ABSTRACT

Title of dissertation:	ESSAYS ON PLATFORM MERGERS AND DYNAMIC PRICE COMPETITION		
	Xinlu Yao Doctor of Philosophy, 2021		
Dissertation directed by:	Professor Andrew Sweeting Department of Economics		

This dissertation covers two innovative topics related to horizontal mergers. One chapter investigates the welfare effect of mergers in the platform markets; another one studies the impact of mergers where firms conduct dynamic price competition with asymmetric information.

In Chapter 1, I analyze the welfare effects of lowering the costs to buyers of searching and multihoming in a setting with multiple two-sided platforms. The analysis is motivated by observed changes following the 2017 acquisition of IronPlanet, which is an online auction marketplace for used heavy equipment by Ritchie Brothers Auctioneers, which operates the largest offline auction marketplace. As is quite common after platform mergers, RBA maintained both platforms but made it easier for buyers of equipment to search across the platforms (multihoming), which has the potential to render the allocation of equipment more efficient, benefitting both buyers and sellers. These efficiencies could offset the market power created by the merger. I use pre- and post-merger transaction data to estimate a new model of search and auction entry by buyers and quantify the increase in welfare effects of the observed changes. Depending on the specification, the proportion of multihoming buyers increases substantially (by 50% in the baseline specification), and the total surplus can increase by more than 8%, although heterogeneity exists in the welfare impact on different market participants. I also consider several additional counterfactuals involving changes in commission and changes in equipment allocation across the marketplaces.

In Chapter 2, my coauthors and I model differentiated product pricing by firms that possess private information about serially-correlated state variables, such as their marginal costs, and can use prices to signal information to rivals. In a dynamic game, signaling can raise prices significantly above static complete information Nash levels even when the privately observed state variables are restricted to lie in narrow ranges. We calibrate our model using data from the beer industry, and we show that our model can explain changes in price levels and price dynamics after the 2008 MillerCoors joint venture.

ESSAYS ON PLATFORM MERGERS AND DYNAMIC PRICE COMPETITION

by

Xinlu Yao

Dissertation submitted to the Faculty of the Graduate School of the University of Maryland, College Park in partial fulfillment of the requirements for the degree of Doctor of Philosophy 2021

Advisory Committee: Professor Andrew Sweeting, Chair Professor Maureen Cropper Professor Ginger Jin Professor Chenyu Yang Professor Jie Zhang © Copyright by Xinlu Yao 2021

Acknowledgments

First of all, I would express my deep and sincere gratitude to my main advisor Professor Andrew Sweeting for his dedicated support and guidance during my P.h.D study. It was my great privilege and honor to work and study under his guidance. Through many research projects, Professor Sweeting was always enthusiastic about providing encouragement and assisting in any way he could. The talks with him help me build my skill set and encourage me to learn new things continuously. His rigorous academic attitude, sincerity, and kindliness have deeply inspired me. Without his help, I cannot complete my dissertation.

I would also like to thank Professor Maureen Cropper, Professor Ginger Jin, and Professor Chenyu Yang for providing advice to my study and help to my job search process. By working with them, I have got the chance to learn many interesting topics and develop my skills. Besides, they provide me with many useful instructions, which helps me to overcome the obstacles in completing my dissertation.

I want to express many thanks to all participants who took part in my P.h.D study and enabled this dissertation to be possible. Professor Laurance Ausubel, Professor John Shea, Professor Daniel Vicent and Professor Jie Zhang offered me helpful feedback about my research. Also, I get many supports from senior students in my department. Particularly, I want to thank Yongjoon Park for his patience and help with my study.

Finally, I would like to thank the National Science Foundation for their sup-

port. The paper in Chapter 1, titled "Platform Mergers in Search Markets: An Application in the U.S. Used Heavy Truck Market," was supported by NSF Grant SES-1949348. The paper in Chapter 2, titled "Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Horizontal Mergers," was supported by NSF Grant SES-1658670 (Sweeting).

Table of Contents

Acknow	edgements	ii
Table of	Contents	iv
List of T	bles	viii
List of H	gures	X
1 Plat Truc 1.1 1.2	rm Mergers in Search Markets: An Application in the U.S. Used Heavy Market. ntroduction	$ \begin{array}{c} 1\\ 1\\ 10\\ 10\\ 11\\ 12\\ 13\\ 15\\ 18\\ 20\\ 23\\ 24\\ 25\\ 28\\ 28\\ 28\\ 28\\ 28\\ 28\\ 28\\ 28\\ 28\\ 28$
	1.3.6 Search Choice Stage 1.3.7 Economics of Search 1.3.7.1 Search Frequency 1.3.7.2 Homing Choices	32 35 35 38
1.4	identification	40 42 42 44 46
		тυ

	1.5	Estima	tion		47
		1.5.1	Paramet	ric Assumptions and Normalizations	47
		1.5.2	Algorith	m	48
			1.5.2.1	Solving Equilibrium Bidding and Auction Entry Stra	te-
				gies	49
			1.5.2.2	Estimating the Distributions of WTP and Equilib-	
				rium Search Choice	50
			1.5.2.3	Search Cost	54
	1.6	Estima	tion Resu	ılts	55
		1.6.1	Distribut	tion of WTP and Distribution of Search Costs	55
		1.6.2	Model F	it and Sensitivity Analysis	60
	1.7	Counte	erfactuals		62
		1.7.1	Roadma	D	62
		1.7.2	Welfare	Analysis with Changes in Search Costs	64
		1.7.3	Welfare	Analysis with Changes in Commission	68
		174	Welfare	Analysis with Changes in Supply Side	71
	18	Discus	sion		77
	1.0	181	No Addi	tional Data About Trucks in the Outside Market	77
		182	With Da	ta About Transacted Trucks in the Outside Market	80
	19	Conclu	ision	a mout manuacteu marks in the outside market .	83
		Contere			00
2	Dyn	amic Ol	igopoly P	ricing with Asymmetric Information: Implications for	or
	Hor	izontal l	Mergers		85
	2.1	Introd	uction		85
	2.2	Model			90
		2.2.1	Outline.		90
		2.2.2	Within-F	Period Timing.	91
		2.2.3	Assumpt	ions	91
		2.2.4	Fully Se	parating Equilibrium in a Finite Horizon and Con-	
			tinuous	Type Game.	92
			2.2.4.1	Final Period (T) .	93
			2.2.4.2	Earlier Periods $(1, \dots, T-1)$.	93
			2.2.4.3	Characterization of Separating Pricing Functions	
				in Period $t < T$.	94
			2.2.4.4	Existence and Uniqueness of a Fully Separating	
				Equilibrium.	97
	2.3	Examp	oles	*	98
		2.3.1	Continu	ous-Type Duopoly Example.	98
			2.3.1.1	Specification.	98
			2.3.1.2	Equilibrium Outcomes and Strategies.	99
		2.3.2	Merger A	Analysis.	105
		2.3.3	Signalin	g Incentives and the Existence of Separating Equi-	
			libria.	· · · · · · · · · · · · · · · · · · ·	108
		2.3.4	Addition	al Examples.	109
	2.4	Empiri	ical Appli	cation: The MillerCoors Joint Venture	110
		r.			

		2.4.1	The JV and Its Effects.	110
		2.4.2	Data	111
		2.4.3	Calibration of the Dynamic Asymmetric Information Model.	113
			2.4.3.1 Products	115
			2.4.3.2 Demand	115
			2.4.3.3 Marginal Costs	118
			2.4.3.4 Objective Function, Matched Statistics and Identi-	
			fication.	118
			2.4.3.5 Parameter Estimates and Model Fit.	123
			2.4.3.6 Predicted Effects of the JV	124
		2.4.4	Testing Alternative Explanations for the Post-JV Price In-	
			creases	127
	2.5	Conclu	usion	134
A	App	endix fo	or Chapter 1	136
	A.1	Supple	ement to Market and Data	136
		A.1.1	Market	136
		A.1.2	Cross-sectional Variation	136
	A.2	Supple	ement to Model	139
		A.2.1	Markets with at Most Two Quality Levels	139
			A.2.1.1 One Platform	139
			A.2.1.2 Two Platforms	144
		A.2.2	Markets with More Than Two Quality Levels	145
			A.2.2.1 One Type of Buyers	146
		- I	A.2.2.2 Two Types of Buyers	146
	A.3	Supple	ement to Identification	150
		A.3.1	Baseline	150
		A.3.2	Extensions	151
	A.4	Supple	ement to Estimation	153
		A.4.1	Algorithm for Solving Equilibrium Payoffs of an Auction	153
	A.5	Supple	ement to Estimation Results	154
		A.5.1	Estimation Results for Two-type Model	154
		A.5.2	Model Fits	154
		A.5.3	Sensitivity Analysis	156
	A.6	Supple	ement to Counterfactuals	159
		A.6.1	Algorithm Used to Solve the New Equilibrium in the Coun- terfactuals	159
				107
В	App	endix fo	or Chapter 2	160
	В.1	Comp	utational Algorithms	160
		B.1.1	Finite Horizon Model.	160
			B.1.1.1 Preliminaries.	160
			B.1.1.2 Period T	163
			B.1.1.3 Period $T = 1,,$	163
			B.1.1.4 Earlier Periods	- 167

	B.1.2	Infinite l	Horizon Model.	167
		B.1.2.1	Speeding Up Solutions By Interpolating Pricing Func-	-
			tions	168
		B.1.2.2	Tolerances and Updating Rules Used for the Esti-	
			mation of the Cost Parameters Using the Infinite	
	D 1 0		Horizon Model.	169
D O	B.1.3	Two-Typ	e Model	170
B.2	Additi	onal Exar	nples. \ldots	173
	B.2.1	Iwo-Iyp		1/3
		B.2.1.1	Outcomes for Alternative Serial Correlation and	174
		D D 1 D	Demand Parameters.	1/4
		B.2.1.2	consistence of the Conditions Required for Existence of	176
		D 9 1 9	Effect of the Number of Eirms on Equilibrium Out	170
		D.2.1.3	comes	170
	B 2 2	Alternat	ive Sources of Asymmetric Information	1/7
ВЗ	D.2.2 Evistor	nce and I	Iniqueness of a Fully Separating Equilibrium in a	101
D.5	Finite	Horizon (Game with Linear Demand	186
	B 3 1	Prelimin	ary Results	188
	B.3.2	Main Re	sult.	189
		B.3.2.1	Proof.	190
B.4	Data			199
	B.4.1	IRI Data		199
		B.4.1.1	Data Selection for the Demand and Conduct Anal-	
			ysis	199
		B.4.1.2	Data Selection for the Calibration of Our Model.	202
	B.4.2	Addition	al Empirical Analyses.	202
		B.4.2.1	Effects of the Joint Venture on Prices	202
		B.4.2.2	Market Shares Around the Joint Venture	205
		B.4.2.3	Price Correlations Across Brands Before and After	
			the JV	207
		B.4.2.4	Price Dynamics in Los Angeles and Seattle Around	
			the JV (Real Prices)	209
B.5	Testin	g the Sup	ermarkup Model	210
	B.5.1	Testing t	the Supermarkup Model Version 1	211
	B.5.2	Testing t	The Supermarkup Model Version 2	212
Bibliogr	aphy			214

List of Tables

1.1	Commission Pre- and Post-Merger	15
1.2	Characteristics and Prices of the Trucks Sold on IP and RBA	17
1.3	Hedonic Regression Used to Construct the One-Dimensional Quality	
	Index	18
1.4	Sensitivity of Price on IP to the Trucks Available on Each Platform .	23
1.5	Estimation Results of the Model with One Type of Buyers	56
1.6	New Equilibrium Under Alternative Changes in Search Costs	64
1.7	Welfare Under Alternative Changes in Search Costs	65
1.8	Average Trading Surplus and Split at Different Quality Levels (1) .	67
1.9	New Equilibrium Under Reserve Prices and Alternative Changes in	
	Commission Rates	70
1.10	Welfare Under Reserve Price and Alternative Changes in Commis-	
	sion Rates	71
1.11	New Equilibrium Under Alternative Changes in Supply Side	72
1.12	Welfare Under Alternative Changes in Supply Side	73
1.13	Average Trading Surplus and Its Split at Different Quality Levels (2)	75
21	Fauilibrium Prices and Welfare in the Duopoly Game	100
2.1	The Effects of Signaling on Mergers and Merger Analysis When	100
2.2	Firms Use Infinite Horizon Signaling Strategies	106
2.3	Equilibrium Pricing in a Finite Horizon Game with Alternative Cost	100
2.0	Specifications	109
2.4	Highest-Selling Beer Brands in 2007 with Ownership. Share and	107
	Average Prices.	112
2.5	Estimates of Demand	116
2.6	AR(1) Price Regressions Using Flagship Market-Pack Size-Week or	-
	-Month Data	121
2.7	Parameter Estimates for Six Specifications	123
2.8	Model Fit for Three Specifications Using Weekly Data, Average Price	
	Elasticity of -3 and Flagship Diversion of 85%	125
2.9	Predicted Average Prices Before and After the MC JV	126
2.10	Observed and Predicted Changes in Price Dynamics	128

2.11	Testing Alternative Models Using a Generalized Conduct ParameterFrameworkFramework	131
A.1	Characteristics of Trucks in the TX Transaction Data and Auction	105
	Transaction Data	137
A.2	Estimation Results for Two-type Model	155
A.3	Model Fitness of Targeted Moments	157
B.1	Price Effects in Models Where Alternative Elements of Firm Objec-	
	tive Functions are Private Information	182
B 2	Estimates of the Effects of the Joint Venture on Prices	204
D 0	Cross Drand Correlations in Drisss for 12 Desla	207
в.3	Cross-Brand Correlations in Prices for 12-Packs	207

List of Figures

1.1 1.2 1.3 1.4 1.5	Auction Locations of RBA Definition of Regions Quality Distribution Across Platforms Price Distributions of Trucks in Different Quality Groups Quality Distributions of Trucks Purchased by Different Types of Buyers Timpling of the Came	12 14 19 21 22 25
1.0 1.7	Net Expected Dayoffs $W(m \text{ home}^*)$ When $M = 5$ and $M = 10$	25 26
1.7	Implied Distribution of WTP and distribution of equilibrium search choices	56
1.9	Cumulative Distribution Functions of Search Cost in the One-type	
1.10 1.11	Model	58 59
	Post-merger	61
1.12	Number of trucks and Average Trading Surplus in Different Quality Groups	76
1.13	Average Trading Surplus in Different Markets	76
1.14	Timeline of the Game With Endogenous Platform Entry of Sellers (1)	78
1.15	Timeline of the Game With Endogenous Platform Entry of Sellers (2)	81
2.1	Period <i>T</i> and <i>T</i> – 1 Pricing Strategies in the Finite Horizon, Continuous Type Signaling Game	101
2.2	Expected $T - 1$ Period Profit Function: $c_{1,T-1} = 58.025$ and $c_{1,T-2} = c_{2,T-2} = \8	103
	Game and Various Periods of the Finite Horizon Game.	104
2.4	Average Nominal Prices (excluding sales) of 12-Packs of the Domes-	
2.5	tic Flagship Brands in Two Regional Markets Around the JV Estimated Pre-JV Price Dynamics and the Combined Market Shares	114
0.6	of AB, Miller and Coors.	122
2.6	(1) of Table 2.7).	127
A.1	Pictures of Truck Tractors	136

A.2	Web Pages of RBA in 2016 and 2018	137
A.3	Price Distribution Across Platforms	137
A.4	The Distribution of Truck Quality and Prices Paid by Different Types	
	of Buyers in Texas	138
A.5	Variation in the Number of Trucks Across Markets	138
A.6	Expected Payoffs from Entering an Auction $U^*(q, v)$ (One Type of	
	Buyers)	147
A.7	Auction Entry Behavior of Buyers	147
A.8	Expected Payoffs from Entering an Auction $U(q, v)$ (Two types of	
	Buyers)	149
A.9	Measurements and Assumptions for the Identification of Different	
	Models	152
A.10	CDF of Search Costs in Two-type Model	156
A.11	Sensitivity Analysis of the One-Type Model	158
р 1	Fourilibrium Arrange Drive Increases in the Infinite Herizon Two	
B.1	Equilibrium Average Price increases in the infinite Horizon Two-	
	tion of Costs	175
ро	LIOH OF COSts	1/5
Б.2	Equilibrium Prices in the two-type Marginal Cost Model (parame-	176
рр	Deriod T_{i} (c. Drofit Europtions in the Two Type Come	170
D.J	Ferrou $T = 0$ Front Functions in the Two-Type Game	1/0
D.4	Equilibrium Effects on Average Prices in the minine Horizon Two-	
	Firme	100
ЪГ	Plillis	100
D.Э D.6	Dialiu Warket Shares Around the John venture	205
D.0	Average Real Plices (excluding sales) of 12-Packs of the Domestic	200
	Flagsnip Brands in Two Regional Markets Around the JV.	209

Chapter 1: Platform Mergers in Search Markets: An Application in the U.S. Used Heavy Truck Market.

1.1 Introduction

Transaction platforms, especially digital ones, that link buyers and sellers play an increasingly important role in the economy. As of August 2019, digital platforms exceed more than four trillion dollars in market capitalization.¹ When different goods are listed for sale on different platforms, multihoming by buyers (i.e., searching multiple marketplaces) can raise the efficiency of the ultimate allocation. One efficiency that might be considered when analyzing a merger of platforms is that the merged firm could facilitate multihoming by developing cross-platform search tools after the merger. This type of cognizable efficiency could potentially offset any market power created by the merger.²

The partial platform integration, in which both platforms remain distinct but cross-platform searching is facilitated, is a common outcome of platform mergers. For example, after acquiring the resale ticket platform StubHub, eBay reports the

¹https://www.chicagobooth.edu/research/stigler/events/ antitrust-competition-conference

²The Supreme Court decision regarding *American Express vs. Ohio* indicates it is important to evaluate welfare effects on both sides of a transaction platform.

listings available on StubHub when a buyer searches on eBay. Similarly, after their acquisition by CoStar Group, Apartment.com and ApartmentFinder.com use a cross-search mechanism to help users search for apartments on both platforms.

While two competing platforms may be reluctant to facilitate multihoming since it might intensify competition for single-homing sellers (Caillaud and Jullien (2003), Armstrong (2006), Armstrong and Wright (2007)), a merged platform might seek to encourage multihoming to increase the surplus the platforms generate, which it may have a greater ability to extract. In this paper, I quantify several welfare effects of this type of integration by analyzing the effects of the 2017 acquisition of Iron Planet (IP), an entirely online auction marketplace for used heavy equipment, by Ritchie Brothers Auctioneers (RBA), the largest traditional auction marketplace that operated in 31 physical auction sites across the United States. The transaction was subject to a second request investigation by the U.S. Department of Justice's Antitrust Division (DOJ), but was not ultimately challenged.

I use data of sales of used truck tractors from both platforms before and after the acquisition to estimate a model of buyers' choices of whether to multihome, how many listings to search and which auction to enter to quantify how buyer behavior changed after the acquisition and the effects of this change on the surplus of different market participants. I find that multihoming by buyers substantially increased, and there was some increase in the average number of searches. Despite the small increase in commission, I find that the changes in buyer behavior increased the combined welfare of buyers and sellers and the platform's revenues. This finding is consistent with the DOJ's decision. As I will discuss, the platforms accounted for a relatively small share of the overall used truck market, even though they accounted for up to 60% of the truck auction market. Therefore, the fact that the merger created only limited market power in this setting is unsurprising. However, the welfare benefits due to the lower search costs are substantial and suggesting that these types of efficiencies should also be seriously considered in settings in which a merger may create more market power.

This paper makes two contributions. The first contribution is providing the first evaluation of a change in platform design that is commonly associated with platform mergers in which both platforms exist post-merger. Most empirical existing platform literature focuses on traditional media markets (Rysman (2004), Argentesi and Filistrucchi (2007) and Jeziorski (2014)) and the mechanisms used on one online platform (Arnosti et al. (2014), Fradkin (2017) and Horton (2019)). The trade-off between changes in search costs and other policy changes has not been clearly discussed while considering the competition and mergers between platforms.

To the best of my knowledge, the only study that also empirically analyzes the effects of online platform mergers is Farronato et al. (2020). These authors consider the acquisition in pet-sitting services in which one platform was shut down post-merger and provide a reduced-form analysis of how platform use changes. In contrast, I estimate a structural model of buyer search to evaluate the welfare effects of changes that facilitated search across platforms.

The second contribution is to develop and to estimate a new buyer search

model that combines choices regarding whether to multihome, how many truck auctions to search, and which auction to enter. This model, combining search with endogenous entry, can be applied to other transaction platforms where allows products to be vertically differentiated and sold by auction.

My model works in the following way. Given their private draws of marginal search costs for additional search and fixed search costs to multihome, buyers decide whether to search both platforms and how many auctions to search. In the used truck auction market, the marginal search costs include an effort to search for trucks online, consult with sales representatives, etc., while the fixed costs of multihoming include the time to learn two systems, register two accounts to monitor trucks, etc. Searching allows buyers to discover individual trucks' characteristics and their private values. Then, buyers simultaneously decide which searched auctions to enter. Finally, buyers bid in the entered auctions. I develop a numerical method to solve the equilibrium in the auction entry stage. I consider two variants of the model. In one variant, all buyers have the same preference for quality, and all single-homing buyers follow the same random entry rule when choosing platforms, and in the other, buyers can have two different preferences for quality and can follow different rules while choosing platforms when single homing. The two-type model is motivated by the fact that buyers who are trucking companies can operate locally or interstate.

The model provides several predictions regarding the market outcome under certain assumptions. It predicts that lower search costs will encourage buyers to perform more searches across the two platforms. Then, buyers can access more information about the trucks before making their auction entry choices. Highquality trucks can be sold to a set of buyers with higher willingness to pay (WTP). Meanwhile, the lower search costs provide buyers access to more trucks across the platforms, resulting in the allocation results of one platform more sensitive to the quality of trucks available on the other platform. These predictions are consistent with the patterns in the pre- and post-merger transaction data.

I estimate a parametric version of the model, although I can prove the nonparametric identification in some cases when there is enough variation in the sets of available trucks across markets. I estimate the model using a two-step procedure. In the first step, I adopt a Nested Fixed Point Algorithm to estimate the distribution of WTP and equilibrium search choices. In the second step, I estimate the bounds of the distribution of search costs by calculating the benefits of different search choices and using the equilibrium conditions in the model.

The estimation results show that buyers' WTP depends on the observed quality of trucks (which I reduce to a single index) and that buyers discount the quality of trucks sold online. Regarding the distribution of search costs, the marginal search cost per search and the fixed search cost of performing multihoming both significantly decrease after the merger. Although buyers' searches are strategic substitutes, the lower search costs still encourage buyers to search for more trucks on average (in the model with one type of buyers, the average number of searches increases from 5.8 to 6.3). Following the merger, more than 50% of buyers shifting from single-homing to multihoming in the baseline model. The increase in the search frequency of buyers with a high-quality preference choosing single-homing offline is more significant than that of buyers with a low-quality preference.

Based on the estimation results, I quantify the welfare effects of the changes associated with the merger. I focus on the following three types of changes: search costs, commission fees and supply side (numbers and types of trucks available on each platform). I capture the partial effect of different elements in different counterfactuals by controlling for other changes.

In the first and main counterfactual, I look at the effect of changes in search costs, keeping the commission rates and supply-side fixed. This comparison shows that the merger can increase the total surplus by more than 8%, among which 6% comes from better matches, and the rest comes from saving in search costs. The total surplus of the buyers and sellers from trading significantly increases. However, the split of the trading surplus among the participants is uneven. While sellers with high-quality trucks can always benefit, it is more ambiguous for other groups. For example, buyers' trading surplus is lower post-merger, considering the fiercer competition among buyers. The cost decomposition shows there is efficiency gain from lowering the cost to multihome alone.

The second counterfactual discusses how the changes in commission structures can impact social welfare. When there is no reserve price, the observed change in commission rate transfers a share of the surplus from sellers to platforms, but the total surplus is the same if we treat the supply side exogenous. Suppose the platforms use auctions with reserve prices and set significantly higher commission rates. In that case, buyers may be discouraged from conducting more searches because buyers' expected payoffs from more searches become lower. However, based on my calculation, it requires a more significant increase in commission rate to offset the efficiency gain from lower search costs in this merger, given the supply side fixed.

Finally, I analyze the additional welfare effect from a possible change on the supply side following the merger. The way to construct the possible change is motivated by the data: post-merger, high-quality goods are more likely to be listed offline, and low-quality goods are more likely to be listed online. I consider a model of two types of buyers where the interstate companies have a higher estimated quality preference and are assumed to choose the offline platform if they conduct single-homing. The results show that this change can generate additional benefits: specifically, with the post-merger search costs, it can increase the total trading surplus by about 2%. This is because it can help buyers with different search strategies to target the goods they prefer more easily.

The reader should be aware of the limitation of my analysis. This paper focuses on buyers' behavior while considering sellers and platform decisions exogeneous because of limited data of other auction platforms and the computation burden. I analyze the changes in sellers and platforms in the counterfactual part. Therefore, the indirect network effect in this platform market cannot be properly analyzed. I discuss the plan to endogenize sellers' platform entry and incorporate indirect network effect in the paper. Additionally, actual buyers' search data are not publicly available; thus, the distribution of search costs is estimated based on the observed transaction and bidding data.

7

Related Literature. This paper builds upon the literature concerning multi-sided platforms, search, and auction.

Most theoretical papers (Caillaud and Jullien (2003), Rochet and Tirole (2003), Rochet and Tirole (2006), Armstrong (2006), and Weyl (2010)) have focused on prices when discussing mergers and assume that no search costs exist in the market. These papers construct models to analyze the competition between multi-sided platforms. They focus on the number of users on two sides and the indirect network effect rather than the composition of users and matching distortion between users on different sides. Bardey and Rochet (2010) is among the very few papers that consider vertically differentiated users in the health insurance market.

In addition to the literature concerning traditional media markets, recently, more studies have focused on online platforms', Arnosti et al. (2014) mention the potential congestion in the matching market with costly screening and uncertain availability. Fradkin (2017) discusses transaction costs and potential congestion in the Airbnb market. As a typical format of online platforms, online auction markets are analyzed in several papers. Krasnokutskaya et al. (2020) study the role of an online procurement market. These authors develop a way to estimate primitives when unobserved seller heterogeneity exists. Bodoh-Creed et al. (2016) discuss efficiency in decentralized auction platforms. Marra et al. (2019) show how the careful design of a commission structure can improve welfare in a wine auction platform with network effect exits. However, most of the analysis focuses on mechanism design within one online platform.

The search model used in this paper is related to Allen et al. (2014) and Salz

(2020). Allen et al. (2014) point out the importance of considering the search costs in the market when analyzing the merger effect in the Canadian Mortgage industry. Salz (2020) discusses the function of intermediation in New York City's trade waste market. Both papers use a non-sequential search model (De los Santos et al. (2012)) and introduce a competition stage to determine the price rather than posted price setting (Hortaçsu and Syverson (2004)). In their papers, with more searches, consumers can access more lenders/carters (corresponding to sellers in this paper), and lenders/carters compete in the auctions by offering the lowest prices to consumers. Differently, in my model, by searching for more goods, buyers can observe the quality and private values of these goods. Based on this information, buyers make their auction entry choices and compete in the auctions.

Therefore, my model includes a buyers' endogenous auction entry stage. Levin and Smith (1994) and Athey et al. (2011) study the endogenous auction entry model, where Athey et al. (2011) compare the sealed bid and open formats in the U.S. Forest Service timber auctions. They assume bidders have information about their private value after entering the auctions, so bidders are not selective. Different from these papers, my model involves buyers choosing which auction to enter, assuming that they know their values of the goods being sold in each auction that they search.³

This paper proceeds as follows. Section 1.2 introduces the market and data. I illustrate several descriptive findings in the data. Section 1.3 describes the game

³There is an extensive literature on entry into single auctions under different information assumptions. My model assumes that buyers know their values, as in Samuelson (1985), but more importantly, they are choosing which, of several auctions, to enter rather than considering an "in/out" entry choice into a single auction.

played by buyers and discusses the economics of search choices. Sections 1.4 and 1.5 describe the identification and estimation strategies. Section 1.6 presents the structural estimates. Section 1.7 is the counterfactual part and shows the welfare analysis of different policy changes from the merger. Section 1.8 talks about the plan to relax the assumption of exogenous sellers. Section 1.9 concludes the paper.

1.2 Market and Data

This section first provides an overview of the market and acquisition. Then, I discuss the data used in this paper and summarize some interesting findings observed in the data.

1.2.1 Market for Used Truck Tractors

In this paper, I study two platforms, i.e., RBA and IP, on which used truck tractors and many other types of heavy equipment are sold through auctions.

1.2.1.1 Channels for Used Truck Tractor Sales

Each month, approximately 20,000 used heavy trucks are sold in the U.S. Among these sales, auctions account for about 10%-15% of used heavy trucks in the U.S. each year. Although the auction channel's market share is less than that of some other intermediaries, such as retailers, it allows sellers to sell "as is, where is," namely, sellers need less certification to sell their trucks via auctions than via other channels. Therefore, the auction market is irreplaceable, and on average, the used trucks sold through auctions are older or have higher mileage. Many different body-style trucks exist, and different trucks are used for different purposes. Truck tractors are among the most popular body-style of heavy trucks (see Figure A.1). The owners of these trucks are usually transportation trucking companies, operating locally or interstate. The inventory of trucks of local trucking companies is much smaller than that of interstate companies. To operate interstate, companies need to register their trucks under an interstate registration plan.⁴

1.2.1.2 Ritchie Bros. Auctioneers

RBA is a primary auction platform that sells heavy industrial equipment through onsite auctions. It has 31 physical auction sites located nationwide (Figure 1.1), and most locations are concentrated in states with large heavy-machinery markets. Post-merger, three auction sites closed.⁵ The dependence of the locations affects the offline auction frequency. Texas has ten auction events every year, but there are only four large auction events each year in Maryland. Although offline auctions are less frequent than online auctions, RBA still accounts for a much larger market share than IP (approximately 4:1). One reason is that many buyers prefer the local inspection opportunities provided by the offline platform.

RBA lists the trucks it will sell in the next two months on its websites. Sellers who choose offline auctions need to transfer the trucks to the auction sites.

⁴The International Registration Plan (the Plan) is a registration reciprocity agreement among the states of the United States and provinces of Canada providing payment for license fees based on the total distance operated in all jurisdictions. https://www.irponline.org/page/ThePlan The trucking companies under this Plan usually operate interstate.

⁵https://www.bizjournals.com/triangle/news/2017/11/10/ the-amazon-effect-ritchie-bros-closes-five.html





Notes: the black points represent the auction sites that closed after the merger.

Since sellers need to pay a high penalty fee if they withdraw the trucks in a short window, I assume that sellers conduct single-homing only. This assumption is consistent with the observation in the data. On the auction days, different auction rings sell various items simultaneously. Buyers can bid in any auction online or in-person with registration. Trucks are sold via English Auctions without reserve prices.

1.2.1.3 IronPlanet

IP is a leading pure online used truck auction platform. There is no location restriction on the online platform. Regardless of their location, all buyers have the same information regarding the trucks sold on IP and can place their bids once allowed online. On the website, information regarding the trucks sold in the next two weeks is available. IP holds auctions every Thursday. A very low starting bid is given by the platform in each auction. Buyers can choose to place a proxy bid before the auction day. On the auction day, several auctions are held almost simultaneously. Since a buyer cannot withdraw her bids in an auction, it is difficult for her to manipulate several auctions if she only has single-unit demand. The auctions proceed very quickly; thus, if a buyer loses in one auction, she loses the opportunity to participate in another auction in which she is interested.

Both RBA and IP sell heavy machinery in addition to trucks. The aggregate market share of these two platforms in the used heavy truck market is no greater than 60% since some other auction platforms exist in the market. ⁶

According to the bidding data and local investigation, buyers also search for and purchase trucks in adjacent states. Therefore, in my analysis, I define "markets" at the region-month level. Figure 1.2 shows the four regions defined in this paper. In each market, the analysis includes all RBA auctions during that month in the region, and all IP auctions are treated as if they occur simultaneously.

1.2.2 Merger

The acquisition was announced in August 2016 and finalized in May 2017. After a second request for additional information under the HSR Act, the case

⁶This information was obtained by consulting with experts in this industry.





received unconditional antitrust clearance from the DOJ.⁷ Some policies related to both price and information changed after the merger.

First, RBA increased the commission rate of its physical auctions to make it more similar to IP. RBA and IP charge buyers different rates for different transaction prices: the rate for a lower-priced truck is proportionately higher than the rate for a higher-priced truck. Table 1.1 shows the commission rates of both platforms before and after the merger. The commission rate for trucks with final prices \$5,000 to \$33,500 increased from 2.5% to 3.85%. All other factors are similar. The second change that I focus on is RBA's integration of the platform media, websites, and support teams to allow users to easily search for and request information regarding trucks on both platforms. First, RBA and IP release news and emails to provide information regarding these two platforms.⁸ Second, both platforms built

⁷https://www.rbauction.com/media/news-releases/archives/2017/ 0170518-rba-ip-secure-antitrust-clearancehttps://www.reuters.com/article/ idCNASC09NTD

⁸The following is an example of news released by RBA after an offline auction in TX: "...,"

	IP		RE	BA	
4	2016 & 2018		2016	2018	
Price	Commission	Price	Commission	Price	Commission
<10	10%	<2.5	10%	< 5	10%
10-33.5	$\min\{3.85\%,\$1,000\}$	2.5-33.5	$\min\{2.5\%, \$950\}$	5-33.5	$\min\{3.85\%, \$500\}$
>33.5	\$1,290	>33.5	\$1,290	>33.5	\$1,290

Table 1.1: Commission Pre- and Post-Merger

Notes: unit of the price is \$1,000.

new websites with a cross-listing mechanism. Both RBA and IP changed their Websites. Before the merger, buyers could only find the trucks sold on the platform they entered; after the merger, buyers can easily find some information regarding the trucks sold on the other platform regardless of which website they enter (Figure A.2). Finally, the platforms share the same customer service. After the merger, the customer service team of either platform can help buyers obtain information regarding the trucks sold on both platforms.

These cross-platform mechanisms enable buyers to easily find information regarding the trucks on both platforms, potentially decreasing the search costs in the market.

1.2.3 Data

The primary data set contains transaction data collected from IP's and RBA's websites. From both platforms, I obtain all transactions of used trucks from 02/01/2016 to 09/30/2016 (the pre-merger period) and from 02/01/2018 to

said Alan McVicker, Regional Sales Manager, Ritchie Bros. "Bidders were very active, competing on a great selection of equipment consigned from more than 650 owners. For those buyers unable to get what they needed in Houston, we have an online Iron Planet auction today (Thursday, April 18) with close to 1,000 items available."https://www.rbauction.com/news-releases/ 20190418-ritchie-bros-sells-us47-million-of-equipment-in-houston-tx-this-week

09/30/2018 (the post-merger period). The transaction data includes the transacted prices, truck characteristics (VIN, age, mileage, make and model), and listing characteristics (platform, transaction price, location, date, and commission fees paid to the platforms). I do not have data of other small truck auction sites. The following additional information of IP after the merger is considered: bids in auctions, each bidder's location (states), and bidding time. To obtain more information regarding the truck models, I match the trucks to the Truck Blue Book to obtain information regarding each model's suggested retail price (MSRP). I assume that the trucks in the transaction data represent an approximation of trucks available in the market. Although the data only shows the set of trucks that went transactions, as RBA uses no reserve-price auctions and IP lowers the starting prices in online auctions when there are no bidders, this assumption should be close to satisfied.

Table 1.2 summarizes the characteristics and transaction price of the trucks sold on both platforms before and after the merger. As shown in the table, the trucks' transaction volumes were similar before and after the merger. However, the quality distributions of the trucks sold on these two platforms changed: the trucks sold on IP are significantly older and have a significantly higher average mileage post-merger. I further discuss this analysis in the descriptive findings.

Since many different observed characteristics can be used to measure trucks' quality, it will be hard to do a structural estimation to include all of them directly. Therefore, I use an additional data set to figure out a one-dimensional quality index.

16

	2016				2018	
Measurements	IP	RBA	Combined	IP	RBA	Combined
Total Number	1,423	6,932	8,355	1,690	6,838	8,528
Freightliner	226	567	793	178	1,384	1,562
International	323	2,422	2,745	660	1,884	2,544
Kenworth	125	1,049	1,174	176	806	982
Mack	301	1,058	1,359	184	721	905
Others	67	323	390	216	249	465
Peterbilt	227	1,027	1,254	86	986	1,072
Volvo	154	486	640	190	808	998
Avg. Price	14,024	17,898	17,169	7,492	17,866	15,040
Avg. Age	9.70	9.27	9.35	11.85	9.12	9.66
Avg. Log(Mileage)	12.69	12.66	12.67	13.03	12.94	12.96
Avg. MSRP	72,945	75,618	75,159	71,543	79,870	78,148

Table 1.2: Characteristics and Prices of the Trucks Sold on IP and RBA

Notes: 1. unit of price and MSRP is \$; 2. unit of Mileage is mile.

The data set I use includes the registration data of trucks in Texas.⁹ This data set includes transactions of used trucks in TX through all channels (retailers, wholesalers, large fleets, and auctions) from 01/01/2016 to 08/31/2018. As mentioned above, less friction exists in other channels than the auction channel, so other channels' transaction prices are more closely related to the trucks' quality. Also, the transaction does not contain RBA and IP only. Therefore, I regress the trucks' transaction price in this data set on the observed characteristics listed in Table 1.2 and construct the one-dimensional quality measurement. Table 1.3 shows the hedonic regression used to construct the quality index.

Finally, by combining the registration data with the licensing under the International Registration Plan in Texas,¹⁰ I can identify the trucks purchased by different types of buyers. Using the VIN as the trucks' unique identity, I can match this data set to the auction transaction data. Since Texas has the largest number

⁹Source: DMV of Texas.

¹⁰Source: DMV of Texas.

VARIABLES	logprice
log(mile)	-0.0438***
-	(0.00351)
log(mrsp)	0.228***
	(0.0225)
Constant	3.710***
	(0.122)
Diesel Dummy	-0.0709***
	(0.0130)
Make Dummies	Y
Age Dummies	Y
Observations	46,545
R-squared	0.526
Notes: standard extremes *** p<0.01	rrors in paren- , ** p<0.05, *

Table 1.3: Hedonic Regression Used to Construct the One-Dimensional Quality Index

of buyers according to the online bidding information, this sample is suitable for analyzing two types of buyers' behavior.

p<0.1

1.2.4 Descriptive Findings

Here, I summarize the notable findings in the data. I first illustrate some cross-sectional facts in this market and then discuss the changes pre- and post-merger.

1.2.4.1 Cross-sectional Variation

Cross-platform Facts There is a decrease in the quality of trucks sold in auctions. According to Texas's data, however, the decline in quality is also true for the trucks sold in other channels (Table A.1). Therefore, I assume that the change in general quality is irrelevant to the merger, and I control it in the counterfactuals.





Notes: quality is round to the nearest integers.

Comparing the quality distributions of the trucks sold on IP and RBA in 2016 and 2018 (Figure 1.3), I find that the average quality of the trucks sold offline is higher than that of the trucks sold online. Consistent with Table 1.2, this difference is much more significant after the merger: high-quality trucks are much more likely to be sold on RBA, and low-quality trucks are much more likely to be sold on IP post-merger. The reason for this change is uncertain and may be related to the sellers' platform entry choices or the platforms' re-position policy. I investigate the welfare effect and policy implication of this change in the counterfactual part.

Figure A.3 presents the price distribution (the mean and variance) of trucks belonging to different quality bins. Here, each quality bin is constructed based on the one-dimensional quality measurement and rounded to the nearest integer. First, as shown in the standard deviation figure, considerable heterogeneity exists in trucks' prices with the same quality on the same platform. The price variance increases in quality, indicating that the transaction prices are not linear in the observed quality levels of the trucks. Additionally, the average transaction price of the trucks on RBA is higher than that on IP at all quality bins. One interpretation of this price difference is related to RBA's local inspection opportunity. Although IP posts the inspection reports of the trucks sold online, buyers may doubt the accuracy of the information about the trucks listed online. Therefore, buyers may discount the observed quality of the trucks sold online.

Cross-buyer Facts Using Texas data, I compare the trucks purchased by interstate trucking companies and local trucking companies. As shown in Figure A.4, interstate firms tend to purchase higher quality trucks. Meanwhile, these firms also pay higher prices for these trucks conditional on buying them. This finding indicates that interstate companies may prefer to purchase high-quality trucks more than local companies.

Cross-market Facts According to the definition of the markets, I calculate the number and average quality of the trucks subject to transactions in each market (Figure A.5). The transaction volume is much larger on RBA and more fluctuated across markets than the one on IP. The average quality of trucks varies across markets.

1.2.4.2 Cross-year Variation

Next, I show some interesting changes in market outcomes following the merger.

Transaction Price and Quality of Trucks The first important finding is a



Figure 1.4: Price Distributions of Trucks in Different Quality Groups

Notes: unit of price is \$1,000.

change in the relationship between the transaction price and the trucks' quality. I divide the trucks into different quality groups according to their percentiles. Figure 1.4 shows the histograms of the log(price). The higher quality trucks' prices tend to increase after the merger, although the distribution of the prices of the lower quality trucks remains approximately unchanged. The difference in the quality of the trucks purchased by different types of buyers is more significant postmerger: compared with local trucking companies, interstate trucking companies are more likely to purchase the high-quality trucks post-merger (Figure 1.5).

Transaction Price and Trucks Available on Each Platform The final remarkable change is related to the relationship between the transaction price and trucks available on each platform in the market. To analyze the change in this relationship, I use the number of trucks on each platform and the average quality of these trucks as two main measurements representing the available trucks. Since the quality and number of trucks fluctuated more on RBA, I regress an on-





line truck's price on its quality and these measurements. I conduct this regression pre-merger and post-merger separately (Table 1.4). One robust finding is that a truck's price on IP is more sensitive to the average quality of the trucks on RBA. When the average quality of the trucks sold on RBA is high, trucks' transaction price on IP tends to be low post-merger if everything else remains the same.

Summary In summary, we can observe that after the merger: (1) compared with the low-quality trucks, the probability of selling high-quality trucks at high prices increases after the merger; (2) compared with local trucking companies, interstate trucking companies are more likely to purchase high-quality trucks; (3) the price distribution on IP is more sensitive to the quality of trucks available on RBA in the same market.

I show that the change in search costs can potentially explain these three findings.
	2016	2016	2018	2018
VARIABLES	log(price)	log(price)	log(price)	log(price)
quality	1.086***	1.164***	0.790***	-2.689***
	(0.0390)	(0.298)	(0.0373)	(0.349)
$quality^2$		-0.0115		0.547***
		(0.0434)		(0.0545)
$\overline{quality_{-i}^{IP}}$	0.430***	0.433***	0.285*	0.0855
5	(0.119)	(0.119)	(0.162)	(0.158)
$quality^{RBA}$	0.0400	0.0359	-0.657***	-0.450**
	(0.105)	(0.106)	(0.221)	(0.215)
N^{IP}	-0.001*	-0.001*	-0.001**	-0.002***
	(0.001)	(0.001)	(0.001)	(0.001)
N^{RBA}	0.000	0.000	-0.000**	-0.000*
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	-2.284***	-2.409***	1.325*	6.661***
	(0.619)	(0.779)	(0.802)	(0.940)
Observations	1,354	1,354	1,470	1,470
R-squared	0.456	0.456	0.258	0.305

Table 1.4: Sensitivity of Price on IP to the Trucks Available on Each Platform

Notes: standard errors in parentheses * * * p < 0.01, * * p < 0.05, * p < 0.1.

1.3 Model

This section presents a model of the platform markets that endogenizes buyers' search, auction entry, and auction bidding. I assume that trucks with different quality levels are sold in single-unit auctions on two platforms. My model treats the supply-side as exogenous while developing an equilibrium model of buyers' behavior. Given their private draws of search costs, buyers simultaneously choose whether to search both platforms and how many auctions to search. Searching allows buyers to discover the characteristics and private values of individual trucks. Then, buyers simultaneously decide which of the searched auctions to enter. Buyers have a unit demand and can enter exactly one auction. The model is static because I do not allow buyers to consider the possibility of entering subsequent auctions if they fail to purchase a truck in the auction market.

1.3.1 Sellers

The realized set of trucks available on each platform is drawn from the observed sets of trucks in the data across regions and time periods. The information regarding the trucks available on each platform includes the realized quality levels (q), namely, the one-dimensional quality index, of the trucks on each platform and the number of trucks at each quality level on each platform.

1.3.2 Buyers

All buyers are ex-ante symmetric but distinguished by their i.i.d draws of fixed search cost ($fc \sim H^{fc}(\cdot)$), marginal search cost ($mc \sim H^{mc}(\cdot)$), and private value of each truck (v). A buyer's WTP of a truck depends on the quality of the truck and the her private value for that truck. I make the following assumption regarding the buyers' WTP.

Assumption 1. Distribution of WTP A buyer's WTP for a truck follows a log-normal distribution as follows: $V = exp(\theta q + v), v \sim N(\mu, \sigma)$.

Assumption 1 places the restriction on the WTP that q and v are not additively separable. The multiplicative structure $\frac{\partial V^2(q,v)}{\partial q \partial v} > 0$ indicates that the values are always positive and consistent with the fact that the variance of the realized prices increases with quality. A log-normal distribution is also a type of distribu-



tion commonly used in the auction literature (e.g., Laffont et al. (1995)). The specific form that I use also allows for a WTP discount ($\alpha < 0$) for online trucks, specifically, the WTP of a buyer with private value v for a truck with quality q online is $V = exp(\theta(q + \alpha) + v)$.

In addition to this baseline model, I also consider a model with two types of buyers who can have different quality preferences, where the draws are i.i.d. within types, and the mix of types fit the interstate/local data from TX. I allow different types of buyers to have different coefficients of θ^H and θ^L for quality.

1.3.3 Timing

Figure 1.6 shows the timeline of the game.

- Search Choice Stage During this stage, buyers draw private search costs independently from the cost distributions. Based on the private search costs and common knowledge (distribution search costs, distribution of WTP, and distribution of trucks available on each platform), buyers simultaneously determine their search choices. A search choice includes the following two parts:
 - Search frequency ($m \in \{1, ..., M\}$): number of trucks to search under either homing choice.
 - Homing choice (home \in {multi, single}):
 - * Single-homing: a buyer searches trucks randomly on one platform;
 - * Multihoming: a buyer searches trucks randomly on both platforms.

To reduce the computation burden, I assume when buyers make their search choices, they have no information regarding the realized available trucks in the market and only know the distribution of a possible set of trucks available on each platform.

 Market and Platform Entry Buyers enter the market based on the realized number of trucks in the market, and single-homing buyers enter a platform.
 I assume this process is exogenously determined. Specifically, I make the following assumption:

Assumption 2. Market and Platform Entry When the realized numbers of trucks sold on Platform A and Platform B in a market are N^A and N^B ,

- Market Entry: $\gamma \times (N^A + N^B)$ buyers enter the market, where γ is a scalar parameter and exogenously determined.
- Platform Entry: the probability of a single-homing buyer entering one platform is given according to a random entry rule which depends on the number of goods on each platform in that market $Prob(A) = \frac{N^A}{N^A + N^B}$, $Prob(B) = \frac{N^B}{N^A + N^B}$.

When doing estimation, γ is calculated from bidding data. Under the random entry rule, all trucks on both platforms have the same probability to be chosen. The only difference between single-homing buyers and multihoming buyers is the composition of the searched trucks they can choose when making their auction entry choices. In reality, an implication of this assumption is that single-homing buyers' platform choice is closely related to the scarcity of offline auctions in that market. In a market, there are many offline auctions, buyers tend to single-homing offline; otherwise, they are more likely to single-homing online. Additionally, I consider alternative rules where buyers can target a specific platform when they conduct single-homing in the estimation part.

- Auction Entry Stage After randomly searching in the market, buyers simultaneously make their auction entry choice (*e*). They determine which of the searched auctions to enter according to the information regarding the searched trucks.
- Auction Bidding Stage After entering auctions, buyers submit their bids (*b*) in the auctions. The trucks are sold via English auctions. They leave the market regardless of whether they win.

When describing the equilibrium, I will use superscripts $\{A, B\}$ to denote different platforms, subscripts $i \in \{1, ..., N_{buyer}\}$ to represent a buyer and $j \in \{1, ..., N\}$ to represent a truck.

1.3.4 Equilibrium

The equilibrium of this game is defined as follows.

Definition 1. (A Symmetric Bayesian Nash Equilibrium for Buyers)

A symmetric Bayesian Nash equilibrium in the market for buyers with common knowledge is a set of search strategies $\{m^*(\cdot), home^*(\cdot)\}\)$, auction entry strategies $e^*(\cdot)$, and bidding strategies $b^*(\cdot)$ such that any buyer

- bids optimally
- enters a searched auction according to an optimal rule
- decides how to search based on an optimal rule

given the equilibrium strategies of the other buyers in each stage.

The game can be solved by backward induction. I describe the equilibrium starting from the bidding stage.

1.3.5 Auction Bidding and Entry Stage

In the bidding stage, a buyer's private information is her WTP for the truck sold in the auction she enters: V(q, v). Given this private information, she makes

her bidding decision. Since the auction is an IPV English auction, truthfully bidding is the dominant strategy for all buyers, we have $b^*(V(q, v)) = V(q, v)$.¹¹

In the auction entry stage, after searching, buyers have private information regarding the searched trucks. I denote the private information of buyer *i* who searched *m* trucks under homing choice home as $(\mathbf{q}_i^{m,\text{home}}, \mathbf{v}_i^{m,\text{home}})$, where $\mathbf{q}_i^{m,\text{home}}$ and $\mathbf{v}_i^{m,\text{home}}$ are $m \times 1$ vectors, $\mathbf{q}_i^{m,\text{home}}$ includes the quality information regarding these *m* trucks and $\mathbf{v}_i^{m,\text{home}}$ includes the private values associated with these trucks. Given the private information, buyers make their entry choice. Buyer *i*'s entry strategy is a $m \times 1$ vector $e_i(\mathbf{q}_i^{m,\text{home}}, \mathbf{v}_i^{m,\text{home}})$, where all the elements are zeros but the chosen one.

The expected payoffs from entering an auction depend on the expected competition in that auction. In the view of other buyers, $Pr_i^e(q_j, v_{ij})$ is the probability that buyer *i* will enter auction *j* with a private value no less than v_{ij} . Similarly, in the view of buyer *i*, $Pr_l^e(q_j, v_{lj})$, $\forall l \neq i$ is the probability that buyer *l* will enter auction *j* with a private value no less than v_{lj} . Then, buyer *i*'s expected payoffs

¹¹When platforms charge a commission from buyers, buyers' bids equal their WTP discount by the commission rate in the English auction. This can completely transfer the burden of the commission from the buyers to sellers if there is no reserve price, we do not consider sellers' platform entry choices, and all buyers have single-unit demand.

from entering an auction with quality q_j and private value v_{ij} is as follows:

$$U_{i}(q_{j}, v_{ij}) = \begin{cases} \int_{\underline{v}}^{v_{ij}} [V^{A}(q_{j}, v_{ij}) - V^{A}(q_{j}, \tilde{v})] d\Pi_{l \neq i} [1 - Pr_{l}^{e}(q_{j}, \tilde{v})] + \dots \\ V^{A}(q_{j}, v_{ij}) \Pi_{l \neq i} [1 - Pr_{l}^{e}(q_{j}, \underline{v})] & \text{if } j \text{ is on } A \\ \int_{\underline{v}}^{v_{ij}} [V^{B}(q_{j}, v_{ij}) - V^{B}(q_{j}, \tilde{v})] d\Pi_{l \neq i} [1 - Pr_{l}^{e}(q_{j}, \tilde{v})] + \dots \\ V^{B}(q_{j}, v_{ij}) \Pi_{l \neq i} [1 - Pr_{l}^{e}(q_{j}, \underline{v})] & \text{if } j \text{ is on } B \end{cases}$$

$$(1.1)$$

In equation (1.1), $\Pi_{l\neq i}[1 - Pr_l^e(q_j, \tilde{v})]$ is the probability that no buyers but *i* enters the auction with quality q_j and random value no less than \tilde{v} .

According to the definition of BNE, an equilibrium entry strategy should maximize buyer *i*'s expected payoffs. A buyer will enter the auction with the highest expected payoff in the set of auctions she has searched. So in equilibrium, the buyer *i*'s entry probability defined above equals

$$Pr_{i}^{e*}(q_{j}, v_{ij}) = \sum_{m, \text{home}} \sum_{\mathbf{q}_{i}^{m, \text{home}}} \left[\int_{\underline{v}}^{\overline{v}} \dots \int_{\underline{v}}^{\overline{v}} \int_{v_{ij}}^{\overline{v}} \mathbf{I}\{q_{j} \in \mathbf{q}_{i}^{m, \text{home}}\} \times \dots \right]$$
$$\mathbf{I}\{U(q_{j}, v) = \max\{U(q_{j'}, v_{ij'})\}_{(q_{j'}, v_{ij'}) \in (\mathbf{q}_{i}^{m, \text{home}}, \mathbf{v}_{i}^{m, \text{home}})}\} \dots$$
$$dF(\mathbf{v}_{i}^{m, \text{home}})Prob(\mathbf{q}_{i}^{m, \text{home}})Prob(m_{i} = m, \text{home}_{i} = \text{home})$$
(1.2)

where $I\{\cdot\}$ are two indicator functions represent the event that truck j is searched and the auction with (q_j, v) has the highest expected payoffs in the set of trucks searched by buyer i.

As equation (1.2) shows, the probability is determined by the distribution of buyer *i*'s search choice $\{Prob(m_i = m, home_i = home)\}$, probability of searching different sets of information including truck j conditional on a search choice and the probability to choose auction j with the searched information.

We can transfer the BNE into probability space.¹² Formally, according to (1.1) (1.2), the symmetry of buyers and the difference between platforms, the problem can be written as

$$\left(\begin{array}{c} Pr^{e^*}(q,v,A) = \Lambda^{e,A}(Pr^{e^*}(q,v))\\ Pr^{e^*}(q,v,B) = \Lambda^{e,B}(Pr^{e^*}(q,v)) \end{array}\right)$$

where $\Lambda^{e,A}$ and $\Lambda^{e,B}$ are the best response probability functions and

$$Pr^{e^*}(q,v) = \left(\begin{array}{c} Pr^{e^*}(q,v,A)\\ Pr^{e^*}(q,v,B) \end{array}\right)$$

Since the best response probability functions are well defined and continuous in the compact convex set of players' probabilities, according to Brouwer fixed-point theorem, at least one equilibrium exists. ¹³

Given the equilibrium probabilities, I can obtain a set of equilibrium expected payoffs $U^*(q, v)$. A buyer's expected payoffs from a set of auctions are the payoffs from the auction with the highest $U^*(q, v)$ in the set. Based on Assumption 2, I can calculate the expected payoffs from different search choices given the realized

¹²There are many other papers, Seim (2006), Aguirregabiria and Mira (2007), etc, treat BNE as being in probability space. Note that here the probabilities come from the different sets of information buyers can get from the random search process, not logit errors.

¹³It is easy to show that the expected payoff functions U is continuous and monotonically increasing in v. In other words, $\forall q$, my expected payoffs to the auction with q will be continuous and increasing in v. Under the optimal rule, buyers choose to entering the auction with the highest expected payoffs. According to equation (1.1)(1.2), as the probability of other people (using their optimal rule) entering auctions increases, my best response will fall continuously.

trucks available in the market.

1.3.6 Search Choice Stage

In the first stage, buyers' private information is their search costs. They decide their homing choices and search frequencies according to their private information.

Since buyers make the choices before they discover the set of trucks available on each platform, the expected payoffs should be average across all possible realizations of available trucks on each platform. I denote buyer *i*'s expected payoffs from a search choice W_i . According to equation (1.2), the expected payoffs depend on the distribution of other buyers' search choice. This distribution is the conditional choice probability (Aguirregabiria and Mira (2007)) associated with a search strategy of a buyer given the distributions of search costs. Define $Pr_i^{m,\text{home}} = Prob(m_i = m, \text{home}_i = \text{home}), \forall i, Pr_{-i}^{m,\text{home}} = \{Prob(m_l = m, \text{home}_l = \text{home}), \forall l \neq i\}$. In equilibrium, buyer *i* choose m_i and home_i that can maximize her net expected payoffs given $Pr_{-i}^{m^*,\text{home}^*}$

$$\max_{m_i, \text{home}_i} [W_i(m_i, \text{home}_i, Pr_{-i}^{m^*, \text{home}^*}) - mc_i \times m_i - \mathbf{I}\{\text{home}_i = \text{multi}\} \times fc_i]$$

where $I{\text{home}_i = \text{multi}}$ is the indicate function that buyer *i* choose multihoming.

So there are two equilibrium conditions that (m^*, home^*) and Pr^{m^*,home^*} should satisfy. The first condition ensures that no one wants to deviate to a different number of searches given their private marginal search costs and the homing choices. Since the expected marginal gain from an additional search decreases with the number of searches,¹⁴ we can find the equilibrium cutoffs at which buyers feel indifferent towards searching different numbers of trucks under the same homing strategy.

$$W_{i}(m^{*}+1, \text{home}^{*}, Pr_{-i}^{m^{*}, \text{home}^{*}}) - W_{i}(m^{*}, \text{home}^{*}, Pr_{-i}^{m^{*}, \text{home}^{*}}) = \underline{mc}(m^{*}, \text{home}^{*})$$
$$W_{i}(m^{*}, \text{home}^{*}, Pr_{-i}^{m^{*}, \text{home}^{*}}) - W_{i}(m^{*}-1, \text{home}^{*}, Pr_{-i}^{m^{*}, \text{home}^{*}}) = \overline{mc}(m^{*}, \text{home}^{*})$$
(1.3)

Second, to ensure that no buyer has an incentive to deviate her homing choice with her search costs, I need to compare the expected payoffs from the proposed equilibrium strategy with the optimal search strategy under the other homing strategy. Therefore, the second set of equilibrium conditions is as follows:

$$W_{i}(m^{*}, \text{single}, Pr_{-i}^{m^{*}, \text{home}^{*}}) - mc_{i} \times m^{*} \ge ...$$

$$\max\{W_{i}(m, \text{multi}, Pr_{-i}^{m^{*}, \text{home}^{*}}) - mc_{i} \times m_{i}\} - fc_{i},$$

$$\max\{W(m, \text{single}, Pr_{-i}^{m^{*}, \text{home}^{*}}) - mc_{i} \times m_{i}\} \le ...$$

$$W_{i}(m^{*}, \text{multi}, Pr_{-i}^{m^{*}, \text{home}^{*}}) - mc_{i} \times m^{*} - fc_{i}.$$
(1.4)

The first inequality is for the case that home^{*} = single and the second inequality is for the case that home^{*} = multi. (1.4) indicates that given a marginal cost, a buyer will choose multihoming only if her fixed cost is lower than a threshold.

Based on (1.3) and (1.4), the conditional choice probabilities can be calcu-

¹⁴Since the expected payoff function is the expectation of the largest order statistics, it is concave in the number of searches (David (1997)).

lated given the distributions of the marginal search cost and fixed search cost. For example, the probability of (m, multi) is as follows:

$$Pr_{i}^{m^{*},\text{home}^{*}}(m, \text{multi}) = \int_{\underline{mc}(m,\text{multi})}^{\overline{mc}(m,\text{multi})} \int_{\underline{fc}}^{fc(mc_{i})} h^{mc}(mc_{i})h^{fc}(fc_{i})dfc_{i}dmc_{i} \qquad (1.5)$$
where $fc(mc_{i}) = W_{i}(m^{*}, \text{multi}, Pr_{-i}^{m^{*},\text{home}^{*}}) - \dots$
 $mc_{i} \times m^{*} - \max\{W_{i}(m, \text{single}, Pr_{-i}^{m^{*},\text{home}^{*}}) - mc_{i} \times m_{i}\}$
and fc is the lower bound of fixed cost.

Therefore, in the symmetric BNE, we have

$$Pr^{m^*, \mathsf{home}^*} = \Lambda^m (Pr^{m^*, \mathsf{home}^*}),$$

where Λ^m is the best response probability function. According to (1.3), (1.4) and (1.5), Λ^m is well defined and continuous in the compact convex set of the buyers' choice probabilities. Based on the Brouwer fixed-point theorem, at least one equilibrium exists.

Given the equilibrium of the game, next, I will discuss the economics of search choices and show how the change in search costs can affect the market outcome.

1.3.7 Economics of Search

1.3.7.1 Search Frequency

Buyers are more likely to find trucks with higher expected payoffs when they search for more trucks. Therefore, a buyer's expected payoffs increase in the numbers of searches. The marginal cost of searching one more truck is the effort required to investigate the quality of the truck and determine the private value of the truck, such as online search and consulting with sales representatives, etc. A buyer decides the number of trucks to search for by trading off between the gain and cost of the marginal search.

When a buyer has lower marginal search cost, she has an incentive to search for more trucks. However, other buyers will also have this motivation when they draw lower marginal search costs. Searching is strategic substitution among buyers because an increase of other buyers' search frequency results in more competition in auctions with high expected payoffs and reduces the gain from more searches. Figure 1.7 shows one example of expected payoff functions from different searches $W(m, \text{home}^*)$ under the same homing choice. The average number of searches chosen by other buyers (λ^m) increases from five to ten, resulting in a flatter expected payoff function. This indirect effect can partly discourage buyers from searching for more goods. However, in the new equilibrium with average lower marginal search costs, the average number of searches among buyers is still higher than the one in the old equilibrium. With more searches in the new equi-



Figure 1.7: Net Expected Payoffs $W(m, \text{home}^*)$ When $\lambda^m = 5$ and $\lambda^m = 10$

Notes: unit of expected payoffs is \$1,000.

librium, the distribution of transaction price will change. Here, I give an example about how the price will change under a simple structure of available trucks in the market.

Proposition 1. (Change in Search Frequency) If (1) there is one platform; (2) there are two types of trucks differentiated by their observed quality levels q^H and q^L , where $q^H > q^L$; (3) buyers in the market choose to search one or two trucks and the equilibrium probability of buyers to search two trucks under different distributions of search costs are Pr^{m^*} and $Pr^{m^{**}}$, where $Pr^{m^*} < Pr^{m^{**}}$, then the difference between the upper tail of the price distribution of high-quality trucks and that of low-quality trucks increases when buyers tend to search two trucks. Formally, using p to denote final prices, we have

$$\begin{aligned} \exists p^*, \forall \tilde{p} \in [p^*, \overline{p}] \\ Prob(p > \tilde{p} | q^H, Pr^{m**}) - Prob(p > \tilde{p} | q^L, Pr^{m**}) &\geq \dots \\ Prob(p > \tilde{p} | q^H, Pr^{m*}) - Prob(p > \tilde{p} | q^L, Pr^{m*}) \end{aligned}$$
(1.6)

Proof. See Appendix A.2.1.1.

In addition to the formal proof in the appendix, here, I briefly present the intuition about this proposition. First, given Assumption 1, I can prove that there is always a threshold of private value above which high-quality trucks are more attractive than low-quality trucks. Without changing the belief about other buyers, when buyers search for two trucks, the auctions with high-quality trucks are more likely to be chosen by buyers with private value above the threshold in the original equilibrium. Then, when all buyers are more likely to search for two trucks, the competition in those high-quality auctions is fiercer. Expecting this, some buyers with moderate private values associated with the high-quality trucks may switch to auctions with low-quality trucks. Namely, the threshold in the new equilibrium will be higher. However, buyers, drawing private values above the new threshold, are more likely to choose auctions with high-quality trucks. Finally, since the transaction price is the second-highest WTP in that auction, relative to low-quality trucks, high-quality trucks are more likely to be transacted with higher prices when

buyers search more intensively.

For markets with more complicated structures of available trucks, the change in buyers' auction entry behavior follows a similar pattern when they tend to search for more trucks. In Appendix A.2.2, I use simulations to illustrate this. Additionally, I give a discussion about a model with two types of buyers and show that buyers with high quality preference are more likely to purchase high-quality trucks than low-type buyers when the marginal search costs are lower.

The analysis above shows that the cross-year change (1)(2) observed in the data can be explained by higher search frequency. This mechanism is also useful to justify identification, which I will discuss later.

1.3.7.2 Homing Choices

When buyers engage in multihoming, they can access a set of trucks from two platforms and choose the platform having auctions with the highest expected payoffs in their choice set, which can smooth the variation in the number and quality across the platforms, and may increase buyers' expected payoffs. The fixed costs of multihoming include learning the two systems, registering two accounts to monitor trucks, etc. Buyers make their homing choices by trading-off between the gain from multihoming and these costs. When the fixed costs decrease, more buyers switch to multihoming.

The increased share of multihoming buyers can change the market outcome. Similar to search frequency, I also show the pattern in a simple setting.

38

Proposition 2. (Change in Homing Choices) If (1) there are two platforms and all the trucks on the same platform have the same quality level; (2) there are two observed quality levels q^H and q^L , where $q^H > q^L$; (3) buyers in the market choose to search one or two trucks and the share of single-homing buyers under different fixed search costs are ω^* and ω^{**} , where $\omega^* > \omega^{**}$, then the difference between the upper tail of the price distribution of a truck on platform A when the average quality of the trucks on platform B is low and that when the average quality of the trucks on platform B is high increases if there are more buyers conduct multihoming. Formally,

$$\begin{split} &\exists p^*, \forall \tilde{p} \in [p^*, \overline{p}] \\ &Prob(p^A > \tilde{p} | \overline{q}^B = q^L, \omega^{**}) - Prob(p^A > \tilde{p} | \overline{q}^B = q^H, \omega^{**}) \geq \dots \\ &Prob(p^A > \tilde{p} | \overline{q}^B = q^L, \omega^*) - Prob(p^A > \tilde{p} | \overline{q}^B = q^H, \omega^*) \end{split}$$

Proof. See Appendix A.2.1.2.

The proposition is based on a similar logic in Proposition 1. The expected payoffs of high-quality trucks are always higher than those of low-quality trucks if the private values drawn by buyers are above some threshold. Therefore, buyers with high private values are more likely to choose Platform B if the trucks on Platform B have high quality. This can explain the change in price distribution shown in Proposition 2.

This proposition shows that the cross-year change (3) observed in the data can be explained by buyers' more multihoming. It will also be called in the iden-

tification.

Finally, the change in the fixed costs may alter the search frequency, and the change in the marginal costs may alter the multihoming choices. For example, under multihoming, there are more variations in the composition of trucks in a choice set. Thus, the expected payoff function is less concave under multihoming when the number of searches is large. Namely, the expected payoffs from searching many trucks are higher under multihoming than the one under single-homing. Therefore, lower fixed search costs can encourage buyers to search for more trucks.

Based on the analysis and propositions above, the cross-year descriptive findings in the data can be explained by the reduction in search costs.

1.4 Identification

In this section, I explain how to identify the two critical components in the model, i.e., the distribution of buyers' WTP and the distribution of search costs. In general, the distribution of search costs is not nonparametrically identified. However, I can still identify the marginal search costs associated with the thresholds between searching different numbers of trucks and the fixed search costs associated with the thresholds between multihoming and single-homing. For simplicity, I refrain from considering post-merger changes and commission and focus on identifying the buyers' WTP and the search cost thresholds.

Definition 2. A model is identified iff $\forall (H^c, F^V, \hat{H}^c, \hat{F}^V)$,

 $(P, B|H^c, F^V, X) = (\hat{P}, \hat{B}|\hat{H}^c, \hat{F}^V, X)$ implies $H^c = \hat{H}^c$ and $F^V = \hat{F}^V$. where

- Exogenous Variables (X): set of realized available trucks on each platform and γ;
- Model Primitives (H^c, F^V) : the distribution of buyers' search costs and WTP;
- Observed Endogenous Outcomes (P, B): the realized transaction price and trucks, and number of bidders in each auction.

As shown in the model, buyers make their search strategies before entering any market. In the auction entry stage, the distribution of equilibrium search choices is a sufficient statistic for a buyer to make her optimal entry strategy. Therefore, I can separate the identification problem into the following two problems: (1) the observed distribution of the transaction prices and number of bidders can identify the distribution of WTP and distribution of buyers' equilibrium search choice ($Pr^{m^*,home^*}$) given the exogenous variables, and (2) the identified distribution of WTP and distributions of equilibrium search choice can identify the bounds of distributions of search costs.

Similar to the discussion about economics of search, I use the case in which all buyers search one or two trucks to show how the model can be identified. The proportion of searching for two trucks is Pr^{m^*} , and the proportion of single homing buyers is ω^* in equilibrium.

1.4.1 Distributions of WTP and Equilibrium Search

1.4.1.1 Baseline: One Platform and One Type of Buyers

I begin with markets with one platform, one type of buyers and trucks in two quality levels (q^H and q^L). Using the observed endogenous outcomes in these markets, I can prove identification.

Given the number of bidders, I can focus on the auctions with a small number of bidders. For example, from the data, I can calculate the price distribution of trucks with quality q^H conditional upon having two buyers in the auctions. This situation includes the following two cases: both buyers in an auction bid lower than p and only one buyer bids lower than p in an auction. Formally, the CDF of the transaction price for these auctions are as follows:

$$\begin{split} F_{2,\text{price}}(p|q^{H},q^{L},N^{H},N^{L}) &= \dots \\ &\{(1-Pr^{m^{*}})F^{v}(log(p)-\theta q^{H}) + Pr^{m^{*}}[\frac{N^{H}-1}{2(N-1)}(F^{v}(log(p)-\theta q^{H}))^{2} + \dots \\ &\underbrace{\frac{N^{L}}{N-1}\int_{\underline{v}}^{log(p)-\theta q^{H}}f^{v}(v)F^{v}(v'|U^{*}(q^{L},v') = U^{*}(q^{H},v))dv]\}^{2} + \dots \\ &\underbrace{N^{L}}_{\text{both buyers bid lower than }p} \\ &2\{(1-Pr^{m^{*}})[1-F^{v}(log(p)-\theta q^{H})] + Pr^{m^{*}}[\frac{N^{H}-1}{2(N-1)}(1-(F^{v}(log(p)-\theta q^{H}))^{2}) + \dots \\ &\underbrace{\frac{N^{L}}{N-1}\int_{log(p)-\theta q^{H}}}^{\overline{v}}f^{v}(v)F^{v}(v'|U^{*}(q^{L},v') = U^{*}(q^{H},v))dv]\} \times \dots \\ &\{(1-Pr^{m^{*}})F^{v}(log(p)-\theta q^{H}) + Pr^{m^{*}}[\frac{N^{H}-1}{2(N-1)}(F^{v}(log(p)-\theta q^{H}))^{2} + \dots \\ &\underbrace{\frac{N^{L}}{N-1}\int_{\underline{v}}}^{log(p)-\theta q^{H}}f^{v}(v)F^{v}(v'|U^{*}(q^{L},v') = U^{*}(q^{H},v))dv]\} \times \dots \\ &\underbrace{\frac{N^{L}}{N-1}\int_{\underline{v}}}^{log(p)-\theta q^{H}}f^{v}(v)F^{v}(v'|U^{*}(q^{L},v') = U^{*}(q^{H},v))dv]\} \\ & (1.7) \end{split}$$

one buyer bids higher than p, one bids lower than p

Here, $N = N^H + N^L$ and $F_{2,\text{price}}(p|q^H, q^L, N^H, N^L)$ is calculated using the data. According to the model, v' is the private value which makes buyers feel indifferent between the auction with q^H and q^L . Under Assumption 1, $F^V(\cdot|q)$ is determined by μ , σ and θ .

Proposition 3. The price distribution function $F_{2,price}(\cdot|q^H, q^L, N^H, N^L), \forall q^H, q^L,$ N^H, N^L can identify the model primitives $F^V(\cdot|q)$ and Pr^{m^*} .

Proof. See Appendix A.3.1.

Here, I describe the basic idea underlying the identification. As shown in the expression of price distribution in (1.7), the price distribution depends on a mixture of the following three distributions: the distributions of WTP for trucks with quality q^H and q^L and the distribution of search frequency, which is simplified as the coefficient Pr^{m^*} . To separately identify these distributions, I need to use the variation in price distributions from markets with different structures of available trucks.

Assume there are two sets of model primitives can generate the same price distributions in markets where all trucks have quality q^H and markets where all trucks have quality q^L . They cannot generate the same price distribution in markets having trucks with both q^H and q^L . As shown in the model part (Proposition 1), when buyers search for more trucks, the trucks with q^H are more likely to be purchased at a high price. If buyers search for two trucks, the difference in the price distribution in these market differs from the difference in the price distribution in the market with one quality trucks. However, the differences are the same when buyers always search for one truck.

This idea can be applied to markets with more quality levels if we can observe the price distribution in markets with various structures of available trucks.

1.4.1.2 Extensions

Two Platforms Assume that buyers' WTP for trucks with the same quality level differs if the goods are listed on different platforms. I denote the distributions of WTP as $F^{V,A}$ and $F^{V,B}$. Buyers also choose to conduct single-homing or multihoming, and the equilibrium probability of conducting single-homing is ω^* . To reduce the number of primitives to identify, here I make the following assumption:

Assumption 3. $Pr^{m^*,single} = Pr^{m^*,multi} = Pr^{m^*}$

According to the baseline model results, I can use the price distributions in markets only with platform A to determine the WTP on platform A and Pr^{m^*} . It is similar to markets only with platform B. An additional initial condition is needed to identify ω^* . One way to get this condition is to use the price distributions in markets with q^H on a platform and q^L on another platform. As discussed in the model part (Proposition 2), the price distribution in these markets and price distribution in markets with one platform are different when the share of multihoming buyers changes.

Two Platforms and Two Types of Buyers In this model, the share of different types of buyers is the same across different markets. In addition to the price distribution of trucks, the types of winners in the auctions are observed. Similarly, I make the following assumption to simplify the identification problem. This assumption and Assumption 3 are kept in estimation.

Assumption 4.
$$Pr_H^{m^*,single} = Pr_H^{m^*,multi} = Pr_H^{m^*}, Pr_L^{m^*,single} = Pr_L^{m^*,multi} = Pr_L^{m^*}$$

As shown in the extended model above, using the price distributions of trucks with different quality levels, it is possible to identify $\{F^{V,A}, F^{V,B}, Pr^{m^*}, \omega^*\}$. This set of statistics can be derived from the underlying model primitives

$$\{F_{H}^{V,A}, F_{L}^{V,A}, F_{H}^{V,B}, F_{L}^{V,B}, Pr_{H}^{m^{*}}, Pr_{L}^{m^{*}}, \omega_{H}^{*}, \omega_{L}^{*}\}.$$

Consider the markets with one platform in each market where the quality of the trucks is the same. Since the share of different types of buyers is the same across different markets, the distribution of equilibrium search choices of both types of buyers is the same in these markets. Given a distribution of search choices $\{Pr_{H}^{m^*}, Pr_{L}^{m^*}, \omega_{H}^{*}, \omega_{L}^{*}\}$, the difference between price distribution of trucks with the same quality purchased by the same type of buyers but on different platforms can identify the difference between the same type of buyers' WTP on different platforms: $\{F_{H}^{V,A} - F_{H}^{V,B}, F_{L}^{V,A} - F_{L}^{V,B}\}$. By combing with the identified $\{F^{V,A}, F^{V,B}\}$, I can express everything as functions of buyers' search choices. Finally, as discussed in the model, using the difference in the quality distribution of the trucks purchased by different types of buyers, $\{\omega_{H}^{*}, Pr_{H}^{m^{*}}, \omega_{L}^{*}, Pr_{L}^{m^{*}}\}$ can be identified separately given $\{\omega^{*}, Pr^{m^{*}}\}$.

Figure A.9 in Appendix A.3.2 gives a summary the measurements and assumptions used to identify different models.

1.4.2 Distribution of Search Costs

Given the distribution of WTP F^V and equilibrium search choice (Pr^{m^*}, ω^*) , I can partially identify the distribution of the search costs according to the equilibrium conditions in the model.

Specifically, given the cutoffs constructed by (1.3) and (1.4), I can map the distribution of equilibrium search choices to the distributions of the search cost. The probability of a buyer searching for two trucks on two platforms equals the probability that the buyer's marginal search cost falls into a range and her fixed search cost is lower than a threshold.

$$(1 - \omega^{*})Pr^{m^{*}} = \int_{\underline{mc}(F^{V}, Pr^{m^{*}}, \omega^{*})}^{\underline{mc}(F^{V}, Pr^{m^{*}}, \omega^{*})} \int_{\underline{fc}}^{fc(mc_{i}, F^{V}, Pr^{m^{*}}, \omega^{*})} h^{mc}(mc_{i})h^{fc}(fc_{i})dfc_{i}dmc_{i}$$

Similarly, I can map the probability of single-homing to a range of the marginal cost and fixed cost. Notably, there is no overlap of the fixed search cost which can support the different homing behavior performed simultaneously given the same marginal cost. For the lowest and highest assumed M, I cannot identify the upper and lower bounds; thus, I make assumptions to identify these mass points.¹⁵

¹⁵I make the following assumptions: (1) the upper bound of the marginal search cost at M = 1 equals the expected payoffs from searching for one truck; (2) the lower bound of the marginal search cost at M = 10 equals zero; and (3) the upper bound of the fixed cost of single-homing buyers equals the highest lower bound of the fixed costs of single-homing buyers.

1.5 Estimation

As shown above, the model can be nonparametrically identified in two steps. I still use a two-step algorithm to estimate the model. To simplify the estimation, I introduce several parametric assumptions and fix certain parameters that are otherwise difficult to estimate.

1.5.1 Parametric Assumptions and Normalizations

- Search Choices: the number of searches performed by buyers follows a Poisson distribution, m_i^{*} ~ Poisson(λ). The share of buyers who choose to conduct single-homing in equilibrium is ω^{*}. I denote the set of all parameters to be estimated as Θ = {θ, ω, α, λ, ω}. When buyers are allowed to have different quality preferences in the model, I obtain the following Θ = {θ^H, θ^L, σ, α, λ^H, λ^L, ω^H, ω^L}. I assume the distribution of WTP is the same before and after the merger, but the distribution of equilibrium search choices can change following the merger.
- Proportion of high/low-type buyers: based on the Texas data, I assume that $share^{H} = 0.6$, $share^{L} = 0.4$;¹⁶
- Ratio of Buyers to Sellers: I fix γ = 4 because the median number of bidders in auctions equals 4;

¹⁶For the trucks transacted in all channels, the share of high type buyers is 0.6 and the share of low type buyers is 0.4. For the trucks transacted in auctions only, the share of high type buyers is 0.35 and the share of low type buyers is 0.65. In Appendix A.5.1, I show the estimation results using both pair of shares.

• I use 6%, 3%, and 5%, which are the weighted average of the observed commission rates, to approximate the commission rates charged by IP, RBA premerger, and RBA post-merger, respectively.

1.5.2 Algorithm

As shown in the identification section, I can first use the observed distribution of prices and bids to estimate the distribution of WTP and the distribution of equilibrium search choices. Then, the distribution of search costs can be estimated based on the estimated distributions. The estimation framework is summarized in the following two steps:

- *Step 1* I use a nested fixed-point algorithm to estimate the distribution of buyers' WTP and the distribution of buyers' search choice in equilibrium based on the observed bidding and transaction data.
 - *Inner Loop* I numerically solve the equilibrium bidding and auction entry strategies of buyers in the inner loop and generate the distribution of prices and the distribution of bids based on simulations.
 - *Outer loop* I use the simulated distributions and observed distributions to construct several moments for estimation. The distribution of WTP and the distribution of equilibrium search choices are estimated in this outer loop.
- *Step 2* I use the estimated distribution of WTP and distribution of equilibrium search choices to nonparametrically estimate the distribution of search

costs based on the equilibrium conditions for the equilibrium search strategies.

Next, I discuss the details in each step.

1.5.2.1 Solving Equilibrium Bidding and Auction Entry Strategies

When there are two types of buyers, there are four equilibrium payoff functions from auctions in each market: $U_{H}^{IP*}, U_{L}^{IP*}, U_{H}^{RBA*}$ and U_{L}^{RBA*} . Since there is no analytical solution to this problem, I propose a numerical way to solve them. To implement the computational method, I assume that the expected payoff function from entering an auction is continuous in the quality of the goods.

Specifically, I use a two-dimension Lagrange interpolation (Judd (1998)) to approximate the equilibrium payoffs. Given the initial guess of the expected payoff functions $U^{(0)}$, I can figure out a set of simulated buyers' entry choices. Then I calculate a new expected payoff $U^{(1)}$ by averaging all the ex-post payoffs of buyers over simulations. A buyer's ex-post payoff from an auction is determined by the equilibrium bidding strategies and set of competitors in that auction. In Appendix A.4.1, I show the details of this computation procedure.

Given a guess of the distribution of WTP and equilibrium search choices, I can simulate a set of bids based on the equilibrium auction entry strategy and bidding strategy. To assign a price to the auctions with only one bidder, I assume that there is always an additional bidder in the auctions who mimics buyers' behavior. They draw WTP from buyers' distribution of WTP, bid truthfully, and discount their bids by the commission. The only difference is they do not make auction entry choices. We can treat these additional bids as bids from sellers or platforms to ensure trucks can be sold with a positive price when there is one buyer in the auctions.

Here I use 100 simulations and denote the sets of bids from simulations $\{\mathbf{b}^s(\Theta)\}_s$.

1.5.2.2 Estimating the Distributions of WTP and Equilibrium Search Choice

As shown in the identification part, the price distributions in markets with different sets of available trucks on each platform are used to identify the model. Using $\{\mathbf{b}^{s}(\Theta)\}_{s}$ and observed data, I can calculate the distribution of bids and the distribution of prices. Based on these distributions, I construct three sets of moments $\{g_{1}(\Theta), g_{2}(\Theta), g_{3}(\Theta)\}$. The estimator can minimize the Wald-type objective function:

$$\hat{\Theta} = argmin_{\Theta} \begin{pmatrix} g_1(\Theta) \\ g_2(\Theta) \\ g_3(\Theta) \end{pmatrix}' W \begin{pmatrix} g_1(\Theta) \\ g_2(\Theta) \\ g_3(\Theta) \end{pmatrix}$$

where W is the weighting matrix.

The underlying justification for using the first two sets of moments is that the differences between the simulated prices and observed prices of the trucks with the

same quality on the same platform faced with the same set of available trucks are independent across auctions and having zero means. The underlying justification for using the third set of moments is that the differences between the simulated quality and observed quality for the trucks purchased by the same type of buyers on the same platform faced with the same set of available trucks are independent across auctions and having zero means.

First Set of Moments ($g_1(\Theta)$): Mean and Standard Deviation of Prices

First, I calculate the mean and standard deviation of the log(price) of IP/RBA before/after the merger. The simulated mean and standard deviation can be achieved by averaging all simulated price across all auctions and sets of available trucks on each platform pre- or post-merger. Using the observed data, I can calculate their correspondence in reality. I construct a set of moments to measure the difference between simulated data and observed data. For example, I consider the mean of the online transaction price pre-merger in the one-type model. The moment is calculated as follows

$$\frac{1}{K} \sum_{k=1}^{K} \frac{1}{N^k} \sum_{j=1}^{N^k} [\frac{1}{N^s} \sum_{s} (\log(p_j^{IP, Pre, k, s})) - (\log(\tilde{p}_j^{IP, Pre, k}))]$$

Here, market 1...K represents K realized sets of available trucks on each platform before the merger, N^s is the number of simulations, and \tilde{p} are the observed prices in the data.

As shown above, the price distribution of trucks in different quality groups can differ. I divide the trucks into two categories according to whether the trucks' quality is above the median. In each category, I calculate the mean and standard deviation of the log(price). For the online auction after the merger, I also construct moments of the price distribution for auctions with one bidder in each auction.

Second Set of Moments ($g_2(\Theta)$): Relationship Among Price, Quality and Trucks Available on Each Platform

The first set of moments includes the aggregate information about price distributions. It is necessary to capture more information about price distributions at different quality levels and different realizations of available trucks. Since there are numerous quality levels and realized sets of available trucks, I use regressions to achieve this goal. Specifically, I regress the transaction price of trucks on the trucks' quality and some measurements of available trucks which are used in the data section. The measurements include the number of trucks available on each platform, average quality of other trucks available on the same platform and the average quality of trucks available on the other platform in the same market. I conduct regressions using both simulated data and observed data. For example, I consider the online transaction price as follows:

$$\begin{split} log(p_{j}^{IP,k,s}) = & \beta_{0}^{IP,s} + \beta_{1}^{IP,s} q_{j}^{IP,k,s} + \beta_{2}^{IP,s} (q_{j}^{IP,k,s})^{2} + \beta_{3}^{IP,s} \overline{q}_{-j}^{IP,k,s} + \beta_{4}^{IP,s} \overline{q}^{RBA,k,s} + \dots \\ & \beta_{5}^{IP,s} N^{IP,k,s} + \beta_{6}^{IP,s} N^{RBA,k,s} + \epsilon_{j}^{IP,k,s} \\ log(\tilde{p}_{j}^{IP,k}) = & \tilde{\beta}_{0}^{IP} + \tilde{\beta}_{1}^{IP} q_{j}^{IP,k} + \tilde{\beta}_{2}^{IP} (q_{j}^{IP,k})^{2} + \tilde{\beta}_{3}^{IP} \overline{q}_{-j}^{IP,k} + \tilde{\beta}_{4}^{IP} \overline{q}^{RBA,k} + \dots \\ & \tilde{\beta}_{5}^{IP} N^{IP,k} + \tilde{\beta}_{6}^{IP} N^{RBA,k} + \tilde{\epsilon}_{j}^{IP,k} \end{split}$$

Above are two regressions based on simulated data and observed data. I conduct

the regressions for the pre-merger period and post-merger period, respectively. Since WTP is assumed to follow a log-normal distribution, I use log(price) instead of price itself. \bar{q}_{-j}^{IP} is the average quality of trucks on IP except for truck j and \bar{q}^{RBA} is the average quality of trucks on RBA in the same market. As I discussed before, the decrease of search cost will encourage buyers with high random values to enter the auctions with high-quality trucks, so I add a quadratic term of trucks' quality and calculate the price sensitivity to quality at the median and third quartile of quality. The coefficient of \bar{q}^{RBA} can capture the cross-platform sensitivity emphasized before.

The moments are used to measure the difference in the relationships between the simulated data and observed data. I can attempt to match all the coefficients directly or match the difference in coefficients pre- and post-merger.

Third Set of Moments ($g_3(\Theta)$): Moments of Different Types of Buyers

The third set of moments is used to estimate the model with two types of buyers. According to the winners' types in auctions won by buyers in Texas, I can divide the data into two subsets. Different types of buyers have different quality preferences, which can result in different quality distributions and price distributions for trucks purchased by different types of buyers. The average quality \bar{q}_T and the average price of \bar{p}_T differ among the trucks in different subsets. The third set of moments is used to measure the difference between simulated data and observed data in these two measurements.

1.5.2.3 Search Cost

After obtaining the estimation results of the distribution of WTP and the distribution of equilibrium search choices, I use the equilibrium conditions mentioned in the model to nonparametrically estimate the bounds two distributions of search costs: $h^{mc}(\cdot)$ and $h^{fc}(\cdot)$.

I first calculate the average expected payoffs from all the possible deviations to other homing strategies and search frequencies, given other buyers follow the estimated equilibrium search. I assume buyers know the difference in distributions of available trucks on each platform pre-merger and post-merger. When calculating the expected payoffs from deviations pre-merger, I average all realized pre-merger markets in the data; when calculating the expected payoffs from deviations post-merger, I average all realized post-merger markets in the data. This assumption about buyers' belief is the same for the model with one type of buyers and the model with two types of buyers.

According to these results and the equilibrium conditions, I can construct a mapping from the distribution of equilibrium search choices to the joint distribution of search costs. The mapping has been shown in the identification part. Since a range of search costs can rationalize a search choice, I can only partially estimate search cost distribution. I denote the joint distribution of the lower bound of search costs as $H_{LB}^{fc,mc}$ and the joint distribution of the upper bound of search costs as $H_{UB}^{fc,mc}$. The corresponding marginal distributions of marginal cost and fixed cost can be calculated and are denoted $\{H_{LB}^{fc}, H_{LB}^{mc}\}$ and $\{H_{UB}^{fc}, H_{UB}^{mc}\}$. Similarly, I

can partially estimate the distributions of search costs in a model with two types of buyers.

1.6 Estimation Results

1.6.1 Distribution of WTP and Distribution of Search Costs

Table 1.5 shows the estimation results of the one-type model. In the Appendix A.5.1, I show the estimation results for the two-type model under different assumptions about single-homing buyers' platform choice. Based on the estimation results, I can draw the implied average WTP, the 25th and 75th percentiles of WTP at different quality levels on different platforms (Figure 1.8). When q = 3, the average WTP is \$8,135 on IP and \$11,172 on RBA; when q = 5, the average WTP is \$35,407 on IP and \$48,626 on RBA . According to the way to construct the quality index, a truck's quality can decrease by two if increasing a truck's age from almost 0 to 15 years old.

Figure 1.8 also shows the implied distribution of search choices pre-merger and post-merger, where the number of searches is truncated at one and ten trucks. By comparing these two distributions, we see that buyers significantly increase the number of trucks to search. The median number of searches increased from 5 trucks to 6 trucks. More than 50% of buyers engage in multihoming, resulting in almost all buyers multihoming after the merger. This is consistent with the fact that the platform presents integrated search results as a default.

Given the estimated distribution of WTP, the change in buyers' equilibrium

Figure 1.8: Implied Distribution of WTP and distribution of equilibrium search choices



Notes: unit of WTP is \$1,000.

Quality Preference	
θ	0.7354
	(0.0027)
Distribution of v	
μ	0.0001
	(0.0205)
σ	0.6015
	(0.0064)
Discount of Quality Online	
α	-0.4314
	(0.0060)
Mean Number of Searches	
λ^{Pre}	5.7907
	(0.0721)
λ^{Post}	6.2964
	(0.0324)
Proportion of Buyers Single-Homing	
$\omega^{*,\mathrm{Post}}$	0.0505
	(0.0134)
$\omega^{*,\operatorname{Pre}}$	0.6392
	(0.0342)

Table 1.5: Estimation Results of the Model with One Type of Buyers

Notes: standard errors are shown in parentheses. They are obtained by numerically calculating the derivatives in $(\frac{\partial g(\Theta)}{\partial \Theta} W \frac{\partial g(\Theta)}{\Theta})^{-1}$.

search choices results from the change in search costs. Using the approach mentioned in the estimation, I obtain the estimates of the distributions of search costs. Figure 1.9 shows the lower bound and upper bound of the cumulative distribution functions of search costs. The lower and upper bound of the median marginal search cost decrease from \$270 to \$224 and from \$328 to \$255, respectively. The lower and upper bound of the median fixed search cost decrease from \$4 to \$0 and from \$10 to \$5, respectively. The estimated costs of additional searches are quite high, but not unreasonable given that buyers need to conduct much searching before determining which trucks to buy. According to an investigation of the used truck market, on average, buyers spend approximately one day to finalize whether to purchase a truck. Considering that the average salary per hour in the U.S. is approximate \$30, the estimation results of the marginal costs are reasonable. One potential explanation for the lower marginal search costs after the merger is that the merged platforms design a better online environment for buyers to search and sales representatives are more familiar with the trucks in the market. This can help buyers to figure out their WTP of trucks with less effort. The estimated multihoming costs are small compared to the marginal costs. The magnitude of fixed costs depends on the assumption about the platform choice of single-homing buyers. I assume buyers know the realized number of trucks in the one-type model before they are randomly allocated to a platform. Therefore, the fixed costs of multihoming are mainly the costs of buyers becoming familiar with the system on these two platforms, the effort they take to register two accounts to monitor the trucks on different platforms, etc. Logically, these fixed costs are not very high.



Figure 1.9: Cumulative Distribution Functions of Search Cost in the One-type Model

Additionally, the integrating policies provided by the merged platforms can lower these costs.

I use an example to illustrate how the change in search costs can change buyers' search choices. This example can also explain why the significant change in marginal search costs only leads to a small change in search frequency. Assume a buyer has mc = \$317, fc = \$19 pre-merger and mc = \$239, fc = \$3 postmerger. Her expected payoff functions are increasing and concave in the number of searches (Figure 1.10). When she searches more than six trucks, her expected payoffs under multihoming are higher than those under single-homing. Given the pre-merger search costs, she will choose to conduct single-homing and search for six trucks to earn her expected payoff \$2, 333.

Notes: unit is \$1,000.


Figure 1.10: How the Change in Search Costs Affects the Buyers' Search Choices

Notes: 1. unit is \$1,000; 2. 1-10 is single-homing, and 11-20 is multihoming.

Next, the buyers' search costs decrease to the post-merger level. If all other buyers still choose the same search choices, this buyer will search for nine trucks under multihoming (Point 2). However, because other buyers are more likely to have lower search costs post-merger, she expects that they will search more aggressively. Because of the strategic substitution of searching among buyers, her expected payoff function shifts downward. In the new equilibrium, she will choose to search for seven trucks on both platforms (Point 3). Compared with her original search, the lower search costs encourage her to search for more trucks on both platforms, which can be applied to other buyers. Thus, the change in the distribution of search costs can explain the estimated change in the distribution of search choices.

1.6.2 Model Fit and Sensitivity Analysis

Table A.3 shows how the model and estimation results fit the targeted moments of the observed prices and bid distributions. In the model with two types of buyers, I show the moments when single-homing high-type buyers choose the offline platform and low-type single-homing buyers follow the random choice rule. The model can fit most of the first and second-order moments of price. Additionally, the estimated results can capture the changes pre-merger and post-merger and fit the quality of trucks purchased by different types of buyers.

Here I show some examples. The observed and simulated average price of offline trucks are \$21,977 and \$22,646 pre-merger and \$21,542 and \$22,198 post-merger, respectively. Both simulated and estimated prices decrease by approximately \$450. According to the observed data, if the average quality offline increases by one and the average quality online decreases by one, we can predict that the price of a truck sold online will decrease by 14% more post-merger than pre-merger; according to the simulated data, if the average quality offline increases by one while the online quality remains the same, I can predict that the price of a truck sold online will decrease by 14.6% more post-merger than premerger. Pre-merger, the average quality of online trucks purchased by low-type buyers is 3.09 pre-merger and 3.12 post-merger in the Texas data. According to the estimates of the two-type model, the average quality is 3.09 pre-merger and 3.15 post-merger.

In addition to the targeted moments, I use the bidding data online post-

Figure 1.11: Observed and Simulated Distributions of Quality and Price Online Post-merger



Notes: unit of price is \$1,000.

merger to show how the estimates fit. In Figure 1.11, I draw the CDF of the quality and price on IP according to the number of bidders in the auction (*n*). For $10 \ge n > 2$, I also divide the quality into two groups. I observe that except for the case $10 \ge n > 2$, $q \le 3.5$, all other price and quality distributions fit well.

Finally, given the estimates, I can test how different moments can be used to identify different model primitives based on the approach proposed by Andrews et al. (2017). I show the results in Appendix A.5.3. Consistent with the model, the parameters of the equilibrium search choices are sensitive to the moments describing how the transaction price is sensitive to the truck's quality and how it is sensitive to the average quality of trucks on the other platform.

1.7 Counterfactuals

1.7.1 Roadmap

As shown above, following the merger, I can observe three types of policy changes. In this section, I analyze each change individually. Here I briefly show the framework of each counterfactual.

- The first change, which is also the focus of this paper, is the change in the buyers' search costs resulting from the change in the integration policies. In this counterfactual, all buyers are ex-ante symmetric. All trucks are sold via auctions without a reserve price. I control for the other two changes. Regarding the supply side, I assume that all observed changes are irrelevant to the merger; thus, when calculating the expected payoffs, I compute payoffs pre- and post-merger search costs across all observed markets; regarding the commission rate, I calculate the welfare change from lower search costs with pre-merger observed commission rates. Except for the estimated change in search costs, I also consider the case where there is a change only in marginal costs and the case there is a change only in fixed costs.
- The second change is the change in commission charged by the platforms and paid by buyers. To make the change affects buyers' search choices, I also consider cases where trucks are sold via auctions with reserve prices. I calculate the welfare under alternative changes in commission rates with one type of buyer and estimated search costs post-merger. In this counterfactual,

I assume that all observed changes on the supply side are irrelevant to the merger.

• The final change is the change in sets of trucks available on each platform as follows: sellers with different quality levels list their trucks across platforms differently. Since this counterfactual focuses on qualities, I allow buyers to have two different preferences for quality. Buyers with high-quality preference will choose the offline platform when they conduct single-homing, and buyers with low-quality preference will follow the random entry rule when they conduct single-homing. All trucks are sold via auctions without a reserve price. I calculate the welfare change from the change on supply side with estimated pre-merger and post-merger search costs separately. The commission rate used in this counterfactual is the same as that observed premerger. Considering the change in the general quality distribution may be irrelevant to the merger, I construct two new sets of available trucks in which both sets have the same population quality distribution and 64 markets. In the first set, the trucks are separated into different platforms according to the rule in the pre-merger data; in the second set, trucks are separated into different platforms according to the rule in the post-merger data.

Next, I discuss how welfare changes in each counterfactual.

1.7.2 Welfare Analysis with Changes in Search Costs

Efficiency Gain Given the estimated search costs of buyers pre- and postmerger, I first re-calculate buyers' equilibrium search choices under different cost structures. The method used to solve the new equilibrium is shown in the Appendix A.6.1. As mentioned above, I control for the change on the supply side by assuming that the possible sets of trucks available on each platform include all realized markets. Table 1.6 shows the results. Regardless of the marginal search costs or the fixed costs are lowered, buyers will search for more trucks and conduct more multihoming. This finding is consistent with the analysis about the economics of search costs in the model part. While lowering marginal costs has a more significant effect on search frequency, the change in fixed costs has a more significant effect on homing choice. In addition, compared with the equilibrium search choices under post-merger costs, lowering both costs can trigger further more searches and more multihoming buyers than simply lowering one.

Table 1.6: New Equilibrium Under Alternative Changes in Search Costs

mc	P	re	Po	ost
fc	Pre	Post	Pre	Post
Mean Number of Searches λ Proportion of Single-homing	5.50	5.66	6.60	6.86
ω^*	0.72	0.50	0.62	0.24

Notes: "Pre" means pre-merger, "Post" means post-merger, "mc" means marginal search costs, and "fc" means the fixed search costs for multihoming.

Next, I calculate the welfare under alternative changes in the search costs. Table 1.7 compares welfare under different cost structures. The search costs of buyers are between \$23.9 million and \$32.2 million and total surplus is between \$231.6 million and \$239.9 million when the search costs follow pre-merger distributions (see the second column in the table). Changes in both the marginal and fixed search costs can result in higher trading surplus and total surplus. Welfare increases more significantly than when only one type of search costs are lowered. Comparing the cost structure pre-merger and post-merger shows that the total trading surplus increases by approximately 6% on average, and the total surplus increases from [\$217.3 million,\$225.6 million] to [\$245.5 million,\$260.1 million]. Among these gains, sellers' welfare gain is derived from higher trading surplus, and buyers' welfare gain comes from lower search costs. Buyers' surplus from trading is lower post-merger.

mc	P	re	Po	ost
fc	Pre	Post	Pre	Post
Buyers	[23.9,32.2]	[21.9,30.5]	[31.7,45.1]	[29.5,44.0]
Trading	133.21	132.7	130.03	129.28
Search Cost	[101.0,109.3]	[102.3,110.9]	[85.0,98.3]	[85.3,99.8]
Sellers	193.4	196.7	212.0	216.0
IP	22.3	22.3	24.3	24.3
RBA	171.1	174.3	187.7	191.7
Platform	14.3	14.4	14.9	15.0
IP	3.1	3.1	3.2	3.2
RBA	11.2	11.3	11.7	11.9
Total Trading	340.9	343.8	356.9	360.3
Total Surplus	[231.6,239.9]	[232.9,241.5]	[258.6,271.9]	[260.5,275.1]

Table 1.7: Welfare Under Alternative Changes in Search Costs

Notes: 1. unit is \$1,000,000; 2. "Search Cost" includes the total search costs generated by all buyers in the market regardless of whether the buyers win the auctions. 3. Sellers' surplus is adjusted by their possible WTP, namely, I randomly draw the WTP of sellers from the distribution of buyers' WTP. For the auctions online, sellers' WTP also has a discount α ; 4."Pre" means pre-merger, "Post" means post-merger.

To explain the observed welfare change, I calculate the average trading sur-

plus of buyers and sellers from choosing trucks at different quality levels on different platforms (Table 1.8). It can be observed that, in general, buyers obtain a lower average surplus from high-quality trucks post-merger. This is especially true for RBA who has a larger market share and relatively more high-quality trucks in the data. With more searches under multihoming, buyers access more information about trucks across platforms before they make auction entry choices on both platforms. As the economics of search choice shows, more buyers with high idiosyncratic values choose the auctions with high-quality trucks and choose the platform that includes those auctions. Therefore, buyers are less likely to win those auctions or pay more to win the auctions. This situation also results in sellers with high-quality trucks obtaining higher surplus. Combined with the increased revenue of platforms, the average surplus from trading is higher post-merger, and the increase is more significant in the auctions with high-quality trucks. This finding is consistent with the economics of search cost analysis mentioned in the model part: when buyers search more trucks under multihoming, it can achieve more assortative matching results across platforms.

Efficiency Loss How close do these changes get us to first-best efficiency conditional on search? Here, I define the first-best efficiency as the case where the trading can generate the highest surplus given buyers' equilibrium search choices.

One way to achieve this is by using a centralized auction mechanism. In this selling mechanism, buyers report their WTP for all the trucks searched. The allocation rule and payment rule in this auction follow the Vickrey–Clarke–Groves (VCG) mechanism (Vickrey (1961)Clarke (1971) and Groves (1973)).The auc-

		$q \in [1,3]$		5	$\epsilon \in (3, 3.5$		6	$i\in(3.5,5)$	
Search Costs	Pre	Post	⊲	Pre	Post	\bigtriangledown	Pre	Post	\bigtriangledown
Total									
IP	8.63	9.02	0.36	13.44	14.02	0.58	22.29	23.34	1.05
RBA	11.87	12.62	0.75	18.79	20.12	1.33	29.32	31.37	2.05
Buyers									
IP	3.69	3.71	0.02	5.43	5.38	-0.05	8.35	8.12	-0.23
RBA	5.04	4.96	-0.06	7.58	7.39	-0.19	11.01	10.57	-0.44
Sellers									
IP	4.31	4.67	0.36	7.04	7.63	0.59	12.31	13.52	1.21
RBA	6.38	6.79	0.41	10.50	11.29	0.79	17.21	18.57	1.36
Platforms									
IP	0.63	0.64	0.01	0.97	1.01	0.04	1.63	1.70	0.07
RBA	0.45	0.87	0.35	0.71	1.44	0.58	1.10	2.23	0.90
Notes: 1. unit is winners from au	\$1,000; 3 ctions on	2. Here, b IP; 3. The	uyers tra	ding surp ce Δ with	lus on IP darker c	means th olor show	le average 's larger n	e trading s umber.	urplus of

Table 1.8: Average Trading Surplus and Split at Different Quality Levels (1)

tioneer (platform) calculates the trading surplus from each possible allocation and picks the one generating the highest total trading surplus.

I assume that buyers' search still follows the estimated distribution postmerger. By calculation, the decentralized auction mechanism's trading surplus accounts for about 98% (\$360.3 million in \$368.7 million) of the trading surplus in the centralized mechanism. The difference mainly comes from the trading surplus of high-quality trucks. This is because, in the decentralized market, buyers have no complete information about other buyers. Some buyers may switch to lowquality trucks when they consider the fiercer competition they might encounter. Therefore, the coordination failure makes some trucks not be allocated to the buyers who have the highest WTP. Nevertheless, the efficiency loss is not big.

1.7.3 Welfare Analysis with Changes in Commission

The merger's main concern is that the merger can increase the market power such that users of the platforms may be harmed. While a larger share of the surplus is transferred from sellers to platforms because of the increased commission rate, sellers also benefit from buyers' search choice change. At least, this situation is true for sellers with high-quality trucks. However, if the change in price policy can generate a counter-search effect, it may partially offset the efficiency gain from lowering search costs.

When the platforms charge the observed higher commission rate from buyers post-merger, the change will not alter buyers' search choices because in the

68

second-price auctions with no reserve price, buyers can completely transfer the burden of commission to sellers by shading their bids if entry choices of sellers are exogenous. Resolving the equilibrium with or without the change in commission shows that buyers' equilibrium search choices keep the same (Table 1.9). The change only affects the welfare split between sellers and platforms (Table 1.10). The platform obtains a larger share of surplus post-merger (from \$15.0 million to \$22.5 million); of this share, approximately \$7.5 million is transferred from sellers to the merged platforms because of the increase in commission on RBA. Given the supply side fixed, the total surplus of buyers and sellers post-merger outweighs the one pre-merger with the observed change in commission rate and estimated change in search costs.

Considering that in reality, some auction platforms, such as eBay, use reserve prices, I construct a counterfactual in which sellers set reserve prices in their auctions. The reserve prices equal the sellers' potential WTP. Then, since buyers cannot completely transfer the burden of commission to sellers by lowering their bids, the change in commission can affect buyers' entry and search choices.

Instead of using the approximation of post-merger commission structure, i.e., 6% for IP and 5% for RBA, I consider alternative higher commission rates. The additional reserve price with high commission works as an additional competitor for all buyers, and the reserve prices are higher for auctions with high-quality trucks. Thus, these commission fees lower buyers' expected payoffs from extensive searching. Some buyers deviate from the original search choices to search for fewer trucks and conduct single-homing. The average equilibrium search frequency will become lower, and the share of single-homing buyers will increase. For example, when the commission rate increases to 50%, the average number of search decreases from 6.86 to 6.11, and the share of single-homing buyers increases from 24% to 33% (see Table 1.9).

When platforms charge buyers a commission rate as high as 70% (see the last column in Table 1.9 and Table 1.10), buyers' average search frequency will decrease to 5.75. The share of single-homing buyers is approximately 39%. The surplus from trading will decrease from that post-merger (\$360.3 million) to a much lower value (\$342.0 million) close to the one pre-merger.

Table 1.9: New Equilibrium Under Reserve Prices and AlternativeChanges in Commission Rates

Search Costs	Pre		Post		
Reserve Price	No	N	0	Ye	es
Commission Rate	6%, 3%	6%, 3%	6%, 5%	50%	70%
Mean Number of Searches		_		_	
λ	5.50	6.	86	6.11	5.75
Share of Single-homing			~ /		
ω^*	0.72	0.	24	0.33	0.39

Notes: "RP" means the case with the reserved price.

In summary, the increased market power may allow the merged platforms to charge a higher commission rate, which can discourage buyers from searching for more trucks on two platforms when they cannot completely transfer the burden to sellers. This will result in a lower total trading surplus. However, at least in this auction setting, the platform needs to charge a higher commission rate to eliminate the welfare gain from lower search costs with the supply fixed. Therefore, when analyzing this merger, if the platform can significantly lower the market's

Search Costs	Pre		Pos	t
Reserve Price	No	N	0	Yes
Commission	6%, 3%	6%, 3%	6%, 5%	70%
Buyers Trading	133.2	129	9.3	103.1
Search Cost	[101.0,109.3]	[85.3,	99.8]	[83.1, 92.9]
Sellers	193.4	216.0	208.5	45.1
Platform	14.3	15.0	22.5	193.7
Total Trading	340.9	360	0.3	342.0
Total Surplus	[231.6,239.9]	[260.5,	275.1]	[249.1,258.9]

Table 1.10: Welfare Under Reserve Price and Alternative Changes in Commission Rates

Notes: 1. "RP" means the case with reserved price and I solve for commission fees which eliminate the welfare gain from lowering search costs; 2. unit is \$1,000,000; 3. "Search Cost" includes the total search costs generated by all buyers in the market regardless of whether the buyers win the auctions. 4. Sellers' surplus is adjusted by their possible WTP, namely, I randomly draw the WTP of sellers from the distribution of buyers' WTP. For the auctions online, sellers' WTP also has a discount α ; 5. "Pre" means "Pre-merger" and "Post" means "Post-merger"

search costs, the efficiency gain from integrating policies can be substantial.

1.7.4 Welfare Analysis with Changes in Supply Side

In the data section, I show that sellers with different quality trucks may enter platforms according to a different rule after the merger: sellers with highquality trucks are more likely to list offline, and sellers with low-quality trucks tend to list online. If buyers are ex-ante asymmetric, having different quality and platform preferences, are there any benefits from separating sellers by quality into two different platforms in the observed way? This counterfactual attempts to investigate this issue. As mentioned above, I use two new data sets that follow different rules to separate the trucks into auctions. Then I compute the predicted outcomes in each market based on these two data sets with pre-merger and postmerger search costs. Table 1.11 shows the new equilibrium search choices of buyers under alternative changes. First, regardless of whether pre-merger or post-merger search costs are applied, high-type buyers search for more trucks and tend to singlehome when trucks are separated into platforms in the way post-merger. It is because high-type buyers can easily target high-quality trucks by single-homing. The change is more significant when the search costs of multihoming are high since, in that case, buyers cannot easily access the trucks on other platforms and highly count on the composition of trucks on their single-homed platform. Second, faced with fiercer competition from high-type buyers in the offline auctions with highquality trucks, low-type buyers tend to search for fewer trucks if the search costs are the same. Finally, similar to the one-type model, lower search costs can trigger both types of buyers to search for more trucks and conduct multihoming.

Search Costs	P	re	Po	ost
Supply Side	Pre	Post	Pre	Post
Mean Number of Searches				
λ^H	6.28	7.82	7.54	7.93
λ^L	6.98	6.36	7.33	7.23
Share of Single-homing				
ω^{H*}	0.39	0.49	0.23	0.28
ω^{L*}	0.28	0.33	0.05	0.06

Table 1.11: New Equilibrium Under Alternative Changes in Supply Side

Notes: "Pre" means pre-merger, "Post" means post-merger.

Table 1.12 shows the welfare of different groups. Both the change in the supply side and lower search costs can increase the total trading surplus. Specifically, the change in the supply side alone can increase the trading surplus by 4.2%. The supply-side change significantly increases the trading surplus offline and RBA's revenue. This change is more remarkable when the search costs are

high. High-type buyers benefit from this change, but low-type buyers get a lower trading surplus because high-type buyers search more aggressively. The lower search costs post-merger can reduce the loss of low-type buyers. Also, with lower search costs, the surplus of online sellers and IP increases significantly. Similarly,

Search Costs Pre Post Supply Side Pre Pre Post Post High Type 109.1 112.9 110.9 112.2 Low Type 28.5 25.0 26.0 25.4 Search Cost [110.8,127.6] [92.7,109.5] [92.5,110.7] [108.6,121.3] **Total Buvers** [16.3,29.0] [10.3,27.1] [16.5,35.4] [26.9,45.0] Sellers 232.6 248.3 245.8 251.2 Sellers IP 25.9 20.5 29.2 24.6 Sellers RBA 206.7 227.8 216.6 226.6 Platform 15.8 16.0 16.3 16.2 IP 3.7 3.1 3.8 3.3 RBA 12.113.0 12.5 12.9 402.2 399.0 405.0 **Total Trading** 386.1 **Total Surplus** [264.7, 277.4][274.6,291.4] [289.5,306.3] [294.4,312.5]

Table 1.12: Welfare Under Alternative Changes in Supply Side

Notes: 1. unit is \$1,000,000; 2. "Search Cost" includes the total search costs generated by all buyers in the market, regardless of whether the buyers win the auctions. 3. Sellers' surplus is adjusted by their possible WTP, namely, I randomly draw the WTP of sellers from the distribution of buyers' WTP. For the auctions online, sellers' WTP also has a discount α ; 4. "Pre" means pre-merger, "Post" means post-merger.

to determine the reason for the changes in welfare, I calculate the average trading surplus of the different groups at different quality levels under different cases (Table 1.13 and Figure 1.12). The numbers of trucks in the different groups change when separating trucks according to different rules. When it comes to the average trading surplus, except for the high-quality group online, the post-merger rule on the supply side lower the surplus online and increase the surplus offline because online trucks are less likely to be searched by buyers. However, since the number of high-quality offline trucks increases, the total trading surplus significantly increases in Table 1.12.

If the platforms can lower the search costs, some high-type buyers will shift to multihoming and enter online auctions if they are more likely to win those auctions, and some low-type buyers will search more aggressively. The trading surplus from online auctions can notably increase while the trading surplus from offline auctions slightly decreases. The decrease in search costs is more beneficial for the markets with a small number of offline auctions. To be specific, Figure 1.13 illustrates that, in the markets where $\frac{N^{IP}}{N^{RBA}}$ is large, the case with the post-merger costs can generate a significantly higher average trading surplus than the case with the pre-merger costs. Therefore, combining the change in search costs with the change in the number of trucks can considerably increase the total trading surplus.

		$q \in [$	[1,3]			$q \in ($	[3, 3.5]			$q \in (;$	3.5, 5]	
Search Costs	P	re	Pc	st	P	re	Pc	st	P1	e Je	PC	st
Supply Side	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
Number of trucks												
IP	888	1,426	888	1,426	1,384	1,133	1,384	1,133	1,070	654	1,070	654
RBA	3,303	2,765	3,303	2,765	4,440	4,691	4,440	4,691	5,661	6,073	5,661	6,073
Average Surplus												
Total												
IP	8.38	8.10	9.12	8.92	13.80	13.18	15.04	14.49	22.58	22.93	24.75	25.33
RBA	13.25	13.69	13.49	13.67	21.99	22.92	22.45	22.84	36.39	37.55	37.24	37.39
Buyers												
IP	3.75	3.73	3.83	3.75	5.91	5.78	5.94	5.74	9.04	9.34	9.00	9.19
RBA	4.88	4.90	4.90	4.92	7.71	7.71	7.65	7.71	11.68	11.56	11.54	11.52
Sellers												
IP	4.01	3.76	4.62	4.52	6.90	6.46	8.04	7.73	11.94	11.97	14.01	14.37
RBA	7.43	7.82	7.63	7.78	12.74	13.62	13.24	13.55	22.20	23.42	23.14	23.31
Platforms												
IP	0.62	09.0	0.66	0.65	1.00	0.95	1.07	1.03	1.60	1.63	1.74	1.77
RBA	0.94	0.97	0.95	0.97	1.53	1.59	1.56	1.58	2.50	2.57	2.56	2.57
Notes: 1."Pre" means Namely, I randomly discount α .	pre-merg lraw the V	er, "Post" NTP of se	means po illers fron	ost-merge n the dist	er. 2. unit ribution	is \$1,000 of buyers	0; 3. Selle v WTP. Fo	ers' surplu	us is adju tions onli	sted by th ine, selle	neir possi rs' WTP a	ble WTP; llso has a

Table 1.13: Average Trading Surplus and Its Split at Different Quality Levels (2)

Figure 1.12: Number of trucks and Average Trading Surplus in Different Quality Groups



Notes: unit of Average Trading Surplus is \$1,000.

Figure 1.13: Average Trading Surplus in Different Markets



Notes: 1. unit of Average Trading Surplus is \$1,000; 2. the share of $\frac{N^{IP}}{N^{RBA}}$ is round to 0.1; 3. for the markets without offline auctions, $\frac{N^{IP}}{N^{RBA}} = \infty$.

1.8 Discussion

As mentioned many times in the paper, the current paper focuses on the buyers' search behavior and discusses the welfare effect of the merger through altering buyers' search costs. Ideally, we should endogenize both buyers' and sellers' entry decisions to the platform. Although this can provide a complete picture of all the participants in the market, it requires more data and significantly raises the computation burden. In this section, I will discuss the situation where I relax the exogenous supply assumption. While I still treat the buyers' platform entry exogenous, buyers' search choice now will be affected by sellers' platform entry strategy. Therefore, I can partially capture the indirect network effect in this two-sided platform market: a buyer's search choice is indirectly affected by other buyers' search choices through their effect on sellers' platform entry. Following, I will discuss two cases according to the data availability of transacted trucks in the market outside of these two auction platforms.

1.8.1 No Additional Data About Trucks in the Outside Market

Since there is no information about sellers in the outside market, I will still treat the entry of sellers to these two platforms exogenous but consider sellers' choice between these platforms as endogenous.

Consider a new game as follows. After buyers make their search choice, potential sellers draw their private information about the quality of their trucks and entry costs to two platforms (ϵ^A and ϵ^B). Then, they simultaneously make their platform entry choices conditional on their private value and belief about the distribution of buyers' search costs. After that, N sellers are randomly drawn into the market. Sellers enter the market according to the entry choice they have made. The following stages are as the one described in the model section. Figure 1.14 summarizes the new timeline of this game.

Figure 1.14: Timeline of the Game With Endogenous Platform Entry of Sellers (1)



I assume that sellers draw their quality and entry costs randomly from common distributions $F^q(\cdot)$, $F^{A,\epsilon}(\cdot)$ and $F^{B,\epsilon}(\cdot)$. The realized number of trucks in a market is also randomly drawn from a population distribution. Then the realized supply of trucks on each platform is determined by the realized number of trucks in that market, the quality of those trucks, and sellers' platform choices. As we have shown before, given the number of determined trucks on each platform, the

equilibrium in the buyers' auction entry and bidding stages can be solved by backward inductions. Therefore, the expected price of trucks with different quality levels on different platforms in a realized market can be calculated. Further, the expected payoffs of sellers with different quality trucks from different platforms can be calculated by average the expected conditional prices over all possible realization of markets.

In the empirical part, to reduce the computation burden, I can use the realized markets in the data to approximate the distribution of market realization. Similar to the auction entry stage of buyers, the uncertainty when sellers make their platform choices comes from the uncertainty in quality and entry costs of other sellers in the same market. Therefore, I can use a similar algorithm to solve for the equilibrium strategy of sellers. To be specific, in equilibrium, sellers have beliefs about the distributions of other sellers' platform choice $\{Pr_{-j}^A(q), Pr_{-j}^B(q)\}$. The expected payoffs of a seller with (q_j, ϵ_j) is a function of these choice probabilities. On the other hand, these choice probabilities also depend on other sellers' expected payoffs from entering different platforms. Namely, I can show the existence of equilibrium by the fixed-point theorem. In the equilibrium, I can approximate seller *j*'s expected payoffs from different platforms V_S^A and V_S^B as polynomials of the truck's quality and the seller's entry cost. By iterating the expected payoff functions, the equilibrium can be solved.

Instead of estimating the new model with a large amount of time, I will construct some simulations to check how the changes in search costs might affect sellers' equilibrium behavior and the social welfare given reasonable assumptions about the distributions of entry costs. The changes in the distributions of search costs are constructed based on the one estimated in the paper. This exercise is more interesting in a model with two types of buyers where they have different quality and platform preferences. I will use it to check whether, through affecting sellers' platform choices, the change in search costs can explain the observed change in supply on different platforms post-merger.

1.8.2 With Data About Transacted Trucks in the Outside Market

This is an ideal case where I have the number of trucks with different quality levels on each platform, including the outside market. Given more complete data, I can endogenize sellers' entry to the merged platforms. In this case, a buyer's search choice is also indirectly affected by other buyers' search choices through its effect on the aggregate number and quality of sellers on these two platforms. Similar to the game mentioned above, there is an additional stage describing sellers' endogenous strategy. However, now sellers make their choices among three options: platform A, platform B, and the outside market. Still, I assume the number and quality of trucks that appear in a market are given. Figure 1.15 shows the new timeline.

Now sellers' three entry costs are drawn from common distributions $F^{A,\epsilon}(\cdot)$, $F^{B,\epsilon}(\cdot)$ and $F^{O,\epsilon}(\cdot)$. The realized supply on different platforms, including in the outside market, is still determined by the realized number of trucks in that market, their quality, and sellers' platform choices. Assume the trucks in the outside market



Figure 1.15: Timeline of the Game With Endogenous Platform Entry of Sellers (2)

can be sold by giving take-it-or-leave-it offers where a realized price of a truck with quality q_j is $exp(\theta_0 q + v_0)$ and $v_0 \sim N(\mu_0, \sigma_0^2)$. Assume there is always enough potential buyers in the outside market to accept a seller's offer no matter how many other sellers choose the outside option. Then, $V_S^O(q_j, \epsilon_j^O) = E[exp(\theta_0 q_j + v_0)] - \epsilon_j^O$. When sellers make platform choices, they need to compare V_S^A , V_S^B , and V_S^O . In equilibrium, a seller's equilibrium expected payoffs from entering platform A and B are functions of the private (q_j, ϵ_J^A) and (q_j, ϵ_j^B) respectively.

I can solve the equilibrium choice probabilities $Pr^{A}(q)$, $Pr^{B}(q)$ and $Pr^{O}(q)$. Then the elasticities of entry to the quality of trucks can be calculated. For example, $elas^{A} = \frac{\frac{\partial Pr^{A}(q)}{Pr^{A}(q)}}{\frac{\partial q}{q}}$. Without estimating the entire model with endogenous sellers, I will first use pre-merger markets to calculate $elas^{A}$, $elas^{B}$, and $elas^{O}$ in the data. These elasticities can be used to calibrate the model primitives in the outside option, given the distribution of search costs pre-merger. Next, I will simulate the market outcome when buyers draw lower search costs to check how sellers' entry choices are affected.

There are two ways to define the outside market. One is considering all the other channels and other auction platforms as the outside market. Then I can use the data from some states¹⁷, which includes transaction data in all channels, to calculate a reasonable approximation of elasticities and calibrate the model primitives in the outside market. Then I can simulate potential realizations of markets based on the nationwide transaction data in each month and census data of truck inventories in each state. In the nationwide transaction data, I know the number and average quality of trucks transacted each month. In the census data, I have the number of truck inventories in each state. Assuming the turnover rates of trucks are stable across states, I can simulate the number of transacted trucks in each state. Another way is only treating other auction platforms as the outside market. From some websites, such as Truckpaper.com, that collect trucks transacted through different platforms, I have collected data about nationwide trucks transacted pre-merger and post-merger on the other auction platforms. Although this data set might be incomplete, I can still use them to approximate some realized markets, including transactions in the outside market.

¹⁷Transaction data through all channels in Texas and Washington State is collected.

1.9 Conclusion

Using a recent merger case in the U.S. used heavy-truck auction market, this paper investigates how partially integrating two competing platforms can affect market outcomes and social welfare. To clearly analyze the causal effect, I develop a detailed model of buyers' behavior. Based on the model, this paper provides several predictions regarding the market outcome when the search costs are low. To quantify the welfare effect, I structurally estimate the distribution of WTP and distributions of search costs before and after the merger. The estimation results show a significant decrease in the search costs after the merger and reveal that buyers search more extensively across the platforms.

In the counterfactual part, I compare the welfare in the market with and without the change in search costs and observed commission. Buyers' more aggressive search allows them to access more information about the trucks before making auction entry choices. The trading surplus increases with the estimated distributions. While sellers with high-quality trucks always benefit from the change, other participants' welfare is more complicated. For example, buyers' welfare depends on the composition of competitors and the magnitude of the reduced search costs. The cost decomposition shows, lowering the search costs of multihoming can generate efficiency gain. When buyers search more extensively, the decentralized market's allocation results are close to those of a market with a centralized mechanism in which there is no coordination failure among buyers. I also combine the change in search costs with the following two alternative changes: high commission fees when auctions have reserve prices and separation of trucks into platforms according to their quality levels. The analysis shows that (1) the increased commission fees may discourage buyers from extensive searching when buyers cannot completely transfer the burden of commission to sellers; (2) the changes on the supply side may generate additional efficiency gain when considering single-homing buyers' platform preference.

Methodologically, the paper considers the heterogeneity among the participants in many dimensions. Both buyers and trucks are differentiated horizontally and vertically. Building upon the literature, I develop a new model that combines search and endogenous auction entry stages. The model can capture a wide range of transaction markets in which search costs are considerable.

Partial integration or the facilitation of multihoming is a common feature of platform mergers, and my analysis suggests that the welfare benefits may be substantial. To the extent that it is not easy for firms to facilitate multihoming prior to mergers, in the language of the Horizontal Merger Guidelines, these benefits are cognizable efficiencies that could be set against market power created by a merger. In the context of this merger, the benefits to buyers and sellers exceeded the harm caused by the increase in commission fees that followed the merger.

This paper can be extended in several ways. Besides extending the model by endogenizing sellers' platform entry, as discussed above, I will also consider some alternative changes along with the merger. For example, upon the acquisition, RBA forms a strategic alliance with Caterpillar. How this change will affect social welfare is another interesting topic to analyze.

Chapter 2: Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Horizontal Mergers

2.1 Introduction

Theoretical and empirical analyses of differentiated product markets usually assume that firms have complete information (CI) and set prices to maximize their current profits. If an alternative is considered, it is typically tacit collusion with repeated CI stage games. These assumptions provide tractability, but there is surprisingly little evidence that they accurately predict how prices change after events such as mergers. The CI assumption is also inconsistent with how firms closely guard information about the margins of individual product lines and how sensitively this information is treated during merger investigations.

This paper considers what happens when we relax the static and CI assumptions. Specifically, we will assume that each firm has a payoff-relevant state variable, such as its marginal cost, which is imperfectly serially-correlated and unobserved by rivals. In this environment, each firm may want to choose its price strategically to affect its rivals' inferences. We will consider fully separating equilibria where, in equilibrium, a firm's chosen price perfectly reveals its current cost, and beliefs have a simple form. In these equilibria, all firms that do not have the lowest possible marginal cost set prices above static best response levels to credibly signal this information to their rivals. This can, in turn, cause static best response prices to increase, and signaling prices to rise further, a positive feedback that can cause equilibrium prices to be significantly above static CI Nash levels, although, as we discuss, separating equilibria may not exist if prices rise too much. While a small theoretical literature has shown that oligopoly signaling can affect equilibrium prices in two- or three-period models, we provide the first analysis of how large these effects may be, and the first empirical application.

We apply our model to horizontal merger analysis. Signaling is a strategic investment to raise rivals' future prices, and like many strategic investments, the equilibrium incentive to invest can rise when the number of competitors is reduced. We use an example to illustrate how a standard static CI merger simulation can significantly underpredict post-merger price increases if the firms are playing a dynamic signaling game. We then apply the model to data from the U.S. beer market around the time of the 2008 Miller-Coors (MC) joint venture (JV). Miller and Weinberg (2017) (MW) show that, after the JV, domestic brewers' prices increased in a way that is inconsistent with static CI Nash pricing. We calibrate our dynamic signaling model using only data on pre-JV price dynamics and show that it predicts the observed change in the level of prices accurately and that it also predicts directional changes in measures of observed price dynamics. We also extend MW's conduct parameter framework (Bresnahan (1982), Lau (1982), Nevo (1998), Berry and Haile (2014)) to show that the CI tacit collusion explanations for the post-JV price increase advanced by MW and Miller et al. (2020) (MSW) do not fully describe the pricing of domestic brewers, suggesting the need to explore new explanations, such as ours.

Before discussing the related literature, we should be clear about several limitations of our analysis. First, we have to assume that each firm has exactly one privately-known state and can send exactly one signal per period. This imposes restrictions on how firms are modeled after mergers. Second, we only consider fully separating equilibria, even though these may not exist for some parameters and we can only prove existence and uniqueness in special cases. Third, while we can reject some specific tacit collusion models, folk theorems imply that collusive models may exist that could fit the data perfectly.

The rest of this introduction reviews the related literature. Section 2.2 lays out the model and the equilibrium concept. Section 2.3 presents some examples and illustrates the implications for merger analysis. Section 2.4 provides our empirical application. Section 2.5 concludes. The online Appendices detail the computational algorithms; additional examples; a proof of existence and uniqueness for the case of linear demand; and, further details of the data and empirical analysis.

Related Literature. Shapiro (1986) and Vives (2011) examine how equilibrium prices and welfare change when marginal costs are private information in one-shot oligopoly models. Most of our focus will be on models where marginal costs lie in quite narrow intervals and the static effects that these papers identify are

very small. A large theoretical literature has considered one-shot signaling models where only one player has private information. The classic Industrial Organization example is the Milgrom and Roberts (1982) limit pricing model, where an incumbent monopolist may lower its first period price to deter entry in a twoperiod game. Sweeting et al. (2020) develop finite and infinite-horizon versions of this model where an incumbent monopolist's type changes over time, as we will assume in this paper.¹ They estimate the model and show that it can explain why incumbent airlines dropped prices by as much as 15% when Southwest threatened entry on monopoly routes. The oligopoly setting considered here is potentially applicable to many more markets.

The literature on games where multiple players signal simultaneously is much more limited.² Mailath (1988) identifies conditions under which a separating equilibrium will exist in an abstract two-period game with continuous types, and shows that the conditions on payoffs required for the uniqueness of each player's separating best response function are similar to those shown by Mailath (1987) for models where only one player is signaling (Mailath and von Thadden (2013) generalize these conditions). Mailath (1989) applies these results to a two-period pricing game where differentiated firms have static linear demands and marginal costs that are private information but fixed. Firms raise their prices

¹Kaya (2009) and Toxvaerd (2017) analyze one-sided, dynamic signaling games where the informed firm's type is fixed, and, in equilibrium, the informed firm signals until its reputation is established.

²Bonatti et al. (2017) analyze linear signaling strategies in a continuous-time Cournot game where each firm's marginal cost is private information and fixed, but firms cannot perfectly observe the quantities that their rivals choose. We will assume that prices are perfectly observable.

in the first period in order to try to raise their rivals' prices in the second period.³ Mester (1992) extends this approach to a three-period quantity-setting model where marginal costs change over time, and she shows that signaling, which leads to increased output in this case, happens in the first two periods.

We rely on Mailath's results to characterize best response signaling pricing functions, and we will focus on the magnitude, empirical relevance and implications of the equilibrium effects in multi-period settings with more standard forms of differentiated product demand. Fershtman and Pakes (2012) and Asker et al. (2020) develop an alternative approach to discrete state and discrete action dynamic games with asymmetric information. They reduce the computational burden using the concept of Experience-Based Equilibrium (EBE) where firms have beliefs about their payoffs from different actions rather than rivals' types.⁴ Our equilibrium concept is more standard, and the computational burden is reduced by focusing on fully separating equilibria in continuous action games.

We discuss the relationship between our paper and discussions of coordinated effects in horizontal merger analysis (Ordover (2007), Baker and Farrell (ming), Farrell and Baker (2021)) in the conclusion. Our paper is partly motivated by the empirical merger retrospectives literature. Ashenfelter et al. (2014) find that 36 of 49 studies across several industries identify significant post-merger

³Caminal (1990) considers a two-period linear demand duopoly model where firms have private information about the demand for their own product, and also raise prices to signal that they will set higher prices in the final period.

⁴The rest of the literature on dynamic games, following Ericson and Pakes (1995) and Pakes and McGuire (1994), has assumed that players observe all state variables up to iid payoff shocks so that there is no role for signaling.

price increases.⁵ Peters (2009) and Garmon (2017) show that merger simulations and other methods, such as pricing pressure indices, that are derived from static CI first-order conditions often perform poorly at predicting price changes after airline and hospital mergers. This leads naturally to the question of which alternative models can do better.

2.2 Model

In this section, we present our general model. More specific assumptions will be made in our examples and application.

2.2.1 Outline.

There are discrete time periods, t = 1, ..., T, where $T \leq \infty$, with discount factor $0 < \beta < 1$. $\beta = 0.99$ in the rest of the paper. There are a fixed set of N riskneutral firms. Each firm either sells a single-product or sells multiple products, which are symmetric in demand and are produced at the same marginal cost, at a single price. There may be observed and fixed differences in demand and costs across firms, but exactly one dimension of a firm's type is private information. In the text, we will assume that the type is continuous on a known compact interval $[\underline{\theta_i}, \overline{\theta_i}]$, but Appendix B.2.1 uses examples where firms can have two discrete types, $\underline{\theta_i}$ and $\overline{\theta_i}$. Types are assumed to evolve exogenously, and independently, from

⁵Ashenfelter et al. (2014) note that retrospectives have not typically found price increases in banking. Interestingly, the Mester (1992) analysis of a Cournot oligopoly model with asymmetric information was explicitly motivated by a desire to explain why, contrary to the predictions of Nash and tacit collusion models, concentration appeared to lead to more competitive behavior in banking.

period-to-period according to a first-order Markov process, $\psi_i: \theta_{i,t-1} \to \theta_{i,t}$.⁶

2.2.2 Within-Period Timing.

In each period t of the game, timing is as follows. Firms enter period t with their t - 1 types, which then evolve according to ψ_i . Firms observe their own new types, but neither the previous nor the new type of other firms.⁷ Each firm then simultaneously chooses a price, $p_{i,t}$, with no menu costs. Once a firm sets its period t price, it is unable to change it. A firm's profits are given by $\pi_i(p_{i,t}, p_{-i,t}, \theta_{i,t})$ and we assume that $\frac{\partial \pi_i}{\partial p_{-i,t}} > 0$ for all -i. Note that $\pi_i(p_{i,t}, p_{-i,t}, \theta_{i,t})$ only depends on current prices and the firm's type, consistent with static and time-invariant demand. Current and past prices are assumed to be perfectly observed by each firm.

2.2.3 Assumptions.

For continuous types, we make the following assumption.

Assumption 5. Type Transitions for the Continuous Type Model. The conditional $pdf \psi_i(\theta_{i,t}|\theta_{i,t-1})$

1. has full support, so that the type can transition from any value on the support to any other value in a single period.

⁶This assumption seems unrealistic, but it is consistent with how the empirical literature on horizontal mergers and the production function literature that has followed Olley and Pakes (1996) (see Doraszelski and Jaumandreu (2013) for an exception) has treated marginal cost or productivity changes.

⁷Our fully separating equilibria would be unchanged if t - 2 types were revealed.

- 2. is continuous and differentiable (with appropriate one-sided derivatives at the boundaries).
- 3. for any $\theta_{i,t-1}$ there is some θ' such that $\frac{\partial \psi_i(\theta_{i,t}|\theta_{i,t-1})}{\partial \theta_{i,t-1}}|_{\theta_{i,t}=\theta'} = 0$ and $\frac{\partial \psi_i(\theta_{i,t}|\theta_{i,t-1})}{\partial \theta_{i,t-1}} < 0$ for all $\theta_{i,t} < \theta'$ and $\frac{\partial \psi_i(\theta_{i,t}|\theta_{i,t-1})}{\partial \theta_{i,t-1}} > 0$ for all $\theta_{i,t} > \theta'$. Obviously it will also be the case that $\int_{\underline{\theta_i}}^{\overline{\theta_i}} \frac{\partial \psi_i(\theta_{i,t}|\theta_{i,t-1})}{\partial \theta_{i,t-1}} d\theta_{i,t} = 0$.

This assumption implies types are positively, but not perfectly, serially correlated so that a higher type in one period implies that a higher type in the next period is more likely.

Beliefs about rivals' types play an important role in our game. In a fully separating equilibrium, each firm will (correctly) believe that each rival has a particular type in the previous period. For convenience, we assume that beliefs about types in t = 1 have the same structure.

Assumption 6. *Initial Period Beliefs.* Firms know what their rivals' types were in a fictitious prior period, t = 0.

2.2.4 Fully Separating Equilibrium in a Finite Horizon and Continuous Type Game.

We now describe the equilibrium for a game with two ex-ante symmetric single-product duopolists, which we will use in our first example.

2.2.4.1 Final Period (*T*).

In the final period, each firm maximizes its expected payoff given its own type, its beliefs about the types of the other firms and their pricing strategies. Play is therefore consistent with a Bayesian Nash Equilibrium. If firm *j* believes that firm *i*'s period T - 1 type was $\widehat{\theta_{i,T-1}^j}$ and *j*'s period *T* pricing function is $P_{j,T}(\theta_{j,T}, \theta_{j,T-1}, \widehat{\theta_{i,T-1}^j})^8$, then a type $\theta_{i,T}$ *i* will set a price

$$p_{i,T}^*(\theta_{i,T}, \theta_{j,T-1}, \widehat{\theta_{i,T-1}^j}) = \dots$$
$$\arg\max_{p_{i,T}} \int_{\underline{\theta_j}}^{\overline{\theta_j}} \pi(p_{i,T}, P_{j,T}(\theta_{j,T}, \theta_{j,T-1}, \widehat{\theta_{i,T-1}^j}), \theta_{i,T}) \psi(\theta_{j,T} | \theta_{j,T-1}) d\theta_{j,T}.$$

2.2.4.2 Earlier Periods (1, .., T - 1).

In earlier periods, *i* may choose not to set a static best response price in order to affect *j*'s belief about its type. The equilibrium concept that we use is symmetric Markov Perfect Bayesian Equilibrium (MPBE) (Roddie (2012), Toxvaerd (2008)). An MPBE specifies period-specific pricing strategies for each firm *i* as a function of its current type, and its belief about *j*'s previous type, and *j*'s belief about *i*'s previous type; and, each firm's belief about its rival's type given observed histories of prices. Equilibrium beliefs should be consistent with Bayes Rule given equilibrium pricing strategies. If there are multiple rivals, they should all have the same beliefs given an observed history. While only current types and prices are

⁸This notation reflects the fact that we are assuming that player j used an equilibrium strategy in T-1 that revealed its type ($\theta_{j,T-1}$), but we are allowing for the possibility that firm i may have deviated so that j's beliefs about i's previous type are incorrect.

directly payoff-relevant, history can matter in this Markovian equilibrium because it affects beliefs. We will only consider fully separating MPBEs where, in every period, a firm's equilibrium pricing strategy perfectly reveals its current type, and *j*'s belief about *i*'s current type will come from inverting *i*'s pricing function.

2.2.4.3 Characterization of Separating Pricing Functions in Period t < T.

We follow Mailath (1989), which shows that one can apply the results in Mailath (1987) to this problem, in characterizing fully separating pricing functions using a definition of firm *i*'s period-specific "signaling payoff function",

 $\Pi^{i,t}(\theta_{i,t}, \widehat{\theta_{i,t}}, p_{i,t})$. This is the present discounted value of firm *i*'s expected current and future payoffs when its current type is $\theta_{i,t}$, it sets price $p_{i,t}$ and *j* believes, at the end of period *t*, that *i* has type $\widehat{\theta_{i,t}}$. $\Pi^{i,t}$ is assumed to be continuous and at least twice differentiable in its arguments. It is implicitly conditional on (i) *j*'s period *t* pricing strategy, which will depend on beliefs about types at t - 1, and (ii) both players' strategies in future periods. As *j*'s end-of-period *t* belief about *i*'s type enters as a separate argument, $p_{i,t}$ only affects $\Pi^{i,t}$ through period *t* profits. Given conditions on $\Pi^{i,t}$ that will be listed in a moment, the fully separating best response function of firm *i*, which is also implicitly conditioned on *j*'s current pricing strategy and beliefs about previous types, can be uniquely characterized as follows (see Appendix B.3 for a restatement of the Mailath (1987) theorems): *i*'s
pricing function will be the solution to a differential equation where

$$\frac{\partial p_{i,t}^*(\theta_{i,t})}{\partial \theta_{i,t}} = -\frac{\Pi_2^{i,t}\left(\theta_{i,t}, \widehat{\theta_{i,t}^j}, p_{i,t}\right)}{\Pi_3^{i,t}\left(\theta_{i,t}, \widehat{\theta_{i,t}^j}, p_{i,t}\right)} > 0,$$
(2.1)

and a boundary condition. The subscript n in $\Pi_n^{i,t}$ denotes the partial derivative of $\Pi^{i,t}$ with respect to the n^{th} argument. Assuming that lower types want to set lower prices (e.g., a type corresponds to the firm's marginal cost), the boundary condition will be that $p_{i,t}^*(\underline{\theta_i})$ is the solution to

$$\Pi_3^{i,t}\left(\underline{\theta_i}, \widehat{\theta_{i,t}^j}, p_{i,t}\right) = 0, \qquad (2.2)$$

i.e., the lowest type's price maximizes its static expected profits given *j*'s pricing policy. The numerator in (2.1) is *i*'s marginal future benefit from raising *j*'s belief about $\theta_{i,t}$, and the denominator is the marginal effect of a price increase on *i*'s current profit. For prices above a static best response price, the denominator will be negative, and the pricing function will slope upwards in the firm's type.

This characterization of a separating best response will be valid under four conditions on $\Pi^{i,t}$, in addition to continuity and differentiability,

Condition 3. Shape of $\Pi^{i,t}$ with respect to $p_{i,t}$. For any $(\theta_{i,t}, \widehat{\theta_{i,t}^{j}})$, $\Pi^{i,t}\left(\theta_{i,t}, \widehat{\theta_{i,t}^{j}}, p_{i,t}\right)$ has a unique optimum in $p_{i,t}$, and, for all $\theta_{i,t}$, for any $p_{i,t}$ where $\Pi_{33}^{i,t}\left(\theta_{i,t}, \widehat{\theta_{i,t}^{j}}, p_{i,t}\right) > 0$, there is some k > 0 such that $\left|\Pi_{3}^{i,t}\left(\theta_{i,t}, \widehat{\theta_{i,t}^{j}}, p_{i,t}\right)\right| > k$. Condition 4. Type Monotonicity. $\Pi_{13}^{i,t}\left(\theta_{i,t}, \widehat{\theta_{i,t}^{j}}, p_{i,t}\right) \neq 0$ for all $(\theta_{i,t}, \widehat{\theta_{i,t}^{j}}, p_{i,t})$. Condition 5. Belief Monotonicity. $\Pi_{2}^{i,t}\left(\theta_{i,t}, \widehat{\theta_{i,t}^{j}}, p_{i,t}\right)$ is either > 0 for all $(\theta_{i,t}, \widehat{\theta_{i,t}^{j}})$. or < 0 for all $(\theta_{i,t}, \widehat{\theta_{i,t}^j})$.

Condition 6. Single-Crossing. $\frac{\Pi_{3}^{i,t}\left(\theta_{i,t},\widehat{\theta_{i,t}^{j}},p_{i,t}\right)}{\Pi_{2}^{i,t}\left(\theta_{i,t},\widehat{\theta_{i,t}^{j}},p_{i,t}\right)} \text{ is a monotone function of } \theta_{i,t} \text{ for all } \widehat{\theta_{i,t}^{j}} \text{ and for } (\theta_{i,t},p_{i,t}) \text{ in the graph of } p_{i,t}^{*}(\theta_{i,t},\theta_{j,t-1}).$

To interpret these conditions, assume that types correspond to marginal costs. The first condition will be satisfied if, for any marginal cost and distribution of prices that the rival may set, a firm's expected current period profit is quasi-concave in its own price. This will hold for common forms of differentiated product demand such as the multinomial and nested logit models. Type monotonicity requires that, when a firm increases its price, the profit that it loses will be lower if it has higher marginal costs. This will hold for constant marginal costs. Belief monotonicity requires that a firm's expected future profits should increase when rivals believe that it has a higher cost, holding its actual cost fixed. This condition may fail: Appendix B.2.1 discusses in detail a two-type example where j will respond to i having a higher cost by setting a lower price in the next period. The single-crossing condition requires that a firm with a higher marginal cost should always be more willing to raise its price, reducing its current profits, in order to raise its rival's belief about its marginal cost. This condition can also fail.

For completeness, we also need to define beliefs that a firm will have if the rival sets a price that is outside the range of the pricing function (i.e., a price that is not on the equilibrium path). When types correspond to marginal costs, we will assume that when a firm sets a price below (above) the lowest (highest) price in the range of the pricing function, it will be inferred to have the lowest (highest) possible cost type.

2.2.4.4 Existence and Uniqueness of a Fully Separating Equilibrium.

The conditions defined above guarantee the existence and uniqueness of fully separating best responses in any period, but this does not prove the existence or uniqueness of a fully separating equilibrium in the whole game. Mailath (1989) proves existence and uniqueness in a two-period duopoly game with linear demand and there is private information about marginal costs. Appendix B.3 shows the existence and uniqueness in a finite horizon, linear demand duopoly game where marginal costs are private information. The proof requires that the marginal cost interval $(\bar{\theta} - \underline{\theta})$ is small enough so that a single-crossing condition holds when prices rise.

In our application, we will assume nonlinear demand and, to reduce the computational burden, an infinite horizon. We will therefore proceed without proofs of existence or uniqueness. Appendix B.1 details how we compute equilibrium strategies, and verify belief monotonicity and single-crossing as part of the algorithm. We will discuss examples where we cannot find a separating equilibrium below. We have only ever found a single equilibrium in finite horizon games and infinite horizon games with continuous types, but we have found examples of multiplicity in infinite horizon games with two types even when, as we describe

below, we impose a refinement that is needed to guarantee unique best response functions in that case.⁹

2.3 Examples

This section uses examples to illustrate the equilibrium of our game and the effects of a merger. Additional examples described in Appendix B.2 are also discussed.

2.3.1 Continuous-Type Duopoly Example.

2.3.1.1 Specification.

There are two ex-ante symmetric single-product firms. Demand is determined by a nested logit model, with both products in one nest, and the outside good in its own nest. Consumer c's indirect utility from buying from product i is $u_{i,c} = 5 - 0.1p_i + \sigma \nu_c + (1 - \sigma)\varepsilon_{i,c}$ where p_i is the dollar price, $\varepsilon_{i,c}$ is a draw from a Type I extreme value distribution, $\sigma = 0.25$, and ν_c is an appropriately distributed draw for c's nest preferences. For the outside good, $u_{0,c} = \varepsilon_{0,c}$. We will set market size equal to 1, so that our welfare numbers have a "per-consumer" interpretation. We first examine what happens to strategies in a finite horizon game with T = 25periods. The game is solved backwards, starting at the last period.

We assume that marginal cost is private information, and that, for each firm,

⁹In examples where we have found multiplicity, the algorithm that we use elsewhere in the paper appears to consistently pick out an equilibrium that is the limit of the equilibrium in the early periods of a finite horizon game as the number of periods grows.

it lies in the interval $[\underline{c}, \overline{c}] = [\$8, \$8.05]$. Costs evolve independently according to an exogenous truncated AR(1) process where

$$c_{i,t} = \rho c_{i,t-1} + (1-\rho)\frac{\overline{c} + \underline{c}}{2} + \eta_{i,t}$$
(2.3)

where $\rho = 0.8$ and $\eta_{i,t} \sim TRN(0, \sigma_c^2, \underline{c} - \rho c_{i,t-1} - (1-\rho)\frac{\overline{c}+\underline{c}}{2}, \overline{c} - \rho c_{i,t-1} - (1-\rho)\frac{\overline{c}+\underline{c}}{2})$, where TRN denotes a truncated normal distribution, and the first two arguments are the mean and variance of the untruncated distribution, and the third and fourth arguments are the lower and upper truncation points. $\sigma_c = \$0.025$.

Two features of this parameterization are worth highlighting. First, marginal costs are restricted to a narrow range (diverging by less than 0.32% from mean value) and the probability that a firm will switch from a relatively high cost to a relatively low cost across periods is quite high.¹⁰ Therefore, no signal should affect a rival's posterior belief about a firm's next period marginal cost very much. Despite this, we find large signaling effects. Second, the demand parameters imply high margins and limited substitution to the outside good in both static and dynamic equilibria. As we will discuss, these features contribute to the existence of a fully separating equilibrium with large price effects.

2.3.1.2 Equilibrium Outcomes and Strategies.

Table 2.1 shows expected price levels, the standard deviation of prices and various welfare measures when we simulate data using equilibrium strategies in

¹⁰For example, the probability that a firm with the highest marginal cost has a cost in the lower half of the support in the next period is 0.32.

				Expecte	d Welfare M	leasures
				Per I	Market Size	Unit
	Nature of	Mean	Std. Dev.	Cons.	Producer	Total
Period	Equilibrium	Price	Price	Surplus	Surplus	Welfare
T-24	MPBE	\$24.76	\$0.47	\$30.91	\$15.96	\$46.87
T-13	MPBE	\$24.76	\$0.47	\$30.91	\$15.96	\$46.87
T-10	MPBE	\$24.75	\$0.47	\$30.92	\$15.95	\$46.87
T-7	MPBE	\$24.68	\$0.45	\$30.98	\$15.89	\$46.88
T-4	MPBE	\$24.25	\$0.36	\$31.40	\$15.51	\$46.91
T-2	MPBE	\$23.38	\$0.17	\$32.23	\$14.74	\$46.97
T-1	MPBE	\$22.88	\$0.06	\$32.71	\$14.29	\$47.00
Т	BNE	\$22.62	\$0.01	\$32.96	\$14.05	\$47.01
Infinite Horizon	Stationary MPBE	\$24.76	\$0.47	\$30.91	\$15.96	\$46.87

Table 2.1: Equilibrium Prices and Welfare in the Duopoly Game

Notes: except for the last row, all prices are based on equilibrium strategies in a finite horizon model with parameters described in the text. The last line reports results for the stationary strategies in an infinite horizon model with the same parameters.

different periods of the finite horizon game. For comparison, expected jointprofit maximizing prices and static Nash equilibrium prices under CI (given average costs) are \$45.20 and \$22.62, with small standard deviations (\$0.007 and \$0.011). Signaling MPBE prices are higher and significantly more volatile than Nash prices when the game is more than a couple of periods from the end, but they are always much lower than joint profit-maximizing prices. We now describe the strategies that result in these outcomes.

Figure 2.1(a) shows four static BNE period T pricing functions for firm 2, for different values of firm 1's period T - 1 marginal cost $(c_{1,T-1})$, assuming that both firms know/believe that $c_{2,T-1} =$ \$8. Firm 2's price increases with $c_{1,T-1}$ as firm 1's expected period T price rises with $c_{1,T-1}$. However, the variation in firm 1's prior cost affects firm 2's price by less than one cent, and, averaging across



Figure 2.1: Period T and T-1 Pricing Strategies in the Finite Horizon, Continuous Type Signaling Game

all possible cost realizations, average prices and welfare are almost identical to outcomes with CI.¹¹ Therefore the existence of asymmetric information alone (i.e., when not combined with some form of dynamics) does not generate interesting effects given our parameters.

There is an incentive to signal in period T-1 because a firm's price can affect its rival's price in period T. Assuming both firms' period T-2 costs were \$8, Figure 2.1(b) shows firm 1's signaling pricing function (found by solving the differential equation in (2.1) given the boundary condition (2.2) if it expected that firm 2 was using its period T strategy. We reproduce the period T pricing strategy for comparison. The pricing functions intersect for $c_{1,T-1} = \$8$, but signaling may lead firm 1 to raise its price by as much as 20 cents for higher costs. At first blush, this large increase may seem surprising given that we know the effect on firm 2's price can only be small. However, the assumed demand implies that firm 1's profit function, shown in Figure 2.2, is sufficiently flat that, if $c_{1,T-1} = \$8.025$, its expected lost period T-1 profit from using a signaling price of \$22.76, rather than the statically optimal period T-1 price of \$22.61, is only \$0.00070 per consumer, which is less than the (discounted) expected period T profit gain of \$0.00079 from being viewed as a firm with a $c_{1,T-1} = \$8.025$ rather than $c_{1,T-1} = \$8.0001$ (which is how firm 2 would interpret a price of \$22.61).

Figure 2.1(b) assumed that firm 2 was using its period T strategy with no signaling. Figure 2.1(c) shows firm 2's best signaling response when firm 1 uses the strategy in Figure 2.1(b) (repeated in the new figure as a comparison). As firm

¹¹Expected producer and consumer surplus differ by less than \$0.0001 across these models.

Figure 2.2: Expected T - 1 Period Profit Function: $c_{1,T-1} = \$8.025$ and $c_{1,T-2} = c_{2,T-2} = \8



Notes: the profit function is drawn "per potential consumer" for a firm assumed to have a marginal cost of \$8.025, and with a rival using the static BNE pricing strategy when both firms' previous period marginal costs were \$8.

1's expected price has increased, firm 2's static best response pricing function shifts upwards. Of course, this positive feedback will cause firm 1's pricing function to rise as well. Figure 2.1(d) shows the equilibrium period T - 1 pricing functions. The increase in the slope and the dispersion of the pricing functions means that period T - 1 prices will be higher and more volatile than period T prices.

The increased vertical spread also means that period T - 1 prices are more sensitive to perceived period T - 2 costs which increases period T - 2 signaling incentives. Figure 2.3 shows a selection of equilibrium pricing functions for period T - 2 and earlier periods. The pricing functions become more spread out and the level of prices increases, although by successively smaller amounts, in earlier periods. Further back than period T - 15 equilibrium pricing functions and average prices barely change. The figure also plots the stationary pricing strategies Figure 2.3: Equilibrium Pricing Functions for Firm 1 in the Infinite Horizon Game and Various Periods of the Finite Horizon Game.



Notes: all functions drawn assuming that firm 1's perceived marginal cost in the previous period was \$8.

that we compute for an infinite horizon game with the same parameters. They are indistinguishable from the strategies in the early periods of the finite horizon game.¹²

2.3.2 Merger Analysis.

There are many possible applications of our model, but we will focus on its predictions for horizontal mergers. We present a simple motivating example using the infinite horizon, continuous cost model with the same demand and marginal cost parameters that we have just assumed, although we will allow for more firms. We will assume that a merger occurs as an unanticipated one-off shock, i.e., firms signal assuming the prevailing market structure will last forever. As discussed in a two-type example with up to seven firms in Appendix B.2.1.3, signaling tends to have more effects on pricing when there are fewer firms, because each firm's price will tend to have a larger effect on its rivals' next period prices.

Table 2.2 shows the effects of 4-to-3 and 3-to-2 mergers. Before either merger, there are symmetric single-product firms. In the upper panel, we assume that a merger eliminates a product, so that after the merger there are only single-product firms with symmetric demand. If we assume that the firms always use equilibrium signaling strategies, then a 4-to-3 merger with no synergy (implying the firms remain symmetric post-merger) will raise average prices by 8.5%. To prevent the merged firm's price from rising, the merger would need to reduce the

¹²We have consistently found this convergence except in cases when the conditions required for separation are violated or are very close to being violated (in which case the infinite horizon strategies may not converge).

(a) Merger Leads to the Elimination o	f a Product By the M	lerged Firm
	4-to-3 Merger	3-to-2 Merger
Signaling MPBE Pre-Merger Avg. Price	\$18.25	\$19.79
Post-Merger Avg. Price of Merged Firm if No Marginal Cost Synergy	\$19.81 (+8.5%)	\$24.75 (+25.1%)
Post-Merger Avg. Price of Non- Merging Firm if No Marginal Cost Synergy	\$19.81 (+8.5%)	\$24.75 (+25.1%)
Merged Firm Marginal Cost Required to Prevent Merged Firm Avg. Price from Rising	\$5.73	-\$2.20
If Merger Analyzed under CI Implied Pre-Merger Avg. Marginal Cost	\$8.29	\$8.62
Merged Firm Marginal Cost Required to Prevent Prices from Rising	\$7.11	\$5.13
Avg. Merged Firm Price in Signaling Model if Analyst Required Marginal Cost is Realized	\$19.17 (+5.0%)	\$23.25 (+17.4%)
(b) Merging Firm Owns Two	Products Post-Merg	er
	4-to-3 Merger	 3-to-2 Merger
Signaling MPBE		
Pre-Merger Avg. Price	\$18.25	\$19.79
Post-Merger Avg. Price of Merged Firm if No Marginal Cost Synergy	\$21.53 (+18.0%)	\$27.18 (+37.3%)
Post-Merger Avg. Price of Non- Merging Firm if No Marginal Cost Synergy	\$19.12 (+4.8%)	\$23.59 (+19.2%)
Merged Firm Marginal Cost Required to Prevent Merged Firm Avg. Price from Rising	\$2.26	-\$11.92
If Merger Analyzed under CI Implied Marginal Cost	\$8.29	\$8.62
Merged Firm Marginal Cost Required to Prevent Prices from Rising	\$3.43	-\$2.05
Avg. Price in Signaling Model	\$18.85 (+3.2%)	\$23.00 (+16.2%)

Table 2.2: The Effects of Signaling on Mergers and Merger Analysis When Firms Use Infinite Horizon Signaling Strategies

average marginal cost of the merging firm from \$8.025 to \$5.73, a 29% reduction.¹³

We can compare these effects to the predictions of an analyst who knows demand and uses a standard CI merger simulation model.¹⁴ Using average prices and CI first-order conditions, the analyst would infer that average pre-merger marginal costs are equal to \$8.29 (i.e., higher than they really are), and that a 14% synergy (reducing marginal costs to \$7.11) would prevent price increases. If the 14% synergy was achieved but firms play a signaling equilibrium after the merger, then the merged firm's average price would increase to \$19.17 (a 5% post-merger increase). In the case of a 3-to-2 merger, all of the effects seen in the 4-to-3 case become larger, and, in fact, the merged firm's marginal cost would need to be negative to prevent a price increase.¹⁵ The realization of the synergy identified by a CI simulation would not prevent prices from rising by 17%.

The lower panel assumes that, after the merger, the merged firm has two products, which have the same marginal cost and which are sold at the same price. This restrictive assumption preserves the structure that each firm has one piece of private information and can send exactly one signal. Ownership of two products increases incentives to raise prices, and hence the size of required synergies. As in the upper panel, a CI analysis will underpredict price increases and required

¹³We assume that the range of marginal costs, \$0.05, and the process by which marginal costs evolve remain the same after the merger and after any synergy is realized.

¹⁴This characterization follows how merger simulation is used in the academic literature. Agency economists typically calibrate the price and nesting parameters in the demand system to match average margins given CI Nash pricing. An incorrect static CI Nash assumption would then lead to the wrong demand parameters.

¹⁵We assume that the firm cannot freely dispose of products so that it cannot choose to produce an infinite amount if it has negative marginal costs.

synergies.

2.3.3 Signaling Incentives and the Existence of Separating Equilibria.

In our example, signaling incentives are relatively weak because marginal costs are only weakly correlated from period to period. Increasing the AR(1) parameter or $\overline{c} - \underline{c}$, or reducing σ_c tend to increase signaling incentives and raise equilibrium prices. However, when price increases are too large, the conditions for characterizing best responses can fail and we may not be able to find a separating equilibrium.

The first six columns of Table 2.3 show, for different periods, the baseline average prices and average prices when signaling incentives are strengthened. Small parameter changes result in higher equilibrium prices, but larger changes result in the failure of our algorithm as we cannot define best response pricing functions. Pooling or partial pooling equilibria may exist, but we do not know how to characterize them. Appendix B.2.1.2 uses a two-type example to examine the failure of the conditions, including belief monotonicity, in more detail.¹⁶

However, as illustrated in the final column, we can sustain separating equi-

¹⁶The two-type model has a much lower computational burden but requires imposing a refinement just to identify unique separating best responses. Specifically, we always find the best response that achieves separation at the lowest cost to the signaling firm, consistent with the type of "intuitive criterion" (Cho and Kreps (1987)) refinement that has been widely used in one-sided signaling models with two types. However, even with this refinement, we have found examples of multiple separating equilibria in the infinite horizon version of the two-type model. The algorithm that we use to produce the reported results appears to consistently select the equilibrium that corresponds to the limit of a (seemingly unique) equilibrium in a finite horizon game as the number of periods grows large.

					Red	uce	Expand Range
	Baseline	Exp	pand Range	5	Std.	Dev.	&Std. Dev.
$\overline{c} - \underline{c}$	[8,8.05]	[8,8.075]	[8,8.15]	[8,8.3]	[8,8.05]	[8,8.05]	[8,8.50]
σ_c	0.025	0.025	0.025	0.025	0.02	0.01	0.25
T-24	24.76	26.51	-	-	25.71	-	24.90
T-10	24.75	26.59	-	-	25.70	-	24.89
T-9	24.74	26.59	fails	-	25.69	fails	24.89
T-8	24.72	26.57	28.48	-	25.66	28.58	24.89
T-7	24.68	26.50	29.17	fails	25.60	28.76	24.87
T-6	24.61	26.37	29.35	30.40	25.49	28.65	24.85
T-1	22.88	23.05	23.42	23.93	22.93	23.05	23.55
Т	22.62	22.63	22.67	22.74	22.62	22.62	22.84
∞	24.76	26.50	fails	fails	25.71	fails	24.90

Table 2.3: Equilibrium Pricing in a Finite Horizon Game with Alternative Cost Specifications

Notes: unit is \$. Values in all but the last line are based on the duopoly, continuous type, finite horizon model with demand parameters described in the text (cost parameters indicated in the table). The last line reports results for the stationary strategies in the infinite horizon model with the same parameters. "Fails" indicates that the belief monotonicity or single-crossing conditions fail so that we cannot calculate signaling best response pricing functions.

libria if we increase $\overline{c} - \underline{c}$ and increase σ_c simultaneously.¹⁷ This pattern will be relevant for our application.

2.3.4 Additional Examples.

Appendix B.2.1 uses two-type duopoly examples to examine how price effects vary with the number of firms and to examine the relationship between the existence of separating equilibria, the magnitude of price effects, the serial correlation of costs and the extent to which, when a firm's price rises, demand is diverted to the outside good. When there is limited diversion to the outside good we find large increases in prices above static CI Nash levels (an increase of 45% in one case) under duopoly even when there is moderate serial correlation in costs (e.g., $Pr(c_{i,t} = c_{i,t-1})=0.75$). On the other hand, price increases are small with more than three firms, and only small increases can be sustained in separating

¹⁷The probability that a cost goes from one extreme of the support to the opposite half of the support is 0.32, which is the same as in the baseline case.

equilibria when there is more diversion to the outside good even under duopoly. The examples suggest that thinking about the effects of signaling is most relevant when two or three firms dominate a market or a very distinct segment of a market.

Appendix B.2.2 present three simple duopoly examples where marginal costs are fixed and known, but firms have private information about some other element of their payoff function (a feature of demand, the weight managers place on revenues rather than profits, or the weight they place on the profits of rivals). Signaling can raise prices significantly above CI Nash levels in each case.

2.4 Empirical Application: The MillerCoors Joint Venture

In this section, we apply our model to data from the U.S. beer industry around the time of the 2008 MC JV. MW show that, relative to the price of imports, the real prices of brands owned by MC and Anheuser-Busch (AB) increased after the JV.¹⁸ We describe the setting and the data, before explaining the calibration of our model using pre-JV pricing data and reporting how well it predicts observed changes in pricing after the JV. Finally, we examine how well the CI models that have previously been used to explain why price increased fit the data.

2.4.1 The JV and Its Effects.

The MC JV, announced in October 2007, effectively merged the U.S. brewing, marketing and sales operations of SABMiller (Miller) and MolsonCoors (Coors), the second and third largest U.S. brewers. The Department of Justice (DOJ) decided not to challenge the transaction in June 2008 because it expected "large reductions in variable costs of the type that are likely to have a beneficial effect on prices".¹⁹ For example, the JV was expected to lower transportation costs by producing Coors products at Miller breweries around the country. Ashenfelter et al. (2015) provide evidence that transportation efficiencies were realized.

¹⁸Anheuser-Busch was purchased by InBev in 2008. Throughout the paper we will use AB to refer to Anheuser-Busch before 2008 and Anheuser-Busch InBev afterwards.

¹⁹Department of Justice press release, 5 June 2008.

MW show that, at a national level, the real prices (i.e., deflated by the CPI-U price index) of the most popular domestic brands, such as Bud Light (BL), Miller Lite (ML) and Coors Light (CL), increased after the JV, relative to the prices of imported brands, such as Corona Extra and Heineken, which MW use as controls for industry-wide cost shocks. Regressions in Appendix B.4 quantify these price increases to lie between 40 cents and a dollar per 12-pack, or 3%-6%, depending on the specification. We will proceed assuming that MW's interpretation that the relative price increase was a causal anticompetitive effect of the JV is correct.²⁰

An important feature of the relative price change is that AB's prices increased as much as those of Miller and Coors. If AB's marginal costs were unaffected by the JV, this pattern is inconsistent with static CI Nash pricing, as a static best response function would predict that AB should have responded to any JV price increase by raising its prices by a smaller amount.

2.4.2 Data.

We use the same data as MW, which comes from the IRI Academic Dataset (Bronnenberg et al. (2008)) which provides weekly UPC-store-level scanner data for the beer category from an unbalanced panel of grocery stores from 2001 to 2011. Appendix B.4 provides details, but we will note in the text where our treatment differs from MW. We will follow the typical convention of assuming that retail prices are set directly by brewers, and that any retail margin is a fixed component of brewers' marginal costs.²¹

Table 2.4 lists the 20 brands with the largest sales by volume in 2007, together with additional brands that MW include in their analysis. The table lists market shares and average nominal prices (per 144 oz, the volume in a standard 12-pack) in 2007 and 2011. Most domestic brands are differentiated from imports

²⁰This interpretation is complicated by how the Great Recession may have affected demand and the fall in the deflator, from 220.0 in July 2008 to 210.2 in December 2008, at exactly the same time that the merger was being consummated.

²¹MW estimate a model that allows for a monopolist retail margin, and cannot reject a model with fixed retail pass-through.

			2007			2011	
Brand	Company	Packs	% 18 +	Mkt. Share	Price	Mkt. Share	Price
Bud Light*,†	AB	10	72.5%	15.7%	\$8.29	15.7%	\$8.92
Miller Lite ^{*,†}	Μ	10	75.1%	10.0%	\$8.11	8.4%	\$8.73
Coors Light ^{*,†}	U	10	74.8%	8.3%	\$8.36	9.4%	\$8.98
Budweiser [†]	AB	10	70.8%	7.7%	\$8.30	6.5%	\$9.00
Corona Extra ^{†,⇔}	GM	ഹ	15.6%	4.1%	\$13.88	3.9%	\$13.46
Natural Light*	AB	7	68.6%	3.9%	\$6.01	3.2%	\$7.15
Busch Light*	AB	6	78.4%	2.8%	\$6.07	2.5%	\$6.96
Miller High Life [†]	Μ	6	54.1%	2.4%	\$6.33	2.2%	\$7.21
Heineken $^{\dagger,\diamond}$	Η	~	12.8%	2.3%	\$14.06	2.3%	\$13.86
Miller Genuine Draft [†]	Μ	10	67.0%	2.3%	\$8.26	1.3%	\$8.94
Michelob Ultra*; [†]	AB	6	27.4%	2.1%	\$10.05	2.4%	\$10.51
Busch	AB	6	70.0%	1.9%	\$6.08	1.6%	\$7.05
Keystone Light*	U	9	81.4%	1.4%	\$5.83	1.5%	\$7.03
Budweiser Select	AB	6	62.0%	1.3%	\$8.37	0.7%	\$8.76
Milwaukee's Best Light*	Μ	9	66.8%	1.3%	\$5.37	0.8%	\$6.19
Corona Light*,†,◇	GM	с	2.3%	1.2%	\$14.23	1.3%	\$13.79
Tecate^\diamond	Н	7	66.3%	1.2%	\$8.65	1.2%	\$9.04
Natural Ice	AB	7	51.3%	1.1%	\$5.96	0.9%	\$7.19
Pabst Blue Ribbon	SP	6	49.3%	1.0%	\$6.26	1.4%	\$7.53
Milwaukee's Best	Μ	ഹ	61.8%	0.8%	\$5.46	0.4%	\$6.46
Coors†	U	10	73.3%	0.8%	\$8.44	1.0%	\$8.84
Michelob Light ^{*,†}	AB	~	29.3%	0.7%	\$9.76	0.3%	\$10.72
Heineken Prem. Light*,†,◇	Н	ഹ	1.9%	0.6%	\$14.28	0.5%	\$14.18
Notes: market shares and price servings. "Packs" is the number	es are based c r of 2007 bot	n all un tle/can-1	its sold in oack size c	packs equivaler ombinations fo	nt to 6, 12 r 6, 12, 18	, 18, 24, 30 an , 24 and 30 pa	d 36 12oz cks, as 36

packs are rare. "% 18+ " is the percentage of 2007 volume sold in the packs of more than 18 cans or bottles. 2007 companies are: AB=Anheuser-Busch, M=SABMiller, C=MolsonCoors, GM=Grupo-Modelo, H=Heineken,

SP=S&P. Prices are nominal prices per 12-pack equivalent (i.e., total dollars sold in all pack sizes divided by total volume in 144oz. units). *=light beers, † =included in MW's sample, $^{\diamond}$ =imports.

Table 2.4: Highest-Selling Beer Brands in 2007 with Ownership, Share and Average Prices.

by being sold primarily in larger packs and at lower prices. The relative prices of domestic brands increased after 2007, but, although CL gained share at ML's expense, the domestic brewers' market shares remained stable: for example, AB's volume share was 41.3% in 2007, 41.5% in 2009 and 39.6% in 2011, with light beer shares of 50.0%, 50.8% and 50.6% respectively.²²

We calibrate the model to match observed pre-JV dynamics of BL, CL and ML prices. As an example of the dynamics in the data, Figure 2.4 shows monthly average nominal prices of 12-packs for the flagship domestic brands in two large markets for 49 months around the consummation of the JV.²³ Average prices are calculated excluding all sales at prices that IRI indicates are temporary price reductions, as changes in regular prices are more likely to reflect changes in wholesale prices. Within-year volatility is a clear feature of this data, even if we ignore the drop in ML prices during the DOJ's investigation.

2.4.3 Calibration of the Dynamic Asymmetric Information Model.

We calibrate an infinite horizon, continuous marginal cost three-firm/product version of our model using pre-JV data, and then compare its predictions with post-JV data. We say "calibration", even though we estimate five cost parameters, because of the strong assumptions we make to limit the computational burden. The most important simplification is that our calibration will treat data from different markets as data from independent repetitions of the same game, rather than reflecting markets with different demand and cost primitives.

²²Appendix B.4 presents a figure showing the evolution of market shares over this period. The post-JV decline in the shares of several non-flagship domestic brands reflected a continuation of pre-existing trends.

 $^{^{23}}$ We use nominal prices so that they are not distorted by fluctuations in the CPI-U deflator, including the drop referenced in footnote 20. See Appendix B.4 for the same figure plotted using real prices.

Figure 2.4: Average Nominal Prices (excluding sales) of 12-Packs of the Domestic Flagship Brands in Two Regional Markets Around the JV.



Notes: averages are calculated as the total dollar sales of 12-packs at prices not identified as temporary price reductions, divided by the number of 12-packs sold. See Appendix B.4 for the same figure with real prices.

2.4.3.1 Products.

We model the pricing of three brands. We label these brands as BL, ML and CL, and will estimate the cost parameters to match the observed price of dynamics of these flagship products. However, Appendix B.4 shows that the prices of brands in the same portfolio (e.g., Budweiser and BL) are highly correlated, and one can also view the brands as representing the portfolios of AB, Miller and Coors. Products of other brewers, including imports and craft beers, are included in the outside good.²⁴ We will assume that ML and CL are symmetric before the JV, as we will have to assume that MC sets the same price for both products after the JV. Appendix B.4 also shows the correlation of ML and CL prices increased after the JV.

2.4.3.2 Demand.

We assume static, time-invariant nested logit demand, with the three brands in the same nest. The parameters are the nesting and price parameters, and the mean utilities (excluding the effect of price) of BL and ML/CL. Our baseline parameters are chosen so that, at average real prices in the pre-JV data, the average own price elasticity is -3, the market shares of the three products are 28% for BL and 14% each for ML/CL and, on average, if the price of one brand increased, 85% of the demand that it loses would go to the other brands (with the remainder to the outside good).²⁵ When we use weekly data on 6/12/18/24/30-packs and exclude temporary price reductions, the pre-JV cross-market average prices are \$10.09 for BL and \$9.95 for ML/CL, and the implied nesting and price parameters are 0.772 and -0.098, and the BL and ML/CL mean utilities are 1.044 and 0.863 respectively.

As motivation for the assumed elasticity and diversion, Table 2.5 reports five

²⁴In an earlier version, we estimate the model allowing for imports to be a non-signaling fringe that used Bayesian Nash pricing. The model predicted that, after the JV, they would raise their prices by a couple of cents.

²⁵These assumed shares overstate the share of BL relative to ML and CL, but understate the share of AB, relative to Miller and Coors, in the beer market and the light beer segment.

Table 2.5. Estimates of Demand	Table 2.5:	Estimates	of Demand
--------------------------------	------------	-----------	-----------

	(1)	(2)	(3)	(4)	(5)
	Nested	RCNL	RCNL	Nested	Nested
	Logit			Logit	Logit
Nests	All Beer	All Beer	All Beer	Flagship/	Flagship/
				Other	Other
Data Freq.	Monthly	Monthly	Quarterly	Monthly	Weekly
Real Price Coefficient	-0.056	-0.083	-0.099	-0.073	-0.047
(2010 dollars)	(0.017)	(0.014)	(0.014)	(0.018)	(0.011)
Nesting Coefficients					
Single All Brand Nest	0.741	0.838	0.831	-	-
	(0.051)	(0.039)	(0.039)		
Two Nests					
Domestic Flagship	-	-	-	0.838	0.898
				(0.049)	(0.040)
Other Brands	-	-	-	0.634	0.815
				(0.047)	(0.037)
Income Coefficients					
*constant	-	0.014	0.014	-	-
		(0.005)	(0.005)		
*price	-	0.001	0.001	-	-
-		(0.000)	(0.000)		
*calories	-	0.004	0.004	-	-
		(0.002)	(0.002)		
Median Product	-2.31	-4.71	-5.41	-2.51	-3.12
Elasticity					
Mean ML Brand	-1.66	-3.68	-4.22	-3.06	-3.09
Price Elasticity					
Mean Flagshin	0.41	0.48	0.47	0.83	0.00
Diversion	0.41	0.40	0.47	0.05	0.90
	F O O O (0.040/		4.000/	0.000/
% Change in Flagship	-5.20%	-8.24%	-9.65%	-4.30%	-2.20%
Sales Given /5¢					
Domestic Price Kise					
Observations	94,656	94,656	31,777	94,656	405,004

Notes: market size is defined as 50% more than the highest sales observed in the geographic market for monthly and quarterly specifications. For the weekly specifications it is estimated as 50% more than the sum of the highest sales from stores observed in the scanner data that week. ML Brand Elasticity reflects the change in ML sales when the prices of all ML products are increased. Mean Flagship Diversion is the average proportion of lost sales that go to other flagship products (i.e., BL, ML and CL products) when the price of a flagship product is increased. The change in flagship sales after a 75 cent price rise is the average across pre-JV observations change in total flagship sales when the prices of all domestic products are increased by 75 cents. Standard errors, clustered on the geographic market, in parentheses. sets of demand estimates (the first three will be used in Section 2.4.4).²⁶ For these specifications, we follow MW as closely as possible in the choice of data, instruments and controls, except that we use optimal GMM for the nested logit models as doing so affects the estimates.²⁷ The first three columns contain one nested logit specification, using monthly data, and two random coefficients nested logit (RCNL) specifications, where the 13 MW brands are all included in a single inside nest, and preferences vary with income. The remaining columns estimate nested logit models using monthly and weekly data (we will use weekly price changes when estimating the cost parameters) where flagship products are grouped into a flagship nest, and the remaining products are placed in an "other beer" nest with a different nesting coefficient. The flagship nesting coefficients are larger, consistent with these brands being close substitutes.

The table reports several implied statistics for each specification, including the average ML brand elasticity (i.e., the effect on demand when all ML prices increase), the proportion of lost demand that switches to other flagship products when a flagship price is increased, and the average, across pre-JV observations, predicted change in flagship sales when the prices of all domestic products increase by 75 cents, which is within the range of the observed post-JV price change. The statistics vary across the specifications. Given that the limited decline in flagship brand and domestic brewer market shares after the JV, we assume values for elasticity and diversion that are consistent with the estimates in columns (4) and

(5).

²⁶All specifications include time period and product (brand*size) fixed effects, and use data from Jan 2005 to Dec 2011, excluding June 2008 to May 2009. All estimates use two-step optimal GMM. Instruments are the same as in MW for the relevant specification, apart from the two nest models where we define instruments for the number and distance measures for other products based on products in the same nest, and interact instruments with a flagship brand dummy.

²⁷None of the specifications yield exactly the same estimates as MW although the monthly RCNL estimates are almost identical.

2.4.3.3 Marginal Costs.

We assume that the marginal costs of product *i*, $c_{i,t}$, lie on the interval [$\underline{c_i}, \underline{c_i} + c'$], where we estimate $\underline{c_{BL}}$, $\underline{c_{ML/CL}}$ and c'. $c_{i,t}$ evolves according to an AR(1) process with truncated innovations

$$c_{i,t} = \rho c_{i,t-1} + (1-\rho)\frac{c_i + c_i + c'}{2} + \eta_{i,t}$$
(2.4)

where $\eta_{i,t} \sim TRN(0, \sigma_c^2, \underline{c_i} - \rho c_{i,t-1} - (1-\rho)\frac{c_i + c_i + c'}{2}, \underline{c_i} + c' - \rho c_{i,t-1} - (1-\rho)\frac{c_i + c_i + c'}{2})$ and σ_c is the standard deviation of the untruncated innovation distribution. The fit of the model improves only slightly if we allow ρ , σ_c and c' to vary across firms.

2.4.3.4 Objective Function, Matched Statistics and Identification.

The cost parameters are estimated using indirect inference (Smith (2008)). For a given value of the cost parameters, we solve the model (see Appendix B.1.2 for the method) and simulate a time-series of data to calculate six statistics/regression coefficients that we match to ones from the data that we describe below. The estimation problem is

$$\widehat{\theta} = \arg\min_{\theta} g(\theta)' W g(\theta)$$

where $g(\theta)$ is a vector where each element k has the form $g_k = \frac{1}{M} \sum_m \tau_{k,m}^{data} - \widehat{\tau_k(\theta)}$ where $\tau_{k,m}^{data}$ is a statistic estimated using the actual data and $\widehat{\tau_k(\theta)}$ is the equivalent coefficient estimated using simulated data from the model solved using parameters θ . W is a weighting matrix. The reported results use an identity weighting matrix, although the choice of W has little effect on the parameters as we match all of the moments almost exactly. The objective function is minimized using fminsearch in MATLAB (version 2018a). Standard errors are calculated treating different markets before the JV as independent observations on the same game. Estimation takes between 12 and 24 hours.²⁸

For each geographic market, we calculate six statistics using data from January 2001 to the announcement of the JV in October 2007. Our preferred specification uses weekly data and the five most common pack sizes (6, 12, 18, 24 and 30-packs).²⁹ Market-week-brand-size average real prices per 12-pack equivalent are calculated excluding temporary store price reductions, and using only market-weeks where we observe more than five stores.³⁰ The first two statistics that we match are the (unweighted) average prices for BL and ML across pack sizes and weeks. The third statistic is the interquartile range (IQR) of prices for BL. This is calculated as the IQR of the residuals for each market from a regression where, pooling markets, we regress the week-market-size prices of BL products on dummies for the specific set of stores observed in the market-week (interacted with pack size) and week-size fixed effects in order to control for fixed retail price differences across stores and any national promotions. The remaining statistics are coefficients from market-brand-specific regressions of market-week-brand-size prices on the lagged prices of all three brands. Specifically we use the averages of $\rho^{ML,ML}$ and $\rho^{CL,CL}$, $\rho^{BL,CL}$ and $\rho^{BL,ML}$, and $\rho^{ML,CL}$ and $\rho^{CL,ML}$, where $\rho^{i,j}$ is the coefficient on the lagged price of brand *j* when the dependent variable is the price of brand *i*. These AR(1) regressions include dummies for the exact set of stores observed, interacted with pack size, and a linear time trend.

Assuming that the equilibrium is unique, the intuition for identification is straightforward.³¹ Given the assumed demand parameters and the observed price

²⁸Computationally light two-step approaches, which are often used to estimate dynamic games, cannot be used because they require that all serially-correlated state variables, which in our setting would include beliefs, are observed by the researcher.

²⁹Our model does not have different pack sizes, market heterogeneity, varying sets of stores or time trends, so the regressions using simulated data do not control for these factors.

 $^{^{30}}$ See Appendix B.4 for a discussion of the sample selection.

³¹The possibility that our game has multiple equilibria may create two issues for estimation. First, the objective function may be hard to minimize if our solution algorithm jumps between different sections of the equilibrium correspondence. In practice, we can match our moments almost exactly across many alternative parameterizations. Second, another equilibrium supported by different parameters might give similar predictions to the equilibrium that our algorithm finds. This is essentially a potential identification problem. Here we have to rely on the fact that we have never found multiple equilibria in continuous-type games, although we suspect that they may exist for some parameters.

levels, the mark-ups implied by the model will identify the lower bounds on brand marginal costs. The AR(1) coefficients and the dispersion of prices will identify the range of costs and the parameters of the cost innovation process.³² We will compare additional statistics that we do not match during estimation to understand the fit of the model.

To provide a sense of the AR(1) coefficients, Table 2.6 shows the coefficients from similar regressions that pool data from all markets for four alternative samples. Panel (a) reports the results for our preferred specification. The serial correlation parameters for a product's own price are between 0.41 and 0.46, while the cross-product correlations are positive but smaller. If price reductions are included (panel (c)), serial correlations fall, which is consistent with sales lasting one week and being proceeded and followed by higher regular prices. Serial correlation is higher if we use only 12-packs (panel (b)). Panel (d) repeats (a) using monthly prices and market, rather than group-of-store, fixed effects (equivalent regressions will be used in our monthly data specification). In this case, the serial correlation parameters increase, but further investigation reveals that this happens primarily due to the change in the fixed effects.³³

While our calibration does not seek to match cross-market heterogeneity, the serial correlation coefficients show some interesting patterns across markets. Using data simulated from our model, we typically estimate higher serial correlation parameters when we change the parameters to induce larger signaling effects on prices, by, for example, reducing diversion to the outside good. Given any type of logit or nested logit preferences, diversion to other brands will tend to be lower when the market share accounted for by the signaling brands is higher. Figure 2.5(a) shows scatter plots of the estimated market-level serial correlation parameters for BL, ML and CL against the share of all beer sales accounted for AB, Miller

³²Larger cross-brand ρ coefficients imply stronger signaling effects, so that a smaller range of costs may be required to generate the dispersion of prices in the data.

³³We have estimated monthly regressions including set of store fixed effects and dropping market-months where the set of stores changes within months. This causes the number of observations to drop dramatically: for example, the number of observations in the BL regression falls to 2,806, and the estimated coefficient on p_{t-1}^{BL} falls to 0.318. For some individual markets, there is not enough data to estimate serial correlation coefficients.

(a) Week, F	Price Redu	ctions Exc	uded,		(b) Wee	k, Price Red	ductions Excluded,	
Fixed E	All Pack S	Set of Stor	es		Fixe	d Effects fo	or Set of Stores	
	(1)	(2)	(3)		(1)	(2)	(3)	
	$p_{BL,t}$	$p_{ML,t}$	$p_{CL,t}$	_	$p_{BL,t}$	$p_{ML,t}$	$p_{CL,t}$	
$p_{BL,t-1}$	0.451	0.056	0.043		0.489	0.071	0.028	
	(0.033)	(0.017)	(0.010)		(0.032)	(0.026)	(0.018)	
$p_{ML,t-1}$	0.030	0.409	0.016		0.062	0.505	0.028	
,	(0.011)	(0.036)	(0.014)		(0.013)	(0.038)	(0.012)	
$p_{CL,t-1}$	0.027	0.021	0.461		0.004	0.016	0.549	
	(0.012)	(0.015)	(0.040)		(0.012)	(0.015)	(0.043)	
Observations	36,659	36,670	36,700		10,829	10,817	10,828	
R-squared	0.979	0.972	0.978		0.964	0.945	0.957	
Mean Price	10.08	9.95	9.94		10.3	10.22	10.19	
SD residuals	0.184	0.221	0.197		0.144	0.183	0.163	
(c) Week, I	Price Redu	ctions Incl	uded,		(d) Month, Price Reductions Excluded			
	All Pack S	izes,			All Pack Sizes,			
Fixed E	ffects for S	Set of Stor	es		Fixed Effects for Markets			
	(1)	(2)	(3)		(1) (2) (3)			
	$p_{BL,t}$	$p_{ML,t}$	$p_{CL,t}$		$p_{BL,t}$	$p_{ML,t}$	$p_{CL,t}$	
$p_{BL,t-1}$	0.287	0.036	0.020	-	0.646	0.097	0.091	
	(0.027)	(0.013)	(0.013)		(0.025)	(0.015)	(0.012)	
$p_{ML,t-1}$	0.045	0.322	0.010		0.074	0.601	0.066	
	(0.009)	(0.027)	(0.012)		(0.015)	(0.027)	(0.014)	
$p_{CL,t-1}$	-0.023	-0.049	0.267		0.100	0.097	0.682	
	(0.013)	(0.020)	(0.039)		(0.010)	(0.016)	(0.025)	
Observations	37,449	37,431	37,442		13,972	13,973	13,975	
R-squared	0.939	0.941	0.942		0.974	0.971	0.974	
Mean Price	9.79	9.67	9.68		10.08	9.95	9.94	
SD residuals	0.337	0.342	0.336		0.210	0.229	0.216	

Table 2.6: AR(1) Price Regressions Using Flagship Market-Pack Size-Week or - Month Data

Notes: regressions also include time period*pack size interactions and use pack sizes containing volumes equivalent to 6, 12, 18, 24 and 30 12 oz. containers. Market or store fixed effects described in the label to each panel. Standard errors, clustered on the market, are in parentheses. The SD residuals statistic is the standard deviation of the residuals from the regression.



Figure 2.5: Estimated Pre-JV Price Dynamics and the Combined Market Shares of AB, Miller and Coors.

Notes: The estimated univariate regression coefficients, with standard errors in parentheses, for panel (a) are BL: 0.011 (0.226) + $0.558C_3$ (0.288), $R^2 = 0.080$; ML : 0.044 (0.192) + 0.465 C_3 (0.245), $R^2 = 0.077$; CL : -0.025 (0.215) + 0.568 C_3 (0.278), $R^2 = 0.091$; and for panel (b): -0.039 (0.046) + 0.120 C_3 (0.058), $R^2 = 0.088$.

and Coors in 2007 (i.e., the C_3). Figure 2.5(b) shows a similar plot for the average of the six cross-brand coefficients. In both cases there is a positive, and, using a regression analysis, a statistically significant, relationship, consistent with our simulations.³⁴

	(1)	(2)	(3)	(4)	(5)	(6)
Data Frequency	Week	Week	Week	Week	Week	Month
Sizes	All	12 only	All	All	All	All
Price Reductions	Excl.	Excl.	Incl.	Excl.	Excl.	Excl.
Mean Own Price Elasticity	-3	-3	-3	-2.5	-3.5	-3
Mean Flagship Diversion	85%	85%	85%	90%	80%	85%
Lower Bound Cost for BL	\$5.259	\$5.278	\$4.845	\$4.248	\$5.973	\$4.616
(c_{BL})	(0.201)	(0.048)	(0.046)	(0.043)	(0.026)	(0.127)
L.B. Cost for ML/CL	\$6.425	\$6.528	\$5.984	\$5.786	\$6.874	\$5.711
$(c_{ML/CL})$	(0.020)	(0.014)	(0.022)	(0.024)	(0.017)	(0.020)
Width Cost Interval	\$0.625	\$0.752	\$1.246	\$0.556	\$0.672	\$1.793
$(\overline{c_i} - c_i)$	(0.029)	(0.021)	(0.018)	(0.102)	(0.026)	(0.037)
Cost AR(1) Parameter	1.156	0.939	0.850	1.222	0.959	0.742
(ρ)	(0.020)	(0.011)	(0.026)	(0.013)	(0.012)	(0.025)
SD Cost Innovations	\$0.282	\$0.278	\$0.566	\$0.260	\$0.270	\$0.400
(σ_c)	(0.024)	(0.001)	(0.050)	(0.104)	(0.026)	(0.052)

Table 2.7: Parameter Estimates for Six Specifications

Notes: BL = Bud Light, ML = Miller Lite and CL=Coors Light. Standard errors in parentheses. The data specifications using weekly data include group-of-store fixed effects when calculating the data statistics. For the monthly specification, the regression using the data only include market fixed effects

2.4.3.5 Parameter Estimates and Model Fit.

Table 2.7 reports estimates from six specifications, using different data or alternative demand parameters. Estimated marginal costs increase when demand is more elastic, and the range of costs and the standard deviation of the innovations increase when we try to match data that contains temporary price reductions. The estimated marginal cost ranges are much larger than in our examples, but the estimated $\sigma_c s$ imply that the probability that a marginal cost can go from high to

³⁴We also find positive, statistically significant relationships when we look at individual crossbrand coefficients.

low across periods is quite high.³⁵ As we will note below, the volatility of observed prices means that the marginal costs implied by CI Nash or conduct parameter models are also quite volatile.

The upper panel of Table 2.8 reports the fit of the moments that we match during estimation for the column (1), (2) and (3) specifications. The lower part of the table reports moments that are not matched, including the skewness of the innovations from the AR(1) regression. The model systematically underpredicts the standard deviation of price residuals for BL. The other moments are matched quite accurately, except that we cannot match the skewness of the residuals when price promotions are included in the data, consistent with our model having no mechanism to match these types of changes.

2.4.3.6 Predicted Effects of the JV.

Table 2.9 reports predicted prices when we resolve the six models assuming that ML and CL have the same marginal cost and are sold by a single firm at the same price. We assume that MC benefits from a synergy that would have prevented average prices from rising if firms set static CI Nash prices, as this seems consistent with the DOJ's expectation, but the width of the cost interval and the remaining parameters remain the same. The predicted price changes in columns (1)-(5) are all within the estimated 40¢-\$1 or 3-6% ranges.³⁶ We cannot find an equilibrium for the monthly data specification. In this case, the estimated parameters imply marginal costs are more persistent (the probability that a firm with the cost $\overline{c_i}$ will have a cost less than $\frac{c_i+c_i+c'}{2}$ is only 0.067) because, in this case, we are matching coefficients from a regression that does not control for cross-store heterogeneity in retail prices, and signaling incentives raise prices so

³⁵For example, for the specification in column (1) the probability that a firm with marginal cost $\overline{c_i}$ will have a marginal cost in the lower half of the range in the next period is 0.24, similar to 0.32 in our baseline example.

³⁶One might be concerned that our assumed discount factor of $\beta = 0.99$ is too low for weekly data. We have recomputed the column (1) estimates assuming $\beta = 0.998$, implying an annual discount factor of around 0.9. While a higher discount factor increases signaling incentives, the estimated parameters change to rationalize pre-JV dynamics in such a way that the predicted post-JV prices are within 1 cent of those reported in Table 2.9.

toto to dimonstration						
	(1)		(2)		(3)	
Frequency	Week		Week		Week	
Sizes	All		12		All	
Price Reductions	Excl.		Excl.		Incl.	
	Data	Model	Data	Model	Data	Model
		Matched	Moments			
$\operatorname{Mean} p_{BL}$	\$10.09	\$10.09	\$10.30	\$10.30	\$9.81	\$9.81
Mean p_{ML}	\$9.96	\$9.96	\$10.22	\$10.22	\$9.68	\$9.68
Mean $\rho^{ML,ML}$, $\rho^{CL,CL}$	0.402,0.413	0.408	0.468,0.450	0.444	0.330,0.290	0.313
Mean $\rho^{BL,ML}, \rho^{BL,CL}$	0.082,0.066	0.074	0.102, 0.056	0.076	0.070,0.060	0.059
Mean $\rho^{ML,CL}, \rho^{CL,ML}$	0.051, 0.036	0.033	0.065,0.026	0.035	0.049,-0.004	0.028
IQR p_{BL}	\$0.189	\$0.189	\$0.185	\$0.212	\$0.314	\$0.313
		•				
		<u>Unmatche</u>	<u>id Moments</u>			
Mean p_{CL}	\$9.95	\$9.96	\$10.20	\$10.23	\$9.68	\$9.68
$ ho^{BL,BL}$	0.444	0.385	0.442	0.418	0.311	0.296
Mean $\rho^{ML,BL}, \rho^{CL,BL}$	0.059, 0.0.42	0.038	0.065,0.040	0.038	0.076,0.004	0.029
SD of BL Res.	\$0.177	\$0.109	\$0.136	\$0.122	\$0.317	\$0.188
SD of ML/CL Res.	\$0.204,\$0.189	\$0.159	\$0.161,\$0.149	\$0.179	0.322, 0.311	\$0.271
${ m IQR}\ p_{ML}, p_{CL}$	\$0.222,\$0.210	\$0.281	\$0.228,\$0.206	\$0.316	0.335, 0.316	\$0.462
Skewness of BL Res.	-0.361	-0.337	-0.307	-0.314	-0.806	-0.098
ML/CL Res.	-0.100,-0.329	-0.331	-0.296,-0.201	-0.297	-0.717,-0.696	-0.080
Notes: BL = Bud Light, ML	. = Miller Lite and	CL=Coors	Light. $SD = stands$	urd deviatio	on. Res. = residual	s from the
AR(1) regressions. For the	data we report sep	arate value	es for the statistics	for ML and	d CL, but, because	the model
assumes that ML and CL a	re symmetric, and	so predicts	identical statistics	(ignoring	simulation error),	we match
the average of these values	during estimation	and report	a single predictior	į.		

Table 2.8: Model Fit for Three Specifications Using Weekly Data, Average Price Elasticity of -3

A						
	(1)	(2)	(3)	(4)	(5)	(6)
Frequency	Week	Week	Week	Week	Week	Month
Sizes	All	12 only	All	All	All	All
Price Reductions	Excl.	Excl.	Incl.	Excl.	Excl.	Excl.
Average Elasticity	-3	-3	-3	-2.5	-3.5	-3
Flagship Diversion	85%	85%	85%	90%	80%	85%
Pre-JV Mean Prices						
BL	\$10.09	\$10.30	\$9.81	\$10.09	\$10.09	\$10.09
ML/CL	\$9.96	\$10.22	\$9.68	\$9.96	\$9.96	\$9.95
Assumed ML/CL Synergy	-\$1.18	-\$1.20	-\$1.14	-\$1.50	-\$0.94	-\$1.17
Post-JV Mean Prices						
BL	\$10.62	\$10.90	\$10.17	\$10.98	\$10.42	fails
	(+5.3%)	(+5.7%)	(+3.7%)	(+8.7%)	(+3.3%)	
ML/CL	\$10.48	\$10.79	\$10.02	\$10.82	\$10.27	fails
	(+5.2%)	(+5.8%)	(+3.5%)	(+8.5%)	(+3.1%)	

Table 2.9: Predicted Average Prices Before and After the MC JV

Notes: BL = Bud Light, ML = Miller Lite and CL=Coors Light. For the data we report separate values for the statistics for ML and CL, but, because the model assumes that ML and CL are symmetric, and so predicts identical statistics (ignoring simulation error), we report a single prediction.

high that the conditions for separation fail.

Figure 2.6 compares, using the column (1) parameters, BL's equilibrium pricing strategies for the static Bayesian Nash 3-firm model, the estimated signaling 3-firm model and the counterfactual post-JV model. Signaling increases the level and the range of BL prices, which span from the lowest point on the two BL pricing functions to the highest point, especially in the counterfactual.

Table 2.10 compares the cross-market averages of the price dynamic statistics before and after the JV in the data, and the values predicted by the column (1) model. The model correctly predicts the directional change in each statistic except the skewness measures, even if it does not predict which statistics change the most. We view our ability to match qualitative changes in dynamics, as well as the increase in average price levels, even though our model is calibrated using only pre-JV data, as an encouraging result. Figure 2.6: Bud Light Equilibrium Pricing Strategies (for estimates in column (1) of Table 2.7).



Notes: the strategies shown assume that $c_{t-1}^{BL} = \underline{c}^{BL}$ and $c_{t-1}^{ML} = c_{t-1}^{CL} = \underline{c}^{ML/CL}$ (lower line) and $c_{t-1}^{BL} = \overline{c}^{BL}$ and $c_{t-1}^{ML} = c_{t-1}^{CL} = \overline{c}^{ML/CL}$ (upper line). Therefore, for each type of equilibrium, the maximum range of BL's prices spans from the lowest point on the bottom line to the highest point on the upper line.

2.4.4 Testing Alternative Explanations for the Post-JV Price Increases.

Some people have suggested that, even if our model can explain why prices rose after the JV, MW and MSW's CI theories of tacit collusion provide pre-existing and satisfactory explanations. While folk theorems imply that a CI tacit collusion model that fits the data almost perfectly is likely to exist, we can test how well MW and MSW's assumptions fit the data. MSW's supermarkup model of collusion is clearly rejected and, in some specifications, MW's baseline interpretation that there was CI Nash pricing before the JV is also rejected. The estimates also imply that marginal costs are serially correlated and quite volatile, a feature that plays an important role in our model.

Our tests extend MW's conduct parameter framework. The framework as-

		Data		F	'itted Mod	el
	Pre-JV	Post-JV	Change	Pre-JV	Post-JV	Change
IQR of Prices						
BL	\$0.189	\$0.241	+0.052	\$0.189	\$0.353	+0.164
ML	\$0.222	\$0.256	+0.034	\$0.281	\$0.350	+0.069
CL	\$0.210	\$0.244	+0.034	\$0.281	\$0.350	+0.069
AR(1) Regressi	on Coeffic	cients				
$ ho^{BL,BL}$	0.444	0.524	+0.080	0.385	0.415	+0.030
$ ho^{ML,ML}$	0.402	0.483	+0.081	0.408	0.409	+0.001
$ ho^{CL,CL}$	0.413	0.453	+0.040	0.408	0.409	+0.001
$ ho^{BL,ML}$	0.082	0.092	+0.010	0.074	0.122	+0.048
$ ho^{BL,CL}$	0.066	0.095	+0.029	0.074	0.122	+0.048
$ ho^{ML,BL}$	0.059	0.087	+0.028	0.038	0.141	+0.103
$ ho^{CL,BL}$	0.042	0.080	+0.038	0.038	0.141	+0.103
Std. Dev. of AR	(1) regres	ssion resid	luals			
BL regression	\$0.177	\$0.188	+0.011	\$0.109	\$0.203	+0.094
ML regression	\$0.204	\$0.204	+0.000	\$0.159	\$0.204	+0.045
CL regression	\$0.189	\$0.193	+0.004	\$0.159	\$0.204	+0.045
Skewness of AF	R(1) regre	ssion resid	duals			
BL regression	-0.361	-0.181	+0.180	-0.337	-0.504	-0.167
ML regression	-0.100	0.001	+0.101	-0.331	-0.470	-0.139
CL regression	-0.329	-0.104	+0.225	-0.331	-0.470	-0.139

Table 2.10: Observed and Predicted Changes in Price Dynamics

Notes: BL = Bud Light, ML = Miller Lite and CL=Coors Light. The calculation of the statistics is explained in Section 2.4.3.4. Pre-JV averages are calculated for 45 markets, and post-JV averages are calculated for 44 markets, as one market does not have at least 5 stores observed in consecutive weeks after the JV.

sumes that pricing is characterized by stacked static, CI first-order conditions

$$\left(\Omega_{mt} \circ \left[\frac{\partial q_{mt}(p_{mt}, \theta^D)}{\partial p_{mt}}\right]\right) (p_{mt} - c_{mt}) + q_{mt}(p_{mt}, \theta^D) = 0,$$

where p_{mt} , q_{mt} and c_{mt} are vectors of prices, quantities and (constant) marginal costs and $\frac{\partial q_{mt}(p_{mt},\theta^D)}{\partial p_{mt}}$ is a matrix of demand derivatives.

 Ω_{mt} is the "conduct" matrix, with (row *i*, column *j*) element $\Omega_{i,j}$. $\Omega_{i,j} = 1$ if products *i* and *j* are owned by the same firm. Under static Nash pricing, all other elements of Ω_{mt} are zero. MW's baseline specification assumes static Nash pricing before the JV, but allows $\Omega_{i,j} = \kappa$ after the JV if *i* and *j* are owned by different domestic brewers. $\kappa = 1$ is consistent with joint profit-maximization, while $0 < \kappa < 1$ could be interpreted as reflecting partial internalization of pricing

externalities.

Given demand estimates, MW estimate the post-JV κ using equations

$$p_{mt} = W_{mt}\gamma - \left(\Omega_{mt}(\kappa) \circ \left[\frac{\partial s_{mt}(p_{mt}, \theta^D)}{\partial p_{mt}}\right]\right)^{-1} s_{mt}(p_{mt}) + \nu_{mt}.$$
 (2.5)

where $c_{imt} = W_{imt}\gamma + \nu_{imt}$ and W includes time, product (brand-size) and geographic market fixed effects; a "distance measure" that multiplies distance to the brewery or port with real diesel prices; and, a dummy for MC products after the JV to allow for an additional efficiency. The JV is assumed not to affect AB's marginal costs. The instruments are the variables in W and a dummy for domestic products after the JV. The post-JV κ is exactly identified by how much more AB's prices increase than the increase that can be rationalized as a static best response.

MW's single exclusion restriction implies that they cannot estimate separate pre- and post-JV κ s or test whether a change in conduct is the source of the price increase.³⁷ We provide this type of test by adding additional instruments and controls.³⁸ Note, however, that we will only use the model to test MW and MSW's assumptions and we will not interpret positive κ s as evidence of collusion. As shown by Corts (1999), some forms of tacit collusion may be consistent with estimates of κ that are less than or equal to zero, and, as we discuss below, our signaling model tends to imply positive estimates of κ even though there is no collusion.

Our specifications include separate pre- and post-JV product and market fixed effects in W. To understand our choice of instruments, consider the first-order condition for product i owned by AB

$$p_{imt} = W_{imt}\gamma + \frac{q_{imt}}{\frac{\partial q_{imt}}{\partial p_{imt}}} + \sum_{\substack{j \in AB \\ j \neq i}} \frac{\frac{\partial q_{jmt}}{\partial p_{imt}}}{\frac{\partial q_{imt}}{\partial p_{imt}}} (p_{jmt} - c_{jmt}) + \kappa \sum_{k \in M,C} \frac{\frac{\partial q_{kmt}}{\partial p_{imt}}}{\frac{\partial q_{imt}}{\partial p_{imt}}} (p_{kmt} - c_{kmt}) + \nu_{imt}.$$

³⁷MW re-estimate the post-JV κ assuming, but not estimating, different pre-JV $\kappa \leq 0.5$. These estimates imply that κ rose after the JV, although by smaller amounts as the assumed pre-JV κ rises, as a pre-JV κ also implies that AB would increase its prices when MC benefits from an efficiency.

³⁸We continue to assume that imported brands use Nash pricing and that $\Omega_{i,j} = 1$ when *i* and *j* have the same owner.

Valid instruments will be correlated with $\sum_{k \in M,C} \frac{\frac{\partial q_{kmt}}{\partial p_{imt}}}{\frac{\partial q_{imt}}{\partial p_{imt}}} (p_{kmt} - c_{kmt})$ (i.e., the incremental effect of a change in *i*'s price on a rival's profits), and uncorrelated with the cost unobservable ν_{imt} .

The first six columns in Table 2.11 report conduct coefficients for the columns (1)-(3) demand specifications in Table 2.5.³⁹ Columns (1)-(3) use the distance measures of rivals as instruments, as they affect rivals' margins, and, as MW already assume that a product's own distance measure is uncorrelated with ν_{imt} , the additional assumptions required are minimal.⁴⁰ Columns (4)-(9) use additional instruments in the form of the average value of the demand unobservables (ξ s) for rival brewers over either the pre- or post-JV period, and the interactions of these instruments with the distance instruments.⁴¹ These additional instruments are valid if ν_{imt} is uncorrelated with the demand unobservables of rivals' products. This is a stronger assumption, although economists sometimes assume that a product's own demand and marginal costs unobservables are uncorrelated in order to estimate demand (MacKay and Miller (2019)). Columns (7)-(9) include linear domestic-market-fiscal year fixed effects in *W*. These controls allow for possible correlations between local preferences and costs for domestic products as a group, and cause conduct to be identified only from within-market-year cross-

³⁹The specifications in columns (1)-(9) contain time period fixed effects, and separate product and market fixed effects for before and after the JV, as well as the distance measure interacted with combinations of dummies for domestic products and periods after the JV. The specification in column (10) is estimated separately for each fiscal year (e.g., the FY06 year runs October 2005-September 2006), and the specification includes product, city and quarter fixed effects, the distance measure (interacted with a dummy for domestic products) as well as non-linear market fixed effects for the domestic products. We have also estimated specifications using the two nest nested logit models, and specifications that estimate κ s based only on the pricing of the flagship brands. These estimates lead us to reject Nash pricing behavior before the JV, and the pre- and post-JV parameters are closer than those in columns (1)-(6).

⁴⁰There are eight excluded distance instruments. For AB products in market m and time t before the JV, the (m; t) distance measure for Miller and the (m; t) distance measure for Coors are instruments. For pre-JV Miller products, the distance measures for AB and Coors are instruments. For pre-JV Coors products, the distance measures for Miller and AB are instruments. For AB (MC) products in market m and time t after the JV, the (m; t) distance measure for MC (AB) is the instrument.

⁴¹Specifically, we calculate the average value of the demand residuals for the products sold by brewer b in market m either before or after the JV, and then construct eight instruments in the same way that we construct the instruments for distance. We average across periods because the demand unobservables are more variable than the distance measures.
		Te	sts of MW C	onduct Mo	del		Tests of	MW & MSV	<u>N Models</u>	Test of MSW
	(1)	(2)	(3)	(4)	(5)	(9)	(2)	(8)	(6)	(10)
Demand Model	NL	RCNL	RCNL	NL	RCNL	RCNL	NL	RCNL	RCNL	RCNL
(all one nest)	Monthly	Monthly	Quarterly	Monthly	Monthly	Quarterly	Monthly	Monthly	Quarterly	Quarterly
Domestic Firms										FY06, FY07
Pre-JV Conduct	0.274	0.322	0.198	0.340	0.263	0.238	0.958	0.909	0.913	0.977, 0.924
	(0.098)	(0.193)	(0.221)	(0.057)	(0.166)	(0.147)	(0.005)	(0.016)	(0.024)	(0.007),(0.013)
										FYIU, FYII
Post-JV Conduct	0.723	0.651	0.573	0.688	0.767	0.717	0.951	0.914	0.921	0.976, 0.933
	(0.124)	(0.146)	(171.0)	(0.144)	(050.0)	(200.0)	(600.0)	(0.014)	(0.013)	(0.011),(0.015)
p-value diff.	0.004	0.013	0.017	0.000	0.000	0.001	0.483	0.638	0.558	·
Supermarkup Controls			·				Dome	stic*Marke linear	t*Fiscal Yeaı	· Fixed Effects non-linear
Excluded IVs <u>ML 12 Packs</u>	Dome	stic Rival D	istance	Dom. R Rival {	ival Distan s and Inter	ce, Dom. actions	Do	m. Rival D and	istance, Dor Interaction	n. Rival ξs }
Pre-JV: Mean $\widehat{c_{imt}}$	\$2.37	\$5.93	\$6.75	\$1.99	\$6.13	\$6.65	-\$4.61	\$1.73	\$2.79	-\$1.41
Residual ρ	0.414	0.427	0.451	0.410	0.430	0.449	0.239	0.252	0.114	0.030
SD AR(1) res.	\$0.31	\$0.27	\$0.20	\$0.31	\$0.27	\$0.20	\$0.63	\$0.44	\$0.30	\$0.35
Post-JV: Mean $\widetilde{c_{imt}}$	-\$1.34	\$4.20	\$5.40	-\$1.00	\$3.34	\$4.62	-\$4.39	\$1.81	\$2.93	-\$5.74
Residual $ ho$	0.431	0.488	0.403	0.433	0.485	0.417	0.224	0.436	0.050	0.006
SD AR(1) res.	\$0.43	\$0.33	\$0.26	\$0.41	\$0.37	\$0.28	\$0.56	\$0.42	\$0.29	\$0.43
Observations	94,656	94,656	31,777	94,656	94,656	31,777	94,656	94,656	31,777	31,777 total
Notes: specifications est lagged marginal costs (estimated separately be	imated usir imt-1) from ore and aft	12 2-step GN 1 a regression ter the JV. T	IM. Conduct n of ML 12-p. he "SD AR(1	parameters ack margina) res." stati	are reportec l costs on th stics are the	l for four fisc: eir lagged va standard dev	al years. The lues, market viation of the	e "residual f t and time fi e residuals f	" statistics ar xed effects. T from these re	e the coefficients or hese regressions are gressions. Standard

er Ero rnative Models Heing a Generalized Conduct Da Tahle 2 11. Tecting Alter brewer/-product variation. We will also use these specifications to test the MSW model.

We reject Nash pricing after the JV in all nine specifications. This is, of course, consistent with MW's interpretation that there was collusion after the JV. All of the estimated pre-JV κ s are positive, and some are significant. The estimates in columns (1)-(6) are consistent with an increase in κ after the JV, but the estimates with market-year controls suggest that conduct did not change, even though the κ estimates are very precise.

The plausibility of these CI pricing models can also be assessed by looking at what they imply for marginal costs and synergies. Table 2.11 reports average implied marginal costs for ML 12-packs. Less elastic demand and higher κ imply lower marginal costs, and the (1), (4) and (7)-(9) costs are implausibly/impossibly low. The remaining columns imply larger synergies for ML, which was being shipped the same distances before and after the JV in most markets, even than the 17.5% synergy for ML and CL that we assumed for the column (1) specification of our model. Controlling for market and time effects, the implied ν_{imt} s are also serially correlated and quite volatile.⁴² While cost volatility is certainly not inconsistent with CI, we view volatility as suggesting that a collusive interpretation of the data requires a very strong CI assumption: if CI is not satisfied, then, given that prices are volatile, collusion would be hampered by the difficulty of distinguishing cheating from a conforming price set by a low marginal cost firm.

The conduct model is not a fully-specified model of collusion because it does not specify why firms choose not to cheat. Some collusion models cannot be tested using the conduct framework, but the MSW supermarkup model can. MSW assume that, every fiscal year, both before and after the JV, a price leader suggests a "supermarkup" on top of Bertrand Nash prices that domestic brewers should charge. If a domestic firm fails to charge the supermarkup, a punishment phase

⁴²The rich fixed effects in columns (7)-(9) cause the ν_{imt} s to jump across fiscal years, so the estimated serial correlation falls.

ensues, but in a CI subgame perfect equilibrium, the suggested supermarkup will satisfy the incentive-compatibility constraints (ICCs). Prices may increase after a merger if the ICCs are relaxed. We can test this model by using an appropriately defined domestic product market-fiscal year fixed effect to control for the supermarkup. If the "supermarkup on Nash" theory is correct, estimates of conduct κ parameters should be equal to zero once the fixed effects are included.

The columns (7)-(9) include linear domestic-market-fiscal year fixed effects. These specifications are not quite consistent with the MSW's exact theory⁴³, but they are simple to estimate. As already discussed, we can reject $\kappa = 0$ before or after the JV at any significance level. Column (10) tests MSW's exact model by allowing for non-linear domestic-market-fiscal year fixed effects (see Appendix B.5 for details) using the quarterly RCNL model (most favorable to pre-JV Nash pricing in columns (3) and (6)). We estimate the model separately for each fiscal year to reduce the number of coefficients estimated simultaneously. Consistent with column (9), the reported conduct parameters are precisely estimated and are between 0.9 and 1, and, because estimated supermarkups are also positive, most of the implied marginal costs are negative. Therefore, we can clearly reject the MSW formulation of CI collusion, although, as we have emphasized, this does not imply that all models of collusion would be rejected.

While our model implies that the conduct parameter framework is misspecified, because it does not control for beliefs or signaling incentives, we have estimated conduct parameter models using data simulated from two and three-firm versions of our model with cross-firm heterogeneity. The estimated conduct parameters are typically between 0.3 and 1, and the implied marginal costs are usually significantly below their true levels.⁴⁴ The estimated conduct parameters can rise, fall or stay roughly unchanged after a merger. The results of our conduct anal-

⁴³Linear fixed effects would be consistent with a model where the leader suggested domestic firms set Nash prices "as if" all of their marginal costs had been raised by a common fixed amount, rather than suggesting a common dollar per 12-pack equivalent price addition to Nash prices.

⁴⁴If, conditional on controls for costs, firms tend to set higher prices when other firms have higher margins or there is more diversion to those rivals' products, then estimated κ s will be positive. As discussed previously, these features also tend to lead to stronger signaling effects in our model, so we tend to estimate positive κ s using simulated data from our model.

ysis are therefore not inconsistent with what one would expect given our model.

2.5 Conclusion

We have developed a model where oligopolists simultaneously use prices to signal private information that is relevant for their future pricing decisions. Although the possibility that this type of behavior would raise equilibrium prices was identified in the theoretical literature over thirty years ago, we provide the first attempt to quantify the magnitude of these effects, both in examples and in an empirical application. We find that effects can be large, and that they can explain changes in price levels and price dynamics after a large horizontal transaction in the U.S. beer industry. While CI theories of tacit collusion can also explain an increase in price levels, our model provides a natural explanation of the periodto-period price changes observed in this data, and in data from other industries where tacit collusion has been suggested (Ordover (2007)). It is also consistent with how firms treat margin information as highly confidential.

We have often been asked how our model and our empirical analysis relate to theories of "coordinated effects" in merger analysis. There is no standard definition of coordinated effects: the presentation in Ordover (2007) is focused on variants of tacit collusion models, but Baker and Farrell (ming) and Farrell and Baker (2021) use a much broader definition which includes both "purposive" theories of collusion and "non-purposive" theories, a group which includes the non-collusive Markov Perfect theories of Maskin and Tirole (1988). Our model lies within this group. Non-purposive theories are valuable partly because they can explain why it may not be appropriate to rely exclusively on static CI unilateral effects models in industries that do not have the characteristics that economists typically believe favor tacit collusion (Stigler (1964)) or where, before a transaction, prices do not display the rigidity that collusive theories often predict (Athey et al. (2004)). They can also explain why coordinated effects do not raise prices to joint-profit maximizing levels, an outcome that a tacit collusion model will predict if prices are set frequently and firms are patient. However, we also believe that combining tacit collusion and asymmetric information is likely to be a profitable direction for future research, building on the work of Kreps et al. (1982) and Athey and Bagwell (2008) who have examined the links in very stylized theoretical models.⁴⁵ In fact, one of our examples in Appendix B.2.2 illustrates how signaling could exacerbate the impact of small coordination incentives.

One could also ask what our model adds to existing non-purposive theories. Maskin and Tirole (1988) provide examples of price-setting games which lead to both price rigidity and price volatility without any underlying volatility in costs or asymmetries of information. We view our introduction of serially correlated asymmetric information as not only realistic, but, also potentially helpful in solving more complicated models, without assuming price changes are asynchronous or subject to potentially large menu costs (Maskin and Tirole (1988), Nakamura and Zerom (2010)), because it means that firms choose prices against a perceived continuous distribution of rivals' prices. This feature of asymmetric information models has long been appreciated in both the static and dynamic discrete choice games literatures (e.g., Seim (2006)), but there are also benefits when choices are continuous.

⁴⁵Athey and Bagwell (2008) consider an example that is explicitly connected to the Mailath (1989) model.

Appendix A: Appendix for Chapter 1

A.1 Supplement to Market and Data

A.1.1 Market

Figure A.1 shows two examples of truck tractors. Figure A.2 shows an example of the webpages of RBA pre-merger and Post-merger.

A.1.2 Cross-sectional Variation

Table A.1 shows the change in general quality of trucks from TX transaction data and nationwide auction transaction data. Figure A.3 summarizes the price distribution of trucks on different platforms. Figure A.4 shows the price distribution of trucks purchased by different types of buyers. Figure A.5 presents the variation in the quality and transaction volume across different months and states.

Figure A.1: Pictures of Truck Tractors





Figure A.2: Web Pages of RBA in 2016 and 2018

Table A.1: Characteristics of Trucks in the TX Transaction Data and Auction Transaction Data

	2016		2018		
	TX, All	Nationwide, Auctions	TX, All	Nationwide, Auctions	
age	9.1841	9.1056	9.6960	9.6490	
log(mile)	6.0312	5.7716	6.1159	6.0445	

Notes: unit of mileage is 1,000 miles.



Figure A.3: Price Distribution Across Platforms

Notes: 1. unit of Price is \$1,000; 2. Quality is round to the nearest integers.





Notes: quality is round to the nearest integers.

Figure A.5: Variation in the Number of Trucks Across Markets



Notes: markets are sorted according to the market size pre-merger.

A.2 Supplement to Model

This section is used to supplement the model part. It has two subsections. (1) Markets with at most two different quality levels. First, in a one platform setting, I prove some properties of the equilibrium payoff functions in the auction entry stage. Then, based on these properties, I discuss the economics of more searches, which gives proof for Proposition 1. Based on similar intuition, I discuss the economics of more multihoming buyers in simple settings and prove Proposition 2. (2) I extend the markets to include more complicated structures of available trucks and provide more evidence about the economics of different search choices. It includes a simulation showing the pattern of models with two types of buyers.

A.2.1 Markets with at Most Two Quality Levels

Assume there are two possible quality levels $q^H > q^L$ in the market. I use $N^{H,A}, N^{H,B}, N^{L,A}, N^{L,B}$ to represent the number of trucks on different platforms with different quality levels, where $N^A = N^{H,A} + N^{L,A}, N^{H,B} + N^{L,B} = N^B, N^A + N^B = N$. Assume buyers can choose to search for one truck or two trucks on one platform or two platforms. The probability of buyers searching for two trucks is Pr^m and the probability of buyers choosing single-homing is ω . Additionally, for simplicity, I assume buyers' WTP for trucks with the same (q, v) is the same on different platforms, i.e., $V^A(q, v) = V^B(q, v) = V(q, v)$. All buyers are ex-ante symmetric. Assume the platform choice of single-homing buyers follows a rule related to the number of trucks on each platform $g(N^A, N^B)$.

A.2.1.1 One Platform

For now, assume there is only one platform with N goods. The equilibrium payoff function $U^*(q,v)$ is monotonically increasing in v, where $\frac{\partial U^*(q,v)}{\partial v} = \frac{\partial V(q,v)}{\partial v} [1 - Pr^{e^*}(q,v)]^{\gamma N-1} \ge 0$. Intuitively, entry choices of other competitors is

independent of buyer *i*'s private value, higher private value can increase buyer *i*'s WTP for the truck and her chance to win the truck. Therefore, higher private value can increase her expected payoffs from an auction for sure. Furthermore, since V(q, v) is convex in v under Assumption 1 and

$$\begin{split} &\frac{\partial Pr^{e^*}(q^H,v)}{\partial v} = \dots \\ &-\frac{1-Pr^{m^*}}{N}f(v) - Pr^{m^*}[\frac{\binom{1}{N^{H}-1}}{\binom{2}{N}}F(v) + \frac{\binom{1}{N^L}}{\binom{2}{N}}F(v'|U^*(q^H,v) = U^*(q^L,v'))]f(v) < 0, \\ &\frac{\partial Pr^{e^*}(q^L,v)}{\partial v} = \dots \\ &-\frac{1-Pr^{m^*}}{N}f(v) - Pr^{m^*}[\frac{\binom{1}{N^{L}-1}}{\binom{2}{N}}F(v) + \frac{\binom{1}{N^H}}{\binom{2}{N}}F(v'|U^*(q^L,v) = U^*(q^H,v'))]f(v) < 0, \end{split}$$

we have $U^*(q, v)$ is convex in v.

Given these properties, I can proof the following Lemma about the equilibrium expected payoff function in the auction entry stage.

Lemma 1. $\exists v^*$, we have $U^*(q^H, v) \ge U^*(q^L, v), \frac{\partial U^*(q^H, v)}{\partial v} \ge \frac{\partial U^*(q^L, v)}{\partial v}, \forall v \ge v^*.$

Proof. The proof includes two steps.

In the first step I prove that $\exists v^*$ such that $U^*(q^H, v^*) \geq U^*(q^L, v^*)$ and $\frac{\partial U^*(q^H, v^*)}{\partial v} \geq \frac{\partial U^*(q^L, v^*)}{\partial v}$.

The proof is by contradiction.

• Assume $U^*(q^H, v) < U^*(q^L, v), v \in [\underline{v}, \overline{v}].$

Under this assumption, a truck with low-quality level will be chosen if a buyer have the same private values for a high-quality truck and a low-quality truck, so $0 \leq Pr^{e^*}(q^H, v) < Pr^{e^*}(q^L, v), 0 > \frac{\partial Pr^{e^*}(q^H, v)}{\partial v} > \frac{\partial Pr^{e^*}(q^L, v)}{\partial v}, \forall v \in [\underline{v}, \overline{v}].$

Since under Assumption 1,

$$\begin{split} U^*(q,v) &= \dots \\ &- \int_{\underline{v}}^{v} (\gamma N - 1) [V(q,v) - V(q,\tilde{v})] [1 - Pr^{e^*}(q,\tilde{v})]^{\gamma N - 2} \frac{\partial Pr^{e^*}(q,\tilde{v})}{\partial v} d\tilde{v} + \dots \\ V(q,\underline{v}) \times [1 - Pr^{e^*}(q,\underline{v})]^{\gamma N - 1}, \\ V(q^H,v) - V(q^H,\tilde{v}) > V(q^L,v) - V(q^L,\tilde{v}), \forall v > \tilde{v}, \end{split}$$

we have $U^*(q^H, v) > U^*(q^L, v)$, which contradicts the assumption.

• Assume $\exists v', U^*(q^H, v') \geq U^*(q^L, v'), \exists v'', U^*(q^H, v'') < U^*(q^L, v'') \text{ and } \forall v \in [\underline{v}, \overline{v}], 0 < \frac{\partial U^*(q^H, v)}{\partial v} < \frac{\partial U^*(q^L, v)}{\partial v}$

Under this assumption, we have $U^*(q^H, \overline{v}) < U^*(q^L, \overline{v})$. Then, we have $Pr^{e^*}(q^H, \overline{v}) = Pr^{e^*}(q^L, \overline{v}) = 0, 0 > \frac{\partial Pr^{e^*}(q^H, \overline{v})}{\partial v} > \frac{\partial Pr^{e^*}(q^L, \overline{v})}{\partial v}.$

Since

$$\begin{split} \frac{\partial U^*(q,v)}{\partial v} &= \frac{\partial V(q,v)}{\partial v} [1 - Pr^{e^*}(q,v)]^{\gamma N-1},\\ \frac{\partial V(q^H,v)}{\partial v} &> \frac{\partial V(q^L,v)}{\partial v}, \end{split}$$

we have $\frac{\partial U^*(q^H, \overline{v})}{\partial v} > \frac{\partial U^*(q^L, \overline{v})}{\partial v}$, which contradicts the assumption.

• Assume $\frac{\partial U^*(q^H,v')}{\partial v} < \frac{\partial U^*(q^L,v')}{\partial v}$ whenever $U^*(q^H,v') \ge U^*(q^L,v')$ and $\frac{\partial U^*(q^H,v'')}{\partial v} > \frac{\partial U^*(q^L,v'')}{\partial v}$ for some $U^*(q^H,v'') < U^*(q^L,v'')$. If $U^*(q^H,\overline{v}) > U^*(q^L,\overline{v})$, we must have some v satisfies $\frac{\partial U^*(q^H,v)}{\partial v} \ge \frac{\partial U^*(q^L,v)}{\partial v}$ and $U^*(q^H,v) \ge U^*(q^L,v)$; otherwise, we cannot find $U^*(q^H,v) < U^*(q^L,v)$. If $U^*(q^H,\overline{v}) < U^*(q^L,\overline{v})$, according to the above, we must have $\frac{\partial U^*(q^H,\overline{v})}{\partial v} \ge \frac{\partial U^*(q^H,\overline{v})}{\partial v}$. If $U^*(q^L,\overline{v})$. The assumption implies $\frac{\partial U^*(q^H,v)}{\partial v} < \frac{\partial U^*(q^L,v)}{\partial v}$ and $U^*(q^H,\underline{v}) \ge U^*(q^L,\underline{v})$, which cannot be true when $\frac{\partial V(q,\underline{v})}{\partial v} = V(q,\underline{v})$.

In the second step, I prove $\forall v^*$ such that $U^*(q^H, v^*) \geq U^*(q^L, v^*)$ and $\frac{\partial U^*(q^H, v^*)}{\partial v} \geq \frac{\partial U^*(q^L, v^*)}{\partial v}$, we must have $U^*(q^H, v) \geq U^*(q^L, v), \frac{\partial U^*(q^H, v)}{\partial v} \geq \frac{\partial U^*(q^L, v)}{\partial v}, \forall v \geq 0$

If
$$U^*(q^H, v^*) \ge U^*(q^L, v^*)$$
, then $\frac{\partial Pr^{e^*}(q^H, v^*)}{\partial v} \le \frac{\partial Pr^{e^*}(q^L, v^*)}{\partial v} < 0$. Namely,
 $\exists \epsilon \to 0, Pr^{e^*}(q^H, v^* + \epsilon) - Pr^{e^*}(q^L, v^* + \epsilon) \le Pr^{e^*}(q^H, v^*) - Pr^{e^*}(q^L, v^*).$

On the other hand,

 v^* .

$$V(q^{H}, v^{*} + \epsilon) - V(q^{L}, v^{*} + \epsilon) \ge V(q^{H}, v^{*}) - V(q^{L}, v^{*}).$$

Therefore, we have $\frac{\partial U^*(q^H, v^* + \epsilon)}{\partial v} - \frac{\partial U^*(q^L, v^* + \epsilon)}{\partial v} \ge \frac{\partial U^*(q^H, v^*)}{\partial v} - \frac{\partial U^*(q^L, v^*)}{\partial v} \ge 0$ and $U^*(q^H, v^* + \epsilon) \ge U^*(q^L, v^* + \epsilon).$

We can iterate this process by using $v^* + \epsilon$ as the starting point. Therefore, we can show that $\forall v \in [v^*, \overline{v}]$, $U^*(q^H, v) \ge U^*(q^L, v)$ and $\frac{\partial U^*(q^H, v)}{\partial v} \ge \frac{\partial U^*(q^L, v)}{\partial v}$. \Box

Proof of Proposition 1

Proof. Now I will prove how the probability of searching two goods increases from Pr^{m^*} to $Pr^{m^{**}}$ will affect $Pr^{e^*}(q^H, v)$ and $Pr^{e^*}(q^L, v)$. The effect can be analyzed based on the following three equations

$$Pr^{e^{*}}(q^{H},v) - Pr^{e^{*}}(q^{L},v) = Pr^{m} \{ \int_{v}^{\overline{v}} [\frac{\binom{1}{N^{H}-1}}{\binom{2}{N}} - \frac{\binom{1}{N^{L}-1}}{\binom{2}{N}}] F(\tilde{v})f(\tilde{v})d\tilde{v} + \dots \\ \int_{v}^{\overline{v}} [\frac{\binom{1}{N^{L}}}{\binom{2}{N}} F(v'|U^{*}(q^{H},\tilde{v}) = U^{*}(q^{L},v')) - \frac{\binom{1}{N^{H}}}{\binom{2}{N}} F(v'|U^{*}(q^{L},\tilde{v}) = U^{*}(q^{H},v'))] f(\tilde{v})d\tilde{v} \}$$
(A.1)

$$\frac{\partial Pr^{e^*}(q^H, v)}{\partial v} - \frac{\partial Pr^{e^*}(q^L, v)}{\partial v} = -Pr^m \{ [\frac{\binom{1}{N^H - 1}}{\binom{2}{N}} - \frac{\binom{1}{N^L - 1}}{\binom{2}{N}}] F(v)f(v) + \dots \\ [\frac{\binom{1}{N^L}}{\binom{2}{N}}F(v'|U^*(q^H, v) = U^*(q^L, v')) - \frac{\binom{1}{N^H}}{\binom{2}{N}}F(v'|U^*(q^L, v) = U^*(q^H, v'))] f(v) \}$$
(A.2)

$$U^{*}(q^{H}, v) - U^{*}(q^{L}, v) = (\gamma N - 1) \{ \int_{\underline{v}}^{v} [V(q^{L}, v) - V(q^{L}, \tilde{v})] [1 - Pr^{e^{*}}(q^{L}, \tilde{v})]^{\gamma N - 2} \\ \frac{\partial Pr^{e^{*}}(q^{L}, \tilde{v})}{\partial v} - [V(q^{H}, v) - V(q^{H}, \tilde{v})] [1 - Pr^{e^{*}}(q^{H}, \tilde{v})]^{\gamma N - 2} \frac{\partial Pr^{e^{*}}(q^{H}, \tilde{v})}{\partial v} d\tilde{v} \} + \dots \\ V(q^{H}, \underline{v}) [1 - Pr^{e^{*}}(q^{H}, \underline{v})]^{\gamma N - 1} - V(q^{L}, \underline{v}) [1 - Pr^{e^{*}}(q^{L}, \underline{v})]^{\gamma N - 1}.$$
(A.3)

From (A.1)(A.2)(A.3), we see that Pr^m affects $Pr^{e^*}(q^H, v) - Pr^{e^*}(q^L, v)$ and $\frac{\partial Pr^{e^*}(q^H, v)}{\partial v} - \frac{\partial Pr^{e^*}(q^L, v)}{\partial v}$ through two channels: (1) the direct channel, when Pr^{m^*} increase to $Pr^{m^{**}}$, according to Lemma 1, it can attract more competitive buyers with private value above a threshold to the high-quality auctions while keeping the equilibrium expected payoffs the same. This effect will increase $Pr^{e^*}(q^H, v) - Pr^{e^*}(q^L, v)$ and decrease $\frac{\partial Pr^{e^*}(q^H, v)}{\partial v} - \frac{\partial Pr^{e^*}(q^L, v)}{\partial v}$ when v is above the threshold; (2) the indirect channel, lowering the expected payoffs because of the increased competition from (1). When Pr^m increases, according to (A.3), the difference between $U^*(q^H, v)$ and $U^*(q^L, v)$ becomes smaller.

The value of $\frac{\partial Pr^{e^*}(q^H,v) - Pr^{e^*}(q^L,v)}{\partial Pr^m}$ and $\frac{\partial \frac{\partial Pr^{e^*}(q^H,v)}{\partial v} - \frac{\partial Pr^{e^*}(q^L,v)}{\partial v}}{\partial Pr^m}$ can be calculated using the implicit function theorem and (A.1)(A.2)(A.3). In the new equilibrium, $\exists v^{**}$, for $v \in [v^{**}, \overline{v}]$, $\frac{\partial Pr^{e^*}(q^H,v) - Pr^{e^*}(q^L,v)}{\partial Pr^m} \ge 0$. Assume the this is not true. Since $Pr^{e^*}(q^H, \overline{v}) = Pr^{e^*}(q^L, \overline{v}) = 0$, we have $\frac{\partial \frac{\partial Pr^{e^*}(q^H,\overline{v}-\epsilon)}{\partial v} - \frac{\partial Pr^{e^*}(q^L,\overline{v}-\epsilon)}{\partial v}}{\partial Pr^m} > 0, \epsilon \to 0$. Therefore, $\frac{\partial U^*(q^H,\overline{v}-\epsilon) - U^*(q^L,\overline{v}-\epsilon)}{\partial Pr^m} > 0, \epsilon \to 0$, which contradicts the indirect effect mentioned above. Similarly, we must have $\frac{\partial \frac{\partial Pr^{e^*}(q^H,\overline{v}-\epsilon)}{\partial Pr^m} < 0, \epsilon \to 0$. In sum, $\exists v^{**}$, for $v \in [v^{**},\overline{v}]$, $Pr^{e^*}(q^H,v|Pr^{m^{**}}) - Pr^{e^*}(q^L,v|Pr^{m^{**}}) \ge Pr^{e^*}(q^H,v|Pr^{m^*}) - Pr^{e^*}(q^L,v|Pr^{m^*})$ when $Pr^{m^{**}} > Pr^{m^*}$.

The probability of the transaction price to be higher than p is equivalent to

the second highest WTP among bidders is higher than p = V(q, v).

$$\begin{split} &Prob(p > \tilde{p}|q, Pr^{m^*}) = Prob(v^{(n-1:n)} > V^{-1}(\tilde{p}, q)|Pr^{m^*}) = \dots \\ &1 - [1 - Pr^{e^*}(q, V^{-1}(p, q))]^{\gamma N - 1} \\ &Prob(p > \tilde{p}|q^H, Pr^{m^*}) - Prob(p > \tilde{p}|q^L, Pr^{m^*}) = \dots \\ &[1 - Pr^{e^*}(q^L, V^{-1}(\tilde{p}, q^L)|Pr^{m^*})]^{\gamma N - 1} - [1 - Pr^{e^*}(q^H, V^{-1}(\tilde{p}, q^L)|Pr^{m^*})]^{\gamma N - 1} + \dots \\ &[1 - Pr^{e^*}(q^H, V^{-1}(\tilde{p}, q^L)|Pr^{m^*})]^{\gamma N - 1} - [1 - Pr^{e^*}(q^H, V^{-1}(\tilde{p}, q^H)|Pr^{m^*})]^{\gamma N - 1} \end{split}$$

 $\begin{array}{l} \text{When } Pr^{e^*}(q^H, V^{-1}(\tilde{p}, q^L)) - Pr^{e^*}(q^L, V^{-1}(\tilde{p}, q^L)) \text{ and } Pr^{e^*}(q^H, V^{-1}(\tilde{p}, q^H)) - \\ Pr^{e^*}(q^H, V^{-1}(\tilde{p}, q^L)) \text{ increase, } Prob(p > \tilde{p}|q^H) - Prob(p > \tilde{p}|q^L) \text{ increases. Namely,} \\ \text{we have } Prob(p > \tilde{p}|q^H, Pr^{m^{**}}) - Prob(p > \tilde{p}|q^L, Pr^{m^{**}}) \geq Prob(p > \tilde{p}|q^H, Pr^{m^*}) - \\ Prob(p > \tilde{p}|q^L, Pr^{m^*}) \text{ when } \tilde{p} \text{ is above a threshold.} \end{array}$

A.2.1.2 Two Platforms

As mentioned in the main text, the effect of ω^* is smoothing the variation in quality and number of trucks across platforms in different markets. I discuss the effect of lower ω^* (share of single-homing) when there are two types of goods differentiated in their quality.

Proof of Proposition 2

Proof. Under the conditions in Proposition 2 and equation (1.1)(1.2) in the auction entry stage, I can get the following two equations:

$$U^{*}(q^{A}, v, A) = \dots$$

$$\int_{\underline{v}}^{v} [V(q^{A}, v) - V(q^{L}, \tilde{v})] d[1 - Pr^{e^{*}}(q^{A}, \tilde{v}, A)]^{\gamma N - 1} + V(q^{A}, \underline{v})[1 - Pr^{e^{*}}(q^{A}, \underline{v}, A)]^{\gamma N - 1}$$
(A.4)

$$Pr^{e^{*}}(q^{A}, v, A|q^{B} = q^{H}) - Pr^{e^{*}}(q^{A}, v, A|q^{B} = q^{L}) = (1 - \omega^{*})\frac{\binom{1}{N^{B}}}{\binom{2}{N}} \times \dots$$

$$\int_{v}^{\overline{v}} F(v'|U^{*}(q^{A}, \tilde{v}, A) = U^{*}(q^{H}, v', B)) - F(v'|U^{*}(q^{A}, \tilde{v}, A) = U^{*}(q^{L}, v', B))f(\tilde{v})d\tilde{v}.$$
(A.5)

Similar to the one-platform case proved in Lemma 1, there is a threshold of private value, when buyers' private value is above the threshold, we have $U^*(q^H, v', B)) \ge U^*(q^L, v', B))$. According to (A.5), when buyers conduct multihoming, they are less likely to choose the trucks on Platform A if the trucks on platform B have high quality and their private draws associate with those goods are not low. Similar to the indirect effect mentioned above, when ω decrease and $q^B = q^H$, the competition on platform B is fiercer, resulting in some buyers with moderate private values switch to platform A. Similar to Proposition 1, I can prove that $\exists v^{**}$, for $v \in [v^{**}, \bar{v}]$, $Pr^{e^*}(q^A, v, A|q^B = q^H, \omega^{**}) - Pr^{e^*}(q^A, v, A|q^B = q^L, \omega^{**}) < Pr^{e^*}(q^A, v, A|q^B = q^H, \omega^*) - Pr^{e^*}(q^A, v, A|q^B = q^L, \omega^{**}) < \frac{\partial Pr^{e^*}(q^A, v, A|q^B = q^H, \omega^{**})}{\partial v}$. Then since the final price is second highest WTP among buyers in an auction, we can get $Prob(p^A > \tilde{p}|\bar{q}^B = q^L, \omega^{**}) - Prob(p^A > \tilde{p}|\bar{q}^B = q^H, \omega^{**}) \geq Prob(p^A > \tilde{p}|\bar{q}^B = q^L, \omega^{**}) - Prob(p^A > \tilde{p}|\bar{q}^B = q^H, \omega^{**})$ when \tilde{p} is above a threshold.

A.2.2 Markets with More Than Two Quality Levels

For the markets with more complicated structure of available trucks, I do some simulations to present the similar findings about the economics of search choices. Considering the difficulty in proving the two-type model, I also illustrate the findings about two-type model by simulation.

A.2.2.1 One Type of Buyers

To be specific, I show a simulation result under the specification where the WTP follows log-normal distribution:

$$V(q_j, v_{ij}) = \begin{cases} exp(\theta q_j + v_{ij}) & \text{if } j \text{ is on } A\\ exp(\theta(q_j + \alpha) + v_{ij}) & \text{if } j \text{ is on } B \end{cases}$$

$$\theta = 0.8, v_{ij} \sim N(0, 0.5), \alpha = -0.5.$$

The set of realized trucks available on each platform is the same to the first realized market in the data. I consider two distributions of equilibrium searches:

- Case 1: $m^{\text{multi}} \sim Poisson(5), \omega = 0;$
- Case 2: $m^{\text{multi}} \sim Poisson(10), \omega = 0.$

Figure A.6 shows that the expected payoffs from entering an auction U is complementary in (q, v). When all buyers are more likely to search for more trucks in equilibrium, the expected payoffs from entering the popular auctions, i.e., the high-quality auctions on platform B, will decrease significantly.

Figure A.7 shows the entry behavior of buyers who search five trucks across platforms on average and buyers who search ten trucks across platforms. I do 100 simulations and calculate the average v when they choose trucks with different q in the two equilibria. We see that under this specification, the buyers who choose to enter auctions with high-quality trucks have higher v on average when all the buyers search for more trucks. Therefore, the trucks with high quality are more likely to be transacted with high price when all buyers search more trucks.

A.2.2.2 Two Types of Buyers

Buyers with different quality preferences have different expected payoffs from searching, resulting in different equilibrium search choices even if they draw Figure A.6: Expected Payoffs from Entering an Auction $U^{\ast}(q,v)$ (One Type of Buyers)



Notes: unit of expected payoffs \$1,000.





from the same distribution of search costs. Because of the higher WTP for highquality trucks, the expected gain from more searches among high-type buyers is higher than that among low-type buyers. On the other hand, when both types of buyers search for more trucks, low-type buyers are more likely to lose in auctions with high-quality trucks since there are more high-type buyers in those auctions. Notably, some low-type buyers may switch to the auctions with low-quality trucks even if their private value for high-quality truck is high. Given these differences, when the search costs are lower, high-type buyers are more likely to purchase high-quality trucks than low-type buyers. I show this pattern by simulation.

Assume all buyers search on one platform. Here θ can have two different values.

$$V(q_j, v_{ij}) = exp(\theta^T q_j + v_{ij})$$

$$\theta^H = 0.85, \theta^L = 0.6, v_{ij} \sim N(0, 0.5).$$

- Case 1: $m^H \sim Poisson(5), m^L \sim Poisson(5);$
- Case 2: $m^H \sim Poisson(10), m^L \sim Poisson(10).$

Figure A.8 shows while U^{H*} is still complementary in q and v, it is not true for U^{L*} . Buyers with high quality preference are more likely to choose the auctions with high quality trucks relative to low-type buyers when all of them can search more trucks.

Figure A.8: Expected Payoffs from Entering an Auction U(q,v) (Two types of Buyers)



Notes: 1. unit of expected payoffs is \$1,000; 2. "H" represents buyers with higher quality preference and "L" represents buyers with low quality preference.

A.3 Supplement to Identification

A.3.1 Baseline

Proof of Proposition 3

Proof. In the market with $N^H = N$, $N^L = 0$, according to the expression of $F_{2,\text{price}}$ above, at a specified p, I can solve for the $[(1 - Pr^{m^*})F^V(pq^H) + Pr^{m^*}F^V(p|q^H)^2]$ which satisfies the following equation:

$$(1 - Pr^{m^*})F^V(p|q^H) + Pr^{m^*}F^V(p|q^H)^2 = 1 - \sqrt{1 - F_{2,\text{price}}(p|q^H, q^H, N^H, N^L)}$$

Similarly to the case where $N^L = N, N^H = 0$.

Therefore, given the $F_{2,\text{price}}(p|q^H, N)$ and $F_{2,\text{price}}(p|q^L, N)$ at a specified p and Pr^{m^*} , I can get corresponding $F^V(p|q^H)$ and $F^V(p|q^L)$ which satisfy following equations:

$$F^{V}(p|q^{H}) = \frac{\sqrt{(1 - Pr^{m^{*}})^{2} + 4Pr^{m^{*}}[1 - \sqrt{1 - F_{2,\text{price}}(p|q^{H}, N)}]} - (1 - Pr^{m^{*}})}{2Pr^{m^{*}}}$$
(A.6)

$$F^{V}(p|q^{L}) = \frac{\sqrt{(1 - Pr^{m^{*}})^{2} + 4Pr^{m^{*}}[1 - \sqrt{1 - F_{2,\text{price}}(p|q^{L}, N)}]} - (1 - Pr^{m^{*}})}{2Pr^{m^{*}}}$$
(A.7)
$$\forall p \in [\underline{V}, \overline{V}]$$

Note that this means that $F_{2,\text{price}}(p|q^H, N) = F_{2,\text{price}}(p|q^H)$. In the market with one quality level and one platform, the price distribution given a truck has two buyers is independent of the number of trucks available in the market as long as N > 2.

Finally, using the markets with two different quality levels q^H and q^L in the

same market, I can get the condition to pin down Pr^{m^*} :

$$\begin{split} &(1-Pr^{m^*})F^V(p|q^H) + \frac{N^L}{N^H}(1-Pr^{m^*})F^V(p|q^L) + Pr^{m^*}\frac{N^H-1}{N-1}F^V(p|q^H)^2 + \dots \\ ⪻^{m^*}\frac{N^L}{N^H}\frac{N^L-1}{N-1}F^V(p|q^L)^2 + \frac{N^L}{N-1}F^V(V^*(p,q^L,q^H)|q^H)F^V(V^*(p,q^H,q^L)|q^L)] = \\ &1-\sqrt{1-F_{2,\text{price}}(p|q^H,q^L,N^H,N^L)} + \frac{N^L}{N^H}[1-\sqrt{1-F_{2,\text{price}}(p|q^L,q^H,N^H,N^L)}] \\ &\text{Where } U^*(p,q^H) = U^*(V^*(p,q^L,q^H),q^L), \\ &U^*(p,q^L) = U^*(V^*(p,q^L,q^L),q^H) \\ &U^*(p,q^L) = U^*(V^*(p,q^L,q^H),q^L), \\ &U^*(p,q^L) = U^*(V^*(p,q^L,q^H),q^L), \\ &U^*(p,q^L) = U^*(V^*(p,q^L,q^H),q^L), \\ &U^*(p,q^L) = U^*(V^*(p,q^L,q^L),q^L) \\ &U^*(p,q^L) = U^*(V^*(p,q^L,q^L),q^L) \\ &U^*(p,q^L) = U^*(P^*(p,q^L,q^L),q^L) \\ &U^*(p,q^L) = U^*(P^*(p,q^L),q^L) \\ &U^*(p,q^L) = U^*(P^*(p,q^L),q^L) \\ &U^*(p,q^L) = U^*(P^*(p,q^$$

Note that given the mappings (A.6) (A.7), the equilibrium payoffs U^* can be solved as a function of Pr^{m*} and $F_{2,\text{price}}$. Therefore, by solving equation (A.8), I can get the Pr^{m^*} . Namely, (A.6)(A.7)(A.8) can identify $F^V(\cdot|q^H)$, $F^V(\cdot|q^L)$ and Pr^{m^*} . Note that although the explicit part about Pr^{m^*} in (A.8) just in degree 2, it may enter U^* in higher order. However, equation (A.8) should be satisfied for any $\{q^H, q^L, N^H, N^L\}$ where Pr^{m^*} are the same. Then there is enough conditions to pin down a unique Pr^{m^*} .

A.3.2 Extensions

Figure A.9 summarizes the data and assumptions used for identify different models mentioned in the text.

Figure A.9: Measurements and Assumptions for the Identification of Different Models



Algorithm 1: Solving for the Equilibrium Expected Payoffs $\mathbf{U}^*(q, v)$

 $\begin{array}{l} \textbf{Result: } \mathbf{U}^*(q,v) = \{U_T^*(q,v)\}, T \in \{H,L\}, \texttt{platform} \in \{A,B\} \\ \texttt{initialization: } \mathbf{U}^{(0)}(q,v) = \sum_i \sum_j a_{ij}^{(0)} q^i \times v^j, e_{-i}^*(q,v) \\ \textbf{while } \mathbf{a}^{(t)} - \mathbf{a}^{(t-1)} > tol \ \textbf{do} \\ & \left| \begin{array}{c} \mathbf{a}^{(t)} = \mathbf{a}^{(t-2)} + \frac{\mathbf{a}^{(t-1)} - \mathbf{a}^{(t-2)}}{(t-1)^{\frac{2}{3}}} \\ \textbf{for Simulation s } \textbf{do} \\ & \left| \begin{array}{c} \texttt{Calculating the realized payoffs from choosing an auction with } (q,v) \\ & \texttt{when all the other buyers using } e_{-i}^*(q,v) : \tilde{U}_T^{A,s}(q,v), \tilde{U}_T^{B,s}(q,v). \\ \\ \textbf{end} \\ & \texttt{Regress } \tilde{U}_T^{A,s}(q,v|A, Pr_{-i}^{m,home*}) \ and \ \tilde{U}_T^{B,s}(q,v) \ on \ \sum_i \sum_j q^i \times v^j \ to \ get \ \hat{a}. \\ & Update \ \mathbf{a}^{(t)} = \hat{\mathbf{a}} \ from \ the \ regression \\ \end{array} \right.$

A.4 Supplement to Estimation

Algorithm 1 shows the way the details about the algorithm I used to solve for the expected payoff functions in equilibrium, which approximate the payoffs function by two-dimension Lagrange interpolation (Judd (1998)). To speed up the convergence, I update the coefficients for the polynomials "smoothly". This is similar to the way used in Weintraub et al. (2010).

A.4.1 Algorithm for Solving Equilibrium Payoffs of an Auction

A.5 Supplement to Estimation Results

A.5.1 Estimation Results for Two-type Model

In the model with two-type of buyers, I estimate models with different settings. For the single-homing buyers, their platform choice rule can follow one rule listed below:

- Rule 1: Both types of single-homing buyers enter the market according to the market share of the platforms.
- Rule 2: High-type single-homing buyers choose offline trucks and low-type entering the market according to the market share of the platform.
- Rule 3: High-type single-homing buyers choose offline trucks and low-type single-homing buyers choose online trucks.
- Rule 4: High-type single-homing buyers choose online trucks and low-type single-homing buyers choose offline trucks.

The estimation results for different settings are shown in Table A.2. We can see that the changes in buyers' equilibrium search choices have the similar pattern across different settings: buyers search more trucks and the share of buyers doing single-homing increase significantly (for most cases, more than 20%).

Figure A.10 shows the distribution of search costs in the model with two types of buyers under Rule 2 when $share^{H} = 0.6$ and $share^{L} = 0.4$. We see that both marginal and fixed search costs decrease significantly. The level of fixed costs is higher in this model since now high-type buyers only search on the offline platform when they are single-homing, the difference between expected payoffs from mulithoming and single-homing is higher.

A.5.2 Model Fits

Table A.3 summarizes how the estimates fit the observation by comparing the target moments. For second set of moments which includes the regression

	Sharel	H = 0.6	$Share^{H} = 0.35$			
	Share	= 0.0 L = 0.4	$Share^{L} = 0.35$ $Share^{L} = 0.65$			
		- 0.4		$\frac{1}{1}$		
	Rule 2	Rule 3	Rule 1	Rule 3	Rule 4	
Quality Pref.						
$ heta_H$	0.8048	0.7818	0.8375	0.8441	0.8410	
	(0.0046)	(0.0030)	(0.0057)	(0.0051)	(0.0047)	
$ heta_L$	0.7424	0.7568	0.7220	0.7487	0.7213	
5. 6	(0.0065)	(0.0031)	(0.0063)	(0.0062)	(0.0055)	
Dist.of v						
μ	0.0001	0.0001	0.0001	0.0001	0.0001	
	(0.0148)	(0.0116)	(0.0285)	(0.0205)	(0.0238)	
σ	0.5246	0.5463	0.5463	0.5061	0.5487	
	(0.0057)	(0.0072)	(0.0085)	(0.0094)	(0.0089)	
Discount of						
Quality Online						
α	-0.3364	-0.4438	-0.4650	-0.5500	-0.5455	
	(0.0059)	(0.0140)	(0.0080)	(0.0127)	(0.0124)	
Search Freq.						
$\lambda_{II}^{\text{Pre}}$	7.0603	6.6185	6.5483	6.4368	6.3038	
ЧH	(0.0342)	(0.1334)	(0.0460)	(0.0662)	(0.0414))	
$\lambda_{II}^{\text{Post}}$	7.6173	7.3751	7.1978	7.0130	6.4934	
11	(0.0326)	(0.1936)	(0.0397)	(0.0463)	(0.0300)	
λ_I^{Pre}	7.2080	6.3883	6.6464	6.2428	6.0765	
L	(0.0107)	(0.0895)	(0.0363)	(0.0435)	(0.0098)	
λ_L^{Post}	7.4542	6.4519	6.8942	6.4851	6.4961	
	(0.0069)	(0.1598)	(0.0536)	(0.0074)	(0.0257)	
Homing						
"*,Pre	0 4212	0 4030	0 5319	0 5137	0 4402	
ω_H	(0.0160)	(0.0326)	(0.1038)	(0.0922)	(0.0897)	
,,*,Post	0 3178	0.2188	0 1004	0.2136	0.0044	
\sim_H	(0.0264)	(0.0177)	(0.0932)	(0.0563)	(0.0203)	
,,*,Pre	0 210/	0 2446	0 5212	0 2120	0 5275	
ω_L	(0.0104)	(0.0311)	(0 0243)	(0.0120)	(0,0080)	
,,*,Post	0 0001	0 1020	0 0 0 2 10)	0.007/	0 1086	
ω_L	(0,0069)	(0.0173)	(0.007)	(0.0974)	(0.0167)	

Table A.2: Estimation Results for Two-type Model

Notes: standard Errors are in parentheses.



Figure A.10: CDF of Search Costs in Two-type Model

Notes: unit is \$1,000.

parameters, I compare the two most important pattern: the change in price of high quality trucks (75th percentile of quality) and the change in price sensitivity to the quality of trucks on the other platform.

A.5.3 Sensitivity Analysis

I report the results of an analysis of the sensitivity of the model parameters to the moments used in the estimation of the one-type model following the approach proposed by Andrews et al. (2017). This approach can be used to conveniently summarize the identification of the parameters in parametric structural models in which changing a single parameter can affect multiple observed outcomes.

	One	-type	Two-type	e (Rule 2)
Moment	Observed	Estimated	Observed	Estimated
$\overline{p}^{IP,Pre}$	2.8703	2.7155	2.8703	2.7158
$std(p)^{IP,Pre}(q \ge median)$	0.8997	0.7547	0.8997	0.7534
$\overline{p}^{IP,Pre}(q \ge median)$	3.3525	3.0680	3.3525	3.0954
$\overline{p}^{IP,Pre}(q < median)$	2.4882	2.4362	2.4882	2.4150
$\overline{p}^{RBA,Pre}$	3.0917	3.1277	3.0917	3.1842
$std(p)^{RBA,Pre}$	0.7105	0.7487	0.7105	0.7523
$\overline{p}^{RBA,Pre}_{RBA,Pre}(q \ge median)$	3.3811	3.4370	3.3811	3.5248
$\overline{p}^{RBA,Pre}_{IB,Past}(q < median)$	2.0846	2.3749	2.7457	2.7769
$\overline{p}^{IF,Fost}$	2.2634	2.5207	2.2634	2.5176
$std(p)^{II}$, $rost$	0.7547	0.7631	0.7547	0.7253
$\overline{p}^{II,IOSI}(q \ge median)$	2.8631	3.0401	2.8631	3.0546
$p^{II,IOSI}(q < median)$	2.0703	2.3534	2.0703	2.3446
$p_{\mu} p_{\mu} p_{\mu$	3.0/31	3.1092	3.0/31	3.1488
= BBA.Post(a > modium)	0.0409	0.7015	0.0409	0.0984
p^{median} $\overline{p}^{RBA,Post}(q \leq modian)$	3.2013 2.8621	3.3704 2.8162	5.2015 2.8621	5.4550 2 8271
$\frac{p}{1 + 1} IP, Post \qquad (q < mean n)$	2.0021	2.0102	2.0021	2.02/1
bid_1 $\frac{\partial n}{\partial t}$ (75th) IP Post	1.7979	1.8012	1.7979	1.7384
$\frac{\partial p}{\partial q}(q^{75th})^{II,I0St} - \dots$ $\frac{\partial p}{\partial q}(q^{75th})^{IP,Pre}$	0.2208	0.0997	0.2208	0.0153
$\left(\frac{\partial p^{IP,Post}}{\partial \overline{q}^{IP}} - \frac{\partial p^{IP,Pre}}{\partial \overline{q}^{IP}}\right) - \dots$	-0.1384	-0.146	-0.1384	-0.1309
$\left(\frac{\partial \overline{\rho}}{\partial \overline{q}^{RBA}} - \frac{\partial \overline{\rho}}{\partial \overline{q}^{RBA}}\right)$ $= \frac{\partial \overline{\rho}}{\partial \overline{q}^{RBA}}$			3 1568	2 8405
$P_H = \frac{P_H}{\sigma} IP, Pre$			2 2864	2.0403
$\frac{q_H}{\pi RBA, Pre}$			2 2077	2 2 2 0 7
$p_H = \frac{p_H}{\pi RBA, Pre}$			2.5077	2 5004
q_H $_{=}IP,Pre$			2.0011	3.3094
p_L _IP.Pre			3.0927	3.0651
q_L -BBA Pre			3.3659	3.3//9
p_L^{PL}			3.0927	3.085
$\overline{q}_{L}^{\text{normalized}}$			3.3659	3.3779
$\overline{p}_{H}^{II,IOSt}$			2.8206	2.6720
$\overline{q}_{H}^{IF,Fost}$			3.4887	3.4018
$\overline{p}_{H}^{RBA,Post}$			3.2399	3.2405
$\overline{q}_{H_{\text{max}}}^{\text{KBA,Post}}$			3.4887	3.4018
$\overline{p}_L^{IP,Post}$			2.6748	2.5241
$\overline{q}_{L}^{IP,Post}$			3.1200	3.1512
$\overline{p}_{L}^{RBA,Post}$			3.0097	3.0410
$-\vec{R}BA,Post$			3.3016	3.3174

Table A.3: Model Fitness of Targeted Moments

Notes: unit of price is \$1,000.





Notes: reported values are average values of the sensitivity measure for each moment/parameter. Sensitivity is calculated as $(g(\Theta)'Wg(\Theta))^{-1}g(\Theta)W$ where $g(\Theta)$ is the Jacobian of the moments with respect to the parameters evaluated at the parameter estimates and W is the weighting matrix.

Algorithm 2: Solving for the Equilibrium Search Choices in the Counterfactuals Given H^{mc} and H^{fc} : $Pr^{m^*,home^*}(H^{mc}, H^{fc})$

Result: $Pr^{m^*,\text{home}^*}(H^{mc},H^{fc})$ initialization: $Pr^{m^{*(0)},\text{home}^{*(0)}}$; while $Pr^{m^{*(t)},home^{*(t)}} - Pr^{m^{*(t-1)},home^{*(t-1)}} > tol$ do for Simulation s do for Market k do Calculate the realized payoffs from bidding in the centralized auction with the information from using an search strategy $m_i \forall i$ in market k: $U^{*s}(q, v | Pr_{-i}^{m^{*(t)}, home^{*(t)}})$ end end 1. Calculate the average payoffs from choosing a search choice $(m_i \text{home}_i)$ by average over all the markets and simulations: $W_{i}^{m,home}(Pr_{-i}^{m^{*(t)},home^{*(t)}});$ 2. Calculate the range of search costs which can support different search choices according to the equilibrium conditions $\{\overline{mc}(m, home), mc(m, home), \overline{fc}(m, home), fc(m, home)\}_{m,home};$ 3. Update the equilibrium search choices according to the updated thresholds and $\{H^{mc}, H^{fc}\}$: $Pr_i^{m^*,home^*}(m, multi) = \int_{\underline{mc}(m, multi)}^{\underline{mc}(m, multi)} \int_{fc}^{fc(mc_i)} h^{mc}(mc_i) h^{fc}(fc_i) dmc_i dfc_i$ end

A.6 Supplement to Counterfactuals

A.6.1 Algorithm Used to Solve the New Equilibrium in the Counterfactuals

Algorithm 2 presents the approach I used to solve the new equilibrium search choices used by buyers under alternative settings. Given the estimated search costs, I solve the new probability of each search choice based on the fixed-point theorem. Note since I can only identify the bounds of search cost distribution, I assume that the search costs following uniform distributions whose supports are estimated.

Appendix B: Appendix for Chapter 2

B.1 Computational Algorithms

This Appendix describes the methods used to solve our model. We describe the continuous type, finite horizon model in detail, before noting what changes in other cases. Our discussion will assume that there are two ex-ante symmetric duopolists. When firms are asymmetric, all of the operations need to be repeated for each firm.

B.1.1 Finite Horizon Model.

B.1.1.1 Preliminaries.

We specify discrete grids for the actual and perceived marginal costs of each firm, which will be used to keep track of expected per-period profits, value functions and pricing strategies. For example, when each firm's marginal cost lies on [8, 8.05] and we use 8-point equally spaced grids, the points are {8, 8.0071, 8.0143, 8.0214, 8.0286, 8.0357, 8.0429, 8.0500}.¹ We use interpolation and numerical integration to account for the fact that realized types will lie between these isolated points. The discount factor is $\beta = 0.99$.

It is useful to define several functions that we will use below:

• $P_{i,t}\left(\widehat{c_{i,t-1}^{j}}, c_{j,t-1}\right)$ is firm *i*'s pricing function in period *t*. This is a function of the marginal cost that *j* believes that *i* had in the previous period, $\widehat{c_{i,t-1}^{j}}$

¹The examples reported in Section 3 use 12 gridpoints, although we have experimented with as many as 20 gridpoints in each dimension to make sure that this does not have a material effect on the reported results.

(which, when j is forming equilibrium beliefs, will reflect that cost that i signaled in the previous period). It will also depend on the marginal cost that i believes that j had in the previous period, but we solve the game assuming that j is using its equilibrium strategy, so that i assumes that its perception of j's prior cost is correct, so we use the argument $c_{j,t-1}$. The actual price set will depend on $c_{i,t}$, and, when we need to integrate over the values that $p_{i,t}$ may take (e.g., to calculate expected profits) we will include $c_{i,t}$ as an explicit argument in the function.

- π_i(p_{i,t}, p_{j,t}, c_{i,t}) is firm *i*'s one-period profit when it sets price p_{i,t} and has marginal cost c_{i,t}, and its rival sets price p_{j,t}. This function does not depend on t because demand is assumed to be static and time-invariant.
- *V_{i,t}* (*c_{i,t-1}*, *c_{j,t-1}*) is the value function for firm *i* defined at the beginning of period *t*, before firm types have evolved to their period *t* values. It reflects the expected payoffs of firm *i* in period *t* and the discounted value of expected payoffs in future periods given equilibrium play in both *t* and future periods. It depends on the true value of each firm's type in *t* − 1, and the rival's perception of *i*'s *t* − 1 type (reflecting any deviation that *i* made in *t* − 1). In the case of an 8-point grid, *V_{i,t}* is a 512x1 vector.
- $\Pi^{i,t}\left(c_{i,t}, \widehat{c_{i,t}^{j}}, p_{i,t}, \widehat{c_{i,t-1}^{j}}, c_{j,t-1}\right)$ is the intermediate signaling payoff function of firm *i* when it knows its current marginal cost $c_{i,t}$, and is deciding what price to set. It does not know the period *t* type of its rival, but it reflects the pricing function that *i* expects *j* to use, $P_{j,t}\left(c_{j,t-1}, \widehat{c_{i,t-1}^{j}}\right)$. $\widehat{c_{i,t}^{j}}$ is the perception that *j* will have about *i*'s cost at the end of period *t*. When the rival sets price

$$P_{j,t}\left(c_{j,t}, c_{j,t-1}, \widehat{c_{i,t-1}^{j}}\right),$$

$$\Pi^{i,t}\left(c_{i,t}, \widehat{c_{i,t}^{j}}, p_{i,t}, \widehat{c_{i,t-1}^{j}}, c_{jt-1}\right) = \dots$$

$$\int_{\underline{c}_{j}}^{\overline{c}_{j}}\left(\pi_{i}\left(p_{i,t}, P_{j,t}\left(c_{j,t}, c_{j,t-1}, \widehat{c_{i,t-1}^{j}}\right), c_{i,t}\right) + \right)\psi_{j}(c_{j,t}|c_{j,t-1})dc_{j,t}.$$

where we note that $p_{i,t}$ only enters through current profits, and $\widehat{c_{i,t}^j}$ only enters through the discounted continuation value. In practice, our description will make up $\Pi^{i,t}$ into two components: $\Pi^{i,t} = \widetilde{\pi_i} + \widetilde{V_{i,t}}$, where

$$\widetilde{\pi_i}\left(p_{i,t}, P_{j,t}\left(c_{j,t-1}, \widehat{c_{i,t-1}^j}\right), c_{i,t}\right) = \dots$$
$$\int_{\underline{c_j}}^{\overline{c_j}} \pi_i\left(p_{i,t}, P_{j,t}\left(c_{j,t}, c_{j,t-1}, \widehat{c_{i,t-1}^j}\right), c_{i,t}\right) \psi_j(c_{j,t}|c_{j,t-1}) dc_{j,t}$$

and

$$\widetilde{V_{i,t}}\left(c_{i,t},\widehat{c_{i,t}^{j}},c_{j,t-1}\right) = \int_{\underline{c}_{j}}^{\overline{c}_{j}} \beta V_{i,t+1}\left(c_{i,t},\widehat{c_{i,t}^{j}},c_{j,t}\right)\psi_{j}(c_{j,t}|c_{j,t-1})dc_{j,t}.$$

Given a set of fully separating pricing functions $P_{i,t}\left(\widehat{c_{i,t-1}^{j}}, c_{j,t-1}\right)$, the relationship between Π and V is that

$$V_{i,t}\left(c_{i,t-1}, \widehat{c_{i,t-1}^{j}}, c_{j,t-1}\right) = \dots$$
$$\int_{\underline{c}_{i}}^{\overline{c}_{i}} \Pi^{i,t}\left(c_{i,t}, c_{i,t}, P_{i,t}\left(c_{i,t}, \widehat{c_{i,t-1}^{j}}, c_{j,t-1}\right), \widehat{c_{i,t-1}^{j}}, c_{j,t-1}\right) \psi_{i}(c_{i,t}|c_{i,t-1}) dc_{i,t}$$

where we recognize that, in equilibrium, *i*'s period *t* pricing function will reveal its cost to *j*, implying $\widehat{c_{i,t}^j} = c_{i,t}$.

B.1.1.2 Period *T*.

Assuming that play in period T - 1 was fully separating, we solve for BNE pricing strategies for each possible combination of beliefs (on our grid) about period T - 1 marginal costs. A strategy for each firm is an optimal price given the realized value of its own period T cost, given the pricing strategy of the rival, its prior marginal cost and the rival's belief about the firm's period T - 1 cost. Trapezoidal integration is used to integrate over the realized cost/price of the rival using a discretized version of the pdf of each firm's cost transition, and we solve for the BNE prices using the implied first-order conditions (i.e., those associated with maximizing static profits). With symmetric duopolists and 8-point grids, we find 512 equilibrium prices.

We use the equilibrium prices to calculate the beginning of period value function

$$V_{i,T}\left(c_{i,T-1}, \widehat{c_{j,T-1}^{j}}, c_{j,T-1}\right) = \dots$$
$$\int_{\underline{c}_{i}}^{\overline{c}_{i}} \int_{\underline{c}_{j}}^{\overline{c}_{j}} \pi_{i}\left(P_{i,T}^{*}\left(c_{i,T}, \widehat{c_{i,T-1}^{j}}, c_{j,T-1}\right), P_{j,T}^{*}\left(c_{j,T}, c_{j,T-1}, \widehat{c_{i,T-1}^{j}}\right), c_{i,T}\right) \dots$$
$$\psi_{j}(c_{j,T}|c_{j,T-1})\psi_{i}(c_{i,T}|c_{i,T-1})dc_{j,T}dc_{i,T}.$$

B.1.1.3 Period T - 1.

Firms choose prices in period T - 1 recognizing that their prices will affect rivals' prices in period T. We solve for period T-1 strategies, assuming separating equilibrium pricing and interpretation of beliefs in period T - 2, so that each firm has a point belief about its rival's period T - 2 marginal cost. We then use the following steps to compute equilibrium strategies. Step 1. (a) Compute

$$\widetilde{V}_{i,T-1}\left(c_{i,T-1}, \widehat{c_{i,T-1}^{j}}, c_{j,T-2}\right) = \beta \int_{\underline{c}_{j}}^{c_{j}} V_{i,T}\left(c_{i,T-1}, \widehat{c_{i,T-1}^{j}}, c_{j,T-1}\right) \psi_{j}(c_{j,T-1}|c_{j,T-2}) dc_{j,T-1}.$$

 $\widetilde{V}_{i,T-1}$ is the expected continuation value (i.e., not including period T-1 payoffs) for *i* when it is setting its period T-1 price, without knowing the period T-1 realization of c_j (but knowing that, in equilibrium, it will be revealed by $p_{j,T-1}$).

(b) Compute $\beta \frac{\partial \tilde{V}_{i,T-1}(c_{i,T-1},c_{j,T-2})}{\partial c_{i,T-1}^{j}}$ using numerical differences at each of the gridpoints (one-sided as appropriate). This array provides us with a set of values for the numerator in the differential equation (2.1). These derivatives do not depend on period T-1 prices, so we do not repeat this calculation as we look for equilibrium strategies.

(c) Verify belief monotonicity using these derivatives.

Step 2. We use the following iterative procedure to solve for equilibrium fully separating prices.² Use the BNE prices (i.e., those calculated in period *T*) as initial starting values. Set the iteration counter, iter = 0.

(a) Given the current guess of the strategy of firm j, $P_{j,T-1}\left(c_{j,T-1}, c_{j,T-2}, \widehat{c_{i,T-2}^{j}}\right)$, which is equal to the pricing functions solved for in the previous iteration, calculate

$$\frac{\partial \widetilde{\pi_{i,T-1}}\left(p_{i,T-1}, P_{j,T-1}\left(c_{j,T-2}, \widehat{c_{i,T-2}}\right), c_{i,T-1}\right)}{\partial p_{i,T-1}} \text{ for a grid of values } \left(p_{i,T-1}, \widehat{c_{i,T-2}^{j}}, c_{i,T-1}\right)$$

²We do not claim that this iterative procedure is computationally optimal, although it works reliably in our examples. There are some parallels between our problem and the problem of solving for equilibrium bid functions in asymmetric first-price auctions where both the lower and upper bounds of bid functions are endogenous. Hubbard and Paarsch (2013) provide a discussion of the types of methods that are used for these problems.

where

$$\widetilde{\pi_{i,T-1}}\left(p_{i,T-1}, P_{j,T-1}\left(c_{j,T-2}, \widehat{c_{i,T-2}^{j}}\right), c_{i,T-1}\right) = \dots$$
$$\int_{\underline{c}_{j}}^{\overline{c}_{j}} \pi_{i}\left(p_{i,T-1}, P_{j,T-1}\left(c_{j,T-1}, c_{j,T-2}, \widehat{c_{i,T-2}^{j}}\right), c_{i,T-1}\right) \psi_{j}(c_{j,T-1}|c_{j,T-2}) dc_{j,T-1}$$

i.e., the derivative of *i*'s expected profit with respect to its price, given that it does not know what price *j* will charge because it does not know $c_{j,T-1}$. The derivatives are evaluated on a fine grid (steps of one cent) of prices.³ This vector will be used to calculate the denominator in the differential equation (2.1).

For each $(\widehat{c_{i,T-2}^{j}, c_{j,T-2}})$, (b) Solve the lower boundary condition equation $\frac{\partial \widetilde{\pi} \left(p_{i,T-1}^{*}, P_{j,T-1} \left(c_{j,T-1}, c_{j,T-2}, \widehat{c_{i,T-2}} \right), c_{i} \right)}{\partial p_{i,T-1}} = 0$ for $p_{i,T-1}^{*}$, using a cubic spline to interpolate the vector calculated in (a). This gives the static best response price and the lowest price on *i*'s pricing function.

(c) Using this price as the initial point⁴, solve the differential equation, (2.1), to find *i*'s best response signaling pricing function. This is done using ode113 in MATLAB, with cubic spline interpolation used to calculate the values of the numerator and the denominator between the gridpoints.⁵ Interpolation is then used to calculate values for the pricing function for the specific values of $c_{i,T-1}$ on the cost/belief grid $\left(c_{i,T-1}, c_{i,T-2}^{j}, c_{j,T-2}\right)$.

³A fine grid is required because it is important to evaluate the derivatives accurately around the static best response, where the derivative will be equal to zero.

⁴In practice, the exact value of the derivative will be zero at the static best response, so that the differential equation will not be well-defined if this derivative is plugged in. We therefore begin solving the differential equation at the price where $\Pi_3^{i,T-1} + 1e - 4 = 0$. Pricing functions are essentially identical if we add 1e-5 or 1e-6 instead.

⁵See discussion of tolerances in Appendix B.1.2.2.

(d) Update the current guess of *i*'s pricing strategy using

$$P_{i,T-1}^{iter=k+1}\left(\widehat{c_{i,T-1}, c_{i,T-2}^{j}, c_{j,T-2}}\right) = (1-\tau)P_{i,T-1}^{iter=k}\left(\widehat{c_{i,T-1}, c_{i,T-2}^{j}, c_{j,T-2}}\right) + \dots$$
$$\tau P_{i,T-1}^{'}\left(\widehat{c_{i,T-1}, c_{i,T-2}^{j}, c_{j,T-2}}\right) \quad \forall c_{i,T-1}, \widehat{c_{i,T-2}^{j}, c_{j,T-2}}, c_{j,T-2}$$

where $P'_{i,T-1}$ are the best response functions that have just been computed. In the finite horizon case, $\tau = 1$, i.e., full updating, works effectively unless we are close to prices where the conditions required to characterize the unique best response fail to hold, in which case we also try using $\tau = \frac{1}{1+iter^{\frac{1}{6}}}$. See discussion below for how we update in the application where we use an infinite horizon model.

(e) Check if the maximum difference between $P_{i,T-1}^{iter=k}$ and $P'_{i,T-1}$, across all gridpoints, is less than 1e-6. If so, terminate the iterative process, else update the iteration counter to iter = iter + 1, and return to step 2(a).

(f) Verify that the solved pricing functions are monotonic in a firm's own marginal costs, and that, given the pricing functions of the rival, that the singlecrossing condition holds for the full range of prices used in the putative equilibrium.

Step 4. Compute *i*'s value $V_{i,T-1}$,

$$V_{i,T-1}\left(c_{i,T-2}, \widehat{c_{j,T-2}}, c_{j,T-2}\right) = \dots$$

$$\int_{\underline{c}_{i}} \int_{\underline{c}_{j}} \int_{\underline{c}_{j}} \left\{ \pi \left(P_{i,T-1}^{*}\left(c_{i,T-1}, \widehat{c_{i,T-2}}, c_{j,T-2}\right), P_{j,T-1}^{*}\left(c_{j,T-1}, c_{j,T-2}, \widehat{c_{i,T-2}}\right), c_{i,T-1}\right) +\beta V_{i,T}\left(c_{i,T-1}, c_{j,T-1}, c_{i,T-1}\right) \\ \psi_{j}(c_{j,T-1}|c_{j,T-2})\psi_{i}(c_{i,T-1}|c_{i,T-2})dc_{j,T-1}dc_{i,T-1}$$

where we are recognizing that equilibrium play at period T-1 will reveal *i*'s true cost to *j*. Note that this is the case even if, hypothetically, $\widehat{c_{i,T-2}^{j}} \neq c_{i,T-2}$ (i.e., *j* was misled in period T-2) because *i* should find it optimal to use its equilibrium signaling strategy given its new cost $c_{i,T-1}$ in response to *j* using a strategy based
on its $\widehat{c_{i,T-2}^j}$ belief.

B.1.1.4 Earlier Periods.

This process is then repeated for earlier periods, with an appropriate changing of subscripts. Given our assumption that first period beliefs reflect actual costs in a fictitious prior period, this procedure will also calculate strategies in the first period of the game.

B.1.2 Infinite Horizon Model.

We use an infinite horizon model for some of our examples and the empirical application. We find equilibrium pricing functions in the continuous type model using a modification of the procedure described above: in particular, we follow the logic of policy function iteration (Judd (1998)) to calculate values given a set of strategies.

The equilibrium objects that we need to solve for are a set of stationary pricing functions, $P_i^*\left(\widehat{c_{i,t-1}^j}, c_{j,t-1}\right)$ and value functions $V_i\left(c_{i,t-1}, \widehat{c_{i,t-1}^j}, c_{j,t-1}\right)$ which are consistent with each other given the static profit function and the transition functions for firm types.

We start by solving the period T-1 game described previously (i.e., assuming that there is a one more period of play where firms will use static Bayesian Nash Equilibrium strategies) to give an initial set of signaling pricing functions $(P_i^{*,iter=1})$. We then calculate firm values in each state $\left(c_{i,t-1}, c_{i,t-1}, c_{j,t-1}\right)$ if these pricing functions were used in every period of an infinite horizon game. This is done by creating a discretized form of the state transition process and calculating

$$\widehat{V}_{i}^{iter=1} = [I - \beta T]^{-1} \pi'_{i} \left(c_{i,t-1}, \widehat{c_{i,t-1}^{j}}, c_{j,t-1} \right)$$

where

$$\pi'_{i}\left(c_{i,t-1}, \widehat{c_{i,t-1}}, c_{j,t-1}\right) = \int_{\underline{c}_{i}}^{\overline{c}_{i}} \int_{\underline{c}_{j}}^{\overline{c}_{j}} \left\{ \pi_{i} \left(\begin{array}{c} P_{i}^{*,iter=1}\left(c_{i,t}, \widehat{c_{i,t-1}}, c_{j,t-1}\right), \\ P_{j}^{*,iter=1}\left(c_{j,t}, c_{j,t-1}, \widehat{c_{i,t-1}}\right), c_{i,t} \end{array} \right) \right\} \dots$$

 $\psi_j(c_{j,t}|c_{j,t-1})\psi_i(c_{i,t}|c_{i,t-1})dc_{j,t}dc_{i,t}$

and *T* is a transition matrix that reflects the transition probabilities for both firms' types and the behavioral assumption that equilibrium play in *t* (and future periods) will reveal period *t* costs. $P_j^{*,iter=1}\left(c_{j,t-1}, \widehat{c_{i,t-1}^j}\right)$ will reflect $P_i^{*,iter=1}$, applied to the states of the rival, when the firms are symmetric.

 $\hat{V}_i^{iter=1}$ is then used to compute a new set of pricing functions, $P_i^{*,iter=2}$, and the process is repeated until prices converge (tolerance 1e-4). Even though policy function iteration procedures do not necessarily converge, we find they work very well in our setting, when the conditions for separation hold, although it is sometimes necessary to update the pricing function to be a linear combination of the previous guess and the newly calculated best response. As illustrated in Figure 2.3, converged pricing functions found by this method are essentially identical to the pricing functions found for the early periods of long finite horizon games where the exact value of t has almost no effect on equilibrium pricing strategies. The computational advantage of this procedure comes from the fact that we do not perform the iterative procedure described above for every period of the game: instead there is a single iterative procedure where we solve for a single set of pricing strategies for the entire game.

B.1.2.1 Speeding Up Solutions By Interpolating Pricing Functions.

When we consider more than two firms and allow for asymmetries, the solution algorithm laid out above becomes slow, with most of the time spent solving differential equations. For example, with 8-point cost/belief grids, three asymmetric firms and 50 iterations, we would have to solve 25,600 differential equations. This would make estimation of the model using a nested fixed point procedure very slow. On the other hand, reducing the number of gridpoints can lead to inaccurate calculations of expected payoffs, and therefore strategies.

Examination of the equilibrium pricing functions (see, for example, Figure 2.3) shows that as we vary rivals' prior types, a firm's pricing functions look like they are translated without (noticeably) changing shape. We exploit this fact by solving for pricing functions for only a subset of the $(\widehat{c_{i,t-1}^j}, c_{j,t-1})$ gridpoints and using cubic splines to interpolate the remaining values.⁶ This allows us to achieve a substantial speed increase, while continuing to calculate expected values accurately on a finer grid.

B.1.2.2 Tolerances and Updating Rules Used for the Estimation of the Cost Parameters Using the Infinite Horizon Model.

In Section 2.4 we estimate the cost parameters using a nested fixed point algorithm, which means that both speed and accuracy are important. After considerable experimentation, we use the following tolerances:

- for the parameter search using fminsearch we set the tolerance for the parameter values at 1e-5 and the tolerance on changes to the objective function at 1e-5. The value of the minimized objective function is typically less than 0.0002, compared with the initial guess, for which we use estimates of the parameters assuming firms use static Bayesian Nash pricing strategies, which usually gives an objective function value of around 0.2.
- the tolerance for criterion for the pricing functions when solving the model is 1e-6 (i.e., at none of the grid points should the price on the best response pricing function be more than 1e-6 from the current guess).
- for the differential equation solver, the initial step size is 5e-5 and the maximum step size is 0.003 for the first ten iterations of the algorithm, but we

⁶For example, when we estimate our model in Section 2.4, we use a seven-point cost grid ($\{1,..,7\}$) for the profits and values of each firm. We solve for pricing functions for the full interaction of gridpoints $\{1,3,5,7\}$ and then interpolate the pricing functions for the remaining gridpoints.

then use an initial step size of 1e-5 and a maximum step size of 0.001.

• we update the pricing function to be the best response for the first 15 iterations, and then use a linear combination of the best response and the current guess where the weight on the best response changes linearly from 1 (iteration 16) to 0.1 (iteration 115).

When we use these tolerances, the infinite horizon game is typically solved using somewhere between 12 and 45 iterations, taking between 3 and 20 minutes. Estimation of the five parameters usually requires around 250 function evaluations, although the objective function and parameters are usually close to their final values within 100 evaluations.

B.1.3 Two-Type Model.

We use a model where each firm can have one of two types when we want to examine all strategies simultaneously or to consider a large number of alternative demand parameters. An additional advantage is that because prices, profits and values can be calculated for each possible type, we avoid small inaccuracies that result from numerical integration.

The key difference to the solution algorithm is that we no longer solve differential equations to find best response pricing functions. Recall that in the continuous type model, the differential equations characterize the unique separating best response when the signaling payoff function satisfies several conditions. In the discrete type model, one can construct multiple separating pricing functions that can be supported for different beliefs of the rival firm. To proceed we therefore need to choose a particular pricing function. We describe our choice, and the method we use to calculate the best response prices here. This procedure can be embedded within the procedure for solving either a finite horizon or an infinite horizon game.

To be as consistent with the continuous type model as possible, we use the prices that allow the two types to separate at the lowest cost, in terms of foregone

current profits taking the current guess of the pricing function of the rival as given, to the signaling firm (i.e., "Riley" signaling strategies, which would also be those that satisfy application of the intuitive criterion).⁷

The amended computational procedure is as follows (described for the infinite horizon case). Suppose that we are looking to find the pricing strategy of firm *i* in period *t* when it believes that *j*'s previous cost was $c_{j,t-1}$ and *j* believes that *i*'s previous cost was $\widehat{c_{i,t-1}^{j}}$. We will repeat this process for each $\left(\widehat{c_{i,t-1}^{j}}, c_{j,t-1}\right)$ combination, of which there will be four in the duopoly model. We need to solve for two prices: *i*'s price when its cost is $\underline{c_i}$ and its price when its cost is $\overline{c_i}$.

Step 1. Find $p_{i,t}^*(\underline{c_i})$, which will be the static best response, as the solution to $\frac{\partial \tilde{\pi} \left(p_{i,t}, P_{j,t} \left(c_{j,t}, c_{j,t-1}, \widehat{c_{i,t-1}^j} \right), \underline{c_i} \right)}{\partial p_{i,t}} = 0 \text{ where}$ $\widetilde{\pi}_i \left(p_{i,t}, P_{j,t} \left(c_{j,t}, c_{j,t-1}, \widehat{c_{i,t-1}^j} \right), \underline{c_i} \right) = \dots$ $\pi_i \left(p_{i,t}, P_{j,t} \left(\underline{c_j}, c_{j,t-1}, \widehat{c_{i,t-1}^j} \right), \underline{c_i} \right) \Pr(c_{j,t} = \underline{c_j} | c_{j,t-1}) + \dots$ $\pi_i \left(p_{i,t}, P_{j,t} \left(\overline{c_j}, c_{j,t-1}, \widehat{c_{i,t-1}^j} \right), \underline{c_i} \right) \Pr(c_{j,t} = \overline{c_j} | c_{j,t-1}) + \dots$

Step 2. Find $p_{i,t}^*(\overline{c_i})$. This is done by finding the price, p', higher than $p_{i,t}^*(\underline{c_i})$, which would make the low cost firm indifferent between setting price $p_{i,t}^*(\underline{c_i})$ and being perceived as a low cost type, and setting price p' and being perceived as a high cost type, i.e.,

$$\widetilde{\pi}\left(p_{i,t}^{*}(\underline{c_{i}}), P_{j,t}\left(c_{j,t}, c_{j,t-1}, \widehat{c_{i,t-1}^{j}}\right), \underline{c_{i}}\right) + \beta \widetilde{V_{i,t+1}}(\underline{c_{i}}, \underline{c_{i}}, c_{j,t-1}) = \dots$$
$$\widetilde{\pi}\left(p', P_{j,t}\left(c_{j,t}, c_{j,t-1}, \widehat{c_{i,t-1}^{j}}\right), \underline{c_{i}}\right) + \beta \widetilde{V_{i,t+1}}(\underline{c_{i}}, \overline{c_{i}}, c_{j,t-1})$$

⁷Of course, in the game we are considering it could be advantageous to the firms to use higher signaling prices, because of how this raises rivals' prices in equilibrium. This equilibrium consideration is ignored when selecting the Riley best response.

where

$$\widetilde{\beta V_{i,t+1}} \left(c_{i,t-1}, \widehat{c_{i,t-1}^{j}}, c_{j,t-1} \right) = V_{i,t+1} \left(c_{i,t-1}, \widehat{c_{i,t-1}^{j}}, \underline{c_{j}} \right) \Pr(c_{j,t} = \underline{c_{j}} | c_{j,t-1}) + \dots$$
$$V_{i,t+1} \left(c_{i,t-1}, \widehat{c_{i,t-1}^{j}}, \overline{c_{j}} \right) \left(1 - \Pr(c_{j,t} = \underline{c_{j}} | c_{j,t-1}) \right).$$

We verify that, consistent with single-crossing, the $\overline{c_i}$ type prefers to set the price p' rather than setting its static best response price. We also verify belief monotonicity when we calculate the value functions. As illustrated in Section B.2.1, there are parameters for which belief monotonicity fails.

B.2 Additional Examples.

B.2.1 Two-Type Examples.

A model where each firm can have one of two types has a much lower computational burden than the continuous type model. In this Appendix we will consider several parameterizations of a two-type model. In all of them we assume that firms are symmetric and that in any period $c_i = \underline{c} = 8$ or $c_i = \overline{c} = 8.05$. The probability that the cost remains the same as in the last period is $0.5 \le \rho < 1$. There are no signaling incentives when $\rho = 0.5$.

Refinement. A disadvantage of the two-type model is that for a given pricing strategy of firm *j*, firm *i* separating best response pricing function is not unique in the sense that it depends on how firm j will interpret the signal. We therefore impose a refinement that is consistent with the logic of the "intuitive criterion" (Cho and Kreps (1987)), which has often been applied as a refinement in discrete-type signaling games where only one player is signaling. Specifically, we assume that the low cost type's strategy will be the static best response, as in the continuous type model, and, under assumptions that appropriately map Conditions 3-6 to the two-type case, the high cost type's best response price will be the lowest price that the low cost type would be unwilling to set even if this would result in rivals' perceiving it as a high cost type rather than a low cost type. While this does uniquely define the best response, it does not guarantee a unique equilibrium in the oligopoly signaling game, and we have identified several examples in the infinite horizon version of the two-type model where there are multiple equilibria. The results reported in this Appendix use an algorithm which, when an infinite horizon equilibirum exists, appears consistently to select the equilibrium which corresponds to the equilibrium in the early periods of a long finite horizon game.

Method. See Appendix B.1.3 for a description of the method used to solve the two-type model.

B.2.1.1 Outcomes for Alternative Serial Correlation and Demand Parameters.

We assume nested logit demand where the indirect utility function for consumer *c* has the form $u_{i,c} = \beta - \alpha p_i + \sigma \nu_c + (1 - \sigma)\epsilon_{i,c}$. We choose β , α and σ so that, for each combination of parameters that we consider, the CI equilibrium prices (at average cost levels) are \$16 for each firm, the market share of each firm at these prices is 0.25, and the diversion, which measures the proportion of a product's lost demand that goes to the rival's product, rather than the outside good, when its price increases from the CI equilibrium price, has a value that we specify. We focus on diversion because when more demand goes the outside good, which is like a competitor that always offers a fixed utility and does not respond to a signal, firms have less incentive to signal and, as we will show, the belief monotonicity and single-crossing conditions become harder to satisfy.⁸ Given assumed market shares, the lowest possible value of this diversion measure is $\frac{1}{3}$, which corresponds to multinomial logit demand. We vary ρ from 0.5 (in which case there is no incentive to signal) to 0.99. We solve an infinite horizon version of our model.

Figure B.1 shows the results for a fine grid of values of diversion and ρ . The orange crosses indicate combinations where the conditions for characterizing best responses fail and we cannot find a separating equilibrium. For combinations where we can find a separating equilibrium the size and color of the circles indicate the percentage increase in average prices relative to average static Bayesian Nash equilibrium prices with the same demand and serial correlation parameters (these prices are also always very close to \$16). When serial correlation is very low, the price effects are always small whatever the level of diversion, and, for given diversion, the price effects become larger as serial correlation increases. For given serial correlation, higher diversion is associated with larger price effects, as it becomes more beneficial for a firm to increase its rival's price (because more of

⁸The intuition is that when the rival's expected price increases, a firm may have a greater incentive to lower its price, towards a static best response price, to take demand from the outside good. See below for an example.

Figure B.1: Equilibrium Average Price Increases in the Infinite Horizon Two-Type Duopoly Model as a Function of Diversion and Serial Correlation of Costs



Notes: red dots mark outcomes where there is a stationary separating equilibrium with average prices less than 0.5% above static BNE levels. The blue circles mark outcomes where there is a stationary separating equilibrium with larger average price increases relative to static BNE prices, and the size of the circle is linearly increasing in the percentage difference in prices (the largest effect shown has average prices increasing by 44.8%). Orange crosses mark outcomes where the conditions required to solve for best response functions fail and we cannot find an equilibrium. The diversion is measured by the proportion of demand that goes to the rival product when one product experiences a small increase in price at CI Nash equilibrium prices given average costs.

the demand that the rival loses will come to the firm), and the increase in a rival's price has a greater effect on the firm's best response. For moderate diversion, such as 0.6, an equilibrium cannot be sustained once serial correlation increases above 0.66. When diversion to rival products is very high, equilibria can be sustained with very large price effects: we find a maximum price increase of 44.8%.⁹

B.2.1.2 Failure of the Conditions Required for Existence of a Separating Equilibrium.

We now consider in more detail an example where the conditions required for separation fail. Demand is the same as before (i.e., indirect utility is $u_{i,c} =$ $5 - 0.1p_i + 0.25\nu_c + (1 - 0.25)\epsilon_{i,c}$), and each firm's marginal cost is either 8 (low) or 8.05 (high). We assume that $\rho = 0.99$ so a signal is very informative about next period's marginal costs and signaling incentives are strong.

Figure B.2: Equilibrium Prices in the Two-Type Marginal Cost Model (parameters described in the text)



Figure B.2 shows the full set of eight equilibrium prices in each period as we

⁹In the diagram, the highest serial correlation for which we can find an equilibrium falls when we increase diversion above 0.95. This appears to reflect the fact that, at this level, small increases in diversion can increase signaling prices significantly, leading the conditions to fail. For each considered value of diversion above 0.95, we identify a value of ρ where signaling raises prices by more than 43.0% and 44.8%.

move backwards from the end of the game. The legend denotes states by {"the firm's perceived cost in t-1", "its rival's perceived cost in t-1" - "the firm's realized marginal cost in t"} so blue indicates prices for a firm whose perceived marginal cost in the previous period was high, its rival's perceived previous period cost was low, and a cross (circle) indicates that the firm's current cost is low (high).

The green crosses (LL-L) remain almost unchanged across periods, as they represent static best responses when both players know that their rival is very likely to be setting the same price, but, as we move earlier in the game, the remaining prices increase, because they involve either signaling (by a \bar{c} firm) or a static best response to a rival who is likely to be raising its price to signal.

In period T - 6 the order of the prices changes with the HH-H price (red circle) below the HL-H price (blue circle). This implies that in period T - 7, a firm that believes its rival is likely to be high cost, is more likely to increase its rival's next period (T - 6) price if it (the firm) is believed to be *low* cost than if it is believed to be high cost. As profits increase in the rival's price, this will lead belief monotonicity to be violated.

Why does the order of the red and blue circles switch? It reflects changes in both the incentive to signal (i.e., the possible effect on future prices) and the cost of signaling (i.e., the effect on current profits). Recall that in the two-type model the equilibrium price of the \bar{c} type is determined by the lowest price that the low-cost firm would be unwilling to set even if choosing it would lead to it being perceived as high cost. Consider the cost, in terms of foregone period T - 6profit, for a low-cost firm of raising its price. The upper panel of Figure B.3 shows the period T - 6 one-period profit functions for a low cost firm given different beliefs about previous firm types and the expected price of the rival.¹⁰ The lower panel shows the corresponding derivatives of the profit function with respect to the firm's own price. For prices above \$34, the marginal loss in profit from a price increase is greater for a red firm (i.e., a firm likely to face a high cost rival) than

¹⁰For example, an HL firm expects to face a low-cost LH firm (setting a black cross price) with probability 0.99, so the expected rival price is \$29.46.



a blue firm (i.e., a low cost rival) so it is less costly for the blue firm to raise its price.¹¹

Now consider the incentive of a low-cost firm to signal (i.e., to pretend to be high-cost). The incentive of an HL (blue) firm to signal a high cost in period T - 6 is that it is very likely to lead to its rival setting the black cross, rather than the green cross, price in period T - 5. This difference is large, so that the incentive to signal is strong. The incentive of an HH (red) firm to signal is that this will very likely lead to it facing the red, rather than the blue, circle price in period T - 5. These period T - 5 prices are closer together (than the black and green crosses) so the incentive to signal will tend to be weaker. The cost and the incentive effects together lead to a reversal of the order of the period T - 6 equilibrium prices, causing belief monotonicity to fail in period T - 7.

B.2.1.3 Effect of the Number of Firms on Equilibrium Outcomes.

We can also use the two-type model to illustrate the effects of increasing the number of symmetric competitors. Demand is the same as in our example in the text (i.e., indirect utility is $u_{i,c} = 5 - 0.1p_i + 0.25\nu_c + (1 - 0.25)\epsilon_{i,c}$), and each firm's marginal cost is either 8 (low) or 8.05 (high). We assume $\rho = \frac{2}{3}$. Figure B.4 shows average equilibrium prices under CI and in our model in infinite horizon games with between 1 and 7 firms. For comparison, average prices with joint profit maximization under CI are also included. Monopoly prices are (obviously) identical under complete and asymmetric information. Relative to CI, dynamic signaling raises average prices by 7.4% under duopoly, 2.2% under triopoly and 0.9% with four firms. With seven firms, the effect is just 0.1%.

The interpretation of why signaling raises prices more when there are fewer firms is that it is a strategic investment to raise rivals' future prices: a firm sacrifices profits in the current period, in order to raise its profits in the next period. For the

¹¹The crossing of the derivative functions reflects the failure of strategic complementarity (defined as $\frac{\partial^2 \pi_i}{\partial p_i \partial p_j} > 0$) for logit-based demand when prices are significantly above static profitmaximizing levels. The intuition is that, as a rival's price increases, the incentive for a firm to reduce its (high) price towards the static best response price can increase.

Figure B.4: Equilibrium Effects on Average Prices in the Infinite Horizon Two-Type Model with Different Numbers of Symmetric, Single Product Firms



same indirect utility function, a price increase becomes more costly in terms of current lost profits when there are more firms because a firm's residual demand is more elastic. In addition, there is a reduced incentive for a single firm to invest because, all else equal, the expected price of single firm will have less effect on the prices of its rivals.

B.2.2 Alternative Sources of Asymmetric Information.

While it is plausible that, in many industries, firms have some private information about their marginal costs and that whatever is unobserved is likely to be serially correlated, our results are not dependent on assuming that it is marginal costs that are privately observed. In this Appendix we consider three examples where marginal costs are fixed and known and the asymmetric information is embedded in a different part of the profit function. In each case we show that equilibrium prices can be significantly higher, and more volatile, than in the CI or static incomplete information versions of the model. The fact that other formulations generate similar results is not surprising, but we perform the calculations in order to emphasize the point that we are not tied to the marginal cost assumption. In all cases, we assume single-product duopolists, as in Section 2.3, and we solve the continuous type, infinite horizon version of our model. The demand parameters also take on their baseline values from Section 2.3, and marginal cost of each firm is held fixed at 8.

Variant 1: Weights on Profits and Revenues. In the first variant, we allow for there to be uncertainty about the weight that each firm places on profits rather than revenues. A number of theoretical and empirical papers study whether managers want to maximize profits or alternative outcome variables, and whether shareholders might strategically choose to incentivize managers to deviate from profit maximization (e.g., Sklivas (1987), Katz (1991), Murphy (1999), De Angelis and Grinstein (2014)). The empirical literature suggests that managers are affected by a variety of incentives that may be complicated for outsiders to evaluate and which may vary over time, depending on oversight from shareholders or corporate boards, and financial constraints.

Without assuming a particular theory of governance, we suppose that the weight placed on profits by firm *i* in period *t* is $\tau_{i,t}$ and that this variable lies on the interval [0.89, 0.9], with the remaining weight on firm revenues. As before, we suppose that the variable evolves according to a truncated AR(1) process, with

Table B.1: Price Effects in Models Where Alternative Elen	nents of Firm	Objective Functic	ons are Private	e Information
			Infinite	e Horizon
Model	Mean Price	<u>ac bNE</u> Std. Dev. Price	Mean Price	<u>.ng Model</u> Std. Dev. Price
Weight on Firm Profits vs. Revenues is Private Information Weight on Profits on $[0.89, 0.9]$ a = 0.8 std. dev. innovation 0.0044	21.79	0.01	24.45	0.59
(eqm. prices if firms are known to maximize profits, i.e., if weight on profit equals 1, are 22.59)				
Weight on Firm vs. Industry Profits is Private Information Weight on Industry Profit [0.00,0.02] $\rho = 0.8$, std. dev. innovation 0.0088	22.66	0.02	25.34	0.59
(eqm. prices if firms are known to maximize own profits are 22.59)				
Size of Each Firm's Loyal Market is Private Information Size of Loyal Market [0.10,0.12] (as fraction of duopoly market)	25.17	0.07	27.56	0.59
$\rho = 0.8$, std. dev. innovation 0.0088 (eqm. prices when loyal markets known to				
be 11% of duopoly markets are 25.17)				

 $\rho = 0.8$. The standard deviation of the innovations is chosen so that, as for our baseline model where marginal costs are private information, the probability that a type will transition from the highest point of the support in one period to a value in the lower half of the support in the next period is 0.32.

The first panel of Table B.1 reports the average CI price when both firms (are known to) maximize profits is 22.59. When a firm places some weight on revenues, it will tend to set a lower price, and the average static BNE or CI price when the profit weight lies on [0.89, 0.90] is 21.79. However, with signaling, average prices increase significantly: in this example, the average Markov Perfect Bayesian Equilibrium price is 8.2% above the average price level *when both firms are known to maximize profits*, with profits increasing by 18%. This example suggests there may be some advantage to shareholders if they keep managers' incentives opaque to rivals even in markets where firms set prices for differentiated products.¹²

Variant 2: Weight on Profits of Other Firms in the Industry. In the empirical Industrial Organization literature, it is common to model tacitly collusive behavior in a reduced-form way by generalizing static first-order conditions to allow for each firm to place some weight on the profits of other firms in the same market (Porter (1983), Bresnahan (1989), Miller and Weinberg (2017)). This type of formulation could also be rationalized by models where participants in financial markets become more optimistic about a firm's prospects when its rivals announce high profits (Rotemberg and Scharfstein (1990)) or by models where firms maximize the overall returns of shareholders who hold stock in competitors (O'Brien and Salop (1999), Azar et al. (2018)).

We consider a model where rivals have some limited uncertainty about the weight that a firm places on its own profit rather than the profit of the industry. Specifically we assume that each firm places a weight $\tau_{i,t}$ of [0.98, 1] on its own

¹²The usual explanation for why shareholders might want to commit to incentivizing their managers to place some weight on revenues comes from quantity-setting models where other firms will reduce their output when a firm's managers are committed to increase their output. In our model it is uncertainty about what firms are trying to maximize that causes equilibrium prices to rise, through the mechanism of signaling.

profits, and $1 - \tau_{i,t}$ on the profits of the industry as a whole (of course, its own profits also contribute to industry profits). We assume that the transition process has $\rho = 0.8$ and $\sigma = 0.0088$, which means that the probability of a type transitioning from the highest point of the support to below the median is 0.32, as in the first example. As can be seen in the second panel of Table B.1, the effect is, once again, to raise prices substantially in the dynamic game with asymmetric information.

Variant 3: Demand Shocks. Our experience in seminars is that many economists believe it is more intuitive that some aspect of demand will be private information to the firm than marginal costs will be.

Some formulations of demand uncertainty give rise to signaling incentives that would be qualitatively different from the ones in our framework. For example, suppose that demand has a logit structure and that each firm has private information about the serially correlated and unobserved quality of its product. Duopolist firms observe each other's prices but not quantities, so that prices are informative about quality. A firm with higher quality will want to charge a higher price, but its rival's optimal price will likely decrease in the firm's quality, so it is unclear whether a firm will want to be perceived as high quality or as low quality. This is likely to be a case where only some type of pooling equilibrium exists.

Here we consider a simple example where firms do have incentives to raise prices to signal that their demand is high. Suppose that each firm sells its products in two markets. In one market, the firms compete as duopolists, but in the other market the firm is a monopolist (so for example, both firms are in market A, firm 1 is the only firm in market B, and firm 2 is the only firm in market C). Due to the possibility of arbitrage, or some other constraint, each firm can only set one price across the markets. One rationalization of this setup would be that each firm has some loyal or locked-in customers, but that additional consumers are competed for. Product quality is known, but firms are uncertain about the size of their rival's loyal market. Normalizing the size of the common market to 1, the sizes of the loyal markets lie between [0.1, 0.12]. The utility specification is the same as before except loyal customers only choose between a single product and the outside good. The transition assumptions are the same as in variant 2. In this formulation, firms will set prices based on the weighted average marginal revenues from the two markets, and when the size of their monopoly market is larger they will prefer higher prices. A firm will therefore have incentive to raise its price to signal that its monopoly market is larger.

The results are presented in the third panel of Table B.1. The addition of the loyal market, where a firm's demand is less elastic, raises prices under all information structures, but the average signaling equilibrium prices are 10% higher than the prices under CI or in a static game with asymmetric information.

B.3 Existence and Uniqueness of a Fully Separating Equilibrium in a Finite Horizon Game with Linear Demand

As discussed in the text, Mailath (1989) and Mester (1992) provide proofs of the existence and uniqueness of a fully separating equilibrium in a two-period duopoly, linear demand, continuous cost price-setting game and a three-period duopoly, linear demand, continuous cost quantity-setting games respectively. This Appendix presents a theoretical proof of existence and uniqueness of a fully-separating Markov Perfect Bayesian Equilibrium for a finite-horizon duopoly pricing game with linear demand and marginal costs that are private information, under a condition that the range of costs is "small enough" so that the single-crossing condition holds. As explained in the text, we have to rely on computational analysis when assuming nonlinear demand or an infinite horizon, and in our application we assume both.¹³ However, we include our proof for the linear demand and finite horizon case for completeness.

We make the following specific assumptions on the model. There are two firms, and i will index the firm.

Assumptions

A1 (linear demand). $q_{i,t} = a_i - b_{1,i}p_{i,t} + b_{2,i}p_{j,t}$, $b_{1,i} > b_{2,i} > 0$.

A2 (positive demand). The intercepts *a* are large enough that for all of the prices charged on the equilibrium path, both firms will have positive output.

A3 (continuous cost interval). The marginal costs of each firm, $c_{i,t}$, lie on compact intervals where $[\underline{c_i}, \overline{c_i}]$ where $\overline{c_i} > \underline{c_i} > 0$.

A4 (cost transitions). Costs evolve independently according to first-order Markov processes with conditional densities $\Psi_i(c_{i,t}|c_{i,t-1})$, where the conditional density functions are smooth in $c_{i,t}$ and $c_{i,t-1}$ and strictly positive for all $[\underline{c_i}, \overline{c_i}]$. $E(c_{i,t}|c_{i,t-1})$ is continuous and strictly increasing in $c_{i,t-1}$.

A5 (discount factor). There is a common discount factor $0 < \beta < 1$.

The statement of the results and the proof will use the following notation.

¹³Our proof is for two firms that may be asymmetric. Extending the proof to three symmetric firms is straightforward.

- $\pi_{i,t}$ denotes per-period profits in period *t*. $\pi_{i,t} = (p_{i,t} c_{i,t})q_{i,t}(p_{i,t}, p_{-i,t})$.
- V_{i,t}(c_{i,t-1}, c_{i,t-1}, c_{j,t-1}) is i's value at the beginning of period t, before c_{i,t} is revealed, when it is perceived to have cost c_{i,t-1}, and its real cost is c_{i,t-1}, and it believes that j's t − 1 cost was c_{j,t-1}.
- Π^{i,t}(c_{i,t}, c_{i,t}, p_{i,t}, c_{j,t-1}) ("signaling payoff function") represents the expected current and future profits (given equilibrium behavior in future periods) of firm *i* in period *t*, when it sets price p_{i,t}, has cost c_{i,t} and is perceived, at the end of the period, as having cost c_{i,t}. c_{j,t-1} is *i*'s perception of *j*'s cost in period *t* 1. In equilibrium, this perception will be correct so we denote it simply by c_{j,t-1}. Π^{i,t}(c_{i,t}, c_{i,t}, p_{i,t}, c_{j,t-1}) is implicitly conditioned on *j*'s period *t* pricing strategy, which will involve *j* setting a price with an average of p_{j,t} and which *i* assumes will reveal c_{j,t}. Π^{i,t}_k(c_{i,t}, c_{i,t}, p_{i,t}, c_{j,t-1}) denotes the derivative of Π^{i,t}(c_{i,t}, c_{i,t}, p_{i,t}, c_{j,t-1}) with respect to the kth argument.
- Prices (the proof will indicate conditioning arguments where necesary):
 - $p_{i,t}^*$ is *i*'s equilibrium strategy in a fully separating MBPE (i.e., it is a function);
 - $p_{i,t}^{BR}$ is *i*'s separating best response pricing function given some separating strategy (not necessarily the equilibrium strategy) by *j*;
 - p^{**}_{i,t} is a price that is a statically optimal best response (i.e., maximizes
 i's current profits) given j's strategy;
 - $\overline{p_{j,t}}$ is the average price set by j when it uses a particular strategy; and,
 - our description of separating pricing strategies will refer to "initial values", which will reflect a $p_{i,t}^{**}$ price determined as the solution to a static profit maximization problem when $c_{i,t} = \underline{c_i}$, and, the "increment" which refers to the additional price above this initial value that may reflect signaling behavior.

B.3.1 Preliminary Results.

We begin with a useful Lemma.

Lemma 2. In a fully separating Markov Perfect Bayesian Equilibrium, play on the equilibrium path will have the following properties, (L-i) $p_{i,t}^*$ will be a function of $c_{i,t}$ and the costs $c_{i,t-1}$ and $c_{j,t-1}$ revealed by prices at t - 1; (L-ii) the only effect of $c_{i,t-1}$ on $p_{i,t}^*$ is through the effect that it will have on the expected value of $p_{j,t}$; (L-iii) *i*'s period *t* price, and the inference that *j* makes about $c_{i,t}$, based on this price, will affect *i*'s profits in *t* and t + 1 only.

Proof. (L-i) In a fully separating equilibrium, prices at t - 1 will reveal marginal costs at t - 1 and the first-order Markovian assumption on the Ψ_i s implies that costs at t - 1 contain all available information from earlier periods about costs. The Markovian equilibrium assumption implies that strategies depend on payoff-relevant state variables (current costs) and beliefs about those variables, only. This implies that strategies can be functions of $c_{i,t}$ (which is private information to i when $p_{i,t}$ is chosen), $c_{i,t-1}$ and $c_{j,t-1}$ only.

(L-ii) The equilibrium choice of $p_{i,t}^*$ will depend on its effect on expected profits in future periods and expected profits at t. Property (L-i) implies that given $p_{i,t}$, which reveals $c_{i,t}$, $c_{i,t-1}$ will not affect what happens at t + 1. Expected profits in period t are $(p_{i,t} - c_{i,t})(a_i - b_{1,i}p_{i,t} + b_{2,i}\overline{p_{j,t}})$ so $c_{i,t-1}$ can only affect i's payoffs through its effect on $\overline{p_{j,t}}$.

(L-iii) Suppose that instead of equilibrium price $p_{i,t}^*$, *i* sets a price $p'_{i,t}$ in the range of the equilibrium price function. t+1 strategies specify an optimal strategy for *i* given $c_{i,t+1}$, $c_{j,t}$ and the cost implied by $p'_{i,t}$, and it will be optimal to use these strategies at t+1 (because of property (L-ii)), so t+1 strategies will correctly reveal $c_{i,t+1}$. Therefore charging $p'_{i,t}$ not $p^{**}_{i,t}$ only affects profits at t and t+1. \Box

Our results characterizing firm i's separating best response function in period t, given a fully revealing pricing strategy, of any form, by j and the assumed form

of strategies at t + 1, are based on the following theorems which are adapted from Mailath (1987).

Theorem 7. Adapted from Theorems 1 and 2, and the Corollary, in Mailath (1987). If (MT-i) $\Pi^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})$ is smooth in arguments $(c_{i,t}, \widehat{c_{i,t}})$, (MT-ii) $\Pi_2^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) > 0$ [belief monotonicity], (MT-iii) $\Pi_{13}^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})$ > 0 [type monotonicity], (MT-iv) $\Pi_3^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) = 0$ for only one p_i , and for this p_i , $\Pi_{33}^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) < 0$ [strict quasi-concavity], (MT-v) there exists k > 0such that $\Pi_{33}^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) \ge 0$ implies $|\Pi_3^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})| > k$, then a pricing function $p_{i,t}^{BR}(c_{i,t}, c_{j,t-1})$ that solves the differential equation

$$\frac{\partial p_{i,t}^{BR}(c_{i,t},c_{j,t-1})}{\partial c_{i,t}} = -\frac{\Pi_2^{i,t}(c_{i,t},c_{i,t},p_{i,t},c_{j,t-1})}{\Pi_3^{i,t}(c_{i,t},c_{i,t},p_{i,t},c_{j,t-1})}$$
(B.1)

and has a lower initial value condition where $p_{i,t}^{BR}(\underline{c_i}, c_{j,t-1})$ solves $\Pi_3^{i,t}(\underline{c_i}, \underline{c_i}, p_{i,t}^{BR}(\underline{c_i}, c_{j,t-1}), c_{j,t-1}) = 0$ is the unique fully separating best response function if a fully separating best response exists.

Theorem 8. Adapted from Theorem 3 in Mailath (1987). Suppose assumptions $(MT-i)-(MT-\nu)$ in Theorem 7 hold. If $(MT-\nu i)$, for (\widehat{c}_i, p) in the graph of $p_{i,t}^{BR}(c_{i,t}, c_{j,t-1})$, $\frac{\Pi_3^{i,t}(c_{i,t}, \widehat{c}_{i,t}, p_{i,t}, c_{j,t-1})}{\Pi_2^{i,t}(c_{i,t}, \widehat{c}_{i,t}, p_{i,t}, c_{j,t-1})}$ is either strictly increasing or decreasing in $c_{i,t}$ [single-crossing], then the fully separating best response described in Theorem 7 exists.

B.3.2 Main Result.

The following theorem gives our main result.

Theorem 9. If $\overline{c_i}-\underline{c_i}$ is small enough for all *i*, in any finite horizon game there will exist a unique fully separating MPBE where, on the equilibrium path, firm *i*'s equilibrium pricing strategy $p_{i,t}^*(c_{i,t}, c_{i,t-1}, c_{j,t-1})$ in any period t < T has the form of the best response function described in Theorem 7. In period T firms will choose static payoff-maximizing prices given their beliefs about rivals' costs in period T - 1.

In periods t < T, pricing strategies will have the following features: (T-i) (a) the initial values (i.e., static best response prices when $c_{i,t} = \underline{c_i}$) are functions of $c_{j,t-1}$ and $c_{i,t-1}$ only (in the following we will denote the function that determines the initial value $g_{i,t}(c_{j,t-1}, c_{i,t-1})$), and (b) the increment above the initial value (a function $f_{i,t}(c_{i,t}, c_{j,t-1})$) is a continuous function of $c_{i,t}$ and $c_{j,t-1}$ only, and in particular it does not depend on $\overline{p_{j,t}}$; (T-ii) for all $c_{i,t} > \underline{c_i}$ the price charged is always above the static best response price for $c_{i,t}$, (T-iii) the effect of $c_{j,t-1}$ on the increment only comes through its effect on i's belief about the distribution of $c_{j,t+1}$, and (T-iv) (a) i's pricing function is continuous and strictly increasing in $c_{i,t}$, (b) i's pricing function is continuous and strictly increasing in $\overline{p_{j,t}}$, (c) i's pricing function is continuous and strictly increasing in $c_{i,t-1}$ and (i's perception of) $c_{j,t-1}$.

B.3.2.1 Proof.

The proof uses induction, showing that if strategies have this form in periods t + 1,...,T - 1 there will exist a unique MPBE with the required form in any period t < T - 1. We then show that the form of equilibrium strategies in period T will lead to strategies that have the specified form in period T - 1.

Period t < T - 1. The logic of the proof for period t is to show that the conditions required for Mailath's theorems hold given Lemma 2 and the assumed equilibrium form of pricing behavior in t + 1. This shows that there will be a unique best response pricing function for each firm given any separating strategy of the other firm. This will let us show some of the features specified above. We then show that there can be only one pair of pricing functions with these properties that are best responses to each other, and this will allow us to show the remaining features.

Uniqueness, Existence and Form of i's Fully Separating Best Response Function Given *j*'s Strategy.

We go through the conditions required for Mailath's results in turn.

Condition (MT-i): $\Pi^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})$ is smooth in arguments $(c_{i,t}, \widehat{c_{i,t}})$. Lemma 2 implies that

$$\Pi^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) = E\pi_{i,t}(c_{i,t}, p_{i,t}, \overline{p_{j,t}}) + \beta E(V_{i,t+1}(c_{i,t}, \widehat{c_{i,t}}, c_{j,t}|c_{j,t-1}))$$
(B.2)

where the second expectation is over the cost that j reveals in period t.

 $E\pi_{i,t}(c_{i,t}, p_{i,t}, \overline{p_{j,t}}) = (p_{i,t} - c_{i,t})(a_i - b_{1,i}p_{i,t} + b_{2,i}\overline{p_{j,t}})$ which is smooth in $c_{i,t}$. Profits in t + 1 will be equal to $(p_{i,t+1} - c_{i,t+1})(a_i - b_{1,i}p_{i,t+1} + b_{2,i}p_{j,t+1})$ and smoothness of the period-*t* expectation of these profits follows from the assumed smoothness of the Ψ_i conditional densities (A4) and the continuity of the pricing functions (Ti/T-iv). Similar logic (and the results concerning period *T* prices below) implies that the period-*t* expectation of discounted profits in t + 2 and future periods will also be continuous in $c_{i,t}, c_{j,t-1}$ and $\widehat{c_{i,t}}$. Therefore $\beta E(V_{i,t+1}(c_{i,t}, \widehat{c_{i,t}}, c_{j,t}|c_{j,t-1}))$ will be smooth in $c_{i,t}, \widehat{c_{i,t}}$ and $c_{j,t-1}$.

Condition (MT-ii): $\Pi_2^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) > 0$. From Lemma 2, $\widehat{c_{i,t}}$ only affects future profits in period t+1 given equilibrium play from t+1 forwards (L-iii). Denote expected profits in period t+1 when j charges an expected price $\overline{p_{j,t+1}}(\widehat{c_{i,t}}, c_{j,t})$, $E\pi_{i,t+1}(c_{i,t+1}, p_{i,t+1}, \overline{p_{j,t+1}}(\widehat{c_{i,t}}, c_{j,t}))$, so

$$\Pi_{2}^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t},c_{j,t-1}) = \dots$$

$$\int \int \frac{\partial E\pi_{i,t+1}(c_{i,t+1},p_{i,t+1},\overline{p_{j,t+1}}(\widehat{c_{i,t}},c_{j,t})))}{\partial \widehat{c_{i,t}}} \Psi_{i}(c_{i,t+1}|c_{i,t}) \Psi_{j}(c_{j,t}|c_{j,t-1}) dc_{i,t+1} dc_{j,t}$$

Given that $\frac{\partial \overline{p_{j,t+1}}(\widehat{c_{i,t}},c_{j,t})}{\partial \widehat{c_{i,t}}} > 0$ (T-iv (c)), it is sufficient to show that $\frac{\partial E\pi_{i,t+1}(c_{i,t+1},p_{i,t+1},\overline{p_{j,t+1}})}{\partial \overline{p_{j,t+1}}} > 0.$

Express the price that *i* charges in t + 1 as $p_{i,t+1} = p_{i,t+1}^{**}(\overline{p_{j,t+1}}, c_{i,t+1}) + p'$, where $p_{i,t+1}^{**}(\overline{p_{j,t+1}}, c_{i,t+1})$ is the static profit-maximizing best response to $\overline{p_{j,t+1}}$ given $c_{i,t+1}$ and $p' \ge 0$ is an increment above the static best response price. Linear

demand implies that *i*'s expected t + 1 profit is

$$E\pi_{i,t+1}'(c_{i,t+1}, p', \overline{p_{j,t+1}}) = (p_{i,t+1}^{**}(\overline{p_{j,t+1}}, c_{i,t+1}) + p' - c_{i,t+1}) * \dots$$

$$(a_i - b_{1,i}(p_{i,t+1}^{**}(\overline{p_{j,t+1}}, c_{i,t+1}) + p') + b_{2,i}\overline{p_{j,t+1}})$$

$$(B.3)$$

$$= E\pi_{i,t+1}'(c_{i,t+1}, 0, \overline{p_{j,t+1}}) + \int_0^{p'} \frac{\partial E\pi_{i,t+1}'(c_{i,t+1}, x, \overline{p_{j,t+1}})}{\partial x} dx$$

$$= E\pi'_{i,t+1}(c_{i,t+1}, 0, \overline{p_{j,t+1}}) + \int_0^{p'} (-2b_{i,1}x)dx$$
(B.5)

where the last line uses the facts that

$$\frac{\partial E\pi_{i,t+1}'(c_{i,t+1}, x, \overline{p_{j,t+1}})}{\partial x} = a_i - 2b_{1,i}p_{i,t+1}^{**}(\overline{p_{j,t+1}}, c_{i,t+1}) - 2b_{1,i}x + b_{2,i}\overline{p_{j,t+1}} + b_{1,i}c_{i,t+1},$$
(B.6)

and

$$a_i - 2b_{1,i}p_{i,t+1}^{**}(\overline{p_{j,t+1}}, c_{i,t+1}) + b_{2,i}\overline{p_{j,t+1}} + b_{1,i}c_{i,t+1} = 0,$$
(B.7)

as $p_{i,t+1}^{**}(\overline{p_{j,t+1}}, c_{i,t+1})$ is the static profit-maximizing price, so that $\frac{\partial E \pi'_{i,t+1}(c_{i,t+1}, x, \overline{p_{j,t+1}})}{\partial x} = -2b_{1,i}x$.

Therefore,

$$\frac{\partial E\pi'_{i,t+1}(c_{i,t+1}, p', \overline{p_{j,t+1}})}{\partial \overline{p_{j,t+1}}} = \frac{\partial E\pi'_{i,t+1}(c_{i,t+1}, 0, \overline{p_{j,t+1}})}{\partial \overline{p_{j,t+1}}}$$
(B.8)

$$= b_{2,i}(p_{i,t+1}^{**}(\overline{p_{j,t+1}}, c_{i,t+1}) - c_{i,t+1}) > 0.$$
(B.9)

where the final step uses the envelope-theorem as $p_{i,t+1}^{**}(\overline{p_{j,t+1}}, c_{i,t+1})$ is the static profit-maximizing price.

Condition (MT-iii): $\Pi_{13}^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) > 0.$

$$\frac{\partial \Pi^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})}{\partial p_{i,t}} = a_i - 2b_{i,1}p_{i,t} + b_{2,i}\overline{p_{j,t}} + b_{i,1}c_{i,t}$$
(B.10)

as, conditional on $\widehat{c_{i,t}}$, $p_{i,t}$ only affects period t profits. Therefore,

$$\Pi_{13}^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) = \frac{\partial \Pi^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})}{\partial p_{i,t} \partial c_{i,t}} = b_{1,i} > 0$$
(B.11)

Condition (MT-iv): $\Pi_{3}^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) = 0$ for only one $p_{i,t}$, and for this $p_{i,t}$, $\Pi_{33}^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) < 0$.

$$\Pi_{33}^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) = -2b_{1,i} < 0 \ \forall p_{i,t}$$
(B.12)

so $\Pi^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})$ will have a unique maximum in $p_{i,t}$.

Condition (MT-v): there exists k > 0 such that if $\Pi_{33}^{i,t}(c_i, \hat{c}_i, p_i, c_{j,t-1}) \ge 0$ then $\left|\Pi_3^{i,t}(c_i, \hat{c}_i, p_i, c_{j,t-1})\right| > k$. As $\Pi_{33}^{i,t}(c_i, \hat{c}_i, p_i, c_{j,t-1}) < 0$ for all $p_{i,t}$, the condition is trivially satisfied.

Therefore, based on Theorem 9, if a fully separating best response function in period *t* exists, it is uniquely characterized as $p_{i,t}^{BR}(c_{i,t}, c_{j,t-1})$ as the solution to a differential equation

$$\frac{\partial p_{i,t}^{BR}(c_{i,t},c_{j,t-1})}{\partial c_{i,t}} = -\frac{\Pi_2^{i,t}(c_{i,t},c_{i,t},p_{i,t},c_{j,t-1})}{\Pi_3^{i,t}(c_{i,t},c_{i,t},p_{i,t},c_{j,t-1})}$$
(B.13)

with a lower initial condition price $p_{i,t}^{BR}(\underline{c_i}, c_{j,t-1})$ that solves $\Pi_3^{i,t}(\underline{c_i}, \underline{c_i}, p_{i,t}^{BR}(\underline{c_i}, c_{j,t-1}), c_{j,t-1}) = 0.$

Period t Pricing Function Properties, Part I

Before discussing single-crossing, we can now prove some features of period-t pricing functions given this characterization of best responses.

Feature (T-ii): the price charged is always above the static best response price for all $c_{i,t} > c_i$.

Proof: as $\Pi_2^{i,t}(c_{i,t}, c_{i,t}, p_{i,t}, c_{j,t-1}) > 0$, and is independent of the value of $p_{i,t}$, and $\Pi_3^{i,t}(c_{i,t}, c_{i,t}, p_{i,t}, c_{j,t-1}) < 0$ for prices above the static best response price, and $\Pi_3^{i,t}(c_{i,t}, c_{i,t}, p_{i,t}, c_{j,t}) \rightarrow 0$ as $p_{i,t}$ approaches the static best response price for any $c_{i,t}$, the solution to the differential equation for a specific $c_{i,t}$ will be greater than the static best response price given $c_{i,t}$ except at c_i .

Feature (T-i(b)): the increment above the initial value is a function of $c_{i,t}$ and $c_{j,t-1}$ only, and it does not depend on $\overline{p_{j,t}}$.

Proof: the initial value solves $\Pi_3^{i,t}(\underline{c_i}, \underline{c_i}, p_{i,t}^*(\underline{c_i}), c_{j,t-1}) = 0$, i.e., it is a static best response when $c_{i,t} = \underline{c_i}$ to the expected price $\overline{p_{j,t}}$. As the numerator in the differential equation, $\Pi_2^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})$, is independent of $p_{i,t}$ and $\Pi_3^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})$ depends only on the increment of $p_{i,t}$ above the intercept, the increment depends only on $c_{i,t}$ and (possibly) $c_{j,t-1}$.¹⁴

Feature (T-iii): the effect of $c_{j,t-1}$ on the increment only comes through its effect on *i*'s belief about the distribution of $c_{j,t+1}$.

Proof: $c_{j,t-1}$ affects $\overline{p_{j,t}}$ and *i*'s period *t* belief about the distribution of $c_{j,t+1}$, which will affect *i*'s expectation of $\overline{p_{j,t+1}}$. Given T-i(b), $\overline{p_{j,t}}$ does not affect the increment. From Lemma 2 (L-ii), at the start of period t + 1, the expectation of $\overline{p_{j,t+1}}$ will depend only on $c_{j,t}$ (revealed by *j*'s period *t* price) and $\widehat{c_{i,t}}$. Therefore the only

$$p_{i,t}^{**} = \frac{a_i}{2b_{1,i}} + \frac{c_{i,t}}{2} + \frac{b_{2,i}}{2b_{1,i}}\overline{p_{j,t}}$$

¹⁴The proof of (MT-ii) shows that $\Pi_3^{i,t}$ only depends on the increment of $p_{i,t}$ above the static best response price for $c_{i,t}$ (not the initial value which is the best response for $\underline{c_i}$). However, given linear demand, static best responses are given by

so the increment of the static best response price above the static best response for $c_{i,t} = \underline{c_i}$ only depends on $c_{i,t} - \underline{c_i}$.

effect that $c_{j,t-1}$ can have on the period t increment, which is set before $p_{j,t}$ is revealed, is that it affects *i*'s beliefs about the distribution of $c_{j,t+1}$.

Feature (T-iv): (a) the pricing function is increasing and continuous in $c_{i,t}$.

Proof: (a) as $\Pi_2^{i,t}(c_{i,t}, c_{i,t}, p_{i,t}, c_{j,t-1}) > 0$ and $\Pi_3^{i,t}(c_{i,t}, c_{i,t}, p_{i,t}, c_{j,t-1}) < 0$ above the static best response price, the pricing function must be increasing in $c_{i,t}$.

Single-Crossing.

Condition (MT-vi): we need to show that, in the graph of $(\widehat{c_{i,t}}, p_{i,t})$, $\frac{\Pi_3^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})}{\Pi_2^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})}$ is either strictly increasing or decreasing in $c_{i,t}$. This amounts to showing that $\frac{\partial \frac{\Pi_3^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})}{\Pi_2^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})}}$ is either positive or negative within the graph of $(\widehat{c_{i,t}}, p_{i,t})$

$$\frac{\frac{\partial \Pi_{3}^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t},c_{j,t-1})}{\Pi_{2}^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t},c_{j,t-1})}}{\partial c_{i,t}} = \dots}{\Pi_{2}^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t},c_{j,t-1})} - \Pi_{3}^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t},c_{j,t-1})} \frac{\partial \Pi_{2}^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t},c_{j,t-1})}}{\partial c_{i,t}} - \Pi_{3}^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t},c_{j,t-1})} \frac{\partial \Pi_{2}^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t},c_{j,t-1})}}{\partial c_{i,t}} - \Pi_{3}^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t},c_{j,t-1})} \frac{\partial \Pi_{2}^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t},c_{j,t-1})}}{\partial c_{i,t}}$$

The denominator is positive. As $\frac{\partial \Pi_3^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t},c_{j,t-1})}{\partial c_i} = b_1 > 0$, and $\Pi_2^{i,t}(c_{i,t},\widehat{c_{i,t}},p_{i,t}) > 0$ the first term in the numerator is strictly positive, and does not depend on $p_{i,t}$. Recognizing that $\frac{\partial \overline{p_{j,t+1}}}{\partial c_{i,t}\partial \widehat{c_{i,t}}} = 0$,

$$\frac{\partial \Pi_{2}^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})}{\partial c_{i,t}} = \dots$$

$$\beta b_{2,i} \int \int (p_{i,t+1}^{**}(c_{i,t+1}, \overline{p_{j,t+1}}) - c_{i,t+1}) \frac{\partial \overline{p_{j,t+1}}}{\partial \widehat{c_{i,t}}} \frac{\partial \Psi_i(c_{i,t+1}|c_{i,t})}{\partial c_{i,t}} \Psi_j(c_{j,t}|c_{j,t-1}) dc_{i,t+1} dc_{j,t}.$$

 $\frac{\partial \overline{p_{j,t+1}}}{\partial \widehat{c_{i,t}}} \text{ is positive (T-iv). With linear demand, the static mark-up, } p_{i,t+1}^{**}(c_{i,t+1}, \overline{p_{j,t+1}}) - c_{i,t+1}, \text{ will decrease in } c_{i,t+1}, \text{ and given the assumptions on the densities } \Psi_i, \\ \int (p_{i,t+1}^{**}(c_{i,t+1}, \overline{p_{j,t+1}}) - c_{i,t+1}) \frac{\partial \Psi_i(c_{i,t+1}|c_{i,t})}{\partial c_{i,t}} dc_{i,t+1} < 0, \text{ but it will be bounded.}$

For prices at or above the static best response price, $\Pi_3^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1}) \leq 1$

0, but, critically, $\Pi_3^{i,t}(c_{i,t}, \widehat{c_{i,t}}, p_{i,t}, c_{j,t-1})$ must be close to 0 when $p_{i,t}$ is not too far above the static best response price. As the signaling price function is continuous and increasing in $c_{i,t}$, and is equal to the static best response price when $c_{i,t} = \underline{c_i}$, it follows that $\frac{\partial \frac{\Pi_3^{i,t}(c_i,\widehat{c_i}, p_i)}{\Pi_2^{i,t}(c_i,\widehat{c_i}, p_i)}}{\partial c_i} > 0$ when the interval $[\underline{c_i}, \overline{c_i}]$ is small enough.

Therefore, from Theorem 8, the unique fully separating best response function described above exists.

A Unique MPBE in Period t Given the Form of the Best Response Functions.

The proof so far has chosen that, given a separating pricing strategy of j, i will have a unique fully separating best response that takes the required form. We now show that, with linear demand, the pair of separating functions used by i and j, given a pair $c_{i,t-1}$ and $c_{j,t-1}$, as best responses to each other, will be unique (i.e., there cannot be more than one distinct pair of best response functions that are best responses to each other).

Recall that the only effect of a change in $\overline{p_{j,t}}$ is on the intercept of *i*'s pricing function. Therefore, holding fixed strategies in future periods, a change in *j*'s period *t* strategy only translates *i*'s best response pricing function upwards and downwards. It follows that there can only be a unique equilibrium if, for both *i* and *j*, $0 < \frac{\partial p_{i,t}^*}{\partial \overline{p_{j,t}}} < 1$. Proof: $\frac{dp_{i,t}^*(c_i)}{d\overline{p_{j,t}}} = \frac{b_{2,i}}{2b_{1,i}}$, which, given A1, is strictly greater than zero and strictly less than one, as required.

Period t Pricing Function Properties, Part II.

We can now show the remaining features of the equilibrium pricing functions.

Feature (T-i(a)): the initial values (i.e., static best response prices when $c_{i,t} = \underline{c_i}$) are continuous functions of $c_{j,t-1}$ and $c_{i,t-1}$ only.

Proof: this follows directly from the Markovian assumption as $c_{j,t-1}$ and $c_{i,t-1}$

are sufficient to determine both players' beliefs about period *t* costs, and, given Theorem 9, to uniquely determine $\overline{p_{j,t}}$.

In the following, we will denote the function that determines the initial value $g_{i,t}(c_{j,t-1}, c_{i,t-1})$. The increment above the initial value, which we will denote $f_{i,t}(c_{i,t}, c_{j,t-1})$, is a continuous function of $c_{i,t}$ and $c_{j,t-1}$ only. From T-i(b), the increment does not depend on $\overline{p_{j,t}}$.

Feature (T-iv(b)): *i*'s pricing function is continuous and strictly increasing in $\overline{p_{j,t}}$, and feature (T-iv(c)): *i*'s pricing function is continuous and strictly increasing in (*i*'s perception of) $c_{j,t-1}$.

Proof: The equilibrium price functions have the form

$$p_{i,t}^* = g_{i,t}(c_{i,t-1}, c_{j,t-1}) + f_{i,t}(c_{i,t}, c_{j,t-1})$$

where, as already shown, $g_{i,t}(c_{i,t-1}, c_{j,t-1})$ is the solution to

$$g_{i,t}(c_{i,t-1}, c_{j,t-1}) = \frac{a_i}{2b_{1,i}} + \frac{c_i}{2} + \frac{b_{2,i}}{2b_{1,i}}\overline{p_{j,t}}(c_{i,t-1}, c_{j,t-1})$$

which is increasing and continuous in $\overline{p_{j,t}}$. From the perspective of firm *i*, $\overline{p_{j,t}}$ is equal to

$$\overline{p_{j,t}} = \frac{a_j}{2b_{1,j}} + \frac{c_j}{2} + \frac{b_{2,j}}{2b_{1,j}}\overline{p_{i,t}} + \int_{\underline{c_j}}^{\overline{c_j}} f_{j,t}(c_{j,t}, c_{i,t-1})\Psi_j(c_{j,t}|c_{j,t-1})dc_{j,t}$$

where the continuity of the increment f and the conditional density $\Psi_j(c_{j,t}|c_{j,t-1})$, and the properties that (i) $f_{j,t}(c_{j,t}, c_{i,t-1})$ is increasing in $c_{j,t}$, and (ii) the integral is increasing in $c_{j,t-1}$ means that $\overline{p_{j,t}}$ is continuous and increasing in $c_{j,t-1}$, holding $\overline{p_{i,t}}$ fixed. But as $\overline{p_{i,t}}$ is also increasing, and continuous, in $\overline{p_{j,t}}$ and vice-versa, both pricing functions will also be increasing and continuous in both $c_{i,t-1}$ and $c_{j,t-1}$. Strategies in Period T.

It remains to show that strategies in the final period have a form that will lead to the type of separating equilibrium strategies described above in period T - 1. The required features are that:

- the period *T* equilibrium pricing function of firm *i* is continuous in *c*_{*i*,*T*}, *c*_{*i*,*T*-1} and *c*_{*j*,*T*-1}; and,
- the expected value $\overline{p_{i,T}}$ is increasing in $c_{j,T-1}$.

In period T, both firms will use static optimal strategies given their beliefs about their rival's previous price. Therefore

$$p_{i,T}^* = \frac{a_i}{2b_{1,i}} + \frac{c_{i,T}}{2} + \frac{b_{2,i}}{2b_{1,i}}\overline{p_{j,T}}$$

where

$$\overline{p_{j,T}} = \frac{a_j}{2b_{1,j}} + \frac{b_{2,j}}{2b_{1,j}}\overline{p_{i,t}} + \frac{E(c_{j,T}|c_{j,T-1})}{2}$$

and solving these equations simultaneously gives

$$\overline{p_{j,T}} = \frac{\left(\frac{a_j}{2b_{1,j}} + \frac{a_i b_{2,i}}{4b_{1,j} b_{1,i}}\right) + \frac{b_{2,j} E(c_{i,T}|c_{i,T-1})}{4b_{1,j}} + \frac{E(c_{j,T}|c_{j,T-1})}{2}}{\left(1 - \frac{b_{2,i} b_{2,j}}{4b_{1,i} b_{1,j}}\right)}$$

SO

$$p_{i,T}^* = \frac{a_i}{2b_{1,i}} + \frac{c_{i,T}}{2} + \frac{b_{2,i}}{2b_{1,i}} \left(\frac{\left(\frac{a_j}{2b_{1,j}} + \frac{a_ib_{2,i}}{4b_{1,j}b_{1,i}}\right) + \frac{b_{2,j}E(c_{i,T}|c_{i,T-1})}{4b_{1,j}} + \frac{E(c_{j,T}|c_{j,T-1})}{2}}{\left(1 - \frac{b_{2,i}b_{2,j}}{4b_{1,j}b_{1,j}}\right)} \right).$$

Given the form of Ψ_i and Ψ_j (A4), $p_{i,T}^*$ will be continuous in $c_{i,T}$, $c_{i,T-1}$ and $c_{j,T-1}$, and $\overline{p_{j,T}}$ is increasing in $c_{i,T-1}$, as required.

B.4 Data

This Appendix provides additional details on the data used in our empirical analysis, and some additional analyses that are not presented in the text.

B.4.1 IRI Data.

The data comes from the beer category of the IRI Academic Dataset (Bronnenberg et al. (2008)). The underlying data is at the weekly UPC-store-level from 2001 to 2011. We only use data from grocery stores.

We use different samples at different points of our analysis. When performing our demand and conduct parameter analysis, we follow MW as closely as possible (and indeed use their code as the basis for our code). When we are estimating price dynamics to calibrate our model, we use selections that we view as appropriate. For example, MW ignore sales of cans and bottles in 18-packs, which are rare for most brands. However, 18-packs account for more than 20% of sales (by volume) of the three flagship brands (Bud Light (BL), Miller Lite (ML) and Coors Light (CL)) that we use in our calibration so we do not want to exclude them. We also choose to stop our pre-JV sample at the time that the JV was announced, rather than including the period of the DOJ's investigation as, during the investigation, ML prices dropped quite dramatically.

In the following sub-sections, we detail the data selection and definitions used in the two parts of the analysis.

B.4.1.1 Data Selection for the Demand and Conduct Analysis.

We follow MW in using the following selection of data.

• **selection of markets:** 39 geographic (IRI defined) regional markets excluding (e.g., because they lack other types of data that will be used in demand estimation, or are viewed as having too few beer sales) the following markets with some stores selling beer in the data: Harrisburg/Scranton; Philadelphia; Providence RI; Tulsa; Minneapolis-St. Paul; Oklahoma City; Salt Lake City; Kansas City; New England; Pittsfield; Eau Claire, WI.

- brands: 13 brands, which are BL, ML, CL, Budweiser, Miller Genuine Draft, Miller High Life, Coors, Corona Extra, Corona Light, Heineken, Heineken Premium Light, Michelob Ultra, Michelob Light.
- pack sizes: packages of cans and glass bottles containing the equivalent of 6, 12, 24 and 30 12oz. servings. 24 and 30-packs are aggregated into a single "large" size. Prices are calculated as total dollars sold divided by volume in 12-pack equivalents.
- product: a product is a brand × pack size (6-pack, 12-pack, "large") combination.
- **time periods:** for demand and supply estimation, data from January 2005 to December 2011 is used, but months from June 2008 to May 2009, i.e., the period immediately after the JV was consummated, are excluded. Monthly data is created by allocating individual days within a week to their correct month, and assuming that sales within a week are spread equally across the days in the week, before aggregating to the monthly level.
- distances and diesel prices: we use MW's estimated distance from the brewery or port (for Heineken) to the market, measured in thousands of miles. Monthly diesel prices come from the U.S. Energy Information Administration.
- **income data:** the random coefficients models are estimated using data on household income taken from the 2005-2011 PUMS samples of the American Community Survey (ACS). We use the same samples as MW to estimate demand.
- **deflator:** when using real prices, or real diesel prices, they are deflated to January 2010 levels using the CPI-U All Urban Consumers-All Items price

index.

The following additional variables are defined:

- market size: for each market, market size is defined as 150% of the maximum of the total sales, measured in 12-pack equivalents, of all of the brands listed above plus 23 others (including popular brands such as Busch and Busch Light) in the package sizes/types that are being used. When we estimate demand using weekly data, we use an alternative definition that defines demand as 150% of the sum of the maximum sales across the stores observed in the sample that week.
- **distance measure:** the distance measure is constructed by multiplying deflated diesel prices by the driving distance from the brewery, or port in the case of Heineken, to the market.
- **demand instruments:** to estimate demand it is necessary to define instruments for a product's price and its share of volume sold amongst the products in its nest. MW use the following instruments:
 - the product's own distance measure (iv-1)
 - the sum of the distance measures for all of the products in the nest (iv-2)
 - the number of products in the nest (iv-3)
 - a dummy for domestic products after the JV (iv-4)
 - (iv-2) and (iv-3) interacted with a dummy for products produced by Miller, Coors, AB or MillerCoors
 - (iv-2) and (iv-3) interacted with a dummy for products produced by AB

When we estimate demand allowing for a flagship nest and an "other brand" nest, (iv-1) and (iv-4) are interacted with a dummy for flagship products, and the other instruments are defined at the nest level (e.g., adding over all products in the same nest, rather than all products). However, all three package sizes are available for all flagship products in all markets, so, for the flagship nest, the (iv-3) instruments are dropped due to collinearity.

B.4.1.2 Data Selection for the Calibration of Our Model.

For our calibration we depart from this selection in the following aspects.

- selection of markets: we use observations from all market-weeks where we observe the flagship brands being sold in at least 5 stores. This gives us 45 markets before the JV, although some markets do not meet the criteria in some weeks. The markets that are added back are: Eau Claire, Kansas City, Minneapolis, New England, Oklahoma City, Salt Lake City. Boston never meets the 5 store criterion after the JV so it is excluded from our estimates of post-JV price dynamics.
- **pack sizes:** packages of cans and glass bottles containing the equivalent of 6, 12, 18, 24 and 30 12oz. servings. These sizes are treated separately, but prices are converted into 12-pack equivalents.
- **time periods:** we use the months from January 2001 to October 2007 for the pre-JV period. The months after May 2009, until December 2011, are the post-JV period.

B.4.2 Additional Empirical Analyses.

We now describe several additional analyses that support the results presented in the paper.

B.4.2.1 Effects of the Joint Venture on Prices.

MW present estimates of the effects of the joint venture on prices. We present complementary estimates here, which can be compared to the price increases predicted by our calibrated model.
An observation in our analysis is a brand-market-month, where real prices are calculated at the brand level by adding up the total sales in package sizes equivalent to packs of 6, 12, 18, 24, 30 or 36 12oz. containers (we include 36packs in this regression where they are available, although they account for a small proportion of sales). The sample contains the following brands: BL, ML and CL (i.e., the domestic flagship brands), Corona Extra and Heineken which we will treat as providing controls for industry-wide shocks, as MW assume. The sample runs from 2001 to 2011, and includes the period immediately before and following the JV. We consider prices defined using all store-UPC-week observations in the appropriate sizes, and prices that are defined excluding store-UPC-week observations that are identified as being sold at temporary price reduction prices. We use both definitions as our analysis of price dynamics will use price series where price reductions are removed.

Table B.2 presents the results from six specifications that differ depending on whether price reductions are included, we use prices in levels or logs and whether brand-time trends are included. The reported coefficients are the coefficients on Post-JV dummies for the domestic flagship brands, so that they measure the increase in real prices relative to the two imported brands. The estimated price increases vary across the columns, but lie in the range from just over 40 cents to one dollar, or 3% to 6%, and the price increases are smaller when we include brand-specific time trends.

	(1)	(2)	(3)	(4)	(5)	(6)			
	\$ Price/	Log(Price/	\$ Price/	Log(Price/	\$ Price/	Log(Price/			
	12 Pack	12 Pack)	12 Pack	12 Pack)	12 Pack	12 Pack)			
	incl.	incl.	incl.	incl.	excl.	excl.			
Post-JV									
Brand Dummies									
Bud Light	0.853	0.046	0.428	0.046	0.485	0.032			
	(0.049)	(0.005)	(0.064)	(0.005)	(0.080)	(0.007)			
Miller Lite	1.024	0.065	0.415	0.045	0.492	0.034			
	(0.058)	(0.006)	(0.071)	(0.006)	(0.070)	(0.006)			
Coors Light	0.945	0.056	0.438	0.048	0.542	0.040			
	(0.060)	(0.006)	(0.068)	(0.006)	(0.076)	(0.007)			
Brand Time Trends	Ν	Ν	Y	Y	Y	Y			
Observations \mathbf{P}^2	25,740	25,740	25,740	25,740	25,740	25,740			
Γ	0.9/1	0.9/3	0.972	0.9/3	0.970	0.970			

Table B.2: Estimates of the Effects of the Joint Venture on Prices.

Notes: the reported coefficients are on domestic brand \times post-JV interactions. The brands included are those listed, plus Corona Extra and Heineken. Observations at the brand-market-month level, aggregating across packages containing the equivalent of 6, 12, 18, 24, 30 or 36 12oz. containers in cans or glass bottles. All specifications include market-brand and time period fixed effects. Standard errors in parentheses clustered on the market.





Notes: Budweiser, Michelob Ultra and Michelob Light aggregated into "Other AB"; Miller Genuine Draft and Miller High Life aggregated to "Other Miller"; Coors is "Other Coors"; Heineken and Heineken Premium Light are "Heineken" and Corona Extra and Corona Light are "Corona". Shares based on volume sold in packages equivalent to 6, 12, 18, 24, 30 and 36 12oz containers.

B.4.2.2 Market Shares Around the Joint Venture.

Our preferred demand system for the calibration assumes that there is limited substitution between the flagship domestic brands and other brands and the outside good, and that observed post-JV price increases should not reduce demand for the flagship products very much. This is consistent with some of our estimates in Table 2.5, although MW's specifications imply more substitution.

Figure B.5 shows the volume-based market shares of the different brands included in the demand analysis (for this purpose, we define market share based on the shares of all beers in the IRI data). We aggregate the non-flagship brands based on their pre-JV ownership. The main feature of the figure is that while the real prices of the flagship brands and the other domestic brands increase after the JV, the effect on brand market shares is quite limited, except that CL gained market share at the expense of ML (a change that appears unrelated to average price changes). Non-flagship Miller and AB brands do lose share after the JV, but this appears to primarily reflect a continuation of pre-JV trends. The imported brands do not appear to gain share.

		Pre-JV							
		(1)	(2)	(3)	(4)	(5)	(6)		
(1)	Bud Light	1							
(2)	Miller Lite	0.891	1						
(3)	Coors Light	0.891	0.889	1					
(4)	Budweiser	0.994	0.892	0.893	1				
(5)	Miller Genuine Draft	0.872	0.973	0.870	0.872	1			
(6)	Coors	0.804	0.812	0.916	0.807	0.804	1		
			Post-JV						
		(1)	(2)	(3)	(4)	(5)	(6)		
(1)	Bud Light	1							
(2)	Miller Lite	0.857	1						
(3)	Coors Light	0.874	0.967	1					
(4)	Budweiser	0.995	0.856	0.872	1				
(5)	Miller Genuine Draft	0.840	0.957	0.940	0.839	1			
(6)	Coors	0.825	0.934	0.959	0.824	0.916	1		

Table B.3: Cross-Brand Correlations in Prices for 12-Packs

Notes: the correlations are for brand-market-week average prices of 12-packs, before the announcement of the JV and after its consummation. Average prices are calculated including price reductions. Correlations for brands with the same owner are slightly higher if price reductions are excluded.

B.4.2.3 Price Correlations Across Brands Before and After the JV.

A significant limitation of our model is that we can only model each firm setting a single price per period. One can view our model as a representation of a more complicated problem where brewers set prices for portfolios of products, but these products can obviously be sold at different prices and these prices could move in different ways over time. In this Appendix, we report price correlations for six domestic brands and find that the prices of brands sold by the same brewer are especially correlated, and Miller and Coors prices are more correlated after the JV. This provides some comfort that viewing the brewers as choosing a single price is not too misleading.

Table B.3 reports the correlations of market-week prices of 12-packs of the flagship brands, plus Budweiser, Miller Genuine Draft and Coors, before and after the JV. It is noticeable that the prices of products with the same owner (e.g.,

BL and Budweiser) are highly correlated and that the prices of Miller and Coors products become more correlated after the JV.

The reported correlations are high partly because beers retail at different prices in different markets. We can also calculate correlations by regressing the price of one brand on the price of another brand, and market and week fixed effects. These results also show significant increases in correlations of Miller and Coors products after the JV: for example, the coefficient on the CL price when the ML price is the dependent variable increases from 0.68 before the JV to 0.84 after the JV. Patterns in the table and the regressions are similar if we use prices defined to exclude temporary price reductions.

B.4.2.4 Price Dynamics in Los Angeles and Seattle Around the JV (Real Prices).

Figure B.6: Average Real Prices (excluding sales) of 12-Packs of the Domestic Flagship Brands in Two Regional Markets Around the JV.



Notes: Averages are calculated as the total dollar sales of 12-packs at prices not identified as temporary price reductions, divided by the number of 12-packs sold. The text contains the same figure with nominal prices.

Figure B.6 repeats Figure 2.4 but with real prices, rather than nominal prices.

B.5 Testing the Supermarkup Model

MSW assume that each fiscal year, a price leader announces market-specific incentive-compatible markups (m_{mt}) , in dollars, above Nash prices that all domestic brewers should charge. Foreign brands are assumed to use static Nash pricing. This implies that, given m_{mt} , the first-order conditions for an AB product *i* are given in the following expression where $\widetilde{p}^{D} = p_{mt} - m_{mt}$ for domestic products and $\widetilde{p}^{I}(\widetilde{p}^{D})$ are Nash equilibrium prices of imported brands if domestic brewers charged \widetilde{p}^{D} :

$$p_{imt} - m_{mt} = W_{imt}\gamma + \frac{q_{imt}(\widetilde{p^{D}}, \widetilde{p^{I}}(\widetilde{p^{D}}))}{\frac{\partial q_{imt}}{\partial p_{imt}}(\widetilde{p^{D}}, \widetilde{p^{I}}(\widetilde{p^{D}}))} + \sum_{\substack{j \in AB\\ j \neq i}} \frac{\frac{\partial q_{jmt}}{\partial p_{imt}}(\widetilde{p^{D}}, \widetilde{p^{I}}(\widetilde{p^{D}}))}{\frac{\partial q_{imt}}{\partial p_{imt}}(\widetilde{p^{D}}, \widetilde{p^{I}}(\widetilde{p^{D}}))} (p_{jmt} - m_{mt} - c_{jmt}) + \nu_{imt}.$$
(B.14)

The first-order conditions for an imported product k (say a Heineken (H) product) are the standard static first-order conditions

$$p_{kmt} = W_{kmt}\gamma + \frac{q_{kmt}(p)}{\frac{\partial q_{imt}}{\partial p_{imt}}(p)} + \sum_{\substack{l \in H \\ l \neq k}} \frac{\frac{\partial q_{lmt}}{\partial p_{lmt}}(p)}{\frac{\partial q_{kmt}}{\partial p_{kmt}}(p)} (p_{lmt} - c_{lmt}) + \nu_{kmt}$$

To test the model we assume that the imported brands do use static best responses, and we test whether FOCs such as (B.14) describe the pricing of domestic producers. In particular we do this by generalizing the model to allow for a "conduct" parameter, i.e.,

$$p_{imt} - m_{mt} = W_{imt}\gamma + \frac{q_{imt}(\widetilde{p^{D}}, \widetilde{p^{I}}(\widetilde{p^{D}}))}{\frac{\partial q_{imt}}{\partial p_{imt}}(\widetilde{p^{D}}, \widetilde{p^{I}}(\widetilde{p^{D}}))} + \dots$$

$$\sum_{\substack{j \in AB \\ j \neq i}} \frac{\frac{\partial q_{jmt}}{\partial p_{imt}}(\widetilde{p^{D}}, \widetilde{p^{I}}(\widetilde{p^{D}}))}{\frac{\partial q_{imt}}{\partial p_{imt}}(\widetilde{p^{D}}, \widetilde{p^{I}}(\widetilde{p^{D}}))} (p_{jmt} - m_{mt} - c_{jmt}) + \kappa \sum_{k \in M, C} \frac{\frac{\partial q_{kmt}}{\partial p_{imt}}}{\frac{\partial q_{imt}}{\partial p_{imt}}} (p_{kmt} - c_{kmt}) + \nu_{imt}.$$

where, if the supermarkup explanation is correct, $\kappa = 0$. The intuition for the test is that if the supermarkup really is a constant markup on a static Nash price then, *controlling for supermarkup using an appropriately defined fixed effect*, price-setting should not be affected by the incremental effect that a price has on the profits of other domestic brewers. On the other hand, if there is an alternative type of deviation from Nash pricing then the estimated κ may be significantly different from zero.

B.5.1 Testing the Supermarkup Model Version 1.

We use two different implementations of the test. The first is easy-toimplement (which means that we can use it for monthly data) but relies on deviating from the MSW model so that the supermarkups only enter the FOCs linearly. Specifically, suppose that a domestic product *i* in market *m* has marginal cost c_{imt} , and that the collusive plan operates by each domestic product being priced according to static Nash best responses if its marginal costs are $c_{imt} + m'_{mt}$ rather than just c_{imt} .¹⁵ One interpretation would be that the domestic firms act as if they have to pay higher marginal retailing costs, a form of tacit collusion that might be hard to detect. In this case, the MW first-order condition for an AB product is simply

$$p_{imt} = W_{imt}\gamma + m'_{mt} + \frac{q_{kmt}(p)}{\frac{\partial q_{imt}}{\partial p_{imt}}(p)} + \sum_{\substack{j \in AB\\ j \neq i}} \frac{\frac{\partial q_{jmt}}{\partial p_{imt}}(p)}{\frac{\partial q_{imt}}{\partial p_{imt}}(p)} (p_{jmt} - m'_{mt} - c_{jmt}) + \nu_{imt}.$$
(B.15)

and, when we generalize to allow for conduct parameters that should be equal to zero if the supermarkup model is correct, the estimating equation are

$$p_{mt} = W_{mt}\gamma + m_{mt} - \left(\Omega_{mt}(\kappa) \circ \left[\frac{\partial s_{mt}(p_t, \theta^D)}{\partial p_{mt}}\right]\right)^{-1} s_{mt}(p_{mt}) + \nu_{mt}.$$
 (B.16)

The first-order condition has the nice feature that the level of demand and the demand derivatives only depend on observed prices, and the supermarkup enters linearly. This theory can be tested by including domestic market-fiscal year fixed effects to control for m'_{mt} , and testing if conduct parameters equal zero. We

¹⁵The effects on cross-market incentive-compatibility constraints would determine the form of mark-ups that colluding firms prefer to use.

use the domestic rival distance measures, their ξ s (averaged across their portfolios either before or after the JV) and interactions of these variables as excluded instruments that identify the conduct parameters.

B.5.2 Testing the Supermarkup Model Version 2.

Testing version 1 is not the same as testing MSW's supermarkup model because, in that model, m_{mt} enters the first-order conditions non-linearly. Testing the MSW model therefore requires estimating non-linear market-fiscal year fixed effects for domestic products, where, for different values of the fixed effect, we re-evaluate the demand derivative matrix and resolve for the Nash prices that the imported brands would charge in response. This potentially creates a very large computational burden, especially when using the RCNL demand model, even if we use quarterly data. To make estimation feasible, we therefore proceed as follows.

First, we estimate all of the parameters, including the conduct parameters and the linear parameters, separately for each fiscal year, so that we are only estimating 40 (39 supermarkup fixed effects and 1 conduct parameter) nonlinear parameters at a time. We report the conduct coefficients for 2005/6, 2006/7, 2009/10, and 2010/11 fiscal years (i.e., two full fiscal years before the JV and after the JV), but we also estimate them for the partial fiscal years in the sample, and the estimated coefficients are similar, but less precise. We expect separate estimation to reduce the econometric efficiency and power of our test, as will the fact that we do not restrict the supermarkups to be consistent with cross-market incentive compatibility constraints on the domestic brewers. However, in practice, our estimates of the conduct parameters are precise.

Second, and more importantly, rather than recomputing demand derivatives, import best responses prices and inverting matrices to back out implied marginal costs many hundreds of times during estimation, we use interpolation from values that are pre-computed. Specifically, before estimation, we compute implied marginal costs for each observed product-market-quarter observation on a grid of supermarkups ($m_{mt} = \{0, 0.25, 0.50, ..., 6\}$) and conduct parameters ($\kappa = \{0, 0.05, ..., 1.1\}$) then use cubic interpolation to get the required values during estimation (restricting the supermarkups and conduct parameters to lie within these ranges). As a result, the computational burden for each function evaluation involves the computation of around 6,000 cubic interpolations.

As usual, one might be skeptical about a researcher's ability to simultaneously estimate 40 nonlinear parameters. However, in practice, MATLAB's fmincon algorithm works very well on this problem even when it uses numerical derivatives, and it delivers the same estimates from a range of different starting values. The conduct parameter estimates are also comfortingly consistent with those from testing version 1 of the supermarket model.

Bibliography

- Aguirregabiria, V. and Mira, P. (2007). Sequential estimation of dynamic discrete games. *Econometrica*, 75(1):1–53.
- Allen, J., Clark, R., and Houde, J.-F. (2014). The effect of mergers in search markets: Evidence from the canadian mortgage industry. *American Economic Review*, 104(10):3365–96.
- Andrews, I., Gentzkow, M., and Shapiro, J. M. (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics*, 132(4):1553–1592.
- Argentesi, E. and Filistrucchi, L. (2007). Estimating market power in a two-sided market: The case of newspapers. *Journal of Applied Econometrics*, 22(7):1247–1266.
- Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics*, 37(3):668–691.
- Armstrong, M. and Wright, J. (2007). Two-sided markets, competitive bottlenecks and exclusive contracts. *Economic Theory*, 32(2):353–380.
- Arnosti, N., Johari, R., and Kanoria, Y. (2014). Managing congestion in decentralized matching markets. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 451–451.
- Ashenfelter, O., Hosken, D., and Weinberg, M. (2014). Did robert bork understate the competitive impact of mergers? evidence from consummated mergers. *Journal of Law and Economics*, 57(S3):S67–S100.
- Ashenfelter, O. C., Hosken, D. S., and Weinberg, M. C. (2015). Efficiencies brewed: Pricing and consolidation in the us beer industry. *RAND Journal of Economics*, 46(2):328–361.
- Asker, J., Fershtman, C., Jeon, J., and Pakes, A. (2020). A computational framework for analyzing dynamic procurement auctions: The market impact of information sharing. *RAND Journal of Economics*, 51(3):805–839.

- Athey, S. and Bagwell, K. (2008). Collusion with persistent cost shocks. *Econometrica*, 76(3):493–540.
- Athey, S., Bagwell, K., and Sanchirico, C. (2004). Collusion and price rigidity. *Review of Economic Studies*, 71(2):317–349.
- Athey, S., Levin, J., and Seira, E. (2011). Comparing open and sealed bid auctions: Evidence from timber auctions. *The Quarterly Journal of Economics*, 126(1):207–257.
- Azar, J., Schmalz, M. C., and Tecu, I. (2018). Anticompetitive effects of common ownership. *Journal of Finance*, 73(4):1513–1565.
- Baker, J. B. and Farrell, J. (forthcoming). Oligopoly coordination, economic analysis, and the prophylactic role of horizontal merger enforcement. *University of Pennsylvania Law Review*.
- Bardey, D. and Rochet, J.-C. (2010). Competition among health plans: A two-sided market approach. *Journal of Economics & Management Strategy*, 19(2):435–451.
- Berry, S. T. and Haile, P. A. (2014). Identification in differentiated products markets using market level data. *Econometrica*, 82(5):1749–1797.
- Bodoh-Creed, A., Boehnke, J., and Hickman, B. R. (2016). How efficient are decentralized auction platforms? *Becker Friedman Institute for Research in Economics Working Paper*, (2016-22).
- Bonatti, A., Cisternas, G., and Toikka, J. (2017). Dynamic oligopoly with incomplete information. *Review of Economic Studies*, 84(2):503–546.
- Bresnahan, T. F. (1982). The oligopoly solution concept is identified. *Economics Letters*, 10(1-2):87–92.
- Bresnahan, T. F. (1989). Empirical studies of industries with market power. *Handbook of Industrial Organization*, 2:1011–1057.
- Bronnenberg, B. J., Kruger, M. W., and Mela, C. F. (2008). The iri marketing dataset. *Marketing Science*, 27(4):745–748.
- Caillaud, B. and Jullien, B. (2003). Chicken & egg: Competition among intermediation service providers. *RAND journal of Economics*, pages 309–328.
- Caminal, R. (1990). A dynamic duopoly model with asymmetric information. *Journal of Industrial Economics*, 38(3):315–333.
- Cho, I. and Kreps, D. (1987). Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102(2):179–221.
- Clarke, E. H. (1971). Multipart pricing of public goods. *Public choice*, pages 17–33.

- Corts, K. S. (1999). Conduct parameters and the measurement of market power. *Journal of Econometrics*, 88(2):227–250.
- David, H. (1997). Augmented order statistics and the biasing effect of outliers. *Statistics & probability letters*, 36(2):199–204.
- De Angelis, D. and Grinstein, Y. (2014). Performance terms in ceo compensation contracts. *Review of Finance*, 19(2):619–651.
- De los Santos, B., Hortaçsu, A., and Wildenbeest, M. R. (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American economic review*, 102(6):2955–80.
- Doraszelski, U. and Jaumandreu, J. (2013). R&d and productivity: Estimating endogenous productivity. *Review of Economic Studies*, 80(4):1338–1383.
- Ericson, R. and Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of economic studies*, 62(1):53–82.
- Farrell, J. and Baker, J. B. (2021). Natural oligopoly responses, repeated games, and coordinated effects in merger analysis: A perspective and research agenda. *Review of Industrial Organization*, pages 1–39.
- Farronato, C., Fong, J., and Fradkin, A. (2020). Dog eat dog: Measuring network effects using a digital platform merger. *NBER Working Paper*, (28047).
- Fershtman, C. and Pakes, A. (2012). Dynamic games with asymmetric information: A framework for empirical work. *Quarterly Journal of Economics*, 127(4):1611–1661.
- Fradkin, A. (2017). Search, matching, and the role of digital marketplace design in enabling trade: Evidence from airbnb. *Matching, and the Role of Digital Marketplace Design in Enabling Trade: Evidence from Airbnb (March 21, 2017).*
- Garmon, C. (2017). The accuracy of hospital merger screening methods. *The RAND Journal of Economics*, 48(4):1068–1102.
- Groves, T. (1973). Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631.
- Hortaçsu, A. and Syverson, C. (2004). Product differentiation, search costs, and competition in the mutual fund industry: A case study of s&p 500 index funds. *The Quarterly Journal of Economics*, 119(2):403–456.
- Horton, J. J. (2019). Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution. *Management Science*, 65(8):3518–3540.

- Hubbard, T. P. and Paarsch, H. J. (2013). On the numerical solution of equilibria in auction models with asymmetries within the private-values paradigm. In Schmedders, K. and Judd, K. L., editors, *Handbook of Computational Economics*, volume 3, pages 37–115. Elsevier: North-Holland.
- Jeziorski, P. (2014). Estimation of cost efficiencies from mergers: Application to us radio. *The RAND Journal of Economics*, 45(4):816–846.
- Judd, K. (1998). Numerical methods in economics. Technical report, The MIT Press.
- Katz, M. L. (1991). Game-playing agents: Unobservable contracts as precommitments. *RAND Journal of Economics*, 22(3):307–328.
- Kaya, A. (2009). Repeated signaling games. *Games and Economic Behavior*, 66(2):841–854.
- Krasnokutskaya, E., Song, K., and Tang, X. (2020). The role of quality in internet service markets. *Journal of Political Economy*, 128(1):75–117.
- Kreps, D. M., Milgrom, P., Roberts, J., and Wilson, R. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory*, 27(2):245–252.
- Laffont, J.-J., Ossard, H., and Vuong, Q. (1995). Econometrics of first-price auctions. *Econometrica: Journal of the Econometric Society*, pages 953–980.
- Lau, L. J. (1982). On identifying the degree of competitiveness from industry price and output data. *Economics Letters*, 10(1-2):93–99.
- Levin, D. and Smith, J. L. (1994). Equilibrium in auctions with entry. *The American Economic Review*, pages 585–599.
- MacKay, A. and Miller, N. (2019). Estimating models of supply and demand: Instruments and covariance restrictions. Technical report, Harvard University.
- Mailath, G. (1987). Incentive compatibility in signaling games with a continuum of types. *Econometrica*, 55(6):1349–1365.
- Mailath, G. J. (1988). An abstract two-period game with simultaneous signaling: Existence of separating equilibria. *Journal of Economic Theory*, 46(2):373–394.
- Mailath, G. J. (1989). Simultaneous signaling in an oligopoly model. *Quarterly Journal of Economics*, 104(2):417–427.
- Mailath, G. J. and von Thadden, E.-L. (2013). Incentive compatibility and differentiability: New results and classic applications. *Journal of Economic Theory*, 148(5):1841–1861.

- Marra, M. et al. (2019). Pricing and fees in auction platforms with two-sided entry. Technical report, Sciences Po Departement of Economics.
- Maskin, E. and Tirole, J. (1988). A theory of dynamic oligopoly, ii: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica*, pages 571–599.
- Mester, L. J. (1992). Perpetual signalling with imperfectly correlated costs. *RAND Journal of Economics*, 23(4):548–563.
- Milgrom, P. and Roberts, J. (1982). Limit pricing and entry under incomplete information: An equilibrium analysis. *Econometrica*, 50(2):443–459.
- Miller, N. H., Sheu, G., and Weinberg, M. C. (2020). Oligopolistic price leadership and mergers: The united states beer industry. Technical report, Georgetown University.
- Miller, N. H. and Weinberg, M. C. (2017). Understanding the price effects of the millercoors joint venture. *Econometrica*, 85(6):1763–1791.
- Murphy, K. J. (1999). Executive compensation. *Handbook of Labor Economics*, 3:2485–2563.
- Nakamura, E. and Zerom, D. (2010). Accounting for incomplete pass-through. *Review of Economic Studies*, 77(3):1192–1230.
- Nevo, A. (1998). Identification of the oligopoly solution concept in a differentiated-products industry. *Economics Letters*, 59(3):391–395.
- O'Brien, D. P. and Salop, S. C. (1999). Competitive effects of partial ownership: Financial interest and corporate control. *Antitrust Law Journal*, 67:559.
- Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.
- Ordover, J. A. (2007). Coordinated effects in merger analysis: An introduction. *Columbia Business Law Review*, pages 411–436.
- Pakes, A. and McGuire, P. (1994). Computing markov-perfect nash equilibria: Numerical implications of a dynamic differentiated product model. *RAND Journal of Economics*, pages 555–589.
- Peters, C. (2009). Evaluating the performance of merger simulation: Evidence from the u.s. airline industry. *Journal of Law and Economics*, 49(2):627–649.
- Porter, R. H. (1983). A study of cartel stability: The joint executive committee, 1880-1886. *Bell Journal of Economics*, 14(2):301–314.
- Rochet, J.-C. and Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the european economic association*, 1(4):990–1029.

- Rochet, J.-C. and Tirole, J. (2006). Two-sided markets: a progress report. *The RAND journal of economics*, 37(3):645–667.
- Roddie, C. (2012). Signaling and reputation in repeated games, i: Finite games. Technical report, University of Cambridge.
- Rotemberg, J. J. and Scharfstein, D. S. (1990). Shareholder-value maximization and product-market competition. *Review of Financial Studies*, 3(3):367–391.
- Rysman, M. (2004). Competition between networks: A study of the market for yellow pages. *The Review of Economic Studies*, 71(2):483–512.
- Salz, T. (2020). Intermediation and competition in search markets: An empirical case study. Technical report, National Bureau of Economic Research.
- Samuelson, W. F. (1985). Competitive bidding with entry costs. *Economics letters*, 17(1-2):53–57.
- Seim, K. (2006). An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics*, 37(3):619–640.
- Shapiro, C. (1986). Exchange of cost information in oligopoly. *Review of Economic Studies*, 53(3):433–446.
- Sklivas, S. D. (1987). The strategic choice of managerial incentives. *RAND Journal of Economics*, 18(3):452–458.
- Smith, A. (2008). Indirect inference. *The New Palgrave Dictionary of Economics,* 2nd Edition.
- Stigler, G. J. (1964). A theory of oligopoly. *Journal of Political Economy*, 72(1):44–61.
- Sweeting, A., Roberts, J. W., and Gedge, C. (2020). A model of dynamic limit pricing with an application to the airline industry. *Journal of Political Economy*, 128(3):1148–1193.
- Toxvaerd, F. (2008). Strategic merger waves: A theory of musical chairs. *Journal of Economic Theory*, 140(1):1–26.
- Toxvaerd, F. (2017). Dynamic limit pricing. *RAND Journal of Economics*, 48(1):281–306.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37.
- Vives, X. (2011). Strategic supply function competition with private information. *Econometrica*, 79(6):1919–1966.

- Weintraub, G. Y., Benkard, C. L., and Van Roy, B. (2010). Computational methods for oblivious equilibrium. *Operations research*, 58(4-part-2):1247–1265.
- Weyl, E. G. (2010). A price theory of multi-sided platforms. *American Economic Review*, 100(4):1642–72.