relational contracting in imperfectly competitive input markets to validate if (and when) relational contracts cause foreclosure and reduce double marginalization.

[8] If the production process, instead, features *dis*economies of scope, then even relationship specific investments can be anticompetitive through raising rivals' costs.

CHAPTER 3

Quality Incentives and Upgrading in Uganda's Coffee Supply Chain

3.1 Introduction

Quality upgrading to increase participation in Global Value Chains (GVC), and especially to attain quality premia offered on export markets, appeals to many policy makers in low- and middle-income (LMIC) countries. Quality upgrading not only has the potential for large direct effects on incomes through access to higher prices and value-added stages of production but it also can improve productivity in the long-run through knowledge and capital transfers (*e.g.,* as in Atkin et al. (2017a)). For example, the World Bank highlights the importance of "upgrading to higher-value intermediate processes and moving to higher-value-added products" (The World Bank, 2020) and the World Trade Organization emphasizes quality upgrading is necessary to appeal to multinational companies that "rely on their suppliers to meet quality standards" (Park et al., 2023). Many actual policies in LMICs aim to support quality upgrading, as many markets (and firms) in LMICs fail to consistently provide high-quality versions of products naturally (see Verhoogen (2023) for a review on quality upgrading in LMICs). Examples abound in myriad diverse sectors, ranging from the Malaysian Government supporting the Penang Skills Development Center to spur human capital upgrading that facilitates working with technology used by exporters to the Thai government investing in capital and enforcing quality standards throughout the domestic aquaculture supply chain.

We study quality upgrading in the context of coffee supply chains in Uganda; coffee is one of the world's most widely traded agricultural commodities–over 7 billion kilograms of coffee were exported from 2023-2024, with Uganda the second largest producer of Arabica coffee (the highest quality coffee) in Africa and with the highest growth rate.[1] Export market

[1] https://apps.fas.usda.gov/psdonline/circulars/coffee.pdf

quality premiums are large for coffee–specialty coffee retail price index reports that retail prices of high-quality coffee are typically double low-quality coffee.[2] However, overall quality production is low in Uganda, as approximately 97% of total coffee production is low quality.[3]

We focus on the role of intermediation, motivated by evidence that intermediaries appear to capture much of the revenue from commodities as world prices do not reach up supply chains to farmers (Staatz et al., 1989; Coulter and Poulton, 1999; Osborne, 2005). Specifically, intermediaries may shape demand-side incentives for quality upgrading (*i.e.,* the quality premium, or increase in price for high-quality coffee) and quality investments throughout the supply chain. And, even outside coffee supply chains, prior research shows that intermediation is important in LMICs across many settings, such as in consumer goods in Nigeria (Grant and Startz, 2024), domestic agricultural markets in Kenya (Bergquist and Dinerstein, 2020) and Uganda (Bergquist et al., 2024), and Peruvian aquaculture (Hansman et al., 2020b).

First, we assemble novel data on the supply chain structure through a network mapping exercise. This network mapping exercise starts with all suppliers of the largest Ugandan coffee exporter from 2018 to 2022, and identifies their suppliers. This process repeats until the supply chain is traced up to the farm game. We find that 85% of coffee travels through an intermediary rather than directly from farmer to exporter (let alone directly to export markets). The average (and median and modal) domestic supply chain has three actors–a producer (equivalently, farmer or upstream) who sells to an intermediary (equivalently, trader or downstream) who sells to an exporter (equivalently, big company). Second, we collect original transaction-level data, including lab tests of quality, to measure how the quality premium evolves along the supply chain. Our analysis reveals that the quality premium diminishes by more than half for producers relative to intermediaries; however, a statistically significant quality premium exists throughout the supply chain.

What economic mechanisms explain the diminished quality premium upstream relative to downstream, and what are the efficiency, distributional, and policy implications of intermediation? We build a model that emphasizes two mechanisms. First, barriers to entry to high-quality intermediation, such as market access costs, both directly reduce demand and increase buyer market power in the high-quality segment of the market, both of which decrease prices upstream for high-quality coffee relative to low-quality coffee. Second, both upstream and downstream can make investments in quality upgrading (*e.g.,* 73% of upstream

[2] https://www.transactionguide.coffee/reports/specialty-coffee-retail-price-index-2023-q4

[3] Ugandan customs data do not provide transactions records with detailed quality parameters. Therefore, the statistic reported reflects interviews with several leading exporters and refers to all coffee exports (robusta and arabica). Arabica coffee exports account for roughly 20% of total coffee production (in the 2022/23 season).

and 70% of downstream report drying coffee at least once over the course of a season). This *productive substitutability* along the supply chain can also depress the quality premium, as efficient downstream quality upgraders are unwilling to pay a premium over their own low upgrading costs. Productive substitutability is presumably a general feature of main supply chains; for example, quality verification tasks can typically occur anywhere along the supply chain. But there are also settings outside of quality verification that feature substitutable quality investments even in contexts far removed from agricultural supply chains in LMICs; for example; downstream software engineers might not pay a premium for high-quality high-powered computer processing produced upstream when they can cheaply produce efficient and effective software that does not require much processing power.

Additionally, these mechanisms can interact to further dampen the quality premium upstream. When productive substitutability reduces entry into the high-quality segment upstream, then even fewer downstream purchase high-quality coffee from producers. Specifically, because buyer market power distortions are largest when there are few buyers, reductions in the number of downstream buyers of high-quality coffee due to productive substitutability exacerbate market power distortions due to barriers to entry alone.

Although both of these mechanisms–barriers to entry and productive substitutability–can explain the reduced quality premium, these two mechanisms have different implications for efficiency, distribution of surplus, and policies to increase quality upgrading. If the reduced quality premium largely reflects downstream efficiency (*i.e.,* productive substitutability), then efficiency costs in terms of overall production of high-quality coffee are likely small. In that case, downstream upgrading makes up for low levels of investment in quality upstream. However, market power due to barriers to entry distorts prices and quantities of high-quality coffee production through two channels–a standard market power distortion and an additional task distortion that applies only to supply chains. Specifically, when barriers to entry dampen the quality premium upstream, then upstream with upgrading costs between the reduced quality premium upstream and the export market quality premium do not upgrade even though it would be socially optimal for them to do so. By analyzing the separate importance of buyer market power and productive substitutability, this research speaks directly to policy-maker concerns about understanding whether intermediaries largely extract rent from or add value along the supply chain (as in Grant and Startz (2024)).

Barriers to entry and productive substitutability also have different implications for distribution of surplus along the supply chain. Barriers to entry create oligopsony rents for downstream, increasing their surplus relative to upstream; removing barriers to entry increases prices for both low and high-quality coffee upstream, reallocating surplus upstream. And, removing barriers to entry can even increase average surplus downstream by shifting

quality investments (although those who make the same quality investments will lose surplus due to price increases). On the other hand, productive substitutability has ambiguous effects on surplus distribution. Specifically, productive substitutability increases demand for low-quality coffee because efficient downstream will who upgrade also purchase large quantities of coffee, thereby increasing the price *level* for low-quality coffee. However, productive substitutability dampens the quality premium, decreasing the quality premium. Therefore, productive substitutability reallocates surplus from efficient to inefficient upstream agents. It also affects downstream surplus, both through changing input prices and task decisions, such that there are both winners and losers downstream.[4] Because upstream producers are typically among the poorest agents in the supply chain, distribution of surplus along the supply chain matters for poverty alleviation and welfare of the poorest.

Last, we loosely calibrate the model to highlight that these two mechanisms–barriers to entry and productive substitutability–can indeed generate the observed pattern of a diminished quality premium. Specifically, we use the model to quantify how these mechanisms could affect prices and the quality premium, buyer markdowns, efficiency (*i.e.*, production of coffee by quality), surplus distribution, and task allocation along the supply chain.

This paper contributes to a rich literature on quality upgrading in LMICs. Much of the existing literature studies supply-side constraints: firms may lack the information or know-how to improve their product (Bloom et al., 2012; Hanna et al., 2014; Atkin et al., 2017a), access to quality-enhancing equipment or other inputs due to credit market imperfections or trade barriers (Fan et al., 2015; Crinò and Ogliari, 2017); a skilled workforce to maintain and monitor quality (Atkin et al., 2017b); or competitive supply-side pressure to produce quality (Björkman Nyqvist et al., 2021). We focus on demand-side constraints, highlighting how the market and cost structure in domestic supply chains affects incentives for, and investments in, quality.[5]

We also advance a burgeoning literature that studies intermediation and market power in LMICs in two areas. First, we build upon prior literature that studies how intermediation affects prices (*e.g.,* as in Fafchamps and Hill (2008); Bergquist and Dinerstein (2020); Chatterjee (2023); Zavala (2023)) by turning to its effects on incentives for quality production

[4]Downstream who switch into upgrading when increasing productive substitutability generally are winners, as are downstream who continue to source and sell high-quality coffee given that productive substitutability decreases prices for high-quality coffee upstream. However, downstream who source and sell low-quality coffee lose surplus because demand from downstream upgraders increases prices for low-quality coffee.

[5]The limited exceptions focus on settings in which consumers do not reward quality, either because they cannot observe it (as in the case of Chinese watermelon markets (Bai, 2024)), or because they are unwilling to pay for it (as in the case of Ugandan maize markets (Bold et al., 2022)). We study a setting in which both quality is observable and there is demand – at least in export markets – for quality.

and investments in quality upgrading. Second, we highlight that quantitatively important fixed costs related to intermediation extend beyond trade costs (Atkin and Donaldson, 2015; Grant and Startz, 2024) but also can include barriers to entry for high-quality.

Additionally, this paper contributes to our understanding of the role of market power in supply chains in two areas, as prior work has established that buyer market power can affects prices, efficiency, and technology choices in diverse settings ranging from Chinese tobacco markets (Rubens, 2023), to firm-to-firm trade in France (Morlacco, 2020) and the U.S. (Alviarez et al., 2023), to labor markets in Peru (Amodio et al., 2023), the Netherlands (Treuren, 2023), and the U.S. (Azar et al., 2022; Rubens, 2024). Specifically, we show that market power in supply chains can cause not only the standard oligospsony distortion but can also decrease efficiency by creating task misallocation within a supply chain.

This project also relates to work that studies how competition affects quality provision. Both experimental (and quasi-experimental) and theoretical work highlight the ambiguous effects of competition on quality. For example, while Bennett and Yin (2019); Busso and Galiani (2019); Björkman Nyqvist et al. (2021) illustrate that competition improves quality, Macchiavello and Morjaria (2020) show the opposite can occur. Theoretically, Gaynor (2006) emphasizes that firms with market power set quality to correspond to the willingness-to-pay (WTP) for quality of the marginal consumer, whereas competitive markets set quality to the WTP of the average consumer. And structural work indeed shows evidence that market power can improve quality, as Goettler and Gordon (2011) highlight that market power allows firms to extract more of the surplus from quality upgrading and Bodéré (2023) illustrates that increased entry can reallocate market shares to lower-quality firms when buyers have preferences over many product features. This paper provides an additional microfoundation for the ambiguous effects of competition, as market power in our model both reduces quality provision conditional on entry but increases entry into the high-quality segment as firms benefit from the increased profits created by market power in the high-quality segment.

This project also relates to a large trade literature studying global supply chains and how tasks are allocated across countries and firms. While this literature mostly focuses on understanding how relative advantage and contracting frictions influence the supply chain structure (Antràs and Chor, 2013; Costinot et al., 2013; Caliendo and Parro, 2015; Alfaro et al., 2019; Boehm and Oberfield, 2020; Chor and Ma, 2020), we emphasize the role of market and cost structure in explaining task allocation and surplus distribution along the chain.

3.2 Data and Context

3.2.1 Data: Network Mapping

We construct a novel dataset on the supply chain structure through a network mapping exercise. We work with the largest Ugandan coffee exporter from 2018-2022 and survey all of their suppliers on their activity in coffee markets, focusing on their coffee production and purchasing behavior.[6] We then interview each respondent's major suppliers, if any.[7] A chain is considered complete once the agent does not have any additional suppliers.

We then select the most prominent chains to include in our study sample for additional data collection (described below) such that our sample presumably includes the agents along the supply chain that are most likely to increase quality production; these members of the supply chain who are on the margin of quality upgrading are a highly policy relevant group as policies supporting quality upgrading will mostly affect this group and the success of such policies depends on the responses of these agents. Specifically, to find the most prominent chain, we always select the supplier with the largest share of volume (randomly breaking ties). We select 199 chains to include the study–one chain for each agent that sells directly to the big exporter.

3.2.2 Data: High-Frequency Sampling Surveys

We collect original transaction-level data throughout the supply chain to study quality incentives and investment throughout the supply chain; no administrative data (either firm or public) record activity along the supply chain at this level. To do so, we visit intermediaries every two weeks and collect samples of their three largest sales and purchases. As shown in Figure 3.1, respondents record information about the transaction date, price, quantity, and counterparty. Then, we bring the coffee to a lab where we assess its quality. We use a similar process for producers, collecting a sample for the three largest harvests for each producer over the course of the season. To ensure that our data collection captures a large share of overall transactions, we also ask about other transactions beyond the three largest. At the respondent level, on average, we collect these detailed data on 81% of purchased quantities and 86% of sold quantities.[8] Statistics are similar looking at all transactions (*i.e.,* weighting

[6]Coffee production occurs throughout the supply chain, with agents that mostly buy and sell coffee still typically having some production. Therefore, to divide the sample, we define agents purchasing no more than 25% of their coffee sales as producers and the other agents as intermediaries.

[7]In practice, we ask for the largest suppliers until we reach 25% of total supply, although stopping at 10 suppliers even if have not reached 25% of total supply.

[8]These data are only available for intermediaries; however, as producers tend to have vary few transactions, we likely captured an even higher percent of their transactions. Therefore, these statistics are likely lower

by volume rather than by respondent), with our high-frequency sampling data capturing 73% of purchases and 86% of sales volumes.

Figure 3.1: Transaction Data Example



3.2.3 Data: Producer and Intermediary Surveys

In addition to network mapping data and high-frequency sampling surveys, we collect data from producers and intermediaries before the season starts about their typical economic activity around coffee. Questions from the baseline include: an incentivized quality observability test; information about prices, quantities, and profits by quality for transactions from the prior season (August to December, 2021) and expectations for them over the upcoming season (August to December, 2022); coffee production practices and activities (both harvest and post-harvest); capital assets used for coffee production; and knowledge about quality production.

bounds for the total percent of quantity included in our data.

3.2.4 Context: Coffee Quality and Production in the Rwenzori Region of Uganda

Over 90% of unwashed arabica coffee (locally known as Drugar) is produced in the Western region of Uganda. In the main crop season of 2022-2023 (*i.e.,* from roughly September to December) close to 120,000 60-kg bags of drugar were produced in Uganda – of which close to 95% were produced in the Rwenzori region. About half of these total bags were directly purchased by large exporters, with the remaining were exchanged through various intermediaries and traders, presumably going to Kampala and then to other countries that purchase low-quality coffee internationally (*e.g.,* South Africa).

Coffee quality in the Rwenzori region is closely tied to the physical properties of the green coffee beans that big exporters ship to world markets.[9] The three physical quality measures are: *i)* moisture; *ii)* defects; and *iii)* foreign matter. Moisture refers to the moisture level of the coffee beans, with overly moist coffee turning moldy and unusable for consumption (as well as being difficult to roast properly). Defects refer to coffee cherries that are defective for any reason, such as underripe, insect-damaged, or moldy. Foreign matter is the amount of non-coffee, such as twigs, pebbles, and dirt, that are included in the coffee. All three of the physical properties are measured in weight and are easily observable.[10]

We use a standard for high-quality coffee employed by the largest coffee exporter in Uganda from 2018-2022. Specifically, this exporter requires moisture to be no more than 14.5%, defects to be no more than 33%, and foreign matter to be no more than 4%. Based on our understanding, the importance of this threshold reflects the exporter's ability to export coffee to international buyers who impose strict binarized quality standards.

[9] As roasted coffee only remains in optimal condition for roughly two weeks, most exported coffee is unroasted with roasters local to consumers typically providing the most premium options.

[10] Defects and foreign matter are easily visible–it does not take long to be able to visually identify which beans are defective and it is even easier to distinguish between coffee and non-coffee. Moisture is also observable (although likely not the exact percentage), as participants perceive signals of moisture from the color of the beans and how tough they are between their teeth.

Table 3.1: Summary Statistics: Producers

	Average	Median	SD
Age	43.96	42	13.05
Profits per Season (1000 UGX)	1789.71	750	2855.42
Number of Buyers	1.55	1	1.34
Sales Quantity (kg)	760.56	75	2231.18
Sale Price (UGX)	9235	9048.33	1540.76
Sales Share of High-quality Coffee (%)	27.86	0	36.23

Note: Data from the networking mapping (number of buyers), high-frequency sampling surveys (quantities, prices, and share of high-quality coffee), and the producer surveys (age and profits per season). Averages and median are at the respondent level.

Table 3.2: Summary Statistics: Intermediaries

	Average	Median	SD
Age	43.99	42	11.51
Profits per Season (1000 UGX)	2559.89	1340	3458.54
Number of Suppliers	19.94	7	27.8
Purchase Quantity (kg)	1155.69	650	1373.86
Purchase Price (UGX)	8262.09	8336.36	1024.57
Purchase Share of High-quality Coffee (%)	12.17	0	20.6
Number of Buyers	.91	0	1.44
Sales Quantity (kg)	1317.38	467	2216.77
Sale Price (UGX)	9196.76	9000	1248.84
Sales Share of High-quality Coffee (%)	19.95	5.59	28.88

Note: Data from the networking mapping (number of suppliers and buyers), high-frequency sampling surveys (quantities, prices, and share of high-quality coffee), and the intermediary surveys (age and profits per season). Averages and median are at the respondent level.

Tables 3.1 and 3.2 provide descriptive information about producers and intermediaries in our sample, combining data across the networking mapping, high-frequency sampling surveys, and the producer and intermediary surveys. Four key patterns emerge. First, intermedaries operate at a larger scale than producers–the median intermediary purchases

nine times as much as the median producer harvests and sells. Second, we replicate a stylized fact in the literature (*e.g.,* as in Dominguez-Iino (2023)) that LMIC agricultural value chains are "seive shaped," meaning that they are more atomistic upstream. Specifically, we observe that the median intermediary sources from many different suppliers (7) whereas the median producer sells to only one intermediary. Third, the median intermediary has roughly twice the profit of the median producer (although the profit ratio is much smaller than the quantity ratio, which suggests that intermediaries might specialize in more costly activities and/or use most costly production processes for the same activities). Last, quality investments occur throughout the supply chain, as upstream sells high-quality coffee and the share of high-quality coffee sold downstream is higher than the share of high-quality purchased, which means that downstream must make quality investments.

3.3 Intermediation, Quality Incentives, and Quality Upgrading

3.3.1 Intermediation

First, we study the empirical relevance of intermediation using our network mapping data. We have two key findings. First, the vast majority–76%–of supply chains include intermediation between the farm gate and the big exporter (and the big exporter itself could reasonably be considered as an additional intermediary). The average (and modal and median) supply chain from farm gate to the export gate has a single intermediary, although there are some farmers that sell directly to the exporter (24%) and the distribution of chains include a small right tail all the way up to seven chains (see Appendix D.1 for more details). Furthermore, as longer chains tend to process larger volumes, on average, an even higher percent of coffee in terms of quantity–85%–passes through an intermediary before reaching the big exporter.

3.3.2 Quality Incentives: The Quality Premium

Second, we use our transaction data from the high-frequency sampling surveys to explore how the quality premium varies along the supply chain. Specifically, we compare prices by quality as a function of supply chain position, where downstream are agents who sell directly to the exporter in the network mapping survey and upstream are agents whose output travels through an intermediary before reaching the exporter (based on the network mapping data). We estimate a regression of price for transaction i in subcounty s in month t on indicators for high quality, sale by downstream, and their interaction, with controls for quantity and

subcounty-by-month fixed effects (*i.e.,* each subcounty-month has a separate fixed effect):

$$\begin{aligned} Price_{i,s,t} = \beta_1 * \text{High Quality}_{i,s,t} + \beta_2 * \text{Downstream}_{i,s,t} \\ + \beta_3 * \left(\text{High Quality}_{i,s,t} * \text{Downstream}_{i,s,t} \right) \\ + \gamma * \text{Quantity}_{i,s,t} + \mu_{s,t} + \epsilon_{i,s,t} \end{aligned} \tag{3.1}$$

In this regression, β_1 represents the quality premium upstream, β_2 represents the average price increase downstream regardless of quality, and β_3 represents the difference in the quality premium downstream (such that the total quality premium downstream is $\beta_1 + \beta_3$). Subcounty-by-month fixed effects are represented by $\mu_{s,t}$ and the residual is $\epsilon_{i,s,t}$. Standard errors are clustered two-ways at the subcounty-by-week and respondent levels.

Table 3.3: Diminished Quality Incentives Upstream

	Price
High quality	244.37**
	(95.46)
Downstream	375.71***
	(104.39)
High quality * Downstream	266.47*
	(149.37)
Observations	1298
Mean DV	8275
Quantity control	Yes
Month X SC FE	Yes
Individual FE	No

Note: Price and quality data from the sampling surveys and downstream is defined based on whether a respondent sells directly to the big exporter in the networking mapping survey. Coefficients are from the regression in 3.1. Standard errors in parentheses are clustered at both the subcounty-by-week and respondent levels. ***, **, and * indicate significance at the 1, 5, and 10 percent levels, respectively.

Note: Price and quality data from the sampling surveys and downstream is defined based on whether a respondent sells directly to the big exporter in the networking mapping survey. Coefficients are from the regression in 3.1. Standard errors in parentheses are clustered at both the subcounty-by-week and respondent levels. ***, **, and * indicate significance at the 1, 5, and 10 percent levels, respectively.

Table 3.3 reports results from the regression in equation 3.1. Estimates illustrate that the quality premium is more than halved upstream (244.37 UGX per kg) relative to downstream (510.84 UGX per kg); however, there still is a statistically significant quality premium upstream. Additionally, the regression analysis highlights that downstream receives higher prices regardless of quality.

We can also measure how the quality premium evolves along the supply chain leveraging variation within respondent, comparing purchases and sales. Specifically, analyses can study whether the quality premium *within respondent* is different for upstream as compared to downstream agents in the supply chain. In other words, we can see if the sales price for high-quality coffee is differentially larger than the purchase price for downstream relative to upstream (rather than comparing the quality premium by layer leveraging variation across respondents, as in 3.3). To do so, we estimate a regression of price for transaction i in subcounty s in month t on indicators for high quality, an indicator for a sale (rather than a purchase) for any layer, the interaction of high quality and the indicator sale for any layer, an indicator for a sale (rather than a purchase) for downstream only, and the interaction of high quality and the indicator for a sale for downstream only, with individual fixed effects as well as controls for quantity and subcounty-by-month fixed effects (*i.e.,* each subcounty-month has a separate fixed effect):[11]

$$\begin{aligned} Price_{i,s,t} = {} & \iota_1 * \text{High Quality}_{i,s,t} \\ & + \iota_2 * \text{Sale by Any Layer}_{i,s,t} + \iota_3 * \left(\text{High Quality * Sale by Any Layer}_{i,s,t} \right) \\ & + \iota_4 * \text{Sale by Downstream}_{i,s,t} + \iota_5 * \left(\text{High Quality * Sale by Downstream}_{i,s,t} \right) \\ & + \mu_i + \gamma * \text{Quantity}_{i,s,t} + \mu_{s,t} + \epsilon_{i,s,t} \end{aligned} \tag{3.2}$$

In this regression model, the key coefficients of interest related to the quality premium are ι_1, ι_3, and ι_5. The quality premium upstream, regardless of whether or not it is a purchase or sale, is represented by ι_1. Heterogeneity in the quality premium when selling is captured

[11] In a regression within individual fixed effects, the subcounty-by-month fixed effects still allow months to have heterogeneous effects across subcounties even though the subcounty fixed effect itself is absorbed by the individual fixed effect.

by ι_3 and ι_5. Specifically, ι_3 captures whether sales by upstream have a different quality premium than purchases and ι_5 measures whether sales by downstream have a different quality premium than purchases. In this regression, we cannot analyze the direct effect of downstream on prices because it is collinear with the individual fixed effects.

Table 3.4: Diminished Quality Incentives Upstream: Within Respondent Variation

	Price
High quality	180.20*
	(108.56)
High quality * Sale	-177.70
	(165.65)
High quality * Sale * Downstream	266.86*
	(158.90)
Observations	1291
Mean DV	8270
Sale control	Yes
Sale * Downstream control	Yes
Quantity control	Yes
Month X SC FE	Yes
Individual FE	Yes

Note: Price and quality data from the sampling surveys and downstream is defined based on whether a respondent sells directly to the big exporter in the networking mapping survey. Coefficients are from the regression in 3.2. Standard errors in parentheses are clustered at both the subcounty-by-week and respondent levels. ***, **, and * indicate significance at the 1, 5, and 10 percent levels, respectively.

Table 3.4 reports results from the regression in equation 3.2. Estimates illustrate that the quality premium is generally 180.20 UGX per kg, which is not statistically different than the estimated average quality premium in 3.3 that leverages comparisons across rather than within respondents. The results also suggest that there is not statistically significant difference in quality premium for purchases and sales within upstream, as the coefficient on ι_3 is not statistically significant. However, the regression highlights that downstream sales have a quality premium that is more than double other transactions at 447.06 UGX per kg, specifically with an increase in quality premium for downstream sales that is almost identical to that in 3.3.

Overall, we show that the result of a diminished (although still significant) quality premium for upstream compared to downstream is robust to leveraging different sources of variation for price comparisons.

3.3.3 Quality Investments

Table 3.5: Transactions by Layer: Quality Investments Along the Chain

	Purchases (kg)		*Sales (kg)*	
Supply Chain Layer	Low Quality	High Quality	Low Quality	High Quality
Downstream	980.1	352.6	1127.9	619.1

Note: Data from sampling surveys. Quantites are not shown for upstream because some supply chains are longer than two (*i.e.,* there are multiple upstream), which would mean that some transactions could be double counted as the coffee flows through the supply chain within upstream agents.

Table 3.5 shows the average purchases and sales per active downstream agent over the course of the season. The table highlights that quality upgrading occurs roughly equally upstream and downstream. Specifically, the quantity of high-quality coffee purchased is roughly half (57%) of high-quality coffee sold, which means that downstream must upgrade the other coffee. Note that the sum of purchases is less than the sum of sales; this gap reflects that downstream tend to also produce some coffee from their own farms (even though the volumes highly that most sales are from aggregation–76%).

3.4 Barriers to Entry, Productive Substitutability, and the Diminished Quality Premium Upstream

We focus on two economic mechanisms, explained in detail below, that can each individually explain the diminished quality premium: barriers to entry to high-quality intermediation and productive substitutability, and show empirical patterns consistent with the importance of these two mechanisms in our empirical setting. Importantly, these two mechanisms need not be mutually exclusive; not only can both mechanisms directly dampen the quality premium upstream relative to downstream, but they can even interact to further reduce the quality premium. Additionally, these mechanisms have different implications for efficiency and distribution, which is relevant as reducing poverty and increasing the welfare of the poorest generally focuses on surplus upstream.

3.4.1 Barriers to Entry to High-Quality Intermediation

Barriers to entry to high-quality intermediation can reduce the quality premium upstream through two channels. First, barriers to entry can lead to fewer downstream purchasing high quality from upstream, which results in fewer competitors for high-quality coffee produced upstream. This reduction in competition results in more buyer power in the high-quality segment, which depress prices for high-quality coffee.

Empirically, barriers to entry also help rationalize low observed equilibrium levels of entry into high-quality intermediation despite seemingly large arbitrage opportunities for downstream to purchase high-quality coffee upstream, where the quality premium is reduced, and sell it to the big exporter at a larger quality premium.[12] Specifically, we find that only one-sixth of active downstream in our empirical setting sell high-quality coffee.

While these barriers to entry may have many different microfoundations, a compelling candidate (for at least part of the cost) is market access to the relatively few (~ 5) exporters that offer a quality premium, which could reflect these exporters having relationships with international buyers of quality or having prior investments in reputation essential for commanding a quality premium. Empirical patterns in D.2 are consistent with this interpretation. Specifically, comparisons of prices based on buyer identity illustrate that big exporters almost triple the quality premium relative to other buyers. Furthermore, analysis of access to big exporters highlights that downstream have superior access both overall and especially for high-quality coffee; compared to upstream, downstream selling high-quality are 26 percentage points more likely to sell to a big company (compared to a mean of only 5% of sales to big companies).

The buyer power attributable to barriers to entry both negatively affects efficiency and redistributes surplus to downstream, which may be undesirable if policy makers prioritize poverty reduction and lifting the poorest individuals out of poverty. Negative effects on efficiency arise from two distinct channels. First, buyer power induces a standard oligopsony distortion, reducing prices and quantities upstream of high quality coffee. Second, buyer power also induces *task* misallocation. Namely, upstream who would upgrade at the quality premium provided by the exporter may not do so at a quality premium that is reduced due to buyer power. Barriers to entry redistribute surplus to downstream (in addition to creating dead-weight loss by decreasing efficiency) because buyer power results in markdowns. These

[12] Based on regression results in 3.3, there is arbitrage opportunity of purchasing high-quality coffee upstream and selling it downstream at 620 UGX per kg, compared to only 375 UGX per kg for purchasing and selling low-quality coffee, for a gain of 245 UGX per kg. Leveraging variation within respondent, as in 3.4, the estimated arbitrage opportunity is 267 UGX per kg. Although these calculations do not incorporate costs, if aggregation costs for downstream are independent of quality, there are large unrealized arbitrage opportunities that could be rationalized by barriers to entry.

markdowns increase surplus for downstream at the expense of upstream.

3.4.2 Productive Substitutability

Productive substitutability can also decrease the quality premium upstream by reducing downstream willingness-to-pay for high-quality coffee upstream. Specifically, when downstream can upgrade at a lower cost than the quality premium upstream, downstream's profit-maximizing task decisions may be to purchase and upgrade low-quality coffee to produce high-quality coffee rather than to aggregate coffee of either quality.

Empirically, both upstream and downstream engage in essential post-harvest quality investments at similar rates. Comparing over the entire season, 73% upstream dry coffee, but so do 70% of downstream. Similarly, 35% upstream but also 39% of downstream sort coffee. Furthermore, analysis of the quality of purchases and sales by downstream in Table 3.5 highlights that roughly half of upgrading (in terms of volume of coffee) occurs at each layer. And although levels of engagement in quality tasks are similar, costs may be different both across and within layer. For example, costs for drying can reflect both the opportunity cost of time to dry coffee as well as prior capital investments. While opportunity cost of time is likely higher for downstream as they are engaged in more economic activity (and, indeed, coordinating times with them for interviews was more challenging!), they may also have superior productive capital for drying coffee–for example, using rakes rather than sticks and drying coffee in wooden tubs elevated from the ground rather than on a tarpaulin on the floor, improving airflow on both sides and reducing risks of increased foreign matter.

While productive substitutability alone does not have any negative effects for efficiency of the supply chain (and actually increases the quantity of high-quality coffee produced overall, *ceteris paribus*), it does shift quality investments (*i.e.,* upgrading tasks) from upstream to downstream by decreasing the quality premium upstream (the "price mechanism"). And, these changed quality investments effectively reallocate surplus to downstream–upstream would have upgraded at the higher quality premium that would occur without productive substitutability gain from upgrading (otherwise they would not change quality production).

3.4.3 Interaction of Barriers to Entry and Productive Substitutability

While barriers to entry and productive substitutability are not mutually exclusive and can both contribute separately to the decreased quality premium upstream, they can also interact to affect the quality premium, efficiency, and distribution of surplus along the chain. Specifically, there are two interactions to highlight. First, productive substitutability can amplify

buyer power distortions created by barriers to entry. Specifically, if some downstream who would have aggregated high quality instead upgrade low quality, then productive substitutability further reduces the number of buyers in the high-quality segment. And, because the *marginal* entrant reduces buyer power more when the *level* of entry is lower, productive substitutability can play an important role in influencing how much buyer power is present in the high-quality segment in settings where the level of high-quality aggregators downstream is low. Therefore, although productive substitutability is not in and of itself inefficient, it can contribute to inefficiency by exacerbating market power distortions through reducing competition for high-quality coffee produced upstream.

3.4.4 Alternative Explanations for a Diminished Quality Premium Upstream

Appendix D.3 evaluates alternative mechanisms that could explain the diminished quality premium. First, we show that quality unobservability is unlikely to be important in this setting by highlighting that study respondents perform very well in an incentivized quality ranking test. Second, we show that relational contracts between upstream and downstream do not appear to empirically relevant in this setting and that quality premiums are similar for respondents most likely to be in relational contracts, which is inconsistent with relational contracts providing insurance (despite evidence of the prevalence of relational contracts in other coffee setting, such as Macchiavello and Morjaria (2020) study in Rwanda).[13] Last, we show that the diminished quality premium (and prices) are unlikely to reflect economically important matching along the chain by studying the average number of distinct buyers as a function of the number of sales transactions by upstream. We find that upstream sell to many distinct buyers, which is inconsistent with matching because matching should lead to sales concentrated among just a few buyers.[14]

3.5 Supply Chain Model of Quality Investments

We then build a parsimonious model to demonstrate how both barriers to entry and productive substitutability can generate the empirically observed reduction in the quality premium

[13]If relational contracts provide insurance against stochastic risk in quality production, then the quality premium should be flatter in relational contracts but with higher prices overall. Intuitively, upstream effectively buys quality insurance from downstream to ensure more stable prices (which may be efficient if downstream have higher risk tolerance).

[14]Another seemingly alternative hypothesis is that search costs drive results; however, through the lens of our model, search costs are grouped with other costs. Therefore, we do not consider search as a possible alternative, distinct mechanism.

upstream. The model takes the supply chain structure seriously, meaning that upstream cannot sell directly to export markets; in other words, coffee must flow through the supply chain. To incorporate barriers to entry, we include a fixed cost for high-quality intermediation. We model agents making endogenous quality investments along the chain (*i.e.,* choosing whether or not to upgrade the coffee); these investments capture productive substitutability because downstream can sell high-quality coffee for the same price to exporters (assuming they overcome the barriers to entry) regardless of which layer of the chain made the quality investment. As a benchmark case, we study the modal, average, and median supply chain: coffee flows from upstream to downstream to the exporter.

3.5.1 Upstream Profit Maximization

Upstream choose the quantity of both low-quality coffee (q_i^L) and high-quality coffee (q_i^H), where q_i^T represents the total quantity of coffee produced across both qualities. They choose quantities to maximize their profits given upstream prices for low-quality coffee (p^L) and high-quality coffee (p^H). Therefore, upstream profit maximization is:

$$\begin{aligned}\max_{q_i^L, q_i^H} &= p^L q_i^L + p^H q_i^H \\ &\quad - \alpha_1^U q_i^T - \alpha_{2,i}^U (q_i^T)^2 - \beta_i^U q_i^H \\ \text{Where: } \alpha_{2,i}^U &\sim \mathcal{N}(\bar{\alpha_2}^U, \sigma^2_{\alpha_2^U}), \quad \beta_i^U \sim \mathcal{N}(\bar{\beta}^U, \sigma^2_{\beta^U})\end{aligned}$$

Therefore, upstream chose to upgrade whenever $p^H - p^L > \beta_i^U$. The notation illustrates that upstream has no seller market power; as the more atomistic layer of the supply chain (as described in 3.2.4), they are not modelled as having any seller power (as in Dominguez-Iino (2023)).

In this model, α_1^U and $\alpha_{2,i}^U$ capture costs that upstream incur regardless of quality, such as the cost of supplying their own labor to harvest coffee cherries from trees. Upstream has increasing marginal costs, which can reflect, for example, capacity constraints related to picking cherries from marginal trees or even extra costs for renting coffee trees to take cherries. Upstream also has upgrading costs that shift the marginal cost of producing high-quality coffee relative to low-quality coffee and can reflect the additional labor and effort to, for example, sort coffee to upgrade quality and store it in a safe, dry location. Random coefficients for both $\alpha_{2,i}^U$ and β_i^U capture unobserved heterogeneity within upstream. This heterogeneity can reflect stochastic features of the environment–such as drier weather that makes drying coffee less costly–as well as other changes in the productive environment, such

as positive labor supply shocks. Additionally, this heterogeneity rationalizes why empirically different quantity levels and quality investments are observed. We do not microfound or distinguish between these sources.[15]

3.5.2 Downstream Profit Maximization

Similarly, downstream also choose the quantity of both low-quality coffee (q_i^L) and high-quality coffee (q_i^H), where q_i^T represents the total quantity of coffee produced across both qualities. These quantity choices map to downstream choosing between three tasks: *i)* aggregating (for purchase and sale) low-quality coffee; *ii)* purchasing low-quality coffee and upgrading it to high-quality coffee; or *iii)* aggregating high-quality coffee. Therefore, downstream profit maximization solves, taking world market prices ($\tilde{p}^L$ for low-quality coffee and $\tilde{p}^H$ for high-quality coffee) as given:

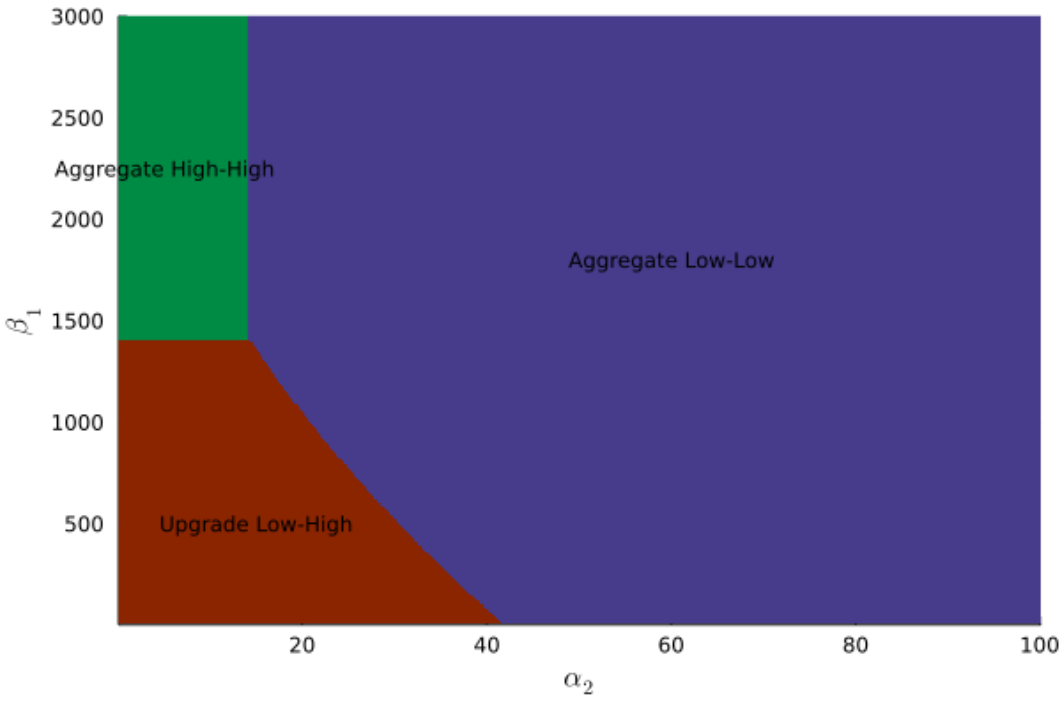
$$
\begin{aligned}
\max_{\mathrm{q}_i^L, \mathrm{q}_i^H} &= \tilde{p}^L \mathrm{q}_i^L + \tilde{p}^H \mathrm{q}_i^H \\
&\quad - \alpha_1^D \mathrm{q}_i^T - \alpha_{2,i}^D (\mathrm{q}_i^T)^2 - p^L(\mathrm{Q}^L, \mathrm{Q}^H)\mathrm{q}_i^L \\
&\quad - 1(\mathrm{q}_i^H > 0)\Bigg(F + \min\{\underbrace{\beta_i^U \mathrm{q}_i^H}_{\text{upgrade low-high}}, \underbrace{(p^H(\mathrm{Q}^L, \mathrm{Q}^H) - p^L(\mathrm{Q}^L, \mathrm{Q}^H))\mathrm{q}_i^H}_{\text{aggregate high-high}}\}\Bigg) \\
\text{Where: } & \alpha_{2,i}^D \sim \mathcal{N}(\bar{\alpha_2}^D, \sigma^2_{\alpha_2^D}), \quad \beta_i^D \sim \mathcal{N}(\bar{\beta}^D, \sigma^2_{\beta^D}) \\
& \mathrm{Q}^L = \sum_{i \in \mathcal{D}} \mathrm{q}_i^L, \quad \mathrm{Q}^H = \sum_{i \in \mathcal{D}} \mathrm{q}_i^H
\end{aligned}
$$

Downstream activity choices depend upon the joint realization of both $\alpha_{2,i}^D$ and β_i^D. Downstream with more advantageous (*i.e.*, lower) draws of $\alpha_{2,i}^D$ are more efficient producers; because their optimal quantity of coffee increases, their higher production makes it worthwhile to pay the fixed cost to enter high-quality intermediation to benefit from the higher world price. Downstream choice of upgrading depends upon realizations of $\alpha_{2,i}^D$ and β_i^D and upstream prices. Specifically, some downstream with low β_i^D may not upgrade, even if $\beta_i^D < \tilde{p}^H - \tilde{p}^L$, when $\alpha_{2,i}^D$ is sufficiently high such that they do not produce large enough quantities to make it worthwhile to upgrade. Importantly, this model captures productive substitutability: either downstream or upstream can perform the quality investment by upgrading coffee. The notation highlights that downstream has no seller market power; as the more atomistic layer of the supply chain relative to exporters (as described in 3.2.4), they

[15] In a reduced-form manner, this heterogeneity may also represent the (stochastic) effects of heterogeneous pre-harvest investments.

are not modelled as having any seller power (again as in Dominguez-Iino (2023)). However, downstream does have buyer market power, as their choices of q_i^L, q_i^H affect Q^L, Q^H and, thereby, affect prices. This buyer power reflects that downstream is less atomistic than upstream (also as in Dominguez-Iino (2023)).

Figure 3.2: Downstream Activity Choices and Cost Function Realization



Note: The heatmap shows downstream activity choices for a given price vector as a function of $\alpha_{2,i}^D$ and β_i^D.

Figure 3.2 illustrates downstream activity choices for a price vector of exporter and upstream prices, highlighting how both parameters $\alpha_{2,i}^D$ and β^D affect choices.[16] Note that the cutoff where agents switch to upgrading low-high occurs at the upstream price quality premium (not explicitly shown); only when β^D us below the upstream market premium is it worthwhile to upgrade. Additionally, because downstream with low upgrading costs have large inframarginal gains from upgrading, they will sell high-quality coffee even when they are less efficient overall (*i.e.,* have a higher draw of $\alpha_{2,i}^D$).

In this model, α_1^D and $\alpha_{2,i}^D$ capture costs that downstream incur regardless of quality, such as the search, storage, and transportation from farm gate to stores where downstream hold

[16]For simplicity, figure 3.2 does not include buyer market power considerations in making choices. In general, buyer market power has ambiguous effects on activity choices. On the one hand, buyer power increases profits from aggregating high-high, which incentivizes entry. On the other hand, buyer power reduces the optimal quantity of high-quality coffee chosen, which makes it less worthwhile to pay the fixed cost.

their coffee. Downstream has increasing marginal costs, which can capture that, for example, sourcing additional coffee may require reaching out to upstream that are more remote. Downstream also has upgrading costs that shift the marginal cost of producing high-quality coffee relative to low-quality coffee and can reflect the additional labor and effort to, for example, sort coffee to upgrade quality and store it in a safe, dry location. Random coefficients for both $\alpha_{2,i}^D$ and β_i^D capture unobserved heterogeneity within downstream. As with upstream, this heterogeneity can reflect stochastic features of the environment–such as drier weather that makes drying coffee less costly–as well as other changes in the productive environment, such as positive labor supply shocks. Additionally, this heterogeneity rationalizes why empirically different quantity levels and task decisions are observed. We do not microfound or distinguish between these sources.

The fixed cost downstream for high-quality intermediation captures barriers to entry to selling high-quality coffee, as described in 3.4.1.

3.5.3 Market Clearing

In our model, upstream prices are determined by market clearing. Specifically, the price vector for low-quality coffee (p^{H*}) and high-quality coffee (p^{H*}) upstream must be such that the quantity of coffee supplied upstream is equal to the quantity of coffee demanded downstream. Formally, market clearing find upstream prices that satisfy:

$$\begin{aligned} Q_S^L(p^{L*}, p^{H*}) &= \mathrm{Q}_D^L(p^{L*}, p^{H*}) \\ Q_S^H(p^{L*}, p^{H*}) &= \mathrm{Q}_D^H(p^{L*}, p^{H*}) \end{aligned}$$

Where:

$$Q_S^L(p^L, p^H) = \sum_{i \in \mathcal{U}} q_i^L(p^L, p^H) \quad , \quad Q_S^H(p^L, p^H) = \sum_{i \in \mathcal{U}} q_i^H(p^L, p^H)$$

$$\mathrm{Q}_D^L(p^L, p^H) = \sum_{i \in \mathcal{D}} \mathrm{q}_i^L(p^L, p^H) \quad , \quad \mathrm{Q}_D^H(p^L, p^H) = \sum_{i \in \mathcal{D}} \mathrm{q}_i^H(p^L, p^H)$$

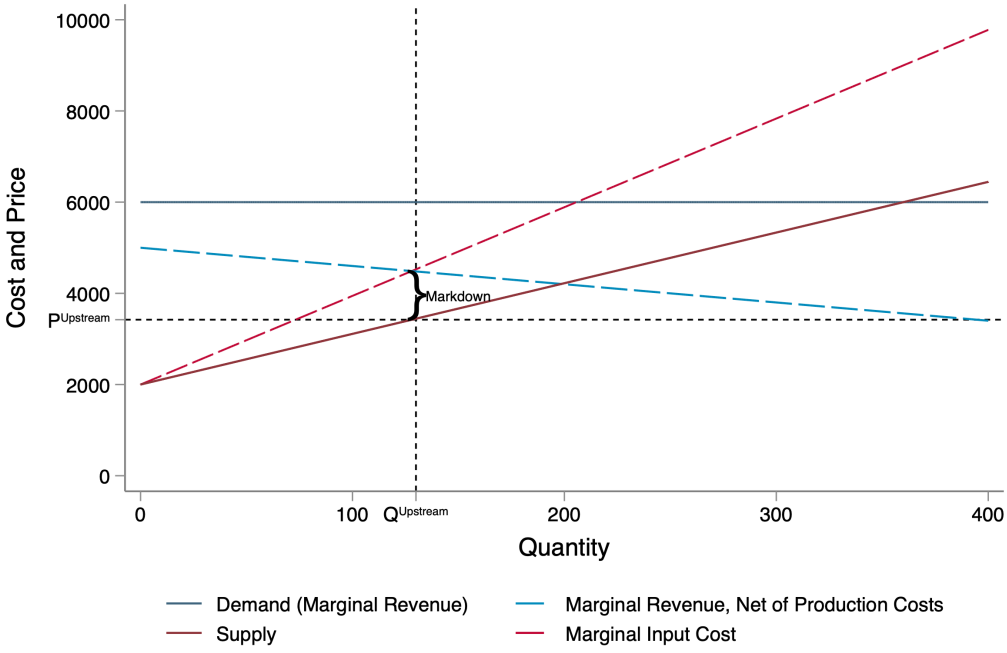
Two brief comments on market clearing are useful. First, barriers to entry, and the resulting differential buyer market power it creates, and productive substitutability both affect market clearing by influencing downstream demand. Specifically, barriers to entry both directly decrease demand for high-quality upstream for any price vector and, conditional on entry into high-quality, reduce quantities of high-quality demanded by downstream through a standard oligosopony framework (explained in more detail in 3.5.4). Second, because both upstream and downstream can change tasks, there can be discontinuities in both supply and demand; specifically, whenever an upstream (downstream) agent switches to upgrade quality,

there is a discontinuous increase in supply (demand) for high quality. Heterogeneity within layer makes agents switch tasks at different price vectors; therefore, such discontinuities tend to be fairly small and for a single agent. Unsurprisingly, then, these discontinuities are quite small in simulations (generally lower than 2.5% of quantities).

3.5.4 Markowns with Value-Added Along the Chain

Analyzing markdowns in the models highlights the importance of downstream market power and downstream productive activity, especially as they relate to efficiency. In a standard oligopsony framework in this environment, without downstream productive activity, markdowns are the difference between marginal input cost (in our empirical setting, the cost of coffee produced upstream) and supply at the quantity where marginal revenue intersects marginal input costs. The markdown distortion refers to the fact that equilibrium quantities (and prices) are not determined by marginal revenue (which is the same as demand in the absence of seller market power) equals supply, as market power creates a wedge between supply and marginal input costs as firms with buyer power increase input prices when they add a marginal unit. This set-up is effectively the opposite of an oligopoly framework (where markups are the difference between marginal revenue and demand at the quantity where marginal cost intersects marginal revenue).

Figure 3.3: Markdowns with Productive Investments by Downstream



Note: In this figure, the demand (marginal revenue) is the price offered by the customer (the big exporter in our empirical setting). In our model, the supply curve is a function of upstream cost parameters , and marginal input cost depends on both supply and market structure. Marginal revenue, net of production costs depends upon both demand (marginal revenue) and downstream cost parameters.

Figure 3.3 highlights how incorporating downstream productive activity affects markdowns. When the firm thinks about its optimal input choice, it must consider its productive costs. Therefore, it picks input quantities at the intersection of marginal revenue *net of production costs* and marginal input costs. Because downstream productive activity incurs costs, marginal revenue *net of production costs* is less than marginal revenue and reduces the quantity of inputs that the firm chooses. Markdowns, then, still are the difference between marginal input cost and supply but at the quantity where marginal revenue *net of production costs* intersects marginal input costs. Through the lens of the model, considering only one quality (*i.e.,* downstream aggregators), the difference between marginal revenue and marginal revenue *net of production costs* is $\alpha_1 + 2\alpha_{2,i} * q_i$. For downstream upgraders in the models, markdowns are different for two reasons. First, the marginal revenue, both overall and *net of production costs*, shifts up because of higher prices from exporters for high-quality coffee. Second, marginal revenue *net of production costs* is shifted downwards by β_i to capture the additional costly productive effort for upgrading.

Importantly, downstream productive activity means that markdowns in the input market are not simply the difference between prices for the same product downstream and upstream. Intuitively, this result reflects that optimal input quantities need to account for the true economic value of a marginal input to downstream, which incorporates any productive costs associated with that selling the production from that input downstream.

Importantly, data alone are insufficient to measure markdowns because marginal revenue *net of production costs* (and marginal input costs) are not observed by the econometrician; the only observed prices are the prices paid by exporter (marginal revenue) and the equilibrium price.

3.6 Calibrated Model Results

We then use a model loosely calibrated to our empirical setting to show how barriers to entry and productive substitutability affect prices (and, thereby, the quality premium), market efficiency, markdowns, distribution of surplus, and task allocation. Specifically, the model is calibrated to have a similar reduction of quality premium upstream (roughly 50%); no other moments are targeted. Therefore, results highlight how barriers to entry and productive substitutability can affect outcomes broadly. However, as the calibration does not target many

features of the empirical setting, we do not interpret magnitudes as seriously quantifying the effects of these mechanisms.

Through the lens of the model, we quantify the effects of removing barriers to entry by setting $F = 0$. Similarly, to quantify the effects of productive substitutability, we can remove downstream's choice to upgrade low-high as an option and require them to choose to aggregate coffee of either quality. Appendix D.5 reports the full list of prices and parameter values used for model simulation.

Table 3.6: Effects of Barriers to Entry & Productive Substitutability on Prices, Quality Incentives, and Markdowns

					Downstream	
	Downstream	Upstream	*Upstream Prices*		*Markdowns*	
	Quality	Quality	Low	High	Low	High
Scenario	Premium	Premium	Quality	Quality	Quality	Quality
Baseline	3000	1594	6994	8588	17	247
No Barriers to Entry For High-Quality Intermediation	3000	2341	7397	9738	30	18
No Productive Substitutability	3000	1616	6961	8576	10	251

Note: Quantification exercises based on model calibrated to match quality premium reduction upstream; prices and parameters are reported in D.5. No barriers to entry corresponds to setting $F = 0$ and no productive substitutability corresponds to eliminating downstream's option of upgrading quality.

Table 3.6 highlights the effects of barriers to entry and productive substitutability on prices, the quality premium (quality incentives), and markdowns. First, the analysis shows that each of barriers to entry and productive substitutability diminish the quality premium upstream. Specifically, a model without barriers to entry for high-quality intermediation but with productive substitutability still does not yield the full quality premium upstream, even though a model without both barriers to entry and productive substitutability will always result in an identical quality premium downstream and upstream, although price

levels may still differ by layer.[17] The results also highlight the importance of the *interaction* of barriers to entry and productive substitutability. Namely, the increase in the quality premium upstream from removing both mechanisms simultaneously ($3000 - 1594 = 1406$; not shown) is meaningfully larger than the sum of changes from removing each of the two mechanisms individually.[18] This result reflects that barriers to entry especially reduce the quality premium in an environment with productive substitutability because productive substitutability decreases competition (and, therefore, increases buyer market power) in the high-quality segment.

The analysis also emphasizes that barriers to entry can result in large markdowns in the high quality segment, although there are still markdowns even for low quality. Markdowns occur for both qualities for two reasons. First, the model holds overall entry into the supply chain fixed. Second, downstream low-high upgraders generally have large markdowns for low quality. These downstream have markdowns because they sell high-quality coffee at a high price; even though they have high productive costs that incorporate upgrading, the increase in price due to the downstream quality premium still results in markdowns. And, it is the downstream who upgrade low-high that drive the increase in markdowns in low quality in the scenario that removes barriers to entry.

[17] Suppose, by contradiction, that in a model without both barriers to entry and productive substitutability that the quality premium is not identical downstream and upstream (and that the equilibrium includes production of both qualities upstream). First, suppose the quality premium is higher upstream than downstream. Then, no downstream will demand high quality because they have would increase their profit by deviating to buying and selling low-quality. Second, suppose that the quality premium is smaller upstream than downstream. Then, downstream has a profitable deviation to purchase and sell high quality. Note that this proof outline does not factor in downstream market power; however, to incorporate downstream market power, the analysis only needs minor additions. Specifically, all profitable deviations highlighted are true for very small quantities. Therefore, define the quality premium difference across layers for both cases as ϵ and consider quantity deviations sufficiently small that the quality premium does not change by more than $\epsilon/2$.

[18] Specifically, $2341 - 1594 = 747; 1616 - 1594 = 22$ and $747 + 22 = 769 < 1406$.

Table 3.7: Effects of Barriers to Entry & Productive Substitutability on Quality Investments and Distrubtion of Surplus

	Upstream Production (kgs)		*Downstream Production (kgs)*		*Producer Surplus (UGX)*	
Scenario	Low Quality	High Quality	Low Quality	High Quality	Upstream	Downstream
Baseline	1894	292	1665	434	7233	4827
No Barriers to Entry For High-Quality Intermediation	1718	1717	0	3388	18380	9955
No Productive Substitutability	1821	300	1796	293	6830	4076

Note: Quantification exercises based on model calibrated to match quality premium reduction upstream; prices and parameters are reported in D.5. No barriers to entry corresponds to setting $F = 0$ and no productive substitutability corresponds to eliminating downstream's option of upgrading quality.

Table 3.7 illustrates the effects of barriers to entry and productive substitutability on production of coffee by quality and layer as well as distribution of producer surplus. First, the results highlight how barriers to entry to high-quality intermediation reduce quality investments throughout the supply chain. Removing barriers to entry pushes all downstream to either upgrade low-high or aggregate high-high. This result reflects that as long as at least one downstream prefers to upgrade low-high rather than aggregate high-high (*i.e.,* as long as there is productive substitutability), then the quality premium upstream is lower than the quality premium downstream. Therefore, if there are no barriers to entry for high-quality intermediation, all downstream will prefer to aggregate high quality over low quality (and some of those may prefer to upgrade). The increase in demand for high quality upstream results in more quality investments both upstream and downstream. As a result, producer surplus increases both upstream and downstream. The increase upstream reflects the higher prices, while the increase downstream occurs due to the task reallocation that maximizes profits. Note, however, that some downstream do experience surplus losses–namely, those downstream that aggregate high-high or upgrade low-high at baseline, as they pay higher input prices when barriers to entry are removed.

The results show that productive substitutability increases efficiency, in terms of total production of high-quality coffee. Specifically productive substitutability allows downstream to make-up for limited quality investment upstream. Additionally, removing productive substitutability does not increase producer surplus anywhere on the chain. Although this result could potentially be different for other parameter values, it reflects that the gains to upstream from task adjustment responding to the higher quality premium do not compensate for losses due to lower overall price levels. The decrease in overall price levels reflects that, once there is no productive substitutability, there are no downstream low-high upgraders. As downstream low-high upgraders tend to demand large quantities of coffee when they upgrade, changing their task to aggregation of either quality meaningfully decreases their demand for low quality. Therefore, the low quality price falls. And, because prices upstream are linked across qualities through market clearing, the price for high quality also falls. Therefore, upstream surplus decreases. Notably, downstream surplus falls even despite the decrease in price levels. This result reflects that downstream upgraders had large surplus that they lose once their task is changed; downstream who complete the same tasks as in the baseline actually experience surplus increases due to the lower prices upstream.

Table 3.8: Effects of Barriers to Entry & Productive Substitutability on Task Allocation

	Upstream		*Downstream*		
	Low Quality	High Quality	Low to Low Aggregate	Low to High Upgrade	High to High Aggregate
Baseline	90.0%	10.0%	97.5%	1.5%	1.0%
No Barriers to Entry For High-Quality Intermediation	59.0%	41.0%	0.0%	34.5%	65.5%
No Productive Substitutability	89.5%	10.5%	98.5%	0.0%	1.5%

Note: Quantification exercises based on model calibrated to match quality premium reduction upstream; prices and parameters are reported in D.5. No barriers to entry corresponds to setting $F = 0$ and no productive substitutability corresponds to eliminating downstream's option of upgrading quality.

Table 3.8 illustrates the effects of barriers to entry and productive substitutability on task

allocation within the supply chain. Barriers to entry have a large effect on task allocation; removing them shifts more upstream and downstream into making quality investments. Conversely, productive substitutability has small effects on task allocation; this result reflects that substitutability at baseline is fairly low, with only 1.5% of downstream upgrading.

3.7 Conclusion

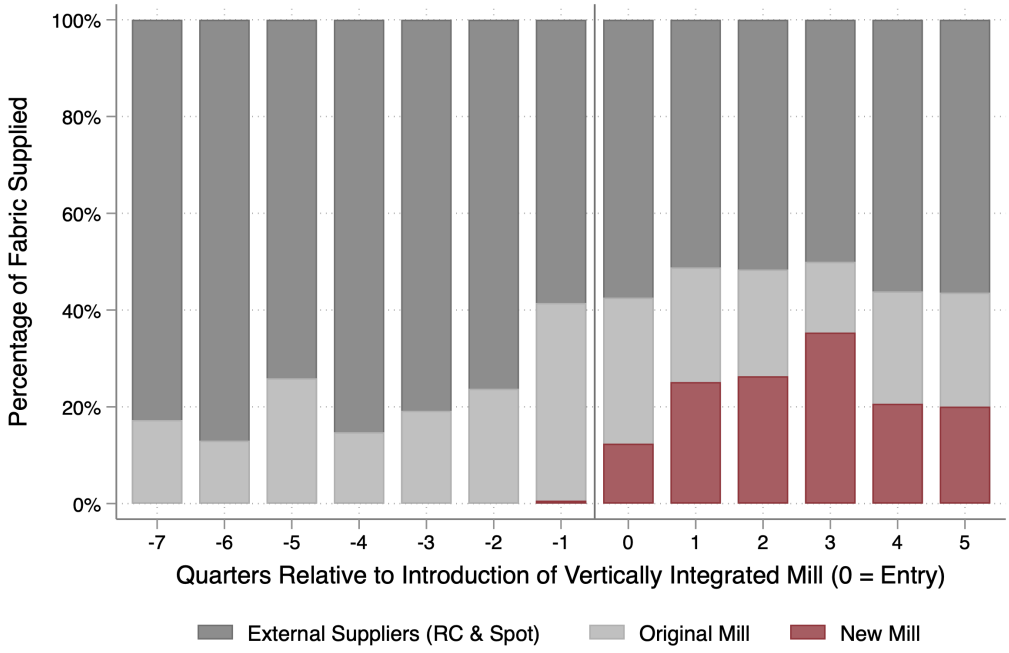
This project highlights the importance of intermediation for shaping quality incentives and investments in a Ugandan coffee supply chain. We show that the vast majority of coffee produced travels through an intermediary, and that intermediaries see a much larger quality premium than producers. We build a model with both barriers to entry and productive substitutability where either of these mechanisms can lead to a diminished quality premium upstream. Simulations verify that indeed these mechanisms can affect outcomes in line with intuition, emphasizing that both market and cost structure of the supply chain play an important role in determining quality incentives and investments. Future work can try to estimate this model (or a similar one) and consider what policy initiatives might be most successful at increasing efficiency and achieving an equitable distribution of surplus by layer.

APPENDIX A

Relational Contracts in Indian Garment Manufacturing

A.1 Input Sourcing by Supplier Type

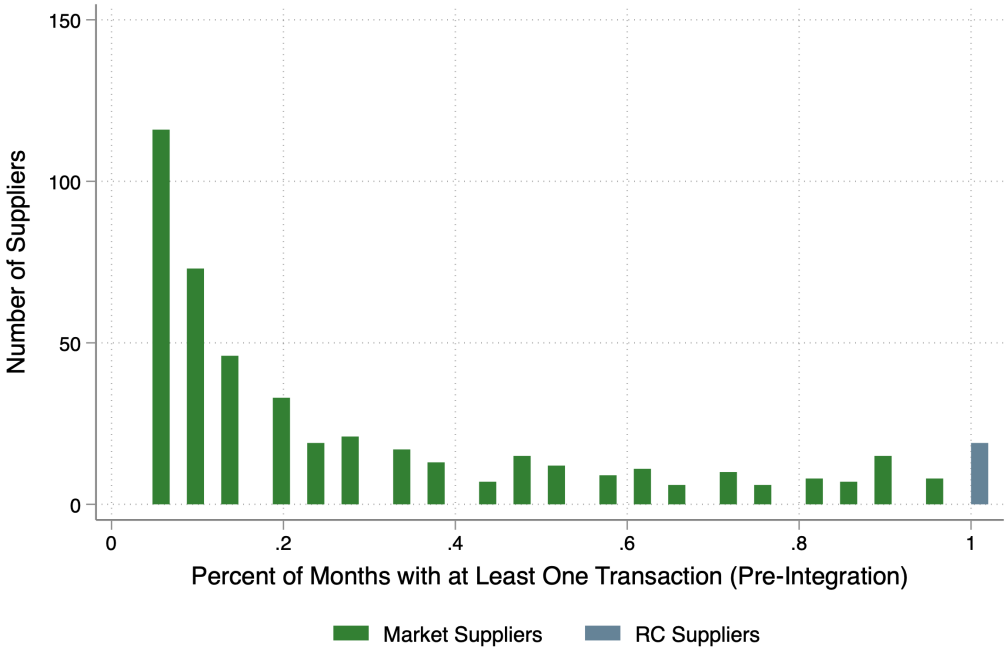
Figure A.1: Fabric Sourcing by Supplier Type



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

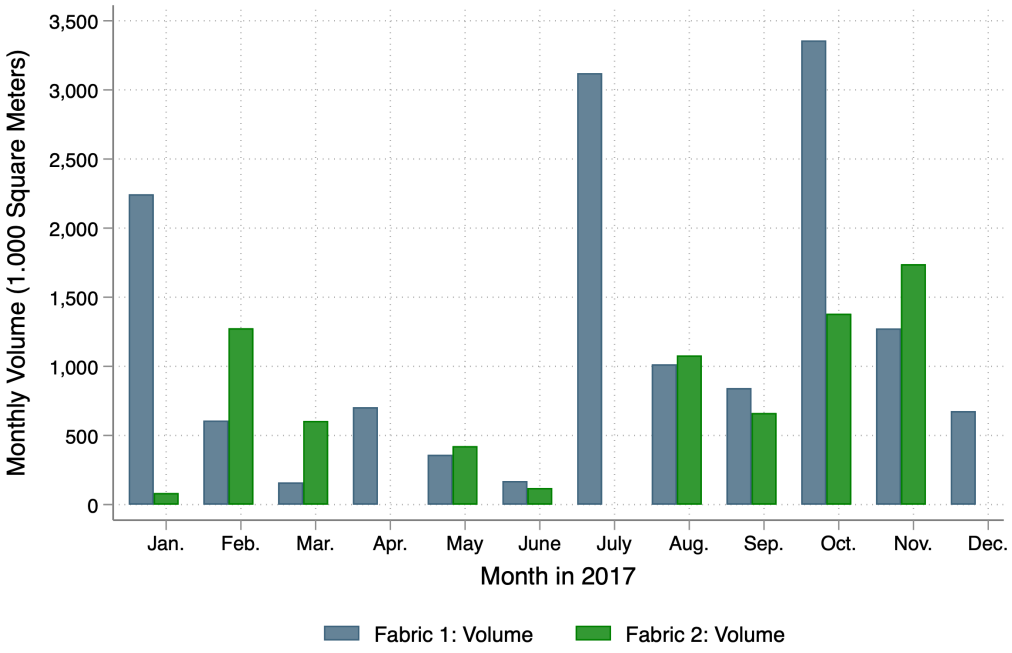
A.2 Relational Contract Discounts

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



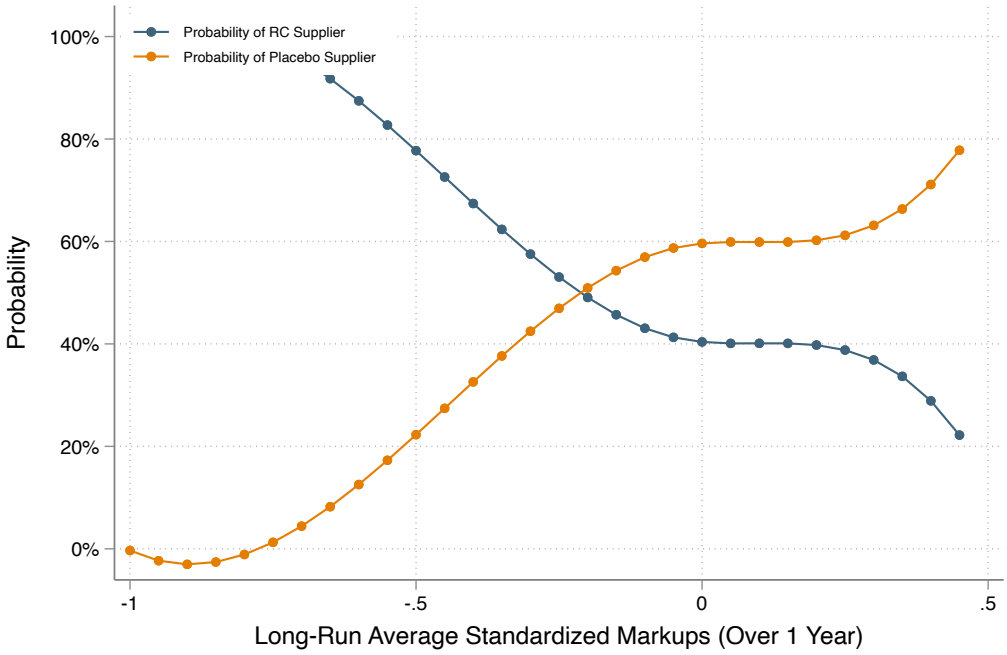
Note: Data from universe of fabric transactions by buyer. Bars represent the number of suppliers with the relevant percent of months with transactions.

Figure A.3: Variation in Monthly Quantity by Fabric



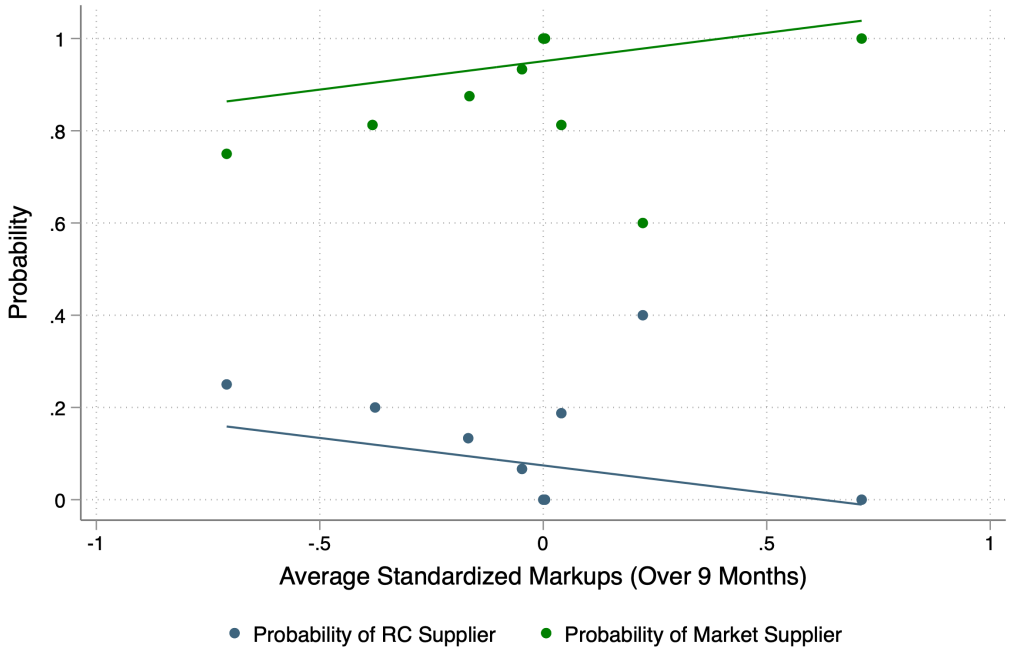
Note: Data from universe of fabric transactions by buyer. Bars represent the quantity of fabric purchased during the month in 2017 for two selected fabrics. Fabrics were selected based on data availability to illustrate variation.

Figure A.4: Discounts and Relational Contract Suppliers



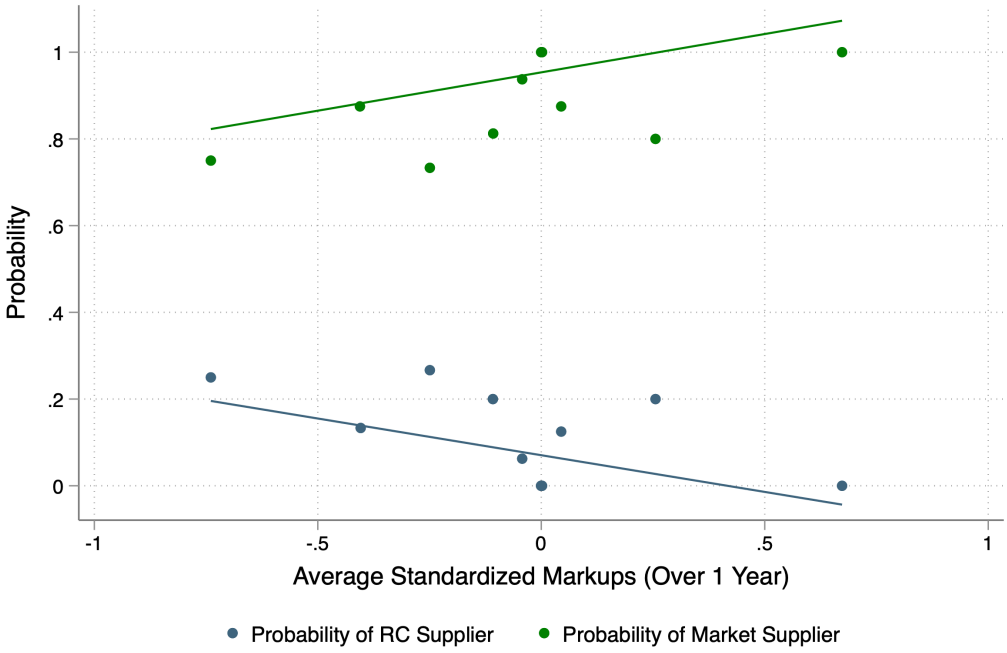
Note: Data from universe of fabric transactions by buyer. Probability estimates represent marginal effects from a non-parameteric series estimator calculated using transactions from the year prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.5: Discounts and Supplier Type



Note: Data from universe of fabric transactions by buyer. Probability estimates from a bin-scatter calculated using transactions from the three quarters prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.6: Discounts and Supplier Type



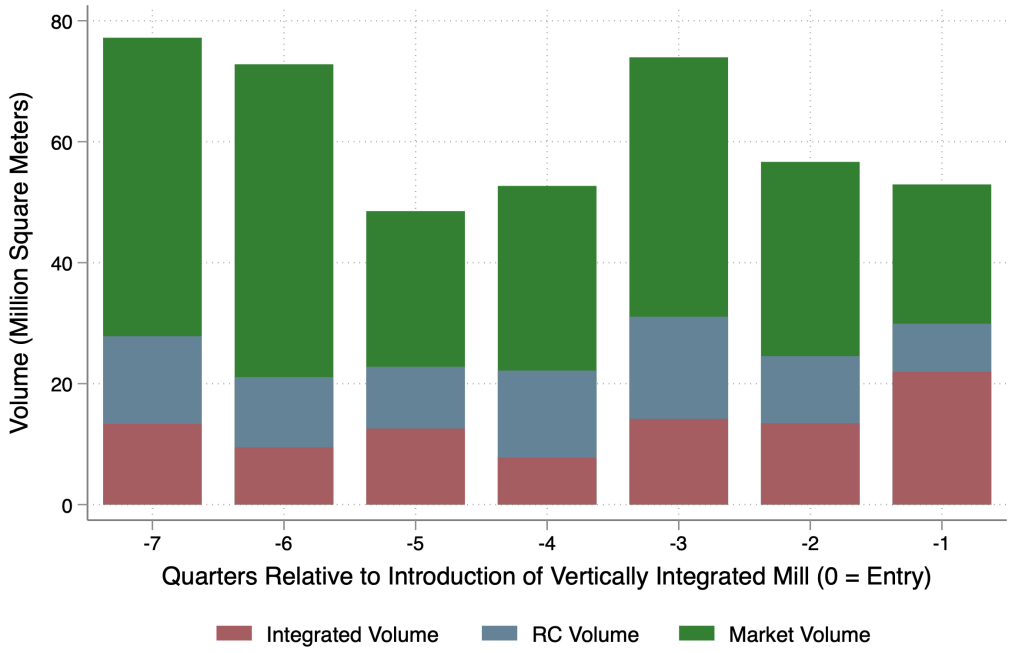
Note: Data from universe of fabric transactions by buyer. Probability estimates from a binscatter calculated using transactions from the year prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.7: Coefficient on Owed Quantity by Supplier Type



Note: Data from universe of fabric transactions by buyer. Results are the coefficient for owed standardized volume from a regression of current volume on owed volue and month fixed effects. Owed standardized volume is calculated for each supplier quarter as the standardize volume in the prior month less the average of monthly standardize volume in the prior year, not including the prior month. Permuted volumes randomly reshuffle standardized volumes, then calculate owed standardized volume using the permuted data across 500 permutations. The value shown for permuted volumes for relational contract suppliers represents the mean across permutations; the 90% empirical CI is also shown.

Figure A.8: Volume by Supplier Type (Pre-Integration)



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

A.3 Owed Quantity Guarantees

In addition to showing that relational suppliers have more stable quantities in general, I explore the association between a proxy for the quantity owed in the relational contract and the volume received in the current period. For relational contract suppliers, owed quantity should be positively correlated with current volume: when the relational contract owes more volume to the supplier, the supplier should receive more volume. To test this hypothesis in the data, I regress volume in a month on a proxy for the volume owed in the relational contract with month fixed effects. A positive association between owed volume and current volume would be consistent with the hypothesized relational contract. Long-run monthly average volume serves as a proxy for the quantity guarantee, as the relational contract delivers the guaranteed quantity to suppliers, conditional on the quantity guarantee not changing in an economically meaningful way over the course of the sample. Owed quantity is then measured as the difference between the proxy for the long-run quantity guarantee and the quantity fulfilled recently (*i.e.,* in the prior month). Alternatively, as there is no

relational contract for market suppliers, there is no reason to expect a relationship between owed volume and current volume for market suppliers.

Figure A.9 illustrates that a one standard deviation increase in owed volume for a relational contract supplier is associated with a .04 standard deviation increase in current volume.[1] Importantly, the same result does not hold for market suppliers, where an increase in owed quantity is associated with a *decrease* in current volume of -.05 standard deviations.

Figure A.9: Coefficient on Owed Quantity and Supplier Type



Note: Data from universe of fabric transactions by buyer. Results are the coefficient for owed standardized volume from a regression of current volume on owed volue and month fixed effects. Owed standardized volume is calculated for each supplier quarter as the standardized volume in the prior month less the average of monthly standardize volume in the three prior quarters, not including the prior month. Permuted volumes randomly reshuffle standardized volumes, then calculate owed standardized volume using the permuted data across 500 permutations. The value shown for permuted volumes for relational contract suppliers represents the mean across permutations; the 90% empirical CI is also shown.

Although this result is consistent with any mean-reverting process, I find that the magnitude of the effect is larger than what would happen if the buyer randomly directed the

[1]The pass-through of less than one could reflect that quantity guarantees operate over a longer time frame than one month, measurement error, or market-level demand shifts that are incorporated into the relational contract but not the analysis.

same quantities to the suppliers over time, which should also be mean reverting by construction. Additionally, the most reasonable alternative mean-reverting contracting strategy distinct from the relational contract described would likely direct volumes based on the manufacturer's knowledge from past transactions. Specifically, if the buyer has directed lower volumes than usual to the supplier, then it should have lower capacity utilization and could potentially pass on some cost savings to the buyer. However, this strategy would likely be ineffective given that supplier capacity can change over time without being observed by the buyer due to transactions with other buyers. Furthermore, if this contracting narrative were correct, the point estimate for the shorter long-run version (nine as opposed to twelve months) should likely be higher, as more recent volume is presumably more relevant for current capacity utilization than older orders (see Appendix Figure A.7).

Empirically, the evidence also suggests that the observed association is neither random nor mechanical. Specifically, a placebo exercise which randomly permutes volume across months for the same supplier, such that the set of volumes associated with each supplier is the same across all simulations but the mapping between month and volume is changed, has a smaller mean point estimate across 500 randomizations than the observed association, although confidence intervals are large given the small number of relational contract suppliers and the restriction to only using pre-integration data to avoid incorporating any effects of integration on the relational contract.

A.3.1 Why Not an Explicit Contract?

There are four factors that presumably lead explicit contracts to fail to resolve the hold-up problem in this setting. First, if the explicit contract is not enforceable because of the institutional environment for contracting, then the explicit contract is effectively irrelevant. Research in India suggests that contracting parties may be concerned that courts will be unable to effectively enforce an explicit contract (Rao (2022); Banerjee and Duflo (2000)). Second, an effective long-term contract would require specifying the price as a function of capacity utilization, as transaction prices shift due to supplier and industry capacity utilization. A contract with constant prices that do not adjust for capacity utilization would result in extreme profit variance, with some orders potentially even unprofitable when supplier capacity usage is sufficiently high; suppliers would be unlikely to agree to any long-term contract that did not condition prices on capacity (at least not without extremely high prices in general). As there are many possible future capacity states due to both the unpredictable stochastic demand from the end client, as well as the unpredictable demand a supplier would face from other buyers, such a contract would likely be difficult (or at least expensive) to

write, even abstracting from enforcement concerns, as it would require agreeing on a different price for a large number of capacity states. The just-in-time nature of the procurement process only exacerbates this problem as the complex contingent contract would need to be written in a short time period. Third, even if a contract could specify a price as a function of the supplier's capacity, it may introduce supplier moral hazard that would undermine the effectiveness of the contract. Specifically, as contracted prices increase in capacity, such a contract reduces supplier incentives to reject sampling opportunities (and, thereby orders), from other buyers to save capacity as the supplier would receive similar markups for all capacity states. It follows that suppliers would pass on costs to the buyer. Note that this concern is relevant as staff at the buyer suggest that suppliers do indeed try to target their capacity to match the buyer's needs. Fourth, capacity utilization is likely not contractible, as it is not observable by the buyer.

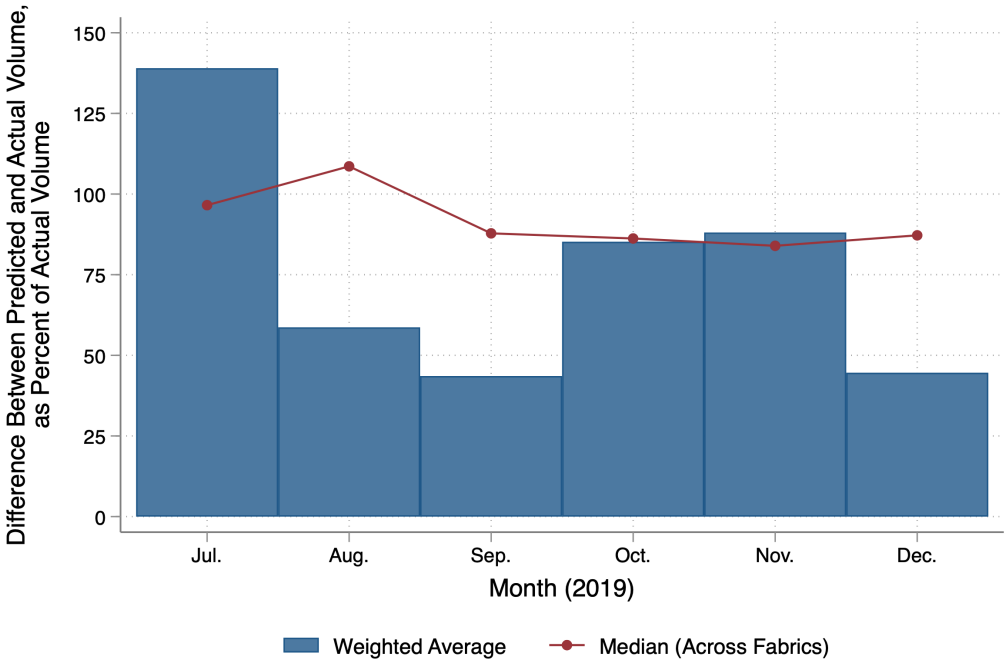
A.4 Why Not Exclusively Relational Contracts?

Given the evidence that relational contracts offer improved pricing relative to market suppliers, the buyer would minimize costs by sourcing exclusively through relational contract suppliers. However, in practice, although relational contract suppliers provide 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 relational contracts reflects that challenges of fulfilling owed quantity guarantees in this setting. Given that the buyer does not know what downstream end client demand will be in any period, it can only credibly commit to the lower bound of estimates of future volume. Additionally, suppliers can typically only produce a subset of fabrics, further decreasing the volume that could credibly be sourced through relational contracts.

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 relational contracts. The data highlight the volatility and unpredictability of volume, especially at the fabric level. Figure A.10 shows the average difference between predicted and actual volume, as a percent of actual volume, weighted by the fabric's share of total volume. These results highlight that errors are, on average, always more than the volume of the typical order. 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 heigher 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 relational contract and integrated suppliers appears to be quite stable pre-integration, with market suppliers absorbing stochastic variation in quantities in A.8.

Figure A.10: Forecast Error in Volume Predictions



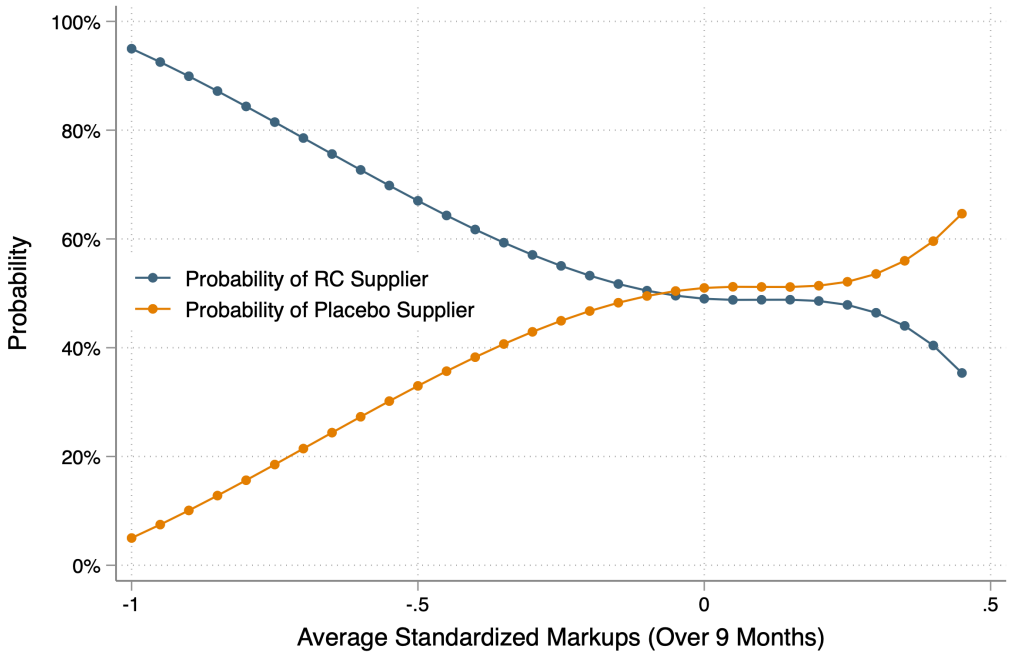
Note: Data from universe of fabric transactions by buyer. Monthly fabric volume forecasts are predictions from a regression of monthly fabric volume on 12 lags of monthly fabric volume, 12 lags of monthly volume across all fabrics, and both fabric and month fixed effects using all but the last six months of data. Errors are the residual fabric volume, expressed as a percent of the monthly fabric volume. Aggregation of errors across fabrics uses the volume-weighted average where volume-weights are the running average percent of volume for the specific fabric over the past year.

A.5 Placebo Analysis: Results

Figure A.11 shows that suppliers that offer discounts are much more likely to be relational contract rather than placebo suppliers, while suppliers that charge high markups are more likely to be placebo suppliers. Additionally, placebo suppliers have more quantity volatitility than relational contract suppliers, and the difference is statistically significant ($p < .01$), as

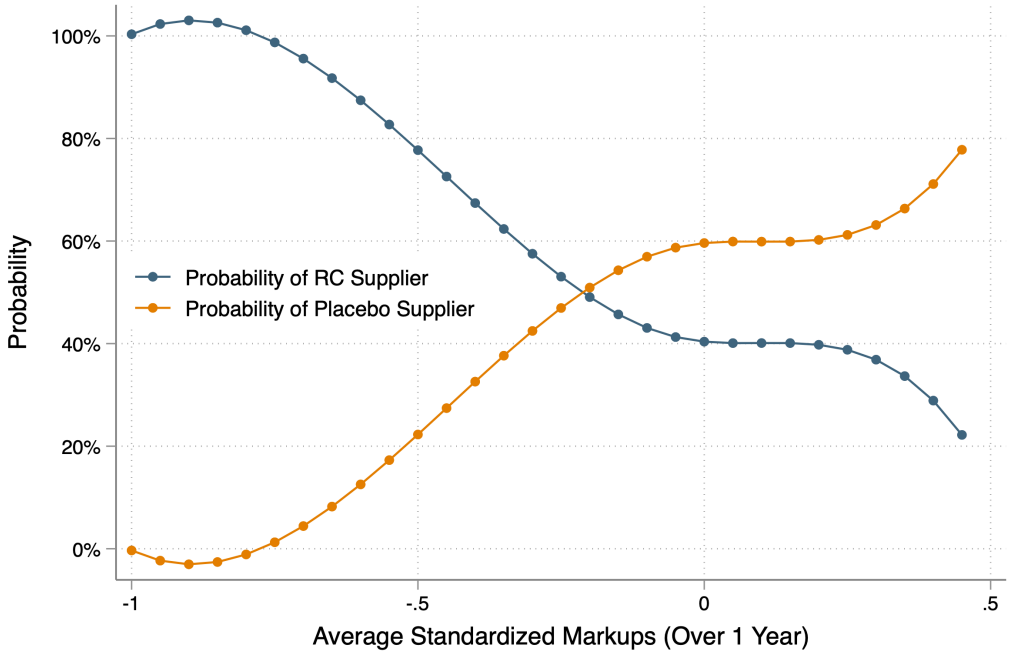
illustrated by Figure A.15. The results for owed volume in Figure A.16 also highlight the difference between relational contract and placebo suppliers, with a one standard deviation increase in owed volume for relational contract suppliers associated with a .04 standard deviation increase in current volume, compared to a -.17 decrease for placebo suppliers. Importantly, given that placebo suppliers and relational contract suppliers have similar monthly average volumes and a similar distribution of owed volumes (see A.17), this difference reflects that current volumes for placebo suppliers are much less associated with owed volume. It follows that it is unlikely that placebo suppliers receive a quantity guarantee.

Figure A.11: Discounts and Supplier Type: Placebo Suppliers



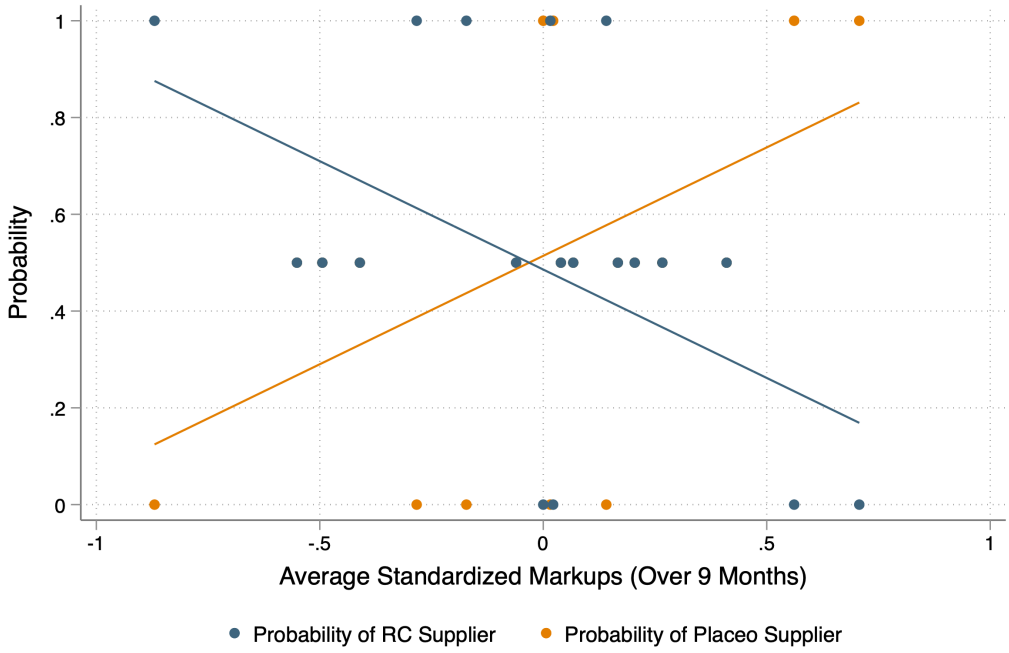
Note: Data from universe of fabric transactions by buyer. Probability estimates represent marginal effects from a non-parameteric series estimator calculated using transactions from the three quarters prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.12: Discounts and Supplier Type: Placebo Suppliers



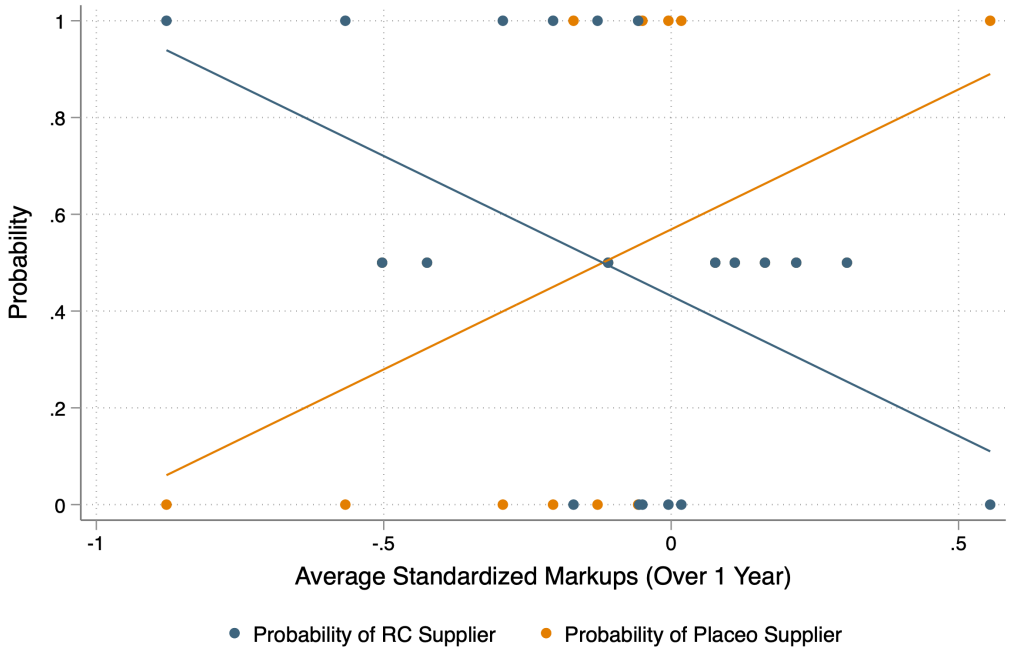
Note: Data from universe of fabric transactions by buyer. Probability estimates represent marginal effects from a non-parameteric series estimator calculated using transactions from the year prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.13: Discounts and Supplier Type: Placebo Suppliers



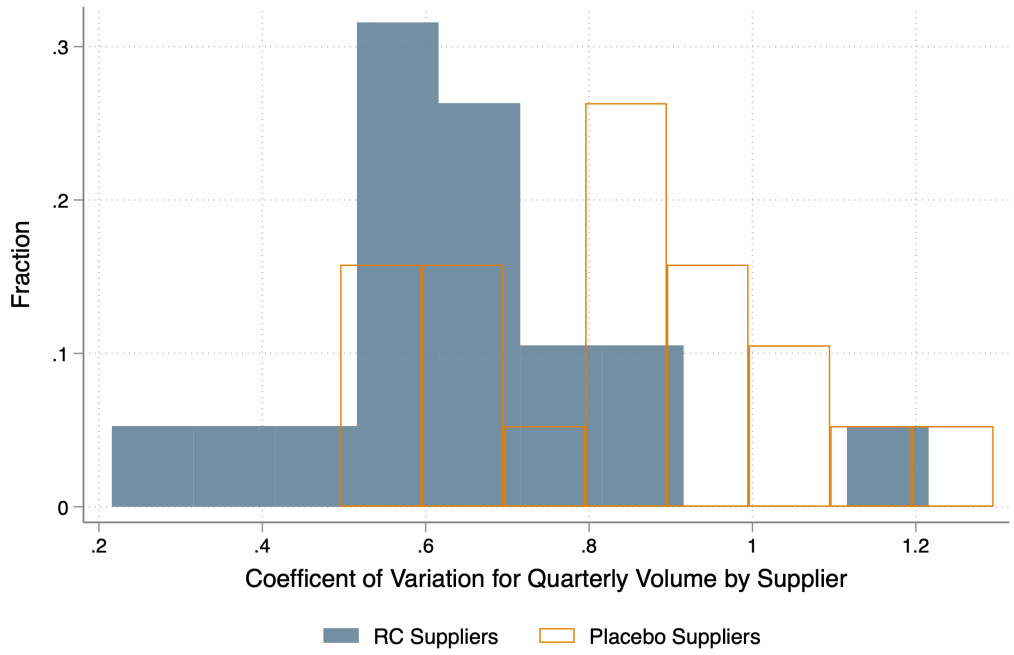
Note: Data from universe of fabric transactions by buyer. Probability estimates from a binscatter calculated using transactions from the three quarters prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.14: Discounts and Supplier Type: Placebo Suppliers



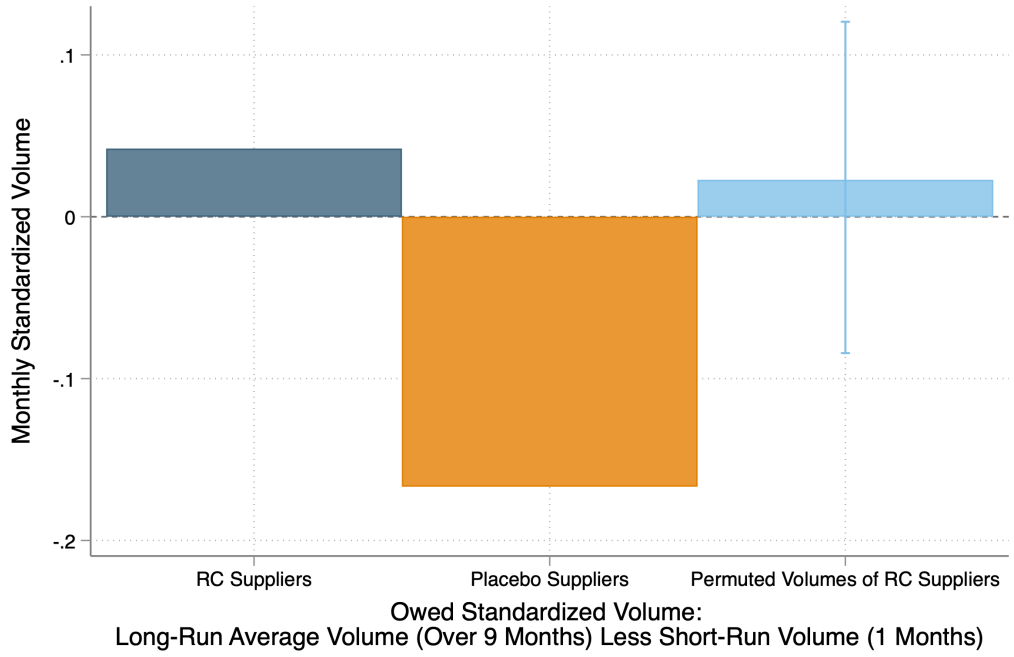
Note: Data from universe of fabric transactions by buyer. Probability estimates from a binscatter calculated using transactions from the year prior to integration. Long-run average standardized markups are the volume-weighted average standardized markup, where markups are standardized by fabric and month. Fabrics with fewer than three observations per month are omitted.

Figure A.15: Coefficient of Variation by Supplier Type: Placebo Suppliers



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

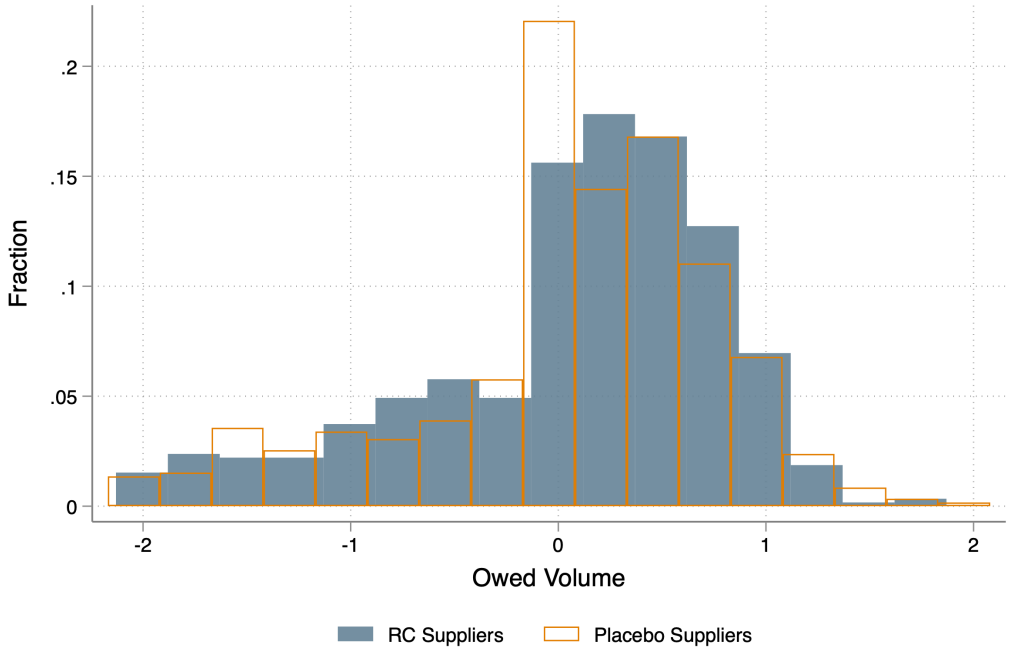
Figure A.16: Coefficient on Owed Quantity by Supplier Type: Placebo Suppliers



Note: Data from universe of fabric transactions by buyer. Results are the coefficient for owed standardized volume from a regression of current volume on owed volue and month fixed effects. Owed standardized volume is calculated for each supplier quarter as the standardize volume in the prior month less the average of monthly standardized volume in the three prior quarters, not including the prior month. Permuted volumes randomly reshuffle standardized volumes, then calculate owed standardized volume using the permuted data across 500 permutations. The value shown for permuted volumes for relational contract suppliers represents the mean across permutations; the 90% empirical CI is also shown.

Rather than a self-enforcing relational contract, the contracting features observed in the data could be enforced by some other feature in the environment, such as social pressure through kin, ethnic, religious, or other social identity or group based identity (as in Sanchez de la Sierra (2021)). However, pricing in the relational contract seems inconsistent with motives based on such social membership, as the discount received by the larger sophisticated firm (*i.e.,* the buyer) means that the contractual enforcement by social group membership also increases in-group inequality, which seems unlikely as such groups presumably have preferences to redistribute to the less well-off members of the group. To be more concrete, it seems unlikely that kin ties effectively pressure the buyer to work with a supplier, but those same kin ties also require the supplier to offer the buyer a discount. Note that it need not be the case that social group membership plays no role–group membership may facilitate the creation of self-enforcing relational contracts or influence which suppliers are selected as relational contract suppliers. Importantly, in these scenarios, the relational contract still operates as a relational contract with its associated impacts on transactions.

Figure A.17: Distribution of Owed Quantity by Supplier Type: Placebo Suppliers



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

APPENDIX B

Vertical Integration and Relational Contracts: The Threat Point Effect

B.1 Robustness Checks for Conceptual Model

B.1.1 Relational Contracts with Dynamics

In a dynamic model that more closely resembles typical models of relational contracts, including a DICC as in Macchiavello and Morjaria (2023), the deviation payoffs change. Note that the on-path payoffs do not change, so I only discuss how the deviation payoffs change in this model. This reflects the interpretation of the relational contract as specifying a pricing rule that is agreed upon before quantities are realized. In the dynamic model, the supplier's outside option is to hold up the buyer making it pay the high market price. Therefore, I search over deviations and find the best static deviation, which occurs when $p^M q^{RC} - p^{RC} q^{RC}$ is the largest:

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

Similarly, the buyer can break the relational contract. The buyer could have an incentive to do so when the relational contract price is higher than the market price, would can occur because the relational contract is designed to smooth profits. For high cost or low quantity states, relational contract prices are high to facilitate supplier profit smoothing. Therefore, I search over static deviations, with the buyer's best deviation occurring when $p^{RC} q^{RC} - p^M q^{RC}$ is the largest. In this case, the buyer's outside option is:

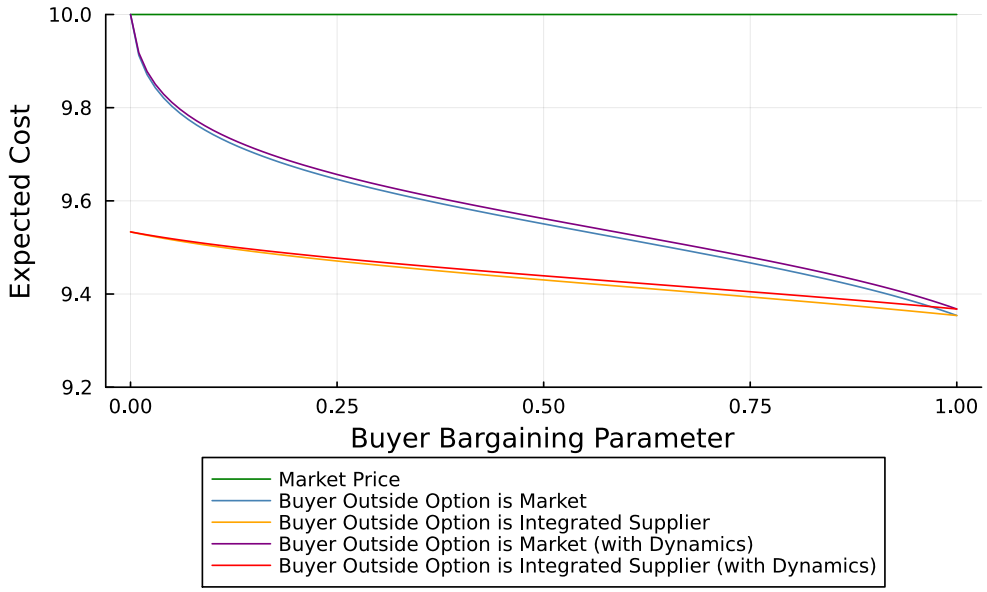
$$O_B = (\delta^B)\Big(-p^M \mathbb{E}_{q^{RC}}[q^{RC}]\Big) + (1-\delta^B)\Big(-p^M q^M\Big) \text{ if not integrated}$$

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

Note that for δ^S, δ^B sufficiently large, the dynamic version approaches the long-run version considered in the main text.

The analysis below shows that results do not change meaningfully when incorporating dynamics.[1]

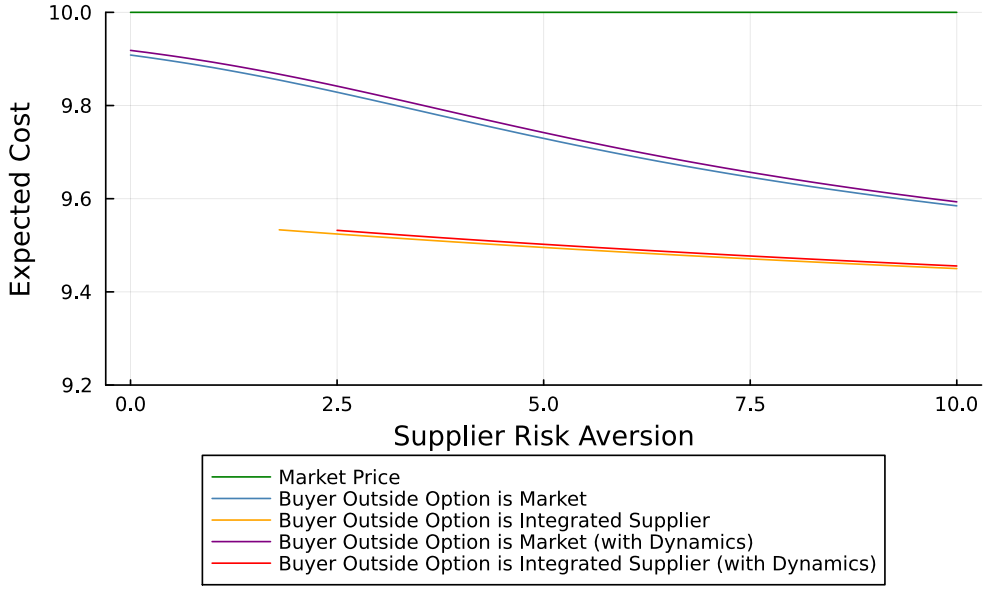
Figure B.1: Relational Contract Prices and the Buyer Bargaining Parameter with Dynamics



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

[1]For this analysis, I set $\delta^S = \delta^B = .9$. As analysis of this model is at the quarter level in general, this 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 should be conservative.

Figure B.2: Relational Contract Prices and Supplier Risk Aversion with Dynamics

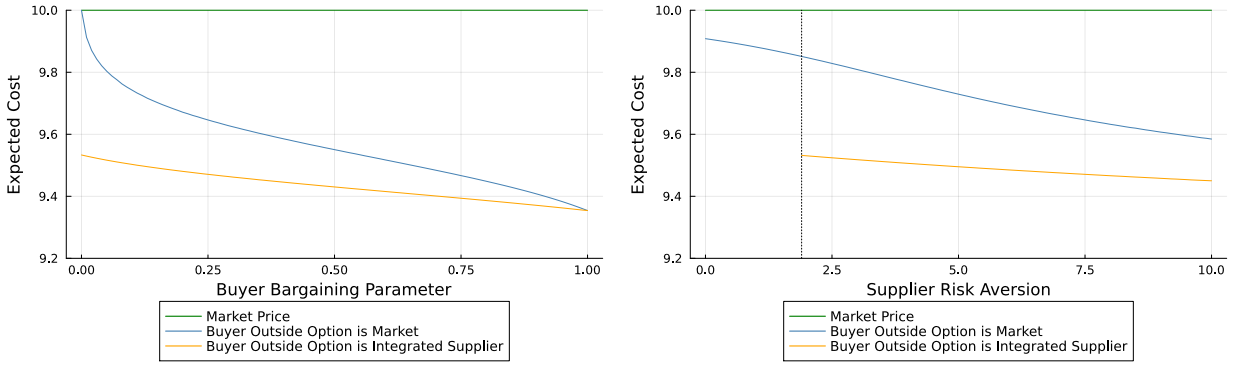


Note: Costs represent volume-weighted average procurement costs, with fabric prices computed from the model with parameterization as described in Appendix B.4.

B.1.2 Cost Heterogeneity Between External and Internal Supplier

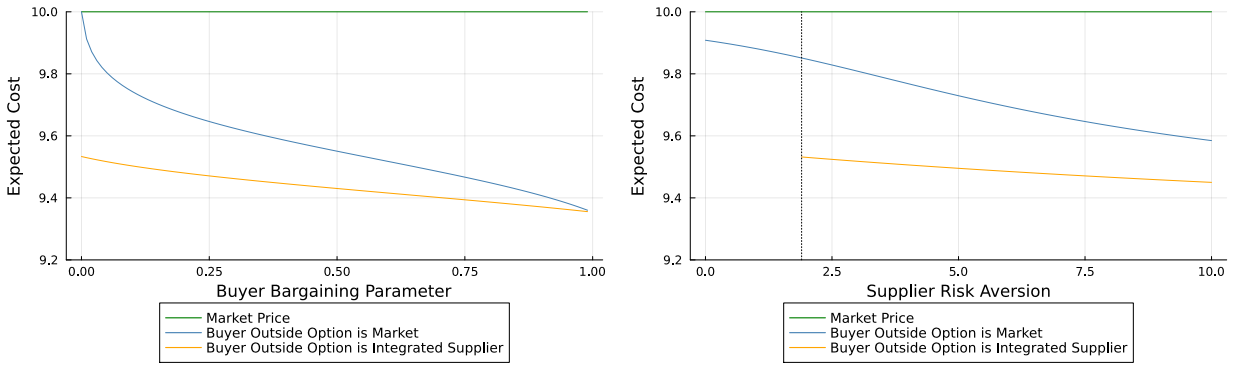
I show that results are robust to the external supplier having lower or higher costs than the integrated supplier, shifting costs by 10%. To ensure that optimal quantity is still $q = 1$ without changing market prices, I only shift the fixed part of costs by 10%, keeping convexity the same. Results (available upon request) shifting all costs are similar. Larger heterogeneity in production costs between the buyer and supplier seem unlikely. If buyer production costs are much lower than supplier production costs, the buyer should fully integrate. On the other hand, if supplier production costs are much lower than buyer production costs, the buyer should never integrate and the integration should not even credibly shift the threat point (assuming the supplier can see the integrated mill and make a reasonable inference about is productivity).

Figure B.3: Relational Contract Prices, the Buyer Bargaining Parameter, and Supplier Risk Aversion when External Supplier has Lower Costs



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

Figure B.4: Relational Contract Prices, the Buyer Bargaining Parameter, and Supplier Risk Aversion when External Supplier has Higher Costs



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

B.2 Credibility of Displacing Relational Contract Supply with Integrated Supply

For relational contract suppliers to credibly believe that the buyer could displace them with integrated supply, the suppliers must believe that the sum of integrated capacity and the relational contract quantity guarantees for other relational suppliers (as the buyer cannot

displace these quantities) must be sufficient to fulfill all demand across the fabrics they produce without going to the spot market. In other words, if the relational contract supplier knows that overall quantities demanded by the end client are large enough for the buyer to need to resort to the spot market even after vertical integration, then relational contract suppliers are unlikely to believe that the integrated supplier is a credible threat because the buyer would prefer to displace spot market supply to integrated supply.[2]

In the context of the model, the buyer and relational contract supplier could bargain over relational contract prices but with the buyer's surplus including total procurement costs for all orders for across all supply sources. Then, replacing relational contract supply with integrated supply would reduce buyer surplus if integrated supply would otherwise displace spot market supply. However, such a contract is clearly infeasible in this setting because the quantities with other suppliers are the buyer's private information and not contractible. And, even if they were public information, the buyer does not know these quantities during relational contract bargaining. However, the intuition still applies–if the supplier knows that integrated supply could instead be used to displace only spot supply, then the supplier should not find the threat of displacement with integrated supply credible.

What are relational contract suppliers likely to believe? The buyer has strong incentives to reduce spot market supply; therefore, relational contract suppliers could credibly believe that the buyer would integrate to the extent that spot market purchases are negligible–recall that the integration is large, increasing total integrated supply by more than six times the largest relational contract supplier. The buyer's website is consistent with this strategic ambiguity–while the buyer highlights overall volumes of garments sold and the sophistication of its mills, it describes the size of the "campus" without providing any specific information about capacity. Furthermore, it need not be the case that the integrated supply alone would displace the spot market, as the buyer also needs to meet quantity commitments for other relational contracts. Therefore, a relational supplier should consider displacement by integrated supply as credible as long as the sum of integrated supply and target quantities for other relational contract suppliers are sufficient to crowd out the spot market. However, relational contract suppliers likely believe that the buyer often purchases some volume from the spot market, at least during peak demand; otherwise, meeting quantity targets until relational contracts would mean leaving much of internal capacity underused. On the other hand, spot market supply is likely most important for fringe fabrics that are not consistently used, as developing relational contracts targeting suppliers specializing in uncommon fabrics

[2]Note that the integrated supplier can also produce quantities over capacity, albeit at high marginal cost, such that the physical presence of an integrated supplier alone may allow the buyer to threaten to displace more quantities than it actually does.

would likely make meeting quantity guarantees challenging. Given the contracting environment in this setting, capacity at the vertically integrated supplier, which fabrics can be made at the vertically integrated supplier, and quantity targets for other relational contract suppliers (let alone the number of other relational contract suppliers)[3] are all presumably private information unknown by any individual relational contract supplier. And, end client demand both overall and for any specific fabric are unknown by both the relational contract supplier and the buyer.

Importantly, relational contract suppliers need not believe that spot would be displaced *for all possible demand realizations*. Rather, believing that spot market supply could be displaced for *some* demand realizations is likely sufficient. Specifically, relational suppliers should find the buyer's threat credible when the probability of integrated supply, combined with other relational contract supply, displacing spot is sufficiently high. Relational contract suppliers would strategically accept the buyer's integrated supplier as the threat point as long as relational contract surplus under the new threat point is higher than the expected utility of keeping the pre-integration relational contract with spot as the outside option for both parties.[4]

Imperfect substitutability in production across fabrics also means that relational contract suppliers need not believe that the buyer has discontinued spot market sourcing for displacement to be credible. For example, the buyer might always purchase uncommonly used fabric from the spot market when no relational contract suppliers nor integrated supplier produce the specific fabric. Therefore, even if relational contract suppliers reasonably suspect that the buyer must purchase some fabric from the spot market, they might still find integration a credible threat for displacement as long as the integrated supplier produces many of the same fabrics as relational supply.

Therefore, I compute the empirical likelihood of displacement for common assumptions in the bargaining literature about beliefs for alternative suppliers: passive beliefs, symmetric beliefs, and wary beliefs (as in Lee et al. (2021)). In the context of my model and setting, relational contract suppliers with passive beliefs assume that other relational contract suppliers behave on-path. Therefore, I compute the residual demand facing any relational

[3]Assuming that the set of firms that have reached an agreement is known is common in much of the empirical literature. While this assumption seems credible in settings such as U.S. insurance markets, where insurance networks become public information, this assumption does not seem believable in the context of my empirical setting of informal contracts in Indian garment manufacturing.

[4]This expected utility is the sum of two products: *i)* the probability of displacement times supplier surplus when trading with the spot market and *ii)* one less the probability of displacement times supplier relational contract surplus when the buyer's threat point is spot market supply. Note that this expected utility assumes that the buyer agrees to keep its threat point as the spot market until a displacement occurs, consistent with optimal behavior on-path.

contract supplier in the period after integration as the total fabric volume sold to end clients after removing supply by other relational contract suppliers and compare it to the residual capacity available at the integrated supplier, where residual capacity is the capacity availabe at the integrated supplier assuming it first uses available capacity to displace spot market suppliers for fabrics that the integrated supplier can produce. I find that displacements occurs in 50% of the four quarters after integration periods across all relational contract suppliers, with all suppliers experiencing displacement during the same quarters because residual capacity at the integrated supplier is correlated across relational contract suppliers since relational contract suppliers fabric sets (imperfectly) overlap. With symmetric beliefs, I model each relational contract supplier as assuming that other relational contract suppliers have an identical relational contract in terms of both fabric and quantity guarantee as their own relational contract. Therefore, residual demand shifts to be the total demand for the fabrics the supplier produces less the the supplier's own quantity guarantee times the number of relational contract suppliers. I find that displacement occurs in all post-integration quarters under symmetric beliefs. Wary beliefs assume that each relational contract supplier has "worst-case" beliefs, meaning that the buyer will use its relational contract suppliers to strategically displace spot volume for the same fabrics as the relational contract supplier. Therefore, I model residual demand with wary beliefs as each relational contract supplier facing residual demand that is the total demand for the fabrics the supplier produces less the quantity guarantees across all other relational contract suppliers. As with symmetric beliefs, displacement occurs in all post-integration quarters with wary beliefs. Note that all of these belief systems implicitly assume that relational contract suppliers know the number of other relational contract suppliers, which is unlikely to be the case in this setting.

B.3 Multilateral Bargaining

The model presented abstracts from the buyer bargaining with all relational contract suppliers rather than just one, as illustrated in the main model. Therefore, I extend the bilateral model to the multilateral case and see what conditions are necessary to reproduce the bilateral result, where I us the same Nash framework for any bilateral negotiation.

B.3.1 Sequential Bargaining

Suppose that the buyer bargains with relational contract suppliers sequentially. As the model illustrates that agreement should always be reached, integrated capacity remains available to use as a threat point for each relational contract supplier. Therefore, as long as the buyer

bargains with relational contract suppliers first before purchasing residual quantities from spot market supply, the result holds. This assumption seems reasonable given that the buyer only has quarterly meetings with relational contract suppliers, and has no incentive to commit any earlier than necessary to spot suppliers. Similarly, if the buyer, the large, sophisticated firm in this setting, gets to choose the order of agreements with suppliers, the bilateral result also holds, as the buyer simply chooses to bargain with relational contract suppliers first. Note that because no relational contracts breakdown, the order within relational contract suppliers is irrelevant.

B.3.2 Nash-in-Nash

Suppose that the buyer bargains with relational contract suppliers, before sourcing from the spot market, in a Nash-in-Nash equilibrium (as described in Lee et al. (2021)))–meaning that "each pair's contract solves the bilateral Nash bargaining problem *taking the contracts agreed to by all other pairs as given.*" If the buyer can use the integrated supplier as the threat point simultaneously with each supplier, then the same equilibrium as the bilateral case prevails. Refinements of Nash-in-Nash that allows bargains between parties to consider alternative parties even not in a contract as possible threats for replacement, such in Nash-in-Nash with Threat of Replacement (NNTR as in Ho and Lee (2019))) or Nash-in-Nash with pairwise stability and replacement (as in Ghili (2022)), do not make sense in this setting, as it is not credible for either party to easily switch to a new relational contract buyer.

It is reasonable to wonder whether it would be credible for the buyer to threaten to use its integrated supplier as the threat point for each relational contract supplier simultaneously. Taking the Nash-in-Nash approach literally would, in principle, allow the buyer to use its integrated supplier as a threat simultaneously for all relational contract suppliers as, in equilibrium, no relational contract supplier breaks the contract. Of course, what threat point is actually used is not observable to any other relational contract supplier (*i.e.,* is private information) and even if other relational contract suppliers could see relational contract outcomes (which they cannot), they could not distinguish between the threat point used in the bargain and the other relational contract supplier having a different type in terms of risk aversion and bargaining parameter. Therefore, it does not seem plausible for any relational contract supplier to claim that the buyer must be using integrated supplier as a threat point for a different relational supplier such that integrated volume is not credible to displace it.

B.4 Parameterization and Solution of Conceptual Model

The model is parameterized as:

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

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

B.4.1 Log Quantities

Table B.1: Log Quantity and Capacity

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

Note: Data from universe of fabric transactions by buyer. Coefficients reported are from a regression of log days between fabric order date and delivery date on log quantity.

Table B.2: Log Quantity and Prices

	Price (Standardized)			
	(1)	(2)	(3)	(4)
Log Quantity	−.07***	−.1203***	−.2322***	−.2335***
	(.0065)	(.009)	(.0145)	(.0145)
Supplier FE	N	Y	Y	Y
Fabric Group FE	N	N	Y	Y
Month FE	N	N	N	Y
N	5456	5438	5396	5396
R^2	.0206	.1081	.2057	.2253

Note: Data from universe of fabric transactions by buyer. Coefficients reported are from a regression of log price on log quantity.

B.5 Structural Model of Relational Contract: Evidence Supporting Modelling Assumptions

B.5.1 Quantity Exogeneity

Although capacity utilization for any supplier is certainly not random, as the buyer provides demand assurance in the relational contract, supplier-period deviations from the targeted capacity in the relational contract are plausibly exogenous. Specifically, these deviations reflect stochastic end client demand, as the buyer knows neither which sample garments will be selected nor what quantities will be conditional on an end client purchasing a sample. It follows that the buyer strategically behaving to reduce capacity variance within a period does not invalidate capacity deviation exogeneity, but, rather, creates the data patterns that explain why capacity variance is lower for relational contract suppliers as compared to other suppliers.

However, buyer behavior that reduces *dynamic* capacity variation would bias estimates, as I would underestimate how much demand assurance the relational contract 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).[5]

To mitigate these concerns, I directly test the most likely examples of such behavior in Table B.3. Specifically, if the buyer provides additional dynamic capacity variance reduction, it would presumably take the form of ensuring that extreme capacity shocks are unlikely to

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

occur repeatedly. I directly test this hypothesis by examining the autocorrelation between extreme capacity realizations, which I define as absolute capacity deviations at 90th percentile of deviations or higher among relational contract suppliers. While I find that prior extreme capacity realizations are associated with lower probability of an extreme capacity realization, the magnitude is small and not statistically significant. Furthermore, the data suggest that the buyer does not even try to ensure that positive capacity shocks are followed by negative capacity shocks (and vice-versa).

Table B.3: Capacity Patterns

	Extreme Capacity (90th Percentile)		Capacity	
	(1)	(2)	(3)	(4)
Lagged Extreme Capacity	−.043	−.056		
	(.059)	(.048)		
Lagged Capacity			.024	.086
			(.059)	(.109)
Cluster Bootstrap 95% CI	[−.132, .064]	[−.117, .114]	[−.163, .173]	[−.202, .173]
Pre-Integration Only	N	Y	N	Y
R^2	.172	.082	.297	.171
N	216	108	216	108

Note: Data from universe of fabric transactions by buyer. Coefficients reported are from a regression using quarter by relational contract supplier panel data of extreme capacity indicator on lagged extreme capacity in columns 1-2 and capacity on lagged capacity in columns3-4. Robust standard errors are in parentheses and the 95% empirical confidence interval bootstrapped by supplier are shown.

Additionally, if parameter estimates are biased, then model prices should be biased. In this case, the untargeted moment validation exercise in 1.4.5.1 should not find that model prices are unbiased.

Another possible concern is that relational contract target quantities are adjusted after the vertical integration; however, this possibility seems unlikely as it is not in the buyer or supplier's interest to decrease the amount sourced from relational contracts given the large volume of fabric still procured from market suppliers that the buyer can displace with integrated supply. Regardless, if the target did decrease, then the estimation approach would systemtically overrestimate capacity shocks in the post-period, which would tend to increase the tuning parameter estimate. And, the increased tuning parameter estimate would lead to underestimates of supplier risk aversion. Although this concern seems unlikely given the small tuning parameter estimate (.002), I nevertheless test to see if relational contract average capacity decreases in the post period in Table B.4 by regressing capacity on an indicator for the post-period and total quantity of fabric procured across all other suppliers. I do not find any evidence that capacities are significantly different post-integration.

Table B.4: Capacity Stability

	Capacity
Post-Period Indicator	−.0099
	(.0141)
	[.0114]
Leave-Out Volume	.0003***
	(.0001)
	[.0001]
Supplier FE	Y
R^2	.338
N	234

Note: Data from universe of fabric transactions by buyer. Coefficients reported are from a regression using quarter by relational contract supplier panel data of capacity on a post-period indicator and total volume at the buyer from all but the supplier. Robust standard errors are in parentheses and standard errors clustered by supplier are in brackets.

Additionally, if parameter estimates are biased, then model prices should be biased. In this case, the untargeted moment validation exercise in 1.4.5.1 should not find that model prices are unbiased. Furthermore, out-of-sample fit should be poor if the model is systematically underestimating capacity utilization in the post-period (unless it is perfectly offset by creating bias in risk aversion of bargaining parameters, in which case the untargeted moment validation should fail). However, I find that strong out-of-sample fit in 1.4.5.2.

B.6 Structural Model of Relational Contract: Comparison of Risk Aversion Estimates

Table B.5: 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

B.7 Confidence Interval for Buyer Bargaining Parameter and Supplier Risk Aversion

Table B.6 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.

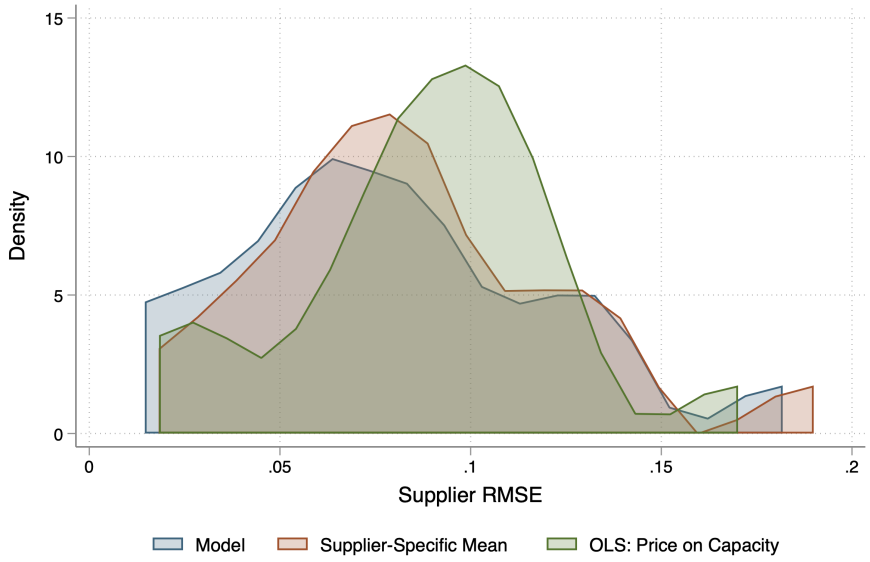
Table B.6: Buyer Bargaining Parameter and Supplier Risk Aversion: 95% CI

Buyer Bargaining Parameter		*Supplier Risk Aversion*	
Estimate	95% CI	Estimate	95% CI
0.8885	[0.0091, 0.9713]	0.0084	[0.0014, 0.0558]
0.9674	[0.0277, 0.9897]	0.0442	[0.0028, 0.0804]
0.5822	[0.0262, 0.8975]	0.0000	[0.0000, 0.0283]
0.0503	[0.0335, 0.9656]	0.0102	[0.0010, 0.0392]
0.0432	[0.0093, 0.4332]	0.0058	[0.0039, 0.0176]
0.0131	[0.0003, 0.0503]	0.0223	[0.0000, 0.0545]
0.0289	[0.0050, 0.9538]	0.0232	[0.0001, 0.0529]
0.6607	[0.0089, 0.9868]	0.0046	[0.0001, 0.0285]
0.9155	[0.0031, 0.9728]	0.0080	[0.0000, 0.0380]
0.0301	[0.0004, 0.9585]	0.0086	[0.0000, 0.0314]
0.0094	[0.0029, 0.9062]	0.0121	[0.0001, 0.0275]
0.9454	[0.3305, 0.9913]	0.0174	[0.0001, 0.1555]
0.8667	[0.0471, 0.9896]	0.0012	[0.0001, 0.0299]
0.9552	[0.0105, 0.9689]	0.0551	[0.0004, 0.1330]
0.3034	[0.0028, 0.9682]	0.0024	[0.0000, 0.0642]
0.9140	[0.0303, 0.9426]	0.0002	[0.0000, 0.0522]
0.0301	[0.0022, 0.9754]	0.0026	[0.0002, 0.0692]
0.9985	[0.9490, 1.0000]	14.6811	[0.1239, 31.2573]

Note: Model estimates.

B.8 Model Validation: Out-Of-Sample Fit RMSE Density

Figure B.5: Distribution of Supplier Squared Residuals by Estimation Method



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

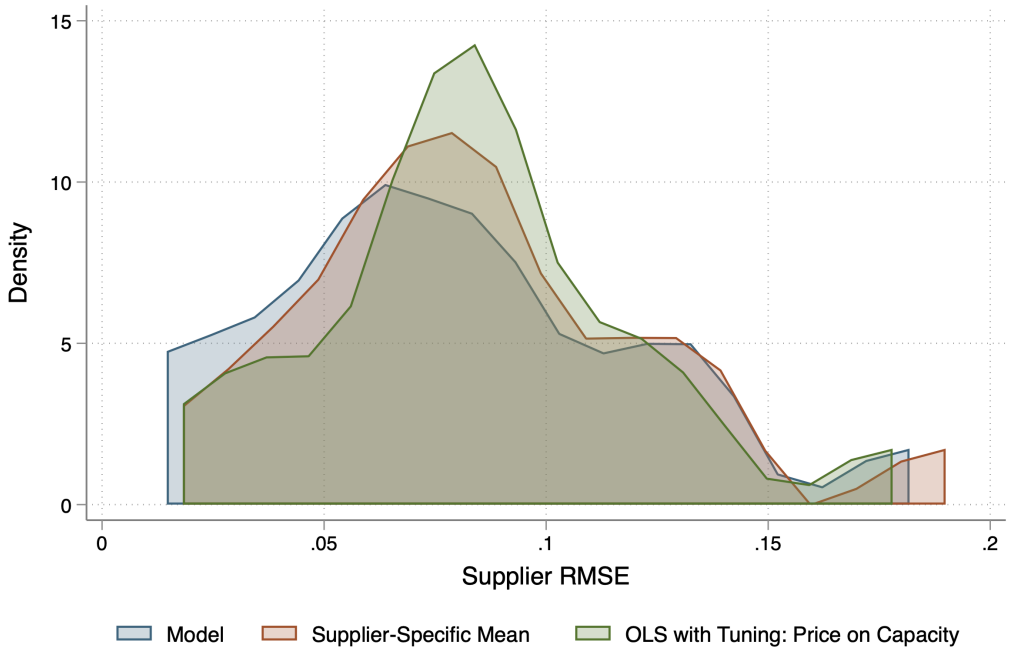
B.9 Model Validation: Out-Of-Sample Fit Comparison with OLS with Tuning Parameter

Table B.7: Out-of-Sample Fit Comparison (OLS with Tuning)

	Structural Model	Supplier Mean	OLS with Tuning
Minimum	.0146	.0182	.0182
25th Percentile	.0482	.0551	.0681
Median	.0741	.0790	.0812
75th Percentile	.1003	.1004	.1014
Max	.1818	.1900	.1780
Mean	.0787	.0844	.0843
Standard Deviation	.0441	.0417	.0379

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

Figure B.6: Distribution of Supplier RMSE by Estimation Method

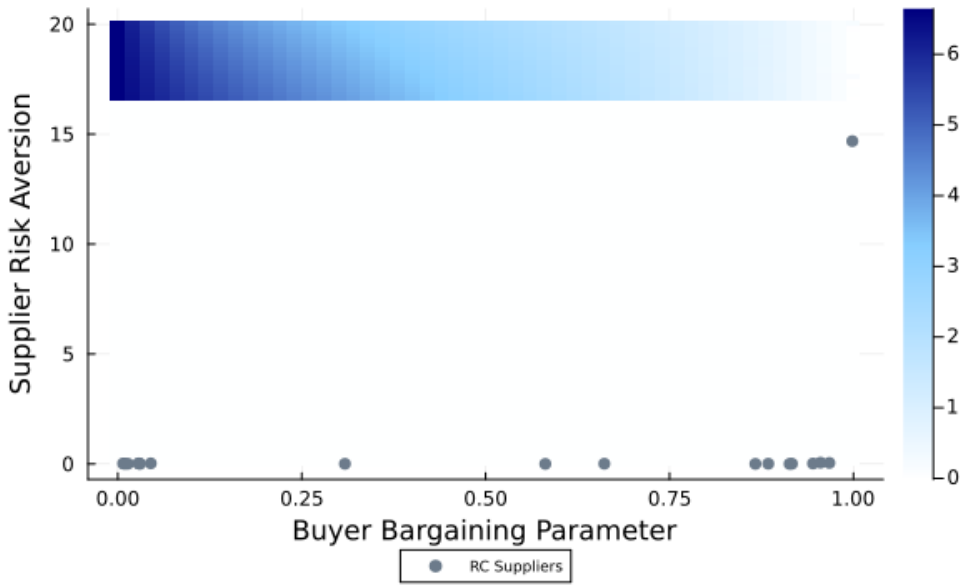


Note: Density of distribution of supplier root-mean squared error in the post- integration error by estimator.

B.10 Validation using Difference-in-Differences Estimates

B.10.1 Model-Implied Threat Point Effect for In-Sample Suppliers

Figure B.7: Threat Point Effect for In-Sample Suppliers



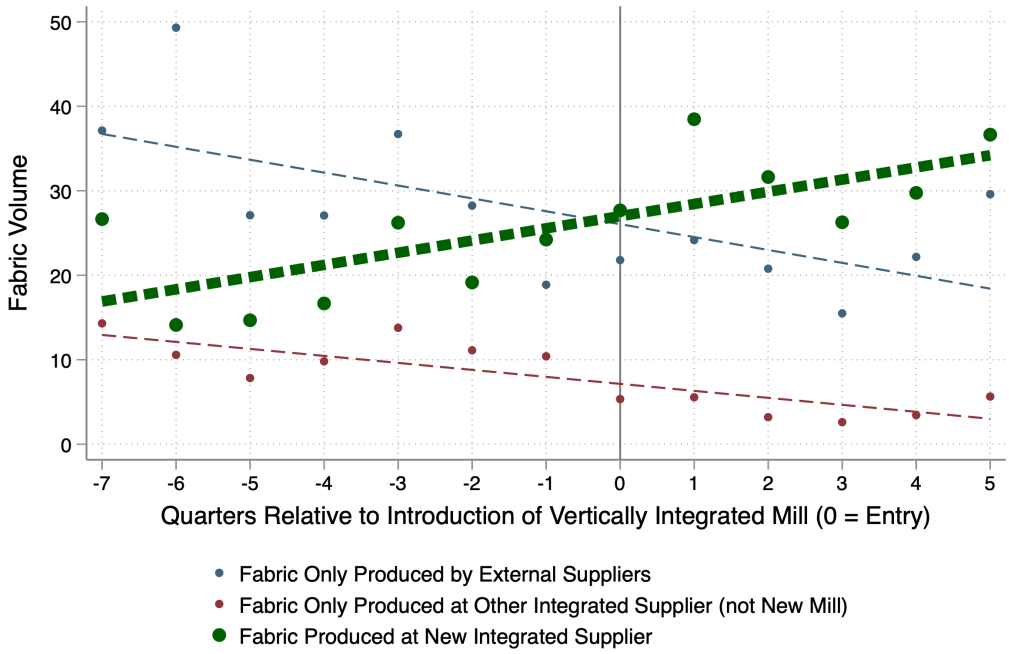
Note: Model estimates.

B.10.2 Doubly Robust Difference-in-Differences Design and Identification

In addition to bringing less production in-house for fabrics with more relational contract suppliers, Figure B.8 illustrates that fabrics with positive trends in demand were also more likely to be brought in-house. Therefore, as suppliers produce different fabrics, some suppliers are more or less likely to face exposure to vertical integration due to industry wide trends in fabric demand rather than supplier performance. To fix ideas using an extreme example, there could be a fantastic supplier which stopped receiving orders after integration due to the fabrics that they produce losing popularity. As trends in fabric demand are correlated with both a supplier's contract quantities and its exposure to vertical integration, counterfactual

outcomes of treated and control suppliers may not evolve in a parallel fashion, violating the parallel trends assumption. Additionally, the importance of trends around construction and initiation time for long-run supply decisions, such as building an integrated mill, is consistent with evidence from other settings, such as initial electricity prices having long-run effects on electricity consumption of manufacturing plants (Hawkins-Pierot and Wagner, 2022).

Figure B.8: Trends in Fabric Demand and Vertical Integration



Note: Data from universe of fabric transactions by buyer. Observations represent the volume of fabric purchased (in millions) across fabric groups, where fabrics are grouped based on whether the fabric is brought in-house or not. Trend lines are from linear trends.

The issue of trends in fabric demand motivates a doubly-robust difference-in-differences approach as in Sant'Anna and Zhao (2020), which weakens the parallel trends assumption to *conditional* parallel trends.[6] In this context, the *conditional* parallel trends assumption is that the change in potential outcomes for suppliers exposed to vertical integration is the same as the change in potential outcomes for suppliers unexposed to vertical integration *after* conditioning on trends in fabric demand that influence supplier exposure to vertical integration.[7] Formally, this assumption can be expressed as:

[6] An alternative option would be to use the method in Freyaldenhoven et al. (2019). Unfortunately, this approach would not work well in this setting, as the covariate adjustment is based on long-running trends rather than short-term contemporaneous movements. In other words, the covariate adjustment is mostly at the supplier, rather than supplier by time, level, which is not ideal variation for this method.

[7] Specifically, linear trend in fabric demand for each fabric from a regression of fabric volume across all

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

In equation B.1, trends in fabric demand are included in $\boldsymbol{X_i}$,[8] and $Y_i^{post}(0)$ denotes the untreated potential outcome in the post period with $D_i = 1$ indicating treatment units and $D_i = 0$ indicating control units.

With the doubly-robust approach, the estimated average treatment effect on the treated is unbiased if at least one of the propensity score model or the regression outcome model is correct. As there is only one treatment event, it is not necessary to adjust for heterogeneity in treatment timing, as in Callaway and Sant'Anna (2021); Sun and Abraham (2021). Importantly, unlike a synthetic control methodology to account for trends, this approach does not use any supplier-level outcome variables when constructing the propensity scores or the regression outcome model. 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 B.8 provides evidence that adjusting for propensity scores alone achieves balance between treatment and control suppliers by showing that treatment is no longer correlated with supplier-level outcome variables after adjusting for the propensity scores using inverse probability weighting. Because the propensity scores use fabric, and not supplier, trends,[8] it is not mechanical that this propensity score approach would achieve balance. I also document in Figure B.9 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.[9]

suppliers on the time period for all quarters in the pre-period. Supplier exposure to trends is calculated as the volume weighted average of the trends for the fabrics the supplier produces. It follows that the supplier's own volume, which is an outcome variable, should not drive results as trends are based on volumes for all suppliers, not just the individual supplier.

[8]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 seven quarters prior to integration that the supplier supplies, an indicator for being a relational contract supplier given that fabrics with relational contracts are less likely to be integrated, and the interaction of the relational contract indicator and the linear term for fabric growth and the count of distinct fabrics.

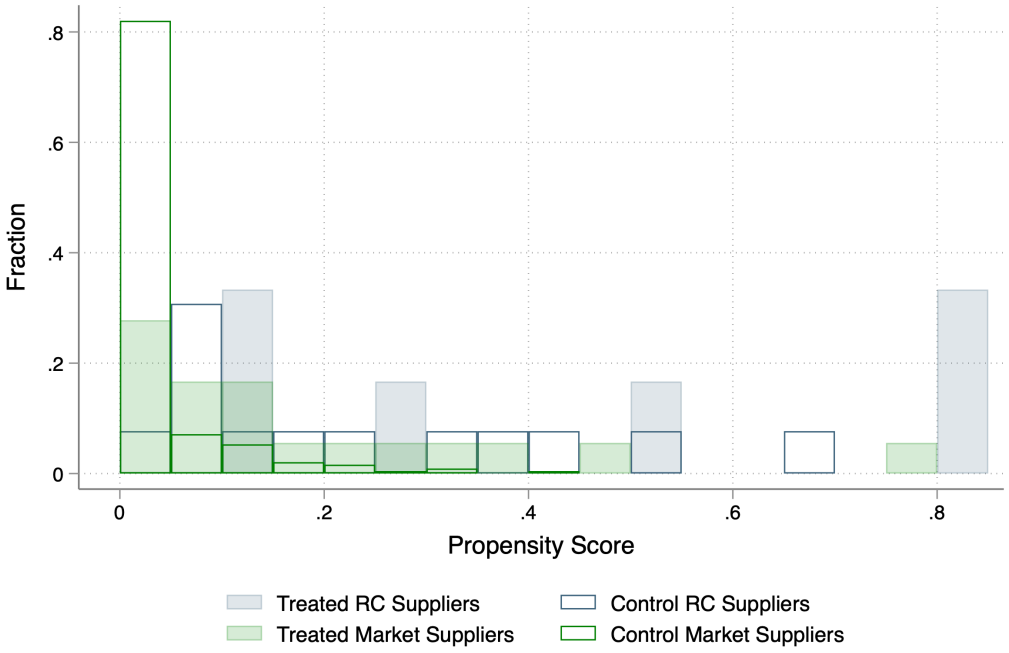
[9]Formally, as stated in Roth et al. (2022), "the conditional probability of belonging to the treatment group, given observed characteristics, is uniformly bounded away from one, and the proportion of treated units is bounded away from zero. That is, for some $\epsilon > 0, \mathbb{P}[D_i = 1|X_i] 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

Table B.8: Balance on Untargeted Outcomes

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

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

Figure B.9: Distribution of Propensity Scores by Supplier Type

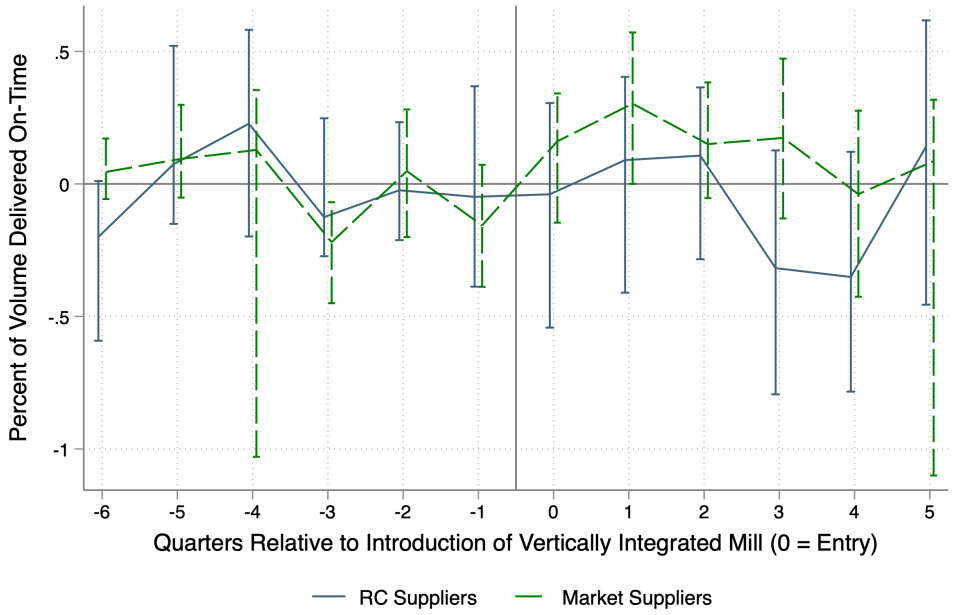


Note: Model estimates using pre-integration data.

B.11 Difference-in-Differences Supplemental Analysis and Robustness Checks

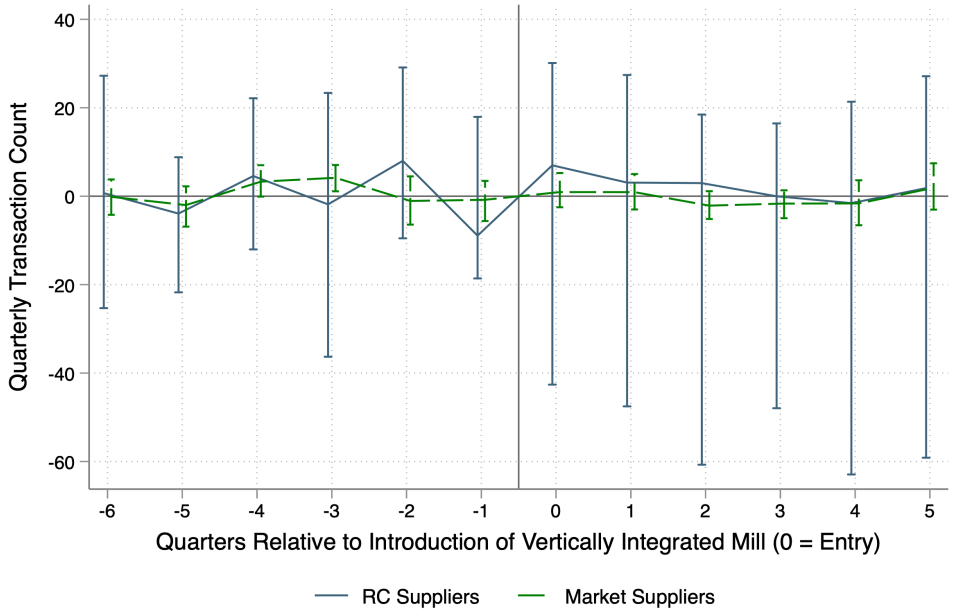
B.11.1 Difference-in-Differences Additional Event Study Plots

Figure B.10: Difference-in-Difference Estimates for Effect of Vertical Integration on Reliability



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

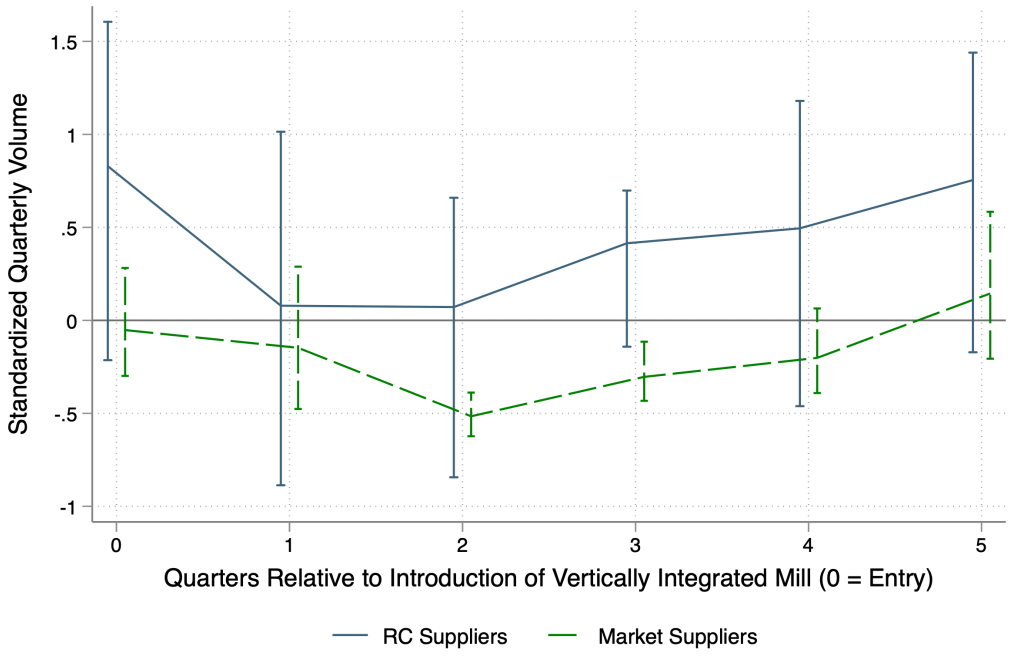
Figure B.11: Difference-in-Difference Estimates for Effect of Vertical Integration on Transaction Count



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

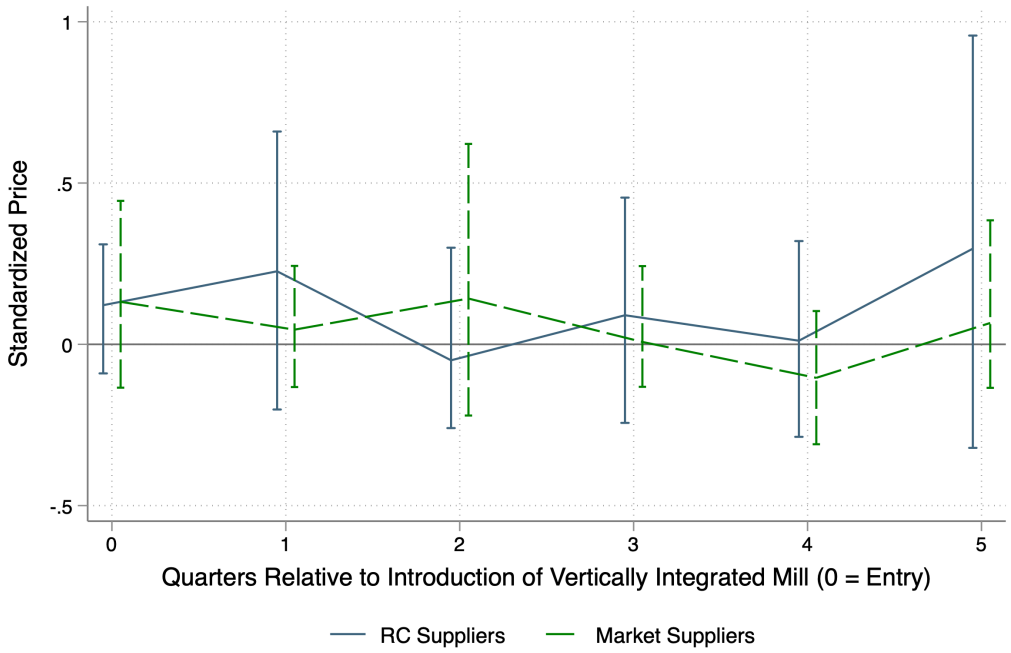
B.11.2 Synthetic Difference-in-Differences

Figure B.12: Synthetic Difference-in-Differences Estimates for Effect of Vertical Integration on Quarterly Volume



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

Figure B.13: Synthetic Difference-in-Differences Estimates for Effect of Vertical Integration on Standardized Prices



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

B.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%.[10]

I also consider the effects of a horizontal acquisition where the acquired firm is a small constrained supplier such that the threat point effect does indeed reduce relational contract prices before the horizontal acquisition. I find that prices increase by 24.6%. Decomposing the change as above, I find that reducing the variance increases price by 23.6% and reducing risk aversion similarly increases prices by 22.5%. These results have similar magnitudes because both of these effects of the merger are sufficient to stop the buyer from threatening to use the integrated supplier as the threat point. This result highlights the role of scale economies in effectively reducing the value of demand assurance, as unfavorable capacity states are less likely to be realized after integration due to the increase in quantities.

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

B.13 Buyer Bargaining Parameter and Supplier Market Power

I measure concentration in supplier's product markets by, first, computing supplier-fabric HHI as the HHI for the fabric across suppliers one quarter prior to integration and, second, aggregating the supplier-fabric HHI to the supplier level as the volume-weighted average of the individual supplier-fabric HHIs. As market power only pertains to relational contract suppliers, I only consider volumes from relational contract suppliers when computing supplier-fabric HHI (results using volumes from all suppliers to compute the supplier-fabric HHI are similar: a point estimate of -1.116 significant at the 5% level).

[10]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.

Table B.9: Buyer Bargaining Parameter and Supplier Concentration

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

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

I find in Appendix Table B.9 that the proxy for supplier market power is indeed strongly, and statistically significantly (using bootstrapped standard errors), correlated with the buyer bargaining parameter. Specifically, higher concentration in a supplier's product markets is negatively correlated with the buyer bargaining parameter, suggesting that suppliers with more market power receive more of the surplus created by the relational contract.

Note that this pattern does not hold for risk aversion, where the estimate is not statistically significant.

Table B.10: Supplier Risk Aversion and Supplier Concentration

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

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

APPENDIX C

Vertical Integration and Relational Contracts: Empirical Evidence on Substitutability from Indian Garment Manufacturing

C.1 Additional Analysis of Substitutability of Vertical Integration and Relational Contracts

Table C.1: Substitutability of VI & RC: OLS

	% Vertically Integrated (Fabric)			
	(1)	(2)	(3)	(4)
Number of RC Suppliers	−.0925***	−.1549**	−.0703***	−.1492***
	(.0327)	(.0624)	(.019)	(.039)
% Volume from Other		−.031		.4395***
Integrated Supplier (Pre-Integration)		(.1217)		(.1334)
Fabric % of Total Volume		16.1952***		7.8367**
		(2.8257)		(3.9553)
Trend in Fabric Demand		0***		0
(Quarterly Growth)		(0)		(0)
Number of All Suppliers		.0113		.013
		(.0268)		(.0131)
Months Pre and Post Integration	3	3	6	6
N	162	162	307	307
R^2	.1391	.647	.0981	.3662
Mean Dependent Variable	.0737	.0737	.0852	.0852

Note: Data from universe of fabric transactions by buyer. Regressions shown consider, alternatively, transactions from three or six months pre- and post-integration. Robust standard errors in parentheses and weights by fabric percent of total volume. ***,** ,* indicate significance at the 1, 5, and 10 percent levels, respectively.

Table C.2: Substitutability of VI & RC: Placebo and Supplier Market Power Robustness

	% Vertically Integrated (Fabric)			
	(1)	(2)	(3)	(4)
Supplier - Fabric HHI	−.3946	−.4236	−.6028	.031
	(.8973)	(.3227)	(.478)	(.517)
Number of RC Suppliers	−.5667***	−.4253***	−.3696***	−.4124***
	(.1817)	(.1303)	(.0756)	(.1142)
Number of Placebo Suppliers	.2398**	.1562	.1381*	.1164
	(.1074)	(.1017)	(.0747)	(.0926)
% Volume from Other		.4163		1.1824***
Integrated Supplier (Pre-Integration)		(.4578)		(.3911)
Fabric % of Total Volume		40.8674***		11.9455
		(8.2633)		(11.8949)
Trend in Fabric Demand		0***		0
(Quarterly Growth)		(0)		(0)
Number of All Suppliers		−.0477		.0223
		(.0627)		(.0605)
Months Pre and Post Integration	3	3	6	6
N	162	162	307	307
R^2	.1822	.5885	.1757	.2745
Mean Dependent Variable	.0737	.0737	.0852	.0852

Note: Data from universe of fabric transactions by buyer. Regressions shown consider, alternatively, transactions from three or six months pre- and post-integration. Robust standard errors in parentheses and weights by fabric percent of total volume. ***,** ,* indicate significance at the 1, 5, and 10 percent levels, respectively.

C.2 Relational Contract Robustness to Buyer Costs to Providing Demand Assurance

I augment the model to incorporate a cost to the buyer of providing demand assurance $\nu^{RC} < \nu$ as $\kappa(\nu)$. While this addition does not change the incentives for the supplier, now the buyer may have a profitable static deviation if they can pay the relational contract price but not provide demand assurance in the future. It follows that the constraints are:

$$
\begin{aligned}
(1-\delta)\Big(p^E q - c^B(q) - p^{RC} q - \kappa(\nu^{RC})\Big) + \delta\Big(p^E q - c^B(q) - p^{RC} q - \kappa(\nu^{RC})\Big) \\
\geq (1-\delta)\Big(p^E q - c^B(q) - p^{RC} q\Big) + \delta\Big(p^E q - c^B(q) - p^F q\Big) \\
\delta\Big((p^F - p^{RC})q - \kappa(\nu^{RC})\Big) \geq (1-\delta)\Big(\kappa(\nu^{RC})\Big) \\
\frac{\delta}{1-\delta}\Big((p^F - p^{RC})q - \kappa(\nu^{RC})\Big) \geq \kappa(\nu^{RC})
\end{aligned}
$$

Plug in for $p^{RC} = p^F - \phi\Delta c^S$:

$$
\frac{\delta}{1-\delta}\Big(\phi\Delta c^S q - \kappa(\nu^{RC})\Big) \geq \kappa(\nu^{RC})
$$

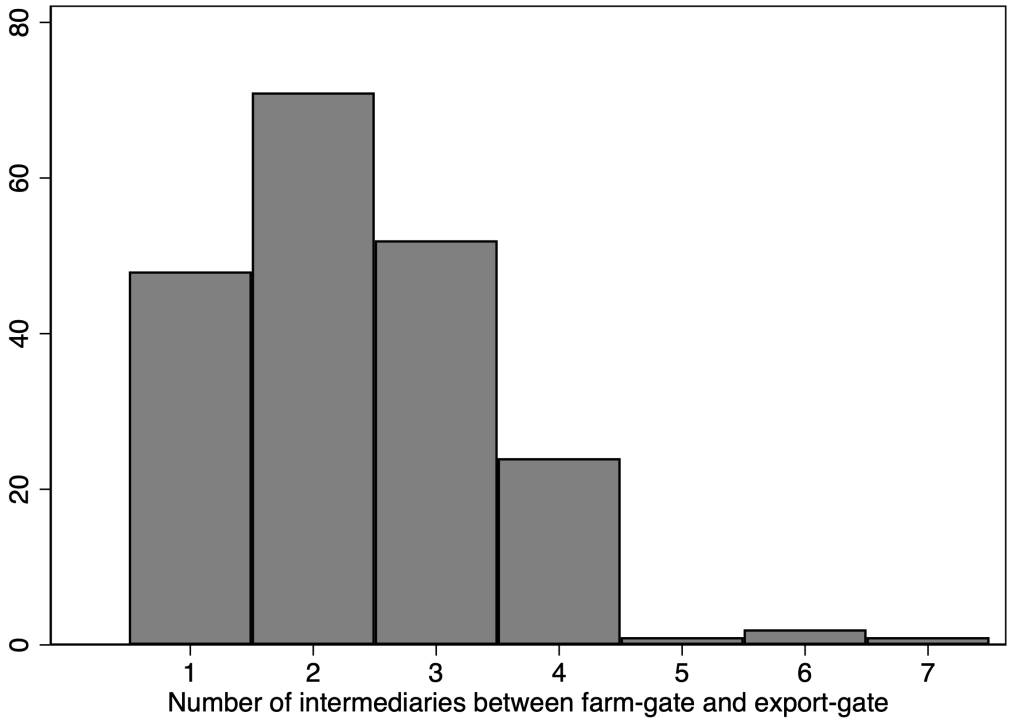
Therefore, as long as the costs to the buyer of providing demand assurance are sufficiently small that $\phi\Delta c^S q > \kappa(\nu^{RC}), \exists\delta \in (0,1)$ such that the relational contract exists. Specifically, as $\delta \to 1$, assuming that $\phi\Delta c^S q > \kappa(\nu^{RC})$, the right-hand side $\to +\infty$ and the relational contract equilibrium exists.

APPENDIX D

Quality Incentives and Upgrading in Uganda's Coffee Supply Chain

D.1 Intermediation: Supply Chain Lengths

Figure D.1: Distribution of Chain Lengths



Note: Data from the network mapping survey. Bars represent the number of chains with the relevant length. Length of 1 indicates that the agent sells directly to the big exporter. Length of 2 indicates that the agent sells to an intermediary which sells to the big exporter. Length of 3 indicates that the agent sells to an intermediary, which sells to another intermediary, which then sells to the big exporter.

D.2 Market Access

Table D.1: Big Exporters More Than Double the Quality Premium

	Price
High quality	283.54***
	(77.76)
Big exporters	625.28***
	(226.11)
High quality * Big exporters	418.64*
	(237.76)
Observations	1297
Mean DV	8275
Quantity controls	Yes
Month X SC FE	Yes

Note: Price, quality, and buyer identity data from the sampling surveys. Coefficients are from a regression of price on indicators for high quality, a sale to a big exporter, and their interaction, as well as controls for quantity and subcounty-by-month fixed effects. Robust standard errors are in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent levels, respectively.

Table D.2: Downstream Have Better Access to Big Exporters, Especially for High Quality

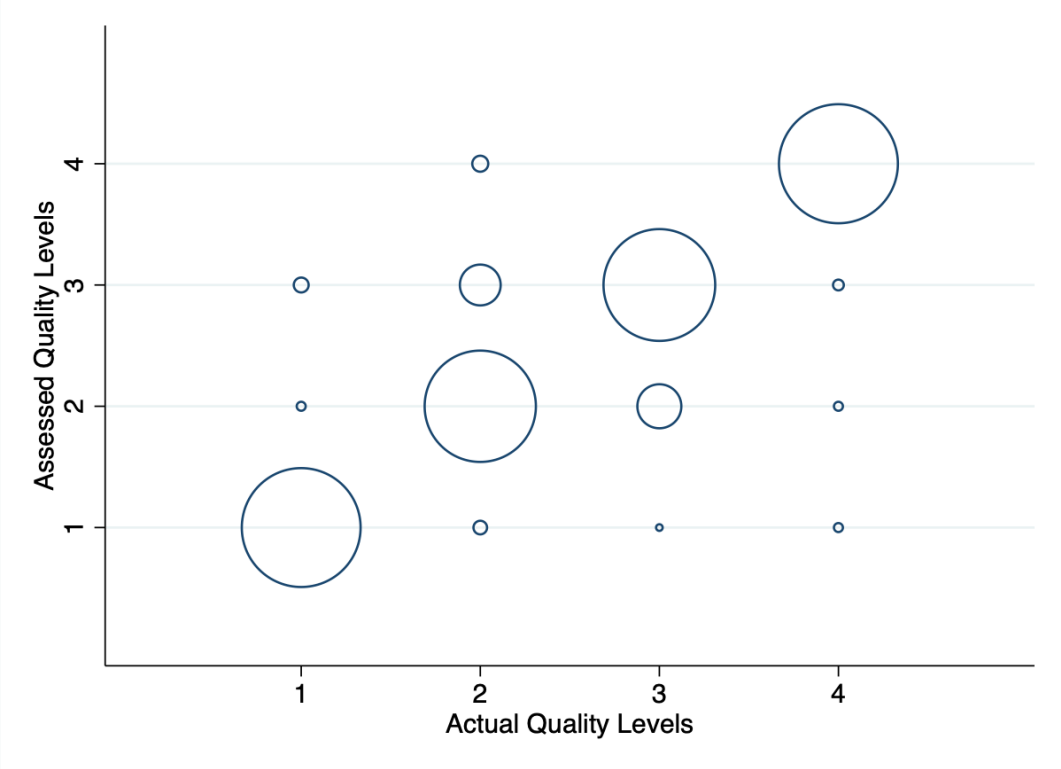
	Big exporters	Big exporters
Downstream	0.12***	0.09***
	(0.03)	(0.03)
High quality		-0.00
		(0.01)
High quality * Downstream		0.17***
		(0.06)
Observations	1450	1303
Mean DV	0.08	0.05
Quantity controls	Yes	Yes
Month FE	Yes	Yes

Note: Price and quality data from the sampling surveys and downstream is defined based on whether a respondent sells directly to the big exporter in the networking mapping survey. Coefficients are from a regression of an indicator for selling to a big exporter on indicators for downstream, high quality, and their interaction, as well as controls for quantity and month fixed effects (fixed effects do not control for subcounty as location is a "bad control" since it can affect an agent's position within the supply chain). Robust standard errors are in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent levels, respectively.

D.3 Alternative Explanations for Reduced Quality Premium Upstream

D.3.1 Quality Observability

Figure D.2: Quality Observability



Note: Data from producer and intermediary surveys. The size of circles correspond to the number of respondents who report the relevant actual and assessed quality levels.

D.3.2 Relational Contracts

Table D.3: Quality Premium Not Associated with Possible Relational Contracts

	Price	Price	Price	Price
High quality	331.22**	60.49	205.87	417.00
	(138.16)	(122.42)	(278.91)	(411.92)
Preference for RC	104.36			
	(215.83)			
High quality * Preference for RC	-240.93	267.95		
	(189.53)	(174.76)		
Interlinked Farmer			-457.66	
			(409.07)	
High quality * Interlinked Farmer			400.75	-683.18
			(531.66)	(528.57)
Downstream	301.81***		124.28	
	(103.02)		(187.20)	
High quality * Downstream	281.01*		822.41	
	(152.34)		(812.58)	
Observations	1030	671	77	47
Mean DV	8349	8420	8562	8789
Quantity control	Yes	Yes	Yes	Yes
Month FE	No	Yes	No	Yes
Subcounty FE	No	Yes	No	Yes
Month X SC FE	Yes	No	Yes	No
Product FE				
Individual FE	No	Yes	No	Yes
Buyer/Supplier FE	No	Yes	No	Yes

Note: Price and quality data from the sampling surveys and downstream is defined based on whether a respondent sells directly to the big exporter in the networking mapping survey. Preference for relational contract is an indicator for whether the respondent indicates that they prefer to trade with a repeated partner or trade for the best price as reported in the producer and intermediary surveys. Interlinked farmer refers to producers that report having borrowed money from downstream in the producer survey. Purchases by upstream are omitted from the sample, as it is not clear whether or not such transactions would be part of relational contracts as upstream may not be able to insure each other. Standard errors in parentheses are clustered at both the subcounty-by-subcounty and respondent levels. ***, **, and * indicate significance at the 1, 5, and 10 percent levels, respectively.

D.4 Search Costs

Table D.4: Upstream Sells to Many Distinct Buyers

# of Sales	Avg. # of Distinct Buyers	Count of Upstream With # of Sales
1	1.00	37
2	1.77	22
3	2.58	19
4	3.75	4
5	3.00	2
6	4.57	7
7	5.00	4
8	5.50	4
9	6.00	1
10	5.50	2
11	9.50	2
12	8.14	7
13	9.50	2
14	9.50	2
15	10.67	6
16	11.00	3
17	11.00	1

Note: Data from sampling surveys. The first column indicates the number of sales by an upstream agent. The second column shows the conditional average number of distinct buyers, where the condition is that the upstream agent had the number of sales as reported in the first column. The third column reports the number of upstream agents with the number of sales as reported in the first column. Upstream definition is based on network mapping data.

D.5 Model Calibration

In the calibration, I use two markets each with 100 agents both upstream and downstream. Prices and parameters take the following values:

Exporter prices:

$$\tilde{p}^L = 9000$$
$$\tilde{p}^H = 12000$$

Upstream cost parameters:

$$\alpha_1^U = 6000$$
$$\bar{\alpha}_2^U = 40$$
$$\sigma_{\alpha_2^U} = 8$$
$$\bar{\beta}^U = 2500$$
$$\sigma_{\beta^U} = 750$$

Downstream cost parameters:

$$\alpha_1^D = 1500$$
$$\bar{\alpha}_2^D = 30$$
$$\sigma_{\alpha_2^D} = 7$$
$$\bar{\beta}^D = 2700$$
$$\sigma_{\beta^D} = 900$$
$$F = 55000$$

BIBLIOGRAPHY

Acemoglu, Daron, Simon Johnson, and Todd Mitton, "Determinants of Vertical Integration: Financial Development and Contracting Costs," *The Journal of Finance*, 2009, *64* (3), 1251–1290.

Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham, "Returns to On-The-Job Soft Skills Training," *Journal of Political Economy*, 2023.

_ , Vittorio Bassi, Anant Nyshadham, and Jorge Tamayo, "No Line Left Behind: Assortative Matching Inside the Firm," *Working Paper*, 2019.

Alfaro, Laura, Davin Chor, Pol Antras, and Paola Conconi, "Internalizing global value chains: A firm-level analysis," *Journal of Political Economy*, 2019, *127* (2), 508–559.

Alviarez, Vanessa I., Michele Fioretti, Ken Kikkawa, and Monica Morlacco, "Two-Sided Market Power in Firm-to-Firm Trade," *Working Paper*, 2023.

Amodio, Francesco, Pamela Medina, and Monica Morlacco, "Labor Market Power, Self-Employment, and Development," *Working Paper*, 2023.

Antràs, Pol and Davin Chor, "Organizing the global value chain," *Econometrica*, 2013, *81* (6), 2127–2204.

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.

Asker, John, "Diagnosing Foreclosure due to Exclusive Dealing," *The Journal of Industrial Economics*, 2016, *64* (3), 375–410.

Atalay, Enghin, Ali Hortaçsu, and Chad Syverson, "Vertical Integration and Input Flows," *American Economic Review*, April 2014, *104* (4), 1120–48.

Atkin, David, Amit K. Khandelwal, and Adam Osman, "Exporting and Firm Performance: Evidence from a Randomized Experiment," *The Quarterly Journal of Economics*, February 2017, *132* (2), 551–615.

_ and Dave Donaldson, "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs," *NBER Working Paper 21439*, 2015.

_ , **Azam Chaudhry, Shamyla Chaudry, Amit K. Khandelwal, and Eric Verhoogen**, "Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan," *The Quarterly Journal of Economics*, March 2017, *132* (3), 1101–1164.

Azar, José, Ioana Marinescu, and Marshall Steinbaum, "Labor Market Concentration," *Journal of Human Resources*, 2022, *57* (S), S167–S199.

Bai, Jie, "Melons as Lemons: Asymmetric Information, Consumer Learning and Seller Reputation," *Working Paper*, 2024.

Baker, George, Robert Gibbons, and Kevin J. Murphy, "Relational Contracts and the Theory of the Firm," *The Quarterly Journal of Economics*, 2002, *117* (1), 39–84.

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.

_ , **Greg Fischer, Dean Karlan, Matt Lowe, and Benjamin N. Roth**, "Does the Invisble Hand Efficiently Guide Entry and Exit? Evidence from a Vegetable Market Experiment in India," *Working Paper*, 2022.

Bassi, Vittorio, Raffaelo Munio, Tommaso Porzio, Ritwika Sen, and Esau Tugume, "Achieving Scale Collectively Achieving Scale Collectively," *Econometrica*, November 2022, *90* (6), 2937–2978.

Bennett, Daniel and Wesley Yin, "The Market for High-Quality Medicine: Retail Chain Entry and Drug Quality in India," *The Review of Economics and Statistics*, March 2019, *101* (1), 76–90.

Bergquist, Lauren Falcao and Michael Dinerstein, "Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya," *American Economic Review*, 2020, *110* (12), 3705–3747.

_ , **Craig McIntosh, and Meredith Startz**, "Search Costs, Intermediation, and Trade: Experimental Evidence from Ugandan Agriculture Markets," *Working Paper*, 2024.

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*, November 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.

Board, Simon, "Relational Contracts and the Value of Loyalty," *American Economic Review*, December 2011, *101* (7), 3349–67.

Bodéré, Pierre, "Dynamic Spatial Competition in Early Education: an Equilibrium Analysis of the Preschool Market in Pennsylvania," *Working Paper*, 2023.

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

Bold, Tessa, Selene Ghisolfi, Frances Nsonzi, and Jakob Svensson, "Market Access and Quality Upgrading: Evidence from Four Field Experiments," *American Economic Review*, August 2022, *112* (8), 2518–52.

Boudreau, Laura, Julia Cajal-Grossi, and Rocco Macchiavello, "Global Value Chains in Developing Countries: A Relational Perspective from Coffee and Garments," *Journal of Economic Perspectives*, August 2023, *37* (3), 59–86.

Breza, Emily and Andres Liberman, "Financial Contracting and Organizational Form: Evidence from the Regulation of Trade Credit," *The Journal of Finance*, 2017, *72* (1), 291–324.

Brugues, Felipe, "Take the Goods and Run: Contracting Frictions and Market Power in Supply Chains," *Working Paper*, 2020.

Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel, "Sell Low and Buy High: Arbitrage and Local Price Effects in Kenyan Markets," *The Quarterly Journal of Economics*, December 2018, *134* (2), 785–842.

Busso, Matias and Sebastian Galiani, "The Causal Effect of Competition on Prices and Quality: Evidence from a Field Experiment," *American Economic Journal: Applied Economics*, January 2019, *11* (1), 33–56.

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.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *The Review of Economic Studies*, 2015, *82* (1), 1–44.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230. Themed Issue: Treatment Effect 1.

Chatterjee, Shoumitro, "Market Power and Spatial Competition in Rural India," *The Quarterly Journal of Economics*, January 2023, *138* (3), 1649–1711.

Chor, Davin and Lin Ma, "Contracting frictions in global sourcing: Implications for welfare," *Working Paper*, 2020.

Ciliberto, Federico and Jonathan W. Williams, "Does multimarket contact facilitate tacit collusion? Inference on conduct parameters in the airline industry," *The RAND Journal of Economics*, 2014, *45* (4), 764–791.

Costinot, Arnaud, Jonathan Vogel, and Su Wang, "An elementary theory of global supply chains," *Review of Economic studies*, 2013, *80* (1), 109–144.

Coulter, Jonathan and C. Poulton, "Cereal Market Liberalization in Africa," Technical Report 1999.

Crinò, Rosario and Laura Ogliari, "Financial imperfections, product quality, and international trade," *Journal of International Economics*, 2017, *104*, 63–84.

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

de la Sierra, Raul Sanchez, "Whither Formal Contracts?," *Econometrica*, 2021, *89* (5), 2341–2373.

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

Fafchamps, Marcel and Ruth Vargas Hill, "Price Transmission and Trader Entry in Domestic Commodity Markets," *Economic Development and Cultural Change*, 2008, *56* (4), 729–766.

Fan, Haichao, Yao Amber Li, and Stephen R. Yeaple, "Trade Liberalization, Quality, and Export Prices," *The Review of Economics and Statistics*, 2015, *97* (5), 1033–1051.

Forbes, Silke J. and Mara Lederman, "Adaptation and Vertical Integration in the Airline Industry," *American Economic Review*, December 2009, *99* (5), 1831–49.

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.

Fujiy, Brian Cevallos, Gaurav Khanna, and Hiroshi Toma, "Cultural Proximity and Production Networks," *Working Paper*, 2023.

Galenianos, Manolis and Alessandro Gavazza, "A Structural Model of the Retail Market for Illicit Drugs," *American Economic Review*, March 2017, *107* (3), 858–96.

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.

Gaynor, Martin, "What Do We Know About Competition and Quality in Health Care Markets?," *Foundations and Trends in Microeconomics*, December 2006, *2* (5), 441–508.

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.

Gil, Ricard, Myongjin Kim, and Giorgio Zanarone, "Relationships Under Stress: Relational Outsourcing in the U.S. Airline Industry After the 2008 Financial Crisis," *Management Science*, 2022, *68* (2), 1256–1277.

Goettler, Ronald L. and Brett R. Gordon, "Does AMD Spur Intel to Innovate More?," *Journal of Political Economy*, 2011, *119* (6), 1141–1200.

Grant, Matthew and Meredith Startz, "Cutting Out the Middleman: The Structure of Chains of Intermediation," *Working Paper*, 2024.

Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartzstein, "Learning Through Noticing: Theory and Evidence from a Field Experiment," *The Quarterly Journal of Economics*, June 2014, *129* (3), 1311–1353.

Hansman, Christopher, Jonas Hjort, Gianmarco León-Ciliotta, and Matthieu Teachout, "Vertical Integration, Supplier Behavior, and Quality Upgrading among Exporters," *Journal of Political Economy*, 2020, *128* (9), 3570–3625.

_ , _ , Gianmarco León-Ciliotta, and Matthieu Teachout, "Vertical Integration, Supplier Behavior, and Quality Upgrading among Exporters," *Journal of Political Economy*, 2020, *128* (9), 3570–3625.

Harris, Adam and Thi Mai Anh Nguyen, "Long-Term Relationships and the Spot Market: Evidence from US Trucking," *Working Paper*, 2023.

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

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

Hortaçsu, Ali and Chad Syverson, "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices," *Journal of Political Economy*, 2007, *115* (2), 250–301.

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.

Macchiavello, Rocco, "Financial Development and Vertical Integration: Theory and Evidence," *Journal of the European Economic Association*, April 2012, *10* (2), 255–289.

_ , "Relational Contracts and Development," *Annual Review of Economics*, 2022, *14* (1), 337–362.

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

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

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

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

Morlacco, Monica, "Market Power in Input Markets: Theory and Evidence from French Manufacturing," *Working Paper*, 2020.

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

_ , "Vertical Integration & Relational Contracts: The Threat Point Effect," *Working Paper*, 2023.

Nyqvist, Martina Björkman, Jakob Svensson, and David Yanagizawa-Drott, "Can Good Products Drive Out Bad? A Randomized Intervention in the Antimalarial Medicine Market in Uganda," *Journal of the European Economic Association*, December 2021, *20* (3), 957–1000.

Osborne, Theresa, "Imperfect competition in agricultural markets: evidence from Ethiopia," *Journal of Development Economics*, 2005, *76* (2), 405–428.

Park, Sang Hyun, Kathryn Lundqist, and Victor Stolzenburg, "Global Value Chain Report 2023: Resilient and Sustainable GVCs in Turbulent Times," Technical Report 2023.

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

Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe, "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *Working Paper*, 2022.

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

_ , "Labor Market Power and Factor-Biased Technology Adoption," *Working Paper*, 2024.

Sant'Anna, Pedro H.C. and Jun Zhao, "Doubly robust difference-in-differences estimators," *Journal of Econometrics*, 2020, *219* (1), 101–122.

Staatz, John M., Josué Dioné, and N.Nango Dembélé, "Cereals market liberalization in Mali," *World Development*, 1989, *17* (5), 703–718.

Startz, Meredith, "The value of face-to-face: Search and contracting problems in Nigerian Trade," *Working Paper*, 2021.

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

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

The World Bank, "World Development Report 2020: Trading for Development in the Age of Global Value Chains," Technical Report 2020.

Treuren, Leonard, "Wage markups and buyer power in intermediate input markets," *Working Paper*, 2023.

Verhoogen, Eric, "Firm-Level Upgrading in Developing Countries," *Journal of Economic Literature*, December 2023, *61* (4), 1410–64.

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

Zavala, Lucas, "Unfair Trade? Monopsony Power in Agricultural Value Chains," *Working Paper*, 2023.