The Value of Bargaining in Online Platform Markets

by

Xu Zhang

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy (Business Administration) in the University of Michigan 2018

Doctoral Committee:

Professor Francine Lafontaine, Co-Chair
Professor Puneet Manchanda, Co-Chair
Professor Daniel Ackerberg
Professor Kanishka Misra

Xu Zhang

xuzh@umich.edu

ORCID iD: 0000-0002-7640-0225

DEDICATION

This dissertation is dedicated to my parents and my husband

for their endless love, support and encouragement,

and to my daughter

for making me stronger, better and more fulfilled
than I could have ever imagined.

ACKNOWLEDGMENTS

I am grateful to everyone who has made a difference in this so far most transformative and defining period of my life - the PhD journey. Specifically, I am eternally thankful for my advisors Francine Lafontaine and Puneet Manchanda, who have provided constant inspiration, held high standards all the time, supported me for both professional and personal growth, and made my completion of this dissertation possible.

The power of a role model is enormous - Francine has served this role for me. I am impressed by her pursuit of rigor in research from our very first meeting to today. Whether it is a tiny germ of idea or a well-developed research project, she always provides constructive feedback and asks incisive questions to make me think deeper and broader. She makes me truly believe that a great scholar is not a person who is only able to do good research, but more importantly, is a person who cares about the impact of the research - to help and contribute to society at a greater level. She helped me to understand the heights and depths of my aspirations and dreams.

There are "random shocks" in life that shape the whole rest of one's life afterwards - getting to know Puneet is no doubt one of the most precious "random shocks" for me. Puneet opened a whole world of new possibilities for me, which allowed me to discover the beauty in marketing and to develop true passion in both career and life. I sincerely admire his mentoring model of "unconditional love" to all of his students - a template which I hope to achieve in my scholarly journey in the future. As my dissertation co-chair, Puneet always knows what I need at any particular moment - whether it be sharing wisdom and experience

during moments of struggle, or standing by me during moments of doubt, or "colorful" encouragements during moments of stress, or "sweet" congratulations after milestone goals, or priceless "babbling" at any time. He inspires me to see the bigger picture in all kinds of questions and pursue what really matters in life.

In addition to my dissertation chairs, I am also truly indebted to Dan Ackerberg and Kanishka Misra for serving on my dissertation committee and providing valuable insights on my research; to Junhong Chu, for her enthusiasm and excitement about ideas, which has always served as energy and motivation boosters to help me overcome research challenges. I am also thankful to a number of other wonderful mentors who shared their perspectives and shaped my development as a researcher and a teacher, including Achyuta Adhvaryu, Anocha Aribarg, Rajeev Batra, Tom Buchmueller, Ying Fan, Fred Feinberg, Tammy Feldman, Kyle Handley, Yan Huang, Tom Lyon, Scott Masten, Sarah Miller, Yesim Orhun, Eric Schwartz, Jagadeesh Sivadasan, S. Sriram, Valerie Suslow, and Marek Zapletal. Since I first entered the PhD program till today, I have benefited greatly from the help and support offered by Brian Jones, Judy Dyer, Chris Feak, and Brenda Imber, who eased my transition between different cultures, different languages, and even different life roles. Earlier in my life, I was fortunate to have some great mentors as well - particularly Miaojie Yu and Ling Li - whose expertise, encouragement, and selfless investment inme guided me on the right path.

Many thanks go to a number of past and present students in the Ross PhD program, their friendship and company made this journey much more delightful. I am grateful for the "secret treasures", ranging from past exams, to mistakes to avoid, to job market advice, passed down by senior students, including Danial Asmat, Sutirtha Bagchi, Sasha Brodsky, Yao Cui, Christina DePasquale, Linda Hagen, Xiaoyang Li, Yao Lu, Michael Olabisi, Mike Palazzolo, Adi Pattabhiramaiah, Anyan Qi, Sarah Stith, Chenyu Yang, Bo Zhao, and KC Zheng. I am also thankful for many valuable experiences, whether happy or sad, shared with fellow students, including Rebecca Chae, Jin Woo Chang, Rahul Chhabra, Yulia Chhabra, Koustav De, Yue Fang, Tong Guo, Zhaohui Jiang, Evgeny Kagan, Murray Lei, Frank Li, Prashant

Rajaram, Sara Ryoo, Paula Suh, Longxiu Tian, Dana Turjeman, Tiffany Vu, Guihua Wang, Chen Zhang, and Dan Zhao. Among them, I would like to give special thanks to Iris Wang for her willingness to share a lovely apartment, a clean office, countless delicious meals, and even a spare car key with me, and to her generous husband Nan Jiang for no complaint on any of the above sharing. I also owe tremendous thanks to my dear long-time and faraway friends - Chunhui Jiao, Ding Han, Sisi Liu, Pan Wang, Xiaonan Wang, Xin Wang, Fan Xiang, and Jing Zhao - for their unwavering trust in me and their 24/7 availability for me.

Lastly, never ending thanks must go to my family, who has always been a pillar of strength for me. I am privileged to be the daughter of inspirational and supportive parents, the wife of a loving and caring husband, and the mother of a happy and lovely daughter. To anyone I have omitted, please accept my humblest apology and know that I truly appreciate your guidance and support.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

Bargaining is an important pricing mechanism, prevalent in both the online and offline worlds. However, little empirical work on the value of bargaining in markets exists, primarily due to the lack of real-world bargaining data. In the first essay, "Meet Me Halfway: The Value of Bargaining", I leverage the availability of rich, transaction-level data on bargaining outcomes on an online platform to quantify the value of bargaining for sellers, buyers, and the platform. I incorporate the decision to bargain, the bargaining realization, and the purchase decision into a structural model, and I perform counterfactual analyses to show the value of allowing bargaining on the platform. I do this by disallowing bargaining, which means that all sellers on the platform must use a fixed-price mechanism. I find that sellers' profits are higher under the fixed-price only regime. The benefits are heterogeneous across sellers, however, with sellers with low reputation levels, high detailed seller ratings, and non-promotion products benefiting more. I also show that buyers' bargaining cost savings are economically significant. Thus, my findings suggest banning bargaining is beneficial for both buyers and sellers, and therefore for a social planner as well. I provide some reasons for why bargaining still exists on the platform despite the above findings. Finally, I show that the results are robust to the assumptions and can be replicated in different product categories.

In the second essay, "Negotiation Pricing on a Health Platform Market: Bringing Hospitals and Patients Together", I study the bargaining between a business-to-consumer platform and the business-side of the platform. The platform connects hospitals with patients who seek health checkups. As a business entity, the platform negotiates with hospitals on the depth of price discount, online or offline payment, and clearing cycle (per transaction, weekly, or monthly). I investigate how the consumer network size and the hospital network size affect the hospitals' willingness to sign contracts and join the platform and the three key negotiation outcomes. I find that both network sizes have positive significant effects on hospitals' willingness to sign contracts and on the three negotiation outcomes, making hospitals more likely to accept terms that are more favorable to the platform. To further investigate the hospital side, I separate public and private hospital network sizes and find their effects on negotiation are heterogeneous. I discuss the findings and explore potential mechanisms behind these findings.

CHAPTER I

Meet Me Halfway: The Value of Bargaining

1.1 Introduction

Bargaining is an important pricing mechanism all over the world. While it has been prevalent so far in offline settings, it has recently been gaining popularity in the online world as well, especially on platforms such as Amazon, eBay, and Alibaba.[1] Much theoretical research has compared the bargaining and the fixed-price mechanisms, with different (market) assumptions yielding different insights (see Arnold and Lippman, 1998; Bester, 1994; Desai and Purohit, 2004; Riley and Zeckhauser, 1983; Wang, 1995). However, there is very little empirical research addressing this question, i.e., whether the bargaining or the fixed-price mechanism is better, partly because it is very hard to observe data on bargaining processes and outcomes due to the nature of the interaction - typically carried out verbally in person over a very short time period. Thus, little is known about the relative benefits and costs of bargaining versus fixed-price mechanisms in real world settings. Intuitively, bargaining offers advantages to sellers as they can price discriminate based on buyers' unobserved willingness to pay, unobserved bargaining intention, and unobserved bargaining ability (I use the terms "buyers" and "consumers" interchangeably throughout the paper). In order to obtain gains from bargaining, sellers are likely to post higher prices, driving away some buyers as a result.

[1] Amazon introduced a "Make an Offer" option on certain product categories in 2014, eBay has employed a "Best Offer" option since 2005, and Alibaba's major e-commerce platform - Taobao - has included a bargaining mechanism since its inception in 2003.

Thus, the gains from allowing bargaining could be offset by the losses in revenue from buyers who do not engage with the seller at all. From a social planner's point of view, the benefits of more completed transactions could be offset by the bargaining cost incurred by buyers. As a result, it is hard to determine ex-ante or in general which pricing mechanism is more beneficial.

In this paper, I attempt to answer this question by quantifying the relative value of bargaining vis-à-vis fixed-price using rich, transaction-level data on bargaining outcomes in a very large digital marketplace (platform) in China, combined with survey data from Chinese buyers on their propensity to bargain and their expectations of success. I do this by building a structural model of consumer demand and sellers' pricing decisions where I allow for the existence of price discrimination and bargaining costs. Specifically, the demand model captures the processes inherent in a transaction where bargaining is possible, including the decision to bargain, the bargaining realization, and the purchase decision. The supply model captures the fact that sellers take the bargaining outcome into consideration when setting the posted price.

The demand model involves three stages. In the first stage, a consumer decides whether to bargain by comparing the expected utilities with and without bargaining. Bargaining cost is a key factor in determining the attractiveness of the bargaining mechanism (Wang, 1995; Jindal and Newberry, 2016). I define the bargaining cost as the cost a buyer expects to incur during the bargaining process, including a psychological cost and expected time and effort costs.[2] The bargaining cost is identified in the data by variation across shopping occasions in the decision to bargain and the expected bargaining outcome. In the second stage, I model the bargaining realization process between a seller and a buyer. While the Nash bargaining equilibrium has been widely used in the literature (e.g., Beckert et al., 2016;

[2]Research in social psychology has shown that the bargaining process can be associated with the possibility of embarrassment or a loss of self-esteem, thus imposing a psychological cost (Rubin and Brown, 1975). In my approach, all costs due to bargaining - psychological, time, etc. - result in an anticipated loss of utility by a buyer prior to initiating bargaining.

Crawford and Yurukoglu, 2012; Draganska et al., 2010; Ellickson et al., 2017; Gowrisankaran et al., 2015; Grennan, 2013; Jindal and Newberry, 2016), its validity hinges on the correct specification of each party's disagreement payoff. An alternative way to model the bargaining process is to use an extensive-form bargaining model (e.g., Keniston, 2011; Larsen, 2015), but that requires detailed alternating-offer data. Given (1) the complications that arise in incomplete information settings (see Ausubel et al. (2002) for a detailed review), (2) the inability to observe each party's disagreement payoff, and (3) the lack of alternating-offer bargaining data, I employ a reduced-form approach involving a two-part model. The first part captures whether bargaining succeeds and the second part captures the discount conditional on success. This approach avoids imposing strong assumptions and can better matches the data. In the third stage, I use a standard consumer discrete choice approach to model the purchase behavior, including the use of a control function to address potential endogeneity issues. Finally, on the supply side, I assume that sellers are playing a static pricing game to maximize profits and obtain the profit-maximizing prices when allowing for buyer bargaining.

The results from my analysis provide rich insights. First, I find the mean bargaining cost to be 9 yuan (or about US $1.5 at the 2012 exchange rate of 6.3 yuan to a dollar), with a range from 3 to 16 yuan across Chinese provinces. To put this into context, this is close to the minimum hourly wage in China that ranges from 11 to 20 yuan. I also find that the variation in the average bargaining costs is related to the economic development with the costs being higher in more developed provinces, reflecting that buyers in these provinces place a higher value on their time. Note that the estimated bargaining cost plays an essential role in analyzing the (dis)advantage of the bargaining mechanism in terms of social welfare. As for the success of bargaining and the realized discount (conditional on bargaining success), I find that both are higher if the posted price is high, no promotion is available, the seller has a low(er) reputation, the buyer has more shopping experience, and the transaction is a repeat purchase. These findings add to the literature on the determinants of bargaining outcomes (e.g. Ayres and Siegelman, 1995; Backus et al., 2016; Draganska et al., 2010; Ellickson et al., 2017; Meza and Sudhir, 2010; Morton et al., 2011; Shelegia and Sherman, 2015), by specifically highlighting several characteristics that are important for online settings but not readily available in offline ones, such as shopping experience and reputation levels. Finally, I find that a 1% increase in posted price on average leads to a 3.4% decrease in conversion rate (the proportion of online traffic to the seller site that results in a purchase). The decrease in the conversion rate comes more from the decrease in transactions made at the posted price than those made at a bargained price. This finding has intuitive appeal as a higher posted price is more likely to scare away consumers who would not bargain. Consumers also value the existence of promotions, high seller reputation levels, and high detailed seller ratings when making the purchase decision.

Using the demand model estimates along with the seller's profit-maximizing price, I perform counterfactual analysis to derive the value of bargaining in this online marketplace. I do this by disallowing bargaining, i.e., "forcing" sellers to move to a fixed-price mechanism. I find that the average posted price decreases by 1.1% but the average transaction price remains about the same after the move. Interestingly, I find that the average conversion rate increases by about 1%. These results are primarily driven by the fact that the gains from the price discrimination are offset by the losses from scaring away consumers by the higher posted price under the bargaining mechanism.

Thus, sellers are overall better off under a fixed-price than a bargaining mechanism, though the average magnitude of improvement is modest. The magnitude of improvement varies across sellers, with "weaker" sellers (e.g., those with low reputation levels) benefiting more by moving to the fixed-price mechanism. For about 1% of sellers, this increase in profits exceeds 5%.

The total annual increase in gross merchandise volume or GMV[3] for the marketplace is

[3]GMV is a term used in online retailing to indicate a total sales dollar value for merchandise sold through a particular marketplace over a certain time frame. Popular online retail sites like eBay, Overstock.com,

around 200 million yuan (30 million dollars) in the focal category - cellphone, representing a 1% increase. The saved bargaining costs for buyers per day are at 6% of the daily GMV. Thus, the results suggest that banning bargaining is beneficial not only from an individual seller's point of view, but also from the buyers' and a social planner's point of view. Note that due to data limitations, I are not able to identify sellers' bargaining costs. Thus, the estimated benefits from the ban on the bargaining mechanism are conservative and the actual benefits can be even larger for both sellers and the social planner. Given these results, a natural question to ask is why sellers have not moved to a fixed-price mechanism in this marketplace. While I discuss this in detail later, the short answer lies in the history of the marketplace's evolution, cultural norms in China, and the very large number of sellers. Finally, I show that the results are robust to my model assumptions and structure and not idiosyncratic to one product category (I am able to replicate them qualitatively for one other product category).

My paper differs from previous empirical work on bargaining (Beckert et al., 2016; Grennan, 2013; Keniston, 2011; Huang, 2012; Jindal and Newberry, 2016) in several important ways, and thus makes the following contributions to both the marketing and economics literature on bargaining. First, previous studies have treated transactions with zero bargained discount as equivalent to no-bargain transactions. However, in this paper, I highlight the difference between the no-bargain transactions and the failed-bargain transactions. This is a more realistic description of the bargaining mechanism and is critical in correctly estimating the effect of bargaining costs. Second, the identification of bargaining costs in the previous literature has been based on functional form assumptions or has depended on stringent data requirements. In the paper, I instead combine secondary (transaction) data with primary (survey) data, which enables me to identify bargaining costs. Third, the two-part model to describe the bargaining realization process makes it easier to not impose the strong assumption of perfect information under the Nash bargaining framework. The advantages of

Google Checkout, Alibaba.com, evenmattress.co.uk, and Avenida.com commonly use this term in place of sales or revenue.

this approach include less stringent data requirements and more flexibility, making it more applicable in most other bargaining settings where the detailed alternating-offer data are unavailable. Lastly, I not only examine the effect of pricing policy change on individual players in the market, but also at the platform level. Given the prevalence of platform markets in e-commerce worldwide, my analysis can serve as a template to investigate the impact of different pricing mechanisms.

The rest of the paper is organized as follows. I describe the institutional setting and the data in §1.2. §1.3 describes the model. The identification and estimation of the demand side model are detailed in §1.4 while the supply side is described in §1.5. The results are discussed in §1.6 and the counterfactual analysis and the robustness checks in §1.7. I conclude in §1.8.

1.2 The Institutional Setting and the Data

Taobao.com is a Chinese leading e-commerce platform founded by the Alibaba Group in 2003. It aims to provide a platform for individual entrepreneurs and small businesses to do business anywhere (Chu and Manchanda, 2016). By 2014, Taobao had more than 500 million registered buyers, over 7 million registered sellers, over 60 million average daily unique visitors, and served customers in more than 200 countries. Its online product listings total over 800 million and there are 72 million transactions each day - more than Amazon and eBay combined. As a result, Taobao.com's market share of e-commerce in China is over 80%.[4]

[4] *About Taobao.* Retrieved from https://www.taobao.com/about/intro.php?spm=1.22376.225943.2 on 11/25/2015. *The Success of Taobao on the Chinese Internet.* Retrieved from http://daxueconsulting.com/success-of-taobao-on-chinese-internet/ on 11/25/2015. *10 Reasons why Alibaba Blows Away Amazon and eBay.* Retrieved from https://www.forbes.com/sites/walterloeb/2014/04/11/10-reasons-why-alibaba-is-a-worldwide-leader-in-e-commerce on 11/25/2015.

1.2.1 The Bargaining Feature

Sellers on Taobao.com typically employ a unique pricing mechanism in that fixed-price and bargaining coexist. Specifically, each product has a posted price - customers can purchase a product immediately at the posted price. However, there is a facility on the seller's site that allows customers to initiate bargaining via a free Skype-like online chatting service called Aliwangwang. All buyers and sellers get an Aliwangwang account automatically when they register on Taobao (and the tool can be used within a web browser or via a mobile app). Figure 1.1 presents a screenshot of a cellphone item page with the price display and the location of the online chatting tool (highlighted). Aliwangwang provides a convenient communication channel between buyers and sellers. It allows users to instantly transmit text, images, and files. Prior to a purchase, a buyer can use the tool to get product information or bargain with a seller. After a purchase, a buyer can inquire about product delivery, exchange and return policy. Taobao buyers are usually accustomed to chatting with sellers as they consider their purchases and are well aware of the possibility of using Aliwangwang to carry out bargaining.[5] If a bargaining process results in a discount, a seller will adjust the price to the agreed-upon price after a buyer adds the product to her shopping cart.

As can be seen from the above, the setup of the bargaining process on Taobao implies that there are no predetermined bargaining rules. Buyers are free to initiate a bargain (or not), carry out alternating-offer bargaining or take-it-or-leave-it bargaining or use any other bargaining heuristic. In addition, I do not observe the Aliwangwang chat history for transactions. As a result, it is hard for me to impose bargaining theory primitives or equilibrium concepts. I therefore specify a flexible model to capture the bargaining outcome as a function of a rich set of related variables.

[5] In the Taobao Consumer Survey, described in §1.2.3, I find that 81% of buyers are aware of the possibility of bargaining. In addition, a search on the biggest Chinese search engine - baidu.com - for the term "Taobao Jiangjia" (Taobao Bargaining) returns around 2,090,000 links. Among the sellers in the sample, 97% have an average online chat waiting time less than 3 minutes.

1.2.2 Transaction Data

The main (secondary) data include a random sample of transactions on cellphones from January 1, 2012 to May 25, 2013 on Taobao.com. The product - cellphone, is particularly suited to study the comparison between the fixed-price and bargaining mechanisms on the e-commerce platform for two reasons. First, bargaining is well accepted in most small brick-and-mortar cellphone stores in China. Thus, the habit of bargaining is likely to be transferred to the online market as well. Second, the cellphone category is very popular and the product is a relatively high-ticket item. The 2012 revenue in the cellphone category was \$6.4 billion, accounting for 5% of the total revenue on Taobao, which makes the cellphone the 3rd largest category out of a total of 113 categories on Taobao. The sample consists of 39,625 transactions made by 24,181 buyers with 8,965 sellers.[6]

For each transaction, I observe detailed attributes of each cellphone e.g., brand, model, memory size, screen size, camera resolutions, carrier compatibility, etc. As Taobao.com does not employ a universal product code (UPC), I extract product attributes from each transaction entry to identify a product. Specifically I identify a product using information on brands and models. For certain products with memory capacity options, like iPhones, I also employ memory sizes to facilitate the product identification process. For those cellphone models with less than 200 transactions, I use brands as an analysis unit. The final sample includes 58 unique products.[7] The top five brands and their market shares are Samsung (18%), Nokia (15%), HTC (9.5%), Apple (9%), and Sony (6.5%).

Table 1.1 presents the summary statistics of the transaction sample. From Figure 1.1,

[6]Less than 5% of transactions show more than one cellphone unit being purchased. As these "bulk" buyers are likely to be different from individual buyers, I exclude them from the sample.

[7]To make the product definition clearer, I provide a few examples. Under Apple, there are five unique products (numbers of transactions are in parentheses) - iPhone4S 32G (1,304), iPhone4 32G (1,094), iPhone5 32G (511), iPhone3GS 32G (336), and a composite product including five less popular models (331). To test the sensitivity of the results to this product definition where I aggregate less popular models to the brand level, I use truncated regressions to predict conditional discount amounts with and without the aggregation. The mean difference between the two predictions is less than 2%. As the bargaining costs are estimated at the province level, I believe the results are not sensitive to this product definition.

I see an item page has a "List Price," a "Posted Price," and a "Promotion Indicator." The Posted Price plays the role of the fixed price in these transactions (note that it is also the highlighted price). If the seller is running any kind of promotion, the Promotion Indicator is highlighted and Posted Price includes the discount. The Posted Price varies a great deal around a mean of 1,263 yuan (about $200 dollars). Sellers employ promotions frequently - on average, 64.5% of the products are offered under a promotion.

A successful bargain occurs if the transaction price is less than the posted price. Based on this metric, 16% of all transactions are associated with a successful bargaining incidence.[8] For these transactions, the bargaining discount amount is captured by the difference between the Posted Price and the transaction price. The mean bargaining discount amount (conditional on success) is 170 yuan, representing 13% of the posted price. The proportion of sellers who have ever offered a bargaining discount is high. Among the sellers with at least 20 observed transactions in the sample, 87% are observed to have offered bargaining discounts in at least one transaction.[9] This suggests that bargaining is a transaction-specific phenomenon rather than a seller-specific phenomenon, leading me to use transactions as the unit of analysis.

Overall, these data patterns suggest that the platform features facilitate bargaining, both buyers and sellers on this platform are aware that they can bargain, and the majority do engage in this activity. The proportion of successful bargained transactions is large enough to enable me to use these data to compare the value of fixed price and bargaining mechanisms.

For each transaction, I also create a repeat purchase indicator, defined as 1 if the focal transaction is at least the second interaction between the seller and the buyer. This information may not be very accurate as I do not have information prior to the sample period.

[8]According to my conversion with the company, the difference between the transaction price and the posted price is an accurate measure of the realized bargaining amount. Buyers could bargain for additional services like free shipping. However, as 95% of the items in the sample are listed with free shipping, this is unlikely to be a concern. Also, most deliveries are within three days, ruling out the concern on shipping speed variation. I also test for this formally by including the distance between the seller and the buyer in the estimation and found that there was no change in the results.

[9]I cannot say for sure if the remaining sellers did not ever offer a discount as they may have done so outside the time window or the random sample of the data.

But this variable is of great interest for the bargaining outcomes and can provide insights to other product categories. About 11% of the transactions are repeat purchases. Additionally, I also calculate the product age, defined as the number of years between the product launch date and the purchase date, to capture the newness of a product and a seller's opportunity cost of not selling the product.

Cellphone sellers on average have three years of selling experience on Taobao. One of the most prominent seller characteristics is the seller "reputation level." Consistent with many other e-commerce platforms, Taobao has adopted a reputation system to alleviate the information asymmetry between buyers and sellers. The seller reputation system on Taobao is based on buyer feedback (and is very similar to the one used on eBay). After each transaction, a buyer has up to 15 days to provide a positive (coded as +1), neutral (coded as 0), or negative (coded as -1) rating, where a positive rating is used when no feedback is provided (see Figure 1.2 for a screenshot). The platform then uses the logarithm of the cumulative feedback to compute the seller's reputation level. However, buyers see a discrete version of this continuous reputation spreading over 20 levels. Each level is represented by a series of well-known symbols (see Figure 1.3) shown on each product page (Figure 1.1). This reputation level plays an important role in conveying information to buyers as it is given a lot of prominence. Given that 99.3% of feedback is positive, a mean reputation level of 8.4 implies roughly 5,000 past transactions for an average cellphone seller.

In addition to the seller reputation level, a seller is also rated on three other dimensions via a five-point scale. These dimensions are "item as described," "service level," and "consignment speed," respectively. There are two differences between the seller reputation level and these detailed seller ratings. First, the seller reputation level is a cumulative measure covering the time period since the seller started selling on the platform, while the detailed seller ratings are rolling measures using only customer feedback in the past four weeks. Second, when a customer does not provide feedback, the default rating for the seller reputation level is recorded as positive but left blank for the detailed seller ratings. Due to these differences, the seller reputation level and the seller detailed ratings provide buyers with information on different aspects of the seller performance. I therefore include both measures in the analysis. Since the three dimensions in the detailed seller ratings are highly correlated, I use the mean of the three measures (mean = 4.8 out of 5). The correlation between the two reputation measures used is -0.16 in the sample. In addition to the two reputation measures, I also have each seller's repeat purchase rate in the past half year. On average, 7.7% of consumers have shopped at a store more than once.

On buyers' side, the most relevant buyer characteristic is the "shopping experience level," which is calculated in a manner similar to the seller reputation level as it is based on the cumulative feedback from the sellers on past transactions. The average buyer shopping experience level is 4.2, suggesting 91-150 past transactions. I also have the location of each buyer (and seller) at the province level and information on gender and age for a subsample of buyers. The average age among buyers is about 29 years old. I see a 70%-30% male-female split in the sample. For those readers who are interested, I report major summary statistics by gender in Table 1.2 and find no statistically significant difference in either the products purchased or the bargaining outcomes across gender.

Finally, for each seller, though I do not observe individual site visits that do not result in purchases by a buyer, I do observe the number of unique visitors and the number of unique purchasers on a four-week rolling basis. I use the ratio between the two to define the conversion rate and to get a measure on the proportion of the no-purchase visits at the seller/product combination level.

1.2.3 Taobao Consumer Survey

One of my goals is to estimate bargaining costs and distinguish failed-bargain transactions from no-bargain transactions. To achieve this, it is critical to have information on bargaining intention and perception of bargaining success likelihood among Taobao consumers for parameter identification. Thus, I supplement the principal transaction data with a Taobao Consumer Survey. I conducted the survey with Taobao customers from mainland China via www.sojump.com, the largest online survey website in China. The survey ran for 10 days (December 20-29, 2016) and I obtained 1,009 responses from 31 mainland provinces.[10] In the survey, in addition to general knowledge about the practice of bargaining on Taobao, I also asked three specific questions related to bargaining: (1) "Would you bargain with a seller if you are going to buy a cellphone priced at 1,500 RMB?", (2) "How likely would you expect to succeed in bargaining?" and (3) "How much discount would you expect to get if bargaining succeeds?"

I use the first question to infer the distribution of consumers' bargaining intention.[11] There are three options to choose from: A. Yes; B. No; C. Maybe. I define a variable indicating a consumer chooses "sure to bargain" as $I(\textit{Certainly Bargain}) = I(\textit{Option} = A)$, and a variable indicating a consumer may bargain as $I(\textit{May} + \textit{Certainly Bargain}) = I(\textit{Option} = A \textit{ or } C)$. Using $I(\textit{Certainly Bargain})$ and $I(\textit{May} + \textit{Certainly Bargain})$, I calculate a lower bound $Pr_l(b = 1)$ and an upper bound $Pr_u(b = 1)$ of the average bargaining intention among Taobao consumers in each province.

The second question provides a measure on the perceived success rate conditional on bargaining among Taobao consumers. There are ten options to choose from, ranging from 10% to 100%. I find that on average consumers believe that the success rate conditional on bargaining is 49%. Since the expected gain from bargaining is the product of the success rate and the expected discount, this number can directly affect a consumer's expected bargaining

[10]I also carried out a survey in 2015 as a pilot with a much smaller number of respondents. The findings of that survey are very similar to those from the 2016 survey, increasing my confidence in the use of the survey data.

[11]In addition to the posted price, the bargaining intention would also depend on the real value of the product, like brands and models, and seller characteristics. However, given the large number of products and continuous measure of seller characteristics in the transaction sample, it is impossible to measure the bargaining intention for each seller/product combination using the survey. As a result, I purposely design the survey question using vague terms without specifying the product and seller characteristics. Under the assumption that respondents' perceived distribution of product and seller characteristics are consistent with the true distribution, the implied bargaining intention at the aggregate level should be consistent with the true bargaining intention.

gain.

I use the third question to test if the survey sample is representative of consumers in the transaction data. By comparing respondents' answers on their expected bargaining discount amount conditional on success in the survey (mean = 165.5 yuan, s.d. = 221.5) and the observed bargaining discount amount conditional on success in the transaction sample (mean = 170.3 yuan, s.d. = 295), I believe the survey respondents represent Taobao consumers well and Taobao consumers on average form reasonable expectations on the conditional bargaining discount amounts. Figure 1.4 also plots a comparison of the age distribution in the transaction data and in the survey. The similarity between the two once again verifies that the two samples are comparable.

Table 1.3 summarizes the survey responses. The mean lower bound of bargaining intention is 54.3%, the mean upper bound is 73.6%, and the perceived success rate conditional on bargaining is 49%. In estimation, both bargaining intention and perceived success rate conditional on bargaining are used at the province level. The province-level lower bounds of bargaining intention range from 35.4% to 81.8%, the upper bounds range from 58.0% to 95%, and the perceived success rates conditional on bargaining range from 31% to 62%.[12] A detailed summary breakdown by gender is provided in Table 1.4. Even though I do not see any difference in the bargaining outcomes in the transaction sample, the survey suggests that men are more optimistic than women in the predicted bargaining outcomes, and are more likely to initiate bargaining, which is consistent with the findings in **?**.

1.3 The Consumer Bargaining and Purchasing Model

I first construct an empirical model of consumers' bargaining and purchasing decisions to describe the demand side of the market. The key objective of my model is to estimate

[12]Note that Ningxia, Qinghai, and Tibet provinces suffer from small sample problems in both the survey and the transaction sample. I group them together to minimize measurement errors.

consumers' bargaining costs and expected bargaining gains. The building block of the model is an individual purchase occasion, which proceeds in three stages, as shown in Figure 1.5. In Stage I, a consumer decides whether to bargain by comparing the expected bargaining gain and her inherent bargaining cost. If she decides to bargain, she proceeds to Stage II (bargaining realization), while if she decides not to bargain, she skips Stage II and proceeds to Stage III (purchase decision). In Stage II, the buyer bargains with the seller to ask for an additional price discount. Subsequently, in Stage III, the buyer decides whether to purchase the product after seeing the realized bargaining discount. Due to the lack of the clickstream data, I are not able to model the search behavior explicitly. Instead, I use proxy measures to control for two common search processes employed by consumers, as detailed later. I describe the model by working backwards from the last stage of the demand model. To facilitate reading, I provide a list of notation definitions in Table 1.5.

1.3.1 Stage III: Purchase Decision

Following Jindal and Newberry (2016), I assume that a consumer knows the exact product she is interested in.[13] On e-commerce platforms, seller characteristics, like seller reputation and seller average ratings, are almost always used to infer product quality by consumers. Thus, I define a shopping occasion as a consumer considering the product that she is interested in being sold by a specific seller, or a product-seller combination. A consumer purchases at most one product in a shopping occasion. The indirect utility consumer i derives from buying product k from seller j is

$$u_{ijk} = \delta_k - \beta_p p_{ijk} + x_{jk}\beta_x + \varepsilon_{ijk} \tag{1.3.1}$$

[13] I assume consumers choose between the focal product and an outside option, instead of specifying a richer consideration set, because of a lack of search data. Given that the defined market shares abstract away from the search process, as detailed in §1.4.4, this assumption should not bias my estimates.

where δ_k is the consumer's average intrinsic preference for product k, and p_{ijk} denotes the transaction price, which can be either the posted price or the bargained price if the bargaining discount amount is positive. x_{jk} is a vector of seller/product characteristics, including seller reputation level, detailed seller rating, site age, and promotion indicator. I include product age and repeat purchase indicator as controls. β_p captures consumers' price sensitivity, and β_x captures their preferences for seller/product characteristics.

With different search processes, a buyer can have different outside options, and therefore different purchase decisions. In order to control for the influence of the search behavior and the resulting competition among sellers, I create two proxy measures to capture two common search processes. The first common search rule is that a buyer searches across sellers for a specific product, for which I use the number of sellers who are selling the same product in that month as a control. The other common search rule is that a buyer has not decided on a product yet but wants to make a purchase from a trustworthy seller (requirements on reputation levels), for which I use the number of sellers with the same reputation level in that month as a control. For brevity, these two proxy measures are part of the vector x_{jk}.

The term ε_{ijk} is the consumer's idiosyncratic shock that is specific to each shopping occasion. As product fixed effects are included, the error term no longer contains the unobserved time-invariant product attributes; thus, it does not suffer from the usual endogeneity problem. However, the error term may include unobserved seller characteristics and unobserved time-varying product attributes, which could still cause an endogeneity problem. To solve this problem, I employ the control function approach (Petrin and Train, 2010). The idea behind the control function approach is to add an extra variable in the utility function to condition out the "bad" variation in the error term that is not independent of the endogenous variable - price.

Specifically, the control function approach posits that the posted price $\bar{p}_{jk}$ can be written as

$$\bar{p}_{jk} = g(x_{jk}, z_j, \mu_{jk}) \tag{1.3.2}$$

where x_{jk} denotes exogenous variables including seller characteristics and product characteristics directly entering the utility function; z_{jk} represents the instrumental variables that affect the price but do not affect utility directly (I discuss the specific variables operationalized in §1.4.4); and μ_{jk} is the unobserved factor that affects the price and is potentially correlated with ε_{ijk}. This correlation is the source of the dependence between $\bar{p}_{jk}$ and the error term ε_{ijk}.

Given that the exogenous variables and the instrumental variables are independent of both μ_{jk} and ε_{ijk}, hence conditional on μ_{jk}, $\bar{p}_{jk}$ is independent of ε_{ijk}. This is the key to use the control function approach to resolve the endogeneity problem; I first estimate a proxy for μ_{jk}, and then estimate parameters conditional on this proxy variable. If the control function is approximated as linear in μ_{jk}, the utility can be formulated as follows:

$$u_{ijk} = \delta_k - \beta_p p_{ijk} + x_{jk}\beta_x + \beta_\mu \mu_{jk} + \tilde{\varepsilon}_{ijk} \tag{1.3.3}$$

where $\tilde{\varepsilon}_{ijk}$ is independent of all the explanatory variables. β_μ is the parameter of the control function to be estimated. The specification of the utility function is completed with the introduction of an outside option which is the no purchase option, i.e., the consumer decides to not purchase the focal product from the seller. The utility of the outside option is normalized to be $u_{ij0} = \tilde{\varepsilon}_{ij0}$.

Consumers make purchase decisions by comparing the utility of the focal product and the outside option. Conditional on μ_{jk}, the probability that consumer i chooses to purchase is equal to

$$Pr_{ijk} = \int I(u_{ijk} > u_{ij0}) f(\tilde{\varepsilon}_{ijk}) d\tilde{\varepsilon}_{ijk} \tag{1.3.4}$$

where $I(\cdot)$ is the indicator function and $f(\tilde{\varepsilon}_{ijk})$ is the density of the error term, assumed to be distributed as i.i.d. Type-I extreme value. If there is no bargaining, the above probability of purchase equals the market share. However, with the inclusion of the bargaining process, consumers face different prices when they make purchase decisions. As a result, the market share is an integral of the realized prices, either the posted price or the bargained price, over the potential pool of consumers. Formally, the market share is a modified version of BLP (Berry et al., 1995), as given by

$$m_{jk} = \int Pr_{ijk}(p_{ijk}) f(p_{ijk}) dp_{ijk} \tag{1.3.5}$$

where $f(p_{ijk})$ is the realized price distribution, which is derived after the bargaining decision and the bargaining realization. Given assumptions on the distribution of $f(\tilde{\varepsilon}_{ijk})$, and after modeling the first two stages, I can compute the integral for the market share.

1.3.2 Stage II: Bargaining Realization

In Stage II, if a consumer decides to bargain, the consumer and the seller enter into a bargaining process through online chatting. A commonly-used concept for bargaining outcome is Nash equilibrium. However, the validity of Nash equilibrium depends on the correct specifications of each party's disagreement payoff, which is unavailable in my context. Also, the prediction power of the Nash equilibrium in the outcomes of buyer-seller bargaining is questionable (**?**). Previous literature also uses extensive form to model the bargaining process under incomplete information settings, but it requires alternating-offer data. Due to the lack of canonical bargaining models with incomplete information and complications with multiple equilibria, I choose to flexibly specify the bargaining outcome as a function of seller, buyer, and product characteristics, as well as price and promotions, without assuming any specific bargaining mechanism, similar to Shelegia and Sherman (2015).

The bargaining outcome has two components: whether it is successful conditional on the decision to bargain and the size of the discount conditional on success. As the discount has a mass point at zero (either no bargain or no success), a natural choice of the model is Tobit I model. However, the Tobit I model is very restrictive because it requires the relative effects of factors to be the same in affecting the two components of the bargaining outcome.

To overcome this restriction, I specify the bargaining outcome as a two-part model (Cragg, 1971). The advantage of this specification is that it is more flexible, fits the data better, and offers clear economic interpretation of the parameters. In the first part, conditional on bargaining ($b_{ijk} = 1$), the success s_{ijk} follows a probit model:

$$p(s_{ijk} = 1|b_{ijk} = 1) = \Phi(\frac{x_{ijk}\gamma}{\sigma_s}) \tag{1.3.6}$$

where Φ is the cumulative distribution function of the standard normal. x_{ijk} represents a vector of seller, buyer, and product characteristics, and the two proxy measures for the two common search processes. Note that x_{ijk} is different from x_{jk} in eq (1.3.1) in the purchase stage in that the subscript ijk represents not only seller/product characteristics but also buyer characteristics, as the bargaining process is realized between a seller and a buyer while the market share is defined for a seller only. x_{ijk} includes the posted price and product fixed effects for brevity, so x_{ijk} essentially represents $[x_{ijk}, \bar{p}_{ijk}, productFE]$. Product fixed effects can partially capture the marginal cost incurred by each seller. The parameter vector γ represents the effects of the explanatory variables on the success rate of bargaining incidence. σ_s is the standard deviation of the unobservable υ_s in the success equation.

In the second part, I use truncated normal regression to model the discount amount conditional on success ($s_{ijk} = 1$). As the discount amounts conditional on success have a positively skewed distribution, I use a logarithm transformation. Specifically, I assume a latent variable d^*_{ijk} following the distribution of:

$$d^*_{ijk} = x_{ijk}\theta + \upsilon_d \qquad \upsilon_d \sim N(0, \sigma_d) \tag{1.3.7}$$

$$d_{ijk} = exp(d^*_{ijk}) - 1 \qquad \text{if } d^*_{ijk} > 0 \tag{1.3.8}$$

I only observe a positive discount amount d_{ijk} when d^*_{ijk} is greater than zero. x_{ijk} is the same set of explanatory variables as in equation (1.3.6), and the parameter vector θ represents the effects of the explanatory variables on the discount amount. σ_d is the standard deviation of the unobservable υ_d in the discount amount equation. In the two-part model, υ_s and υ_d can be correlated and the correlation, though not being explicitly estimated, does not affect the consistency of the estimates (Belotti et al., 2015). If one uses a Type II Tobit model, this correlation can be theoretically estimated but it may not be well identified (Wooldridge, 2010). This model nests the standard Tobit I model as a special case when the truncated normal regression model is assumed for the observations of positive bargaining discounts.

1.3.3 Stage I: bargaining decision

I assume consumers are rational in the market. When a consumer browses the product she is interested in, she first makes a bargaining decision by weighing her expected utility if choosing to bargain and that if choosing not to bargain. Thus, a consumer's bargaining decision b_{ijk} follows:

$$b_{ijk} = \begin{cases} 1 & \text{if } E[\text{utility if bargain}] > E[\text{utility if not bargain}] \\ 0 & \text{otherwise} \end{cases} \tag{1.3.9}$$

where $E[\text{utility if bargain}] = E[max\{u_{ijk}(\bar{p}_{jk} - d_{ijk} + c_i), u_{ij0} - \gamma c_i\}]$, i.e. the possible maximum utility of purchasing at the post-bargain price or the utility of no purchase after bargaining, and $E[\text{utility if not bargain}] = max\{u_{ijk}(\bar{p}_{jk}), u_{ij0}\}$, i.e., the possible maximum utility of purchasing at the posted price or the utility of no purchase. Note that bargaining costs only affect a consumer's bargaining decision but not her purchase decision. For rational consumers, after a bargaining process ends, no matter if it succeeds or fails, bargaining costs are sunk, so they should not affect the purchase decision.

1.4 Identification and Estimation of the Demand

As explained before, to overcome the data limitations, I supplement the transaction sample with an auxiliary survey data set. In order to use the combination of the transaction-level and the survey data to the best advantage, I combine several estimation techniques including truncated regression, control function, generalized method of moments, and simulated maximum likelihood. As a result, the estimation algorithm follows a series of distinct steps (shown in Figure 1.6):

Step one: I first estimate the discount amount conditional on success using the truncated regression model (equation 1.3.7), and then use the model estimates to predict expected discount amount conditional on success for each shopping occasion regardless of the actual bargaining status, $E[d_{ijk}|s_{ijk} = 1], \forall ijk$.

Step two: Using the expected discount amount conditional on success, combined with the survey-based province-level bargaining intention and perceived success rate conditional on bargaining, I recover the lower and upper bounds of the province-level average bargaining costs, denoted as c_{zl}, c_{zu} respectively, using the method of moments (equation 1.3.9).

Step three: Once calculated, the estimated expected discount amount and the estimated bargaining costs imply the probability of bargaining. With the bargaining probabilities and the observed success conditional on bargaining from the survey, I recover the parameters governing the bargaining success function using the simulated maximum likelihood approach (equation 1.3.6).

Step four: For each seller-product combination, I calculate the market share by integrating over potential shopping occasions made by a potential pool of consumers at different transaction prices. The primitives in the utility function (equation 1.3.3) are estimated by minimizing the distance between the predicted and the observed market shares using the generalized method of moments. I use the control function to deal with the endogeneity problem.

A brief summary of identification is provided in Table 1.6. I now describe the estimation and identification strategies for each step in detail.

1.4.1 Step one: discount amount conditional on success

As specified earlier, the logarithm of the discount amount follows a truncated regression model. The observations with positive bargaining discounts are used in the regression. The likelihood function for the truncated normal regression is:

$$l_d = \prod_{\{ijk:d^*_{ijk}>0\}} \frac{\phi(\frac{d^*_{ijk}-x_{ijk}\theta}{\sigma_d})}{1-\Phi(\frac{-x_{ijk}\theta}{\sigma_d})} \tag{1.4.1}$$

where x_{ijk} includes the logarithm of the posted price, seller, buyer, and product time-varying characteristics, and product fixed effects. d^*_{ijk} is the latent bargaining discount amount. As I use consumer-level data, the aggregate unobserved shocks are unlikely to cause any systematic correlation between the price and the error term at the micro level. Thus, following the previous literature using consumer-level data, I assume the price is exogenously given when a consumer initiates a bargaining process (Nevo, 2000). The implicit assumption here is that the discount amount conditional on success is not systematically different between purchasers, whose transactions are observed, and non-purchasers, whose transactions are unobserved. At first glance, this assumption may appear to be a strong one as the discount amount for purchasers on average should be greater than that for non-purchasers. However, the fact that the expected bargaining discount amount in the survey is about the same as the realized bargaining discount amount in the transaction sample suggests that this assumption may not be that strong. Moreover, from an estimation point of view, the consequence of such an assumption is similar to that when buyers have incorrectly calibrated beliefs about bargaining outcomes. In §1.7.2, I show that my results are robust to this assumption.

Once equation (1.4.1) is estimated, I predict the expected discount amount conditional on success for each shopping occasion:

$$E[d^*_{ijk}|s_{ijk} = 1] = x_{ijk}\hat{\theta} + \hat{\sigma}_d\lambda(x_{ijk}\hat{\theta}/\hat{\sigma}_d) \tag{1.4.2}$$

where $\lambda(\cdot)$ is the inverse Mills ratio defined as the ratio of the normal probability density function to the normal cumulative distribution. $\hat{\theta}$ and $\hat{\sigma_d}$ are consistent estimates obtained from the truncated regression model. I calculate the expected discount amount conditional on success by transforming the latent variable $E[d_{ijk}|s_{ijk} = 1] \approx exp(E[d^*_{ijk}|s_{ijk} = 1]) - 1$. Given the standard deviation of d^*_{ijk} is small, this approximation is appropriate.

1.4.2 Step two: bargaining costs

In the first stage of the model, rational consumers make bargaining decisions by weighing the expected utilities between choosing to bargain and not to bargain. However, as described earlier, a data challenge is that I only observe the realized bargaining discounts, not bargaining decisions. That is, I know a consumer chooses to bargain if I see a positive discount, but I do not know if a consumer chooses to bargain when I see a zero discount, which can result from either a no bargain decision or a failed bargain process.

In order to address this gap in the transaction data, I use the consumer survey data (described earlier in §1.2.3). Given how I frame the survey questions, the bargaining decision is simpler in the survey than in a real shopping occasion. In the real shopping occasion, consumers need to compare the expected utilities between choosing to bargain or not, while in the survey, respondents only need to compare the expected bargaining gain and the bargaining costs and put aside the purchase decision. Thus, the bargaining decision becomes $b_{ijk} = I(E[d_{ijk}|b_{ijk} = 1] > c_i)$. As such, bargaining costs are identified by consumers' bargaining decisions across shopping occasions with varying expected bargaining gains. Using the survey, I calculate the *province-level* lower bounds and upper bounds of bargaining intention. As bargaining intention is negatively correlated with the bargaining cost, the lower bound of the bargaining intention $Pr_l(b = 1)$ implies an upper bound for the average bargaining cost and the upper bound of the bargaining intention $Pr_u(b = 1)$ implies a lower bound for the average bargaining cost in each province.

Consumers make bargaining decisions by comparing the expected bargaining gain and the bargaining cost. The expected bargaining gains are calculated by multiplying the perceived success rate conditional on bargaining obtained from the survey and the expected discount amount conditional on success calculated in step one. This implicitly defines the set of shopping occasions that lead a consumer to choose to bargain. Formally, let this set be

$$B(c_i; x_{ijk}) = \{ijk|E[d_{ijk}|x_{ijk}] > c_i\} \tag{1.4.3}$$

Since there are very few multiple transactions made by one consumer and no individual-level bargaining intention is available in the transaction sample, the individual-level bargaining costs are unidentifiable. Instead, I take the advantage of the upper and lower bounds of the province-level bargaining intention among consumers to derive the lower and upper bounds of bargaining costs at the province level. Of course, if I had the individual buyers' bargaining intention and demographic information, I would have been able to estimate the individual-level bargaining cost, and include consumers' demographics into the bargaining realization model. In this sense, my model is flexible enough to incorporate correlations between bargaining costs and the discount amount via consumers' observed characteristics. Further, I use the province-level perceived success rate conditional on bargaining to proxy for that of an average individual in that province. With the assumption that Taobao consumers' bargaining intention and perceived success rate conditional on bargaining in the survey are good proxies for those in the transaction sample, I form the following two moment conditions:

$$Pr[b = 1|c_{zl}] = \int I(ijk \in B(c_{zl}; x_{ijk}))f(x_{ijk})dx_{ijk} \tag{1.4.4}$$

$$Pr[b = 1|c_{zu}] = \int I(ijk \in B(c_{zu}; x_{ijk}))f(x_{ijk})dx_{ijk} \tag{1.4.5}$$

where c_{zl} and c_{zu} are lower and upper bounds of average bargaining costs in province z respectively. They are estimated by matching the above moments calculated based on the transaction sample to those observed in the survey (the sample distribution of x_{ijk} is assumed to be consistent with the population distribution).

The probabilities outlined above are crude frequency simulators with a property of discontinuity, which can cause problems in estimation. To ensure smooth convergence, I use a kernel-smoothed frequency simulator (Mcfadden, 1989) with a univariate survival function as a kernel

$$S(w_{ijk}) = \frac{1}{1 + exp(-hw_{ijk})} \tag{1.4.6}$$

where $w_{ijk} = E[d_{ijk}|x_{ijk}] - c_i$, the difference between the expected bargaining gain and the bargaining cost. h is a tuning parameter in the kernel function, which is set to 0.1. As $w_{ijk} \to \infty$, $S(w_{ijk}) \to 1$. That is, a consumer will for sure bargain if the net expected bargaining gain is large enough. The parameters are estimated using the generalized method of moments.

1.4.3 Step three: bargaining success indicator

After recovering the upper and lower bounds of province-level bargaining costs in step two, I assume that the bargaining cost of consumer i from province z is a random draw from $Uniform(c_{zl}, c_{zu})$. I now briefly discuss the identification of the parameters in the bargaining success function. Assume that I observe multiple shopping occasions where the expected bargaining gains are the same and the bargaining costs follow the same distribution, i.e., the average bargaining intentions are the same. Then the difference between the probabilities of observed success rate can be attributed to observed factors, such as seller characteristics and buyer characteristics.[14] Thus, the effects of the observed factors on the success rate are identified. The likelihood of observing a bargaining success can be written as

$$\begin{aligned} l_s &= \prod_{ijk} l_{s,ijk} \\ &= \prod_{ijk} Pr(s_{ijk}=1|b_{ijk}=1)Pr(b_{ijk}=1) \\ &= \prod_{ijk} \Phi(x_{ijk}\gamma/\sigma_s)Pr(E[d_{ijk}|x_{ijk}]-c_i>0) \\ &= \prod_{ijk} \Phi(x_{ijk}\gamma/\sigma_s)\int I(E[d_{ijk}|x_{ijk}]-c_i>0)f(c_i)dc_i \end{aligned} \tag{1.4.7}$$

The underlying intuition of this likelihood function is as follows. The probability of observing a bargaining success equals the probability of bargaining times the probability of success conditional on bargaining. The fourth equality is an application of a crude frequency simulator. Similar to step two, I use a kernel-smoothed frequency simulator to ensure smoothness of the likelihood function. The parameters (γ, σ_s) are estimated using simulated maximum likelihood.

1.4.4 Step four: purchase utility

In the purchase stage, the unknown parameters are intrinsic product preferences δ_k, marginal utility of income β_p, and seller/product characteristics preference β_x. The parameters are identified by consumers' purchasing patterns between the inside good and the outside good. However, I do not observe individual consumers who visited the site but walked away without purchase. Thus, I use seller-level market shares for each shopping occasion to infer the

[14] In the estimation I treat consumers as myopic, i.e., they only compare the expected bargaining gain with the bargaining cost but not the expected utilities. I check the robustness of my results to this assumption when I compare the results from the analyses with sophisticated consumers versus naive consumers in §1.7.2. There is no substantive change in the main results.

utility preferences (note since there is only one focal product, market share thus equals the conversion rate). Specifically, I observe the number of unique visitors and the number of unique purchasers for each seller in that month. I assume that one in five unique visitors is a serious shopper, thus the market size for each shopping occasion is defined as the number of unique visitors divided by five. Though five is an arbitrarily chosen number here, the results on the price elasticities and the policy simulations would not be affected as these analyses are all in relative terms (the results are same if I choose two or ten). The conversion rate is defined as the ratio of the number of unique purchasers to the market size. The average conversion rate is 10.02%. I assume that the seller-level conversion rate is a reasonable proxy for the seller-product-level conversion rate. Given that less than 5% sellers sell product across categories, I think this is a reasonable assumption. As I use aggregate market shares in estimating the parameters in the utility function, though product fixed effects are included in the model and can absorb any correlation between the price and the unobserved product characteristics, the correlation between the price and the unobserved seller characteristics and time-varying product characteristics may still cause endogeneity problems and yet need to be controlled for. For example, a seller's positive word-of-mouth or an attractive webpage design may drive up the price and the conversion rate at the same time. If I fail to control for those unobserved variables, I may underestimate the price coefficient.

To solve the price endogeneity problem, I use the control function approach. I first regress the endogenous variable (the posted price) on the exogenous and instrumental variables. The instrumental variable is constructed based on a measure of how many pieces of positive feedback (distance) a seller still needs in order to reach the next reputation level, which is denoted as $DistNextLevel$. This instrument is valid for two reasons. First, $DistNextLevel$ is correlated with the posted price and second, it does not directly affect a consumer's purchase behavior. As reputation level generates substantial returns, sellers have a strong incentive to lower their prices to boost transaction volume when they are close to the next reputation level (Fan et al., 2016; Zhong, 2016). This incentive generates a positive correlation between the instrument and the posted price. The first-stage R-squared is 0.68, with F-statistics for the weak IV test greater than 18.5, suggesting the relevance condition is valid (Table 1.7).[15] For a consumer to obtain the information on a seller's $DistNextLevel$, she needs to get information on both the specific threshold for each reputation level and the seller's current cumulative feedback score. As neither piece of information is easily accessible to consumers, $DistNextLevel$ is unlikely to affect a consumer's purchase decision directly.

The probability of making a purchase is:

$$
\begin{aligned}
m_{jk} &= \int Pr_{ijk}(p_{ijk}) f(p_{ijk}) dp_{ijk} \\
&\approx Pr(u_{ijk} > u_{ij0} | b_{ijk} = 0) Pr(b_{ijk} = 0) \\
&\quad + Pr(u_{ijk} > u_{ij0} | s_{ijk} = 0) Pr(s_{ijk} = 0 | b_{ijk} = 1) Pr(b_{ijk} = 1) \\
&\quad + Pr(u_{ijk} > u_{ij0} | s_{ijk} = 1) Pr(s_{ijk} = 1 | b_{ijk} = 1) Pr(b_{ijk} = 1)
\end{aligned} \tag{1.4.8}
$$

which includes three parts (see the bottom of Figure 1.5): the share of consumers who did not bargain and made the purchase, the share of consumers who bargained but failed and made the purchase, and the share of consumers who successfully bargained and made the purchase at a negotiated price. The parameters are estimated by minimizing the distance between the predicted market shares and the observed shares.

1.5 Supply Side

In order to perform policy simulations, I need to model sellers' pricing decisions in addition to the demand side model (see Figure 1.6). I assume that Taobao sellers are playing a static

[15] On a subsample of data where sellers' $DistNextLevel$ is very close to the thresholds, the correlation between $DistNextLevel$ and the price is even stronger.

pricing game to maximize profits.[16] The profit of seller j on product k is given by

$$\begin{aligned}\Pi_{jk} &= (Pr(b_{jk}=0)+Pr(b_{jk}=1)Pr(d_{jk}=0|b_{jk}=1))Pr(a_{jk}(\bar{p}_{jk})=1)(\bar{p}_{jk}-mc_{jk}) \\ &\quad +Pr(b_{jk}=1)Pr(d_{jk}>0|b_{jk}=1)Pr(a_{jk}(\bar{p}_{jk}-\hat{d}_{jk})=1)(\bar{p}_{jk}-\hat{d}_{jk}-mc_{jk}) \\ &= Pr_{jk,zero}(\bar{p}_{jk}-mc_{jk})+Pr_{jk,pos}(\bar{p}_{jk}-\hat{d}_{jk}-mc_{jk}) \qquad (1.5.1)\end{aligned}$$

where $a_{jk}(\cdot)$ is the purchase decision made by a representative consumer. If $u_{jk} > u_{j0}$ at price p_{jk}, then $a_{jk}(p_{jk}) = 1$, otherwise $a_{jk}(p_{jk}) = 0$. The profit function consists of two parts: the profit from transactions completed at the posted price $\bar{p}_{jk}$, and the profit from transactions completed at a bargained price $\bar{p}_{jk}-\hat{d}_{jk}$, where $\hat{d}_{jk}$ is the seller's expected bargaining discount amount for a representative consumer. The above expression of the profit function is a simplification of the real profit function in that a seller uses a representative consumer to proxy for the consumer pool. A representative consumer is defined as a consumer with an average (across all consumers) shopping experience and buyer characteristics of interest. The probabilities and the expected bargaining gains are calculated for this representative consumer. I compute the profit-maximizing price that sellers would set via equating the first-order condition to zero. Thus, the relationship between the marginal cost and the optimal price for seller j on product k can be written as

$$\left(\frac{\partial Pr_{jk,zero}}{\partial \bar{p}_{jk}}+\frac{\partial Pr_{jk,pos}}{\partial \bar{p}_{jk}}\right)(\bar{p}_{jk}-mc_{jk}) = \left(\frac{\partial Pr_{jk,pos}}{\partial \bar{p}_{jk}}\hat{d}_{jk}+Pr_{jk,pos}\frac{\partial \hat{d}_{jk}}{\partial \bar{p}_{jk}}\right)-(Pr_{jk,zero}+Pr_{jk,pos}) \qquad (1.5.2)$$

The intuition behind this equation is as follows. Since bargaining can result in a lower profit margin, a seller needs to consider the conversion rate and the profit margin not only at the posted price but also at the bargained price when setting the optimal posted price. The

[16] It is possible that Taobao sellers set prices in a strategic manner. For example, sellers may lower the posted price in order to boost sales volume as they get closer to the next reputation threshold. In order to model such strategic behavior, I conduct a sensitivity test by adding a weighting function, which is formulated as $(1+\frac{1}{distance\ to\ the\ threshold})$, into the profit maximizing objective. The results are not sensitive to the added weighting function.

above equation nests and reduces to the standard monopoly pricing $(p - mc) = \frac{1}{|e|} \cdot p$ when bargaining was not possible.

1.6 Results

I first discuss the estimates for the bargaining realization and bargaining costs followed by the parameters of the utility function. I report the bootstrapped standard errors for all the estimates.

1.6.1 Bargaining Realization and Bargaining Decision

Table 1.8 shows the estimates of the two-part model in the bargaining realization stage. I estimate the model on both the full sample and a subsample where the information on buyers' age and gender is available. Columns (1) and (2) report the results on bargaining success rates. As expected, I find that a consumer is more likely to bargain successfully under the following scenarios: when the posted price is high, no promotion is available, the seller has a low reputation level, the buyer has more Taobao shopping experience, and this is a repeat purchase between the seller and the buyer. Also, I find that female buyers on average are more likely to succeed in bargaining, but age does not have any statistical significant effect. To make the results more interpretable, I calculate the average partial effect of each explanatory variable. Specifically, I find that a 1% increase in price leads to a 0.6% increase in bargaining success rate. Holding the posted price and everything else constant, the presence of a promotion decreases the bargaining success rate by 0.5%. A one level increase in the seller reputation level decreases the bargaining success rate by 0.1%, while a one level increase in the buyer shopping experience increases the bargaining success rate by 0.2%. If a buyer and the seller have at least one transaction in the past, then the bargaining success rate increases by 1.3%. Lastly, the bargaining success rate of female buyers is greater than male buyers by 0.2%.

Columns (3) and (4) report the results of the discount amount conditional on bargaining success. Most factors have similar signs in the two parts of the bargaining realization model. A 1% increase in the posted price leads the discount amount to increase by 1%. The effect of a presence of a promotion on the discount amount is -56%, a drop of about 20 yuan. A one level increase in the seller reputation decreases the discount amount by 15%, about 6 yuan, suggesting that higher reputation sellers have more bargaining power. Interestingly, I find that though consumers with more shopping experience are more likely to succeed in bargaining, the discount that they get from bargaining is less than those with less shopping experience. A one level increase in consumer shopping experience is associated with 5% decrease, or about 2 yuan, in the bargaining discount amount. Buyer age has a small negative effect on the discount amount, but I see no statistically significant difference between male and female buyers in terms of discount amount conditional on success. Given the results are mostly the same across the full sample and the subsample (with and without buyer age and gender information), I will describe the results from the full sample going forward.

Figure 1.7 shows the estimated province-level bargaining costs. For each province, I report the estimated upper bound, lower bound, and the mean. On average, the bargaining cost is 9.0 yuan, about 1.5 dollars. To put this in context, the minimum hourly wage in China ranges from 11 to 20 yuan. Not surprisingly, the estimated bargaining costs vary across provinces. I attempt to explain this heterogeneity using several factors. Specifically, I use the 2014 China Economic Census data to construct the following variables: province-level disposable income per capita, population density, household size, level of urbanization, and internet penetration rate. I also use the World Value Survey to construct a trust level from the question "Could you tell me whether you trust people you meet for the first time completely (4), somewhat (3), not very much (2) or not at all (1)?" I find that income per capita, population density, level of urbanization, and internet penetration are positively correlated with the province-level bargaining costs, while household size and the trust level are negatively correlated with them. Though the province-level characteristics do not have statistically significant explanatory power in bargaining costs, most likely due to the small sample size (there are only 31 provinces), the above correlations seem to suggest that more developed provinces on average have higher bargaining costs, which is consistent with the economic theory that higher income people have higher time costs and thus may have a lower propensity to bargain.

Using the estimates from the bargaining realization and the estimated bargaining costs, I find that on average bargaining would be initiated in about 78% of shopping occasions. Conditional on bargaining, Figure 1.8 represents the histogram of the estimated success rate across all shopping occasions in the sample. The average success rate conditional on bargaining is 18.8%. Using the estimated bargaining intention and the success rate conditional on bargaining, a back-of-the-envelope calculation of the bargaining success rate equals $78\% * 18.8\% \approx 15\%$ - this is consistent with the 16% success rate observed in the sample.

As I previously highlighted, it is important to distinguish between no-bargain and failed-bargain transactions in the retail settings, as it can directly affect the estimated bargaining cost and further affect the managerial implications. However, this point has been overlooked in previous literature. To assess the importance of this point, I compare the estimation results with the distinction to that without the distinction. Without the distinction, one would incorrectly think the bargaining success rate equals 100% and only those who get a discounted price are the ones who initiate bargaining. Under this inaccurate belief, the bargaining cost would be overestimated by 8 times (76 yuan vs. 9 yuan), and the bargaining intention would be underestimated by 5 times (16% vs. 78%). The big differences between the two results suggest the distinction between no-bargain and failed-bargain transactions is very meaningful and important.

1.6.2 Purchase Decision

Table 1.9 reports results from the purchase model with and without the control function approach. The increase in the magnitude of price and seller characteristics coefficients suggests that the control function approach helps to correct for the endogeneity problem caused by seller unobservables. Without controlling for the unobservables, I would underestimate the price coefficient (less negative) and thus after applying the control function approach, the coefficient on the term to correct for endogeneity is positive and the estimated price coefficient becomes more negative. I find that higher price results in lower conversion rate, while promotions lead to higher conversion rate. Both the seller reputation level and the detailed seller rating affect the conversion rate positively.

To help interpret the results, I numerically calculate the price elasticity on conversion rate for each product/seller combination (Figure 1.9). On average, a 1% increase in price leads to a 3.4% decrease in conversion rate, suggesting the average conversion rate would decrease from 10.02% to 9.68%.[17] As there are two types of transactions, i.e., transactions made at the posted price and those made at a bargained price, I separate the price elasticity on conversion rate into two types. The mean implied price elasticities of the above two types are -3.6 (sd = 0.29) and -2.9 (sd = 0.40) respectively, with a significant difference between the two ($p < 0.000$). Thus, the conversion of transactions at the posted price is more elastic than that after a successful bargain. This result makes intuitive sense as an increase in the posted price will not deter consumers when they can get an additional discount from bargaining as much as when consumers have no choice but to purchase products at the posted price.

I also calculate the mean price elasticities on conversion rates for the top five brands. Nokia has the lowest price elasticity (-3.2) while Apple has the highest (-3.7), consistent with the belief in the Chinese market (at the time of the sample) that Apple products were

[17] The estimated price elasticity implies a 20% margin for the retailers on average. This retail margin seems reasonable compared to the 55% - 75% gross margin reported for manufacturers in the cellphone industry (Techwalls report https://www.techwalls.com/production-costs-of-smartphones/ on 04/21/2017).

considered discretionary products while Nokia products were considered necessary products.

1.7 Counterfactual Analysis: Ban on Bargaining

1.7.1 Counterfactual Simulations

The pricing mechanisms used by the world's top e-commerce platforms have been varied and evolved over time. For example, Amazon started as a fixed-price platform and introduced bargaining on certain product categories in 2014. eBay started as an auction site, then introduced fixed-price, and then bargaining in 2005. Unlike Amazon and eBay, the possibility of bargaining on Taobao arose inadvertently due to the availability of the online chatting tool and became very popular over time. Thus, an interesting question for Taobao is what would happen if it bans the bargaining.[18]

If the platform bans bargaining, i.e., switches from a mixed-price to a fixed-price mechanism, then profit-maximizing sellers would choose their prices according to:

$$p^*_{jk} = \underset{p_{jk}}{\text{argmax}}\{Pr(a_{jk}(p_{jk}) = 1)(p_{jk} - mc_{jk})\} \tag{1.7.1}$$

where the underlying parameters in the above expression are estimated from the baseline model. By setting the first-order-condition to zero, I first estimate sellers' new equilibrium prices after the policy change. Then, using the new prices, I estimate the conversion rate for each seller/product shopping occasion in the sample. I also collect aggregate metrics on the number of transactions and GMV of the Taobao cellphone category in 2012. Using these aggregate measures, I are able to quantify the market-level response to the policy change (banning bargaining).

[18]Though both Amazon and eBay use the bargaining mechanism, their setting is somewhat different from Taobao's in that Taobao sellers do not explicitly show whether they allow for bargaining or not and thus buyers are uncertain about the likelihood of bargaining success. This kind of situation, where bargaining is possible (though that is never stated explicitly), can be seen in stores like BestBuy in the United States and many offline "bazaars" in the developing world.

Table 1.10 reports the result before and after the bargaining is banned. Before the change, the average posted and transaction prices are 1,263 and 1,251 yuan.[19] After the change, as no bargaining is allowed, the average posted price equals the average transaction price, which is 1,249 yuan. The ban on the bargaining mechanism causes the posted price to decrease by 1.1%. This is as expected because sellers would set a higher posted price to leave enough room for bargaining before the policy change. However, due to the additional discounts that consumers could get through bargaining, the final average transaction price between the two scenarios is not very different.

Though the average transaction prices do not seem to differ much, interestingly, I find that when bargaining is banned, the conversion rate becomes higher, increasing from 10.06% to 10.13%, a 0.7% increase. This increase in the conversion rate can arise if either of the following two scenarios unfolds. First, a consumer who has a high bargaining cost would not bargain and would not purchase the product at the high posted price before the change, but she may make the purchase at the lower posted price after the change. Second, a consumer who failed to bargain and refused to purchase the product at the high posted price before the change may buy at the lower posted price after the change. Thus, due to the existence of bargaining costs and the probability that bargaining may fail under the bargaining mechanism, I see that the conversion rate is higher with the pure fixed-price mechanism than the mixed-price mechanism.

Given the benefit of the fixed-price mechanism, a natural question to ask is why sellers do not enforce fixed prices. I offer four potential reasons for this. First, given the existence of the Aliwangwang interface on the website, a commitment to a fixed price is not enforce-

[19] Note that the average transaction price reported here is different from the average transaction price reported in Table 1.1, where the average transaction price is 1,236 yuan. The reason is that the two prices are computed differently. In the summary table, I calculate the average transaction price for the 39,625 transactions in the sample. In contrast, I calculate the average transaction price for the 39,625 shopping occasions where transactions are made by a representative consumer instead of a specific consumer. In other words, the average transaction price in Table 1.1 is the observed transaction price, while here it is the simulated transaction price for a representative consumer under the current mixed-price mechanism. I do this to facilitate a "fair" comparison between the mixed-price and the counterfactual fixed-price mechanisms.

able by the seller and is therefore likely to be deemed as "cheap talk" from the consumers' perspective. I found some evidence to support this in the consumer survey, where when I asked respondents whether they will still bargain with a seller even when the seller posts a "no bargaining" sign on his/her site, 34.5% respondents answered "yes," and 22.2% answered "maybe." Second, based on my informal discussions with Taobao sellers, though most sellers do not like bargaining, they prefer to leave the bargaining decision to buyers due to social and cultural norms that take bargaining as granted. Given this, I see that sellers almost never state "no bargaining" on their item pages.[20] This is also consistent with findings in the prior literature that a fixed-price approach is optimal only if sellers' commitment to it is credible (Riley and Zeckhauser, 1983). Third, given how fragmented the market is, with over 7 million sellers on Taobao, it is nearly impossible for the sellers to take any collective action to move to a fixed-price policy. Finally, the platform has always taken a laissez-faire approach since its inception and has little incentive to interfere with sellers' transaction decisions, as the revenue for the platform primarily comes from advertising sold to sellers, not commissions on transactions.

To study whether the effect of the policy change is heterogenous across sellers, I calculate the profit difference before and after the ban on the bargaining mechanism for each seller/product shopping occasion, i.e., profit difference = profit under the fixed-price - profit under the mixed-price. In order to compare across sellers, I standardize the profit difference by setting the market size to one, that is, assuming there is only one potential customer for each seller/product combination. The mean of the estimated profit increases by 0.2% under the fixed-price mechanism compared with that under the mixed-price mechanism.[21] I regress the estimated profit difference on seller reputation level, detailed seller rating, site age, and the promotion indicator to study the determinants of the difference. The results

[20]Even if there exist stringent fixed-price sellers, my two-part model will capture their preference by setting the probability of successful bargaining rate for these sellers to be close to zero.

[21]The mean of the estimated profit is 23.48 yuan under the mixed-price mechanism and 23.53 yuan under the fixed-price mechanism. As the estimated profits are a function of the market size definition, the absolute magnitude of the estimates cannot be used for a comparison. The relative change, however, is not sensitive to the market size definition and that is why I use it to make the comparison.

are in Table 1.11, where a negative number indicates mixed-price with bargaining is better for sellers with those attributes, and fixed-price is better for sellers lacking those attributes.

I find that sellers with low reputation levels, and thus low bargaining power, benefit more from the fixed-price mechanism. This finding is consistent with the theoretical prediction that a decrease (increase) in a seller's bargaining power favors fixed-price (bargaining) (Wang, 1995) and a relatively higher (lower) buyer's bargaining ability favors fixed-price (bargaining) for the seller (Arnold and Lippman, 1998). Also, for a similar reason, I find that the fixed-price mechanism is more appealing for sellers with high detailed seller rating. Though the site age has a positive effect on the profit difference after the policy change, the size of the effect is negligible. Products with promotion indicator benefit more from bargaining. This result has intuitive explanation in that given the posted price being the same, an indicator of promotion leads to a lower bargaining amount under the mixed-price mechanism but it should not have any effect on the transaction price under the fixed-price mechanism.[22]

Given that the mean number of transactions per day in the Taobao cellphone category is 69,622 (in 2012), the estimated number of transactions per day after banning the bargaining mechanism increases to 70,095. The percentage change in the number of transactions per day equals the percentage change in the conversion rate, provided that the market size remains stable before and after the policy change.[23] Using the average transaction price and the number of transactions per day, I impute the GMV per day as 87.0 million yuan before the change and 87.5 million yuan after the ban, which is about a 0.7% increase.

[22]When performing this exercise on the promotion indicator, I assume that a seller's promotion decision does not vary before and after the policy change. The only thing that changes is the posted price. If I allow the usage of promotion to change, then I need to model sellers' promotion decision explicitly in the model, something that I leave for future research. Here, since I define posted price as list price - promotional discount, I have implicitly accounted for the possible change in sellers' promotion behavior due to the presence of bargaining.

[23]I argue that the market size stability assumption is plausible for two reasons. First, Taobao was a virtual monopoly in China during the sample period, given the fact that about 83% of all the online transactions went through this platform. As a result, for anyone who wants to shop or sell online, Taobao seems to be the only dominant choice. Second, based on the survey, only 4% of Taobao buyers really enjoy bargaining, and over 60% of Taobao buyers do not find bargaining enjoyable. So if there was any change after banning bargaining, I would expect the traffic to increase as the majority do not like bargaining, and this would further reinforce my findings and managerial recommendations.

Comparing the average transaction price and the total sales, it seems that there is not much difference between the two types of pricing mechanisms. However, the existence of bargaining costs gives rise to significant welfare implications. Under the bargaining mechanism, a consumer who initiates bargaining, successfully or not, incurs some bargaining costs. In contrast, under the fixed-price mechanism, as bargaining is not allowed, no bargaining costs are incurred. The incurred bargaining costs before the policy change are roughly calculated as the number of transactions per day divided by the average conversion rate, multiplied by the probability of bargaining and the average bargaining costs. The estimated bargaining costs incurred by buyers per day in the cellphone category alone are 4.9 million yuan, about 6% of the category total sales - an economically significant number. As I do not observe sellers' marginal costs (for products), I am not able to separately identify sellers' bargaining costs. However, given the time and effort sellers spend on dealing with consumers' bargaining, sellers' bargaining costs should also be substantial. As a result, the fixed-price mechanism would be even more favorable (than what I document) than the status-quo from a social planner's and sellers' point of view.

In conclusion, while the platform may not have a strong incentive to switch to a fixed-price mechanism, this move is very beneficial from the social planner's perspective. Such a policy change would greatly help consumers as they will avoid bargaining costs and at the same time may help sellers avoid operational costs if their bargaining costs are high.

1.7.2 Robustness

The goal of this paper is to measure accurately the value of bargaining for sellers, buyers, the e-commerce platform, and the social planner in terms of social welfare. It is important that the counterfactual analysis is robust to the model assumptions and can be generalized. In this section, therefore, I discuss the robustness of my results to several assumptions, particularly those that relate to buyers' bargaining decision and the bargaining realization process. I conclude with a replication of my analysis using data from another product category to judge its generalizability.

Buyers' Bargaining Decision

In the baseline model, I assume that buyers' expected bargaining gain is the product of the perceived success rate conditional on bargaining as revealed in the survey and the expected discount amount conditional on success. In the robustness exercise, I allow Taobao buyers to be either fully sophisticated, i.e., perceive the bargaining success rate to be the same as the realized success rate, or fully naive, i.e., perceive the bargaining success rate as 100% all the time and thus the expected bargaining gain equals the expected bargaining discount amount conditional on success. The sophisticated assumption and the naive assumption provide a lower bound and an upper bound for the bargaining costs, and thus provide two bounds for the counterfactual comparison between the fixed-price mechanism and the mixed-price mechanism.

For each assumption considered, I find that the changes in average transaction price and conversion rate are minimal compared with the baseline counterfactual analysis. In contrast, changes in assumptions have a substantial effect on the saved bargaining cost amount. Under the sophisticated and the naive assumptions, the saved bargaining costs per day equal 1.9 and 6.7 million yuan, respectively. However, changes in bargaining assumptions yield qualitatively similar conclusions.

Bargaining Realization Process

One of the critical assumptions underlying my two-part model for the bargaining realization process is the conditional independence of the error term, that is, the error term is uncorrelated with the explanatory variables, including the posted price. Although the included seller characteristics and the product fixed effects can do a decent job in controlling for the endogenous pricing decision, there could remain a bias caused by potential unobserved seller characteristics in the error term, e.g., sellers' bargaining skills and bargaining willingness. The best way to control for unobserved seller characteristics is to include seller fixed effects. However, to do so, I have to restrict the analysis to sellers with at least two transactions. Such a restriction introduces a sample selection. Given one of my goals is to estimate the value of bargaining for the platform, I want to keep the sample as representative as possible, and thus I did not include seller fixed effects in the baseline model. To evaluate the robustness of my results to the conditional independence assumption, I restrict my analysis to a subsample and include seller fixed effects in the bargaining realization process.

Besides the sample selection issue, given the large number of sellers in the sample and the product fixed effects, it is computationally impossible to estimate the baseline structural model with seller fixed effects. As a compromise, I employ a reduced form approach. Table 1.12 reports the estimated results. Columns (1) and (2) present probit regressions of the bargaining success on the explanatory variables without and with seller fixed effects. Columns (3) and (4) report truncated regressions of the realized bargaining discount amount without and with the seller fixed effects. The sample is restricted to the sellers with at least 20 observed transactions. The estimates are similar across columns, suggesting that my results are robust to the conditional independence assumption. Note that the key difference between the reduced form regressions and the structural model is whether I account for a buyer's bargaining intention. In the structural model, I explicitly estimate the bargaining intention, while the reduced form regressions do not allow to do so. Nevertheless, the comparison with the reduced-form results increases my confidence in my results.

1.7.3 Generalizability

The above analysis uses data from the cellphone category. However, given that the goal of the paper is to evaluate the relative value of bargaining for the platform, it is important to assess whether the findings from the cellphone category can be generalized to other categories. This section considers an additional product category and replicates the findings as a generalizability check.

In addition to cellphones, I were able to obtain data on the women's shoes product category. Following the first two steps in §1.4, I find the the bargaining cost is 8.6 yuan on average in this category.[24] The estimated bargaining cost of 8.6 yuan for women's shoes is very close to my estimate of 9.3 yuan for the cellphone category. Given that the unconditional success rate in women's shoes category is low, I expect similar findings in the counterfactual analysis that both the average transaction price and the conversion rate stay about the same. As before, the major benefits from a ban on bargaining would come from the saved bargaining costs. As women's shoes are lower ticket items than cellphones, the magnitude of the total benefit is likely to be smaller than that for cellphones but in the same direction. Overall, these analyses suggest that my results are not idiosyncratic to the cellphone category.

1.8 Discussion and Conclusion

This paper contributes to a small but growing body of empirical literature on bargaining. I focus on comparing the value of bargaining versus fixed-price mechanism on an online platform. I propose a structural model capturing the stages inherent in a transaction where bargaining is possible - the decision to bargain, the bargaining realization, and the purchase decision. A consumer's bargaining cost, essential for evaluating the impact on social welfare implications of bargaining versus fixed price, is modeled as a critical part in the decision to bargain stage. A two-part model is used to describe the bargaining realization, which is flexible enough to capture the data generating process without onerous data requirements. For the purchase decision, I use a consumer discrete choice model augmented with outcomes

[24] Note that the estimation process did not include product fixed effects (due to the lack of standardized products in this category) and used the same survey data as in my main analysis to compute the average bargaining success rate conditional on bargaining.

from the bargaining realization process and use a control function approach to address the potential endogeneity problem. On the supply side, I model sellers making profit-maximizing pricing decisions that are impacted by expected bargaining outcomes. Using the estimates from the structural model, I perform a counterfactual analysis to derive the value of bargaining by assuming a counterfactual scenario where bargaining is banned in the marketplace. I find that both the marketplace as a whole and the sellers in the marketplace would benefit from this policy change. My results suggest that social welfare will also increase, largely due to the saved bargaining costs for buyers. Further, I investigate the heterogeneity of the benefits from the policy change across sellers. The findings are consistent with theoretical predictions that lower bargaining power favors fixed-price mechanism.

The findings of this paper provide additional evidence with respect to several key components in bargaining. I find that consumers' bargaining costs are comparable to the minimum hourly wage in the geographic setting (China). Further, I find bargaining costs are positively correlated with province-level disposable income per capita, population density, urbanization, and internet penetration, and negatively correlated with household size and trust level. The overall pattern suggests that bargaining costs are higher for people in more developed areas. I also find that sellers' reputation level and buyers' shopping experience have strong effects on bargaining outcomes; this complements the findings of earlier studies on the key determinants for bargaining outcomes. From an applied perspective, online platforms are a very large and growing part of the modern digital economy. Many platforms allow bargaining (in various forms), but there is little research on the value of bargaining versus fixed-price in this setting. My study thus fills the void.

There are several avenues for future research. First, due to data limitations, I define a market at the seller/product level. This allows to abstract away from seller competition and also the consumer search process. In cases where information on consumers' search behavior and simultaneous bargaining with multiple sellers is available, one could get a better assessment of the effect of pricing policy change by incorporating a search model into the framework. Third, though the two-part bargaining model is flexible, an extensive-form bargaining model could be employed if alternating-offer data are available. Finally, in this paper, I implicitly include the seller's bargaining cost as part of the marginal cost. However, if sellers' product marginal costs are observed, it may be possible to separately identify the seller's bargaining cost, which would make the social welfare and sellers' benefits analysis more complete. I hope that future research can carry out these extensions.

1.9 Figures

Figure 1.1: Taobao Item Page

Figure 1.2: Taobao Feedback Page

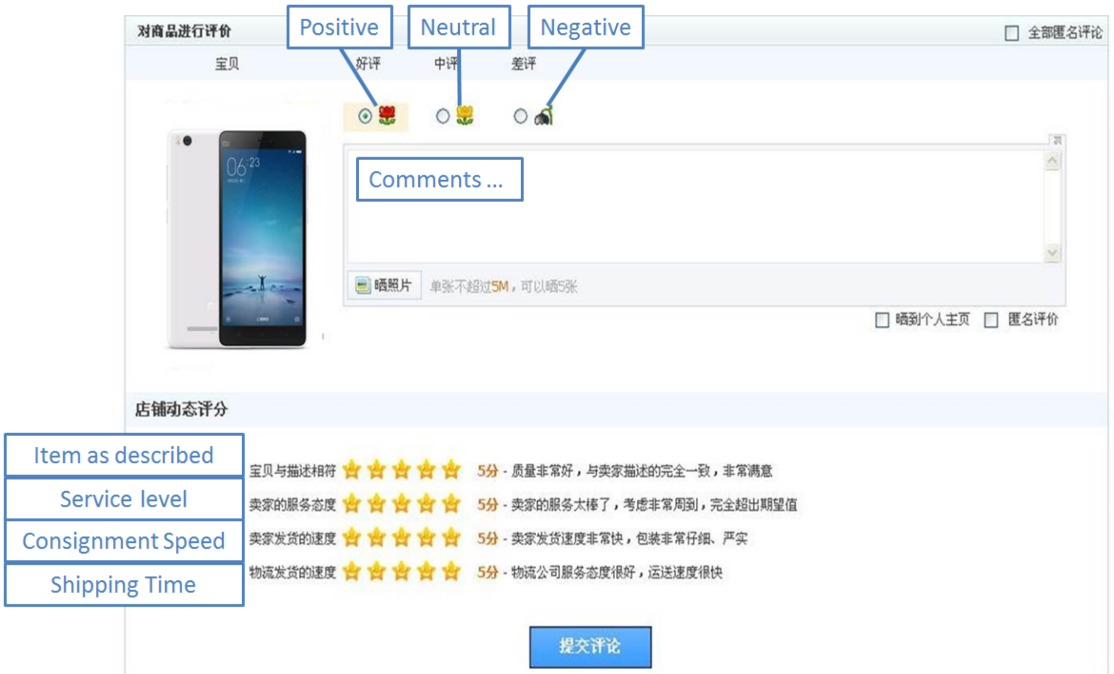

Figure 1.3: Taobao Reputation Rule

Taobao Buyer Reputation Level / **Taobao Seller Reputation Level**

4分-10分	
11分-40分	
41分-90分	
91分-150分	
151分-250分	
251分-500分	
501分-1000分	
1001分-2000分	
2001分-5000分	
5001分-10000分	
10001分-20000分	
20001分-50000分	
50001分-100000分	
100001分-200000分	
200001分-500000分	
500001分-1000000分	
1000001分-2000000分	
2000001分-5000000分	
5000001分-10000000分	
10000001分以上	

Figure 1.4: A Comparison of Age Distribution between the Transaction Sample and the Survey

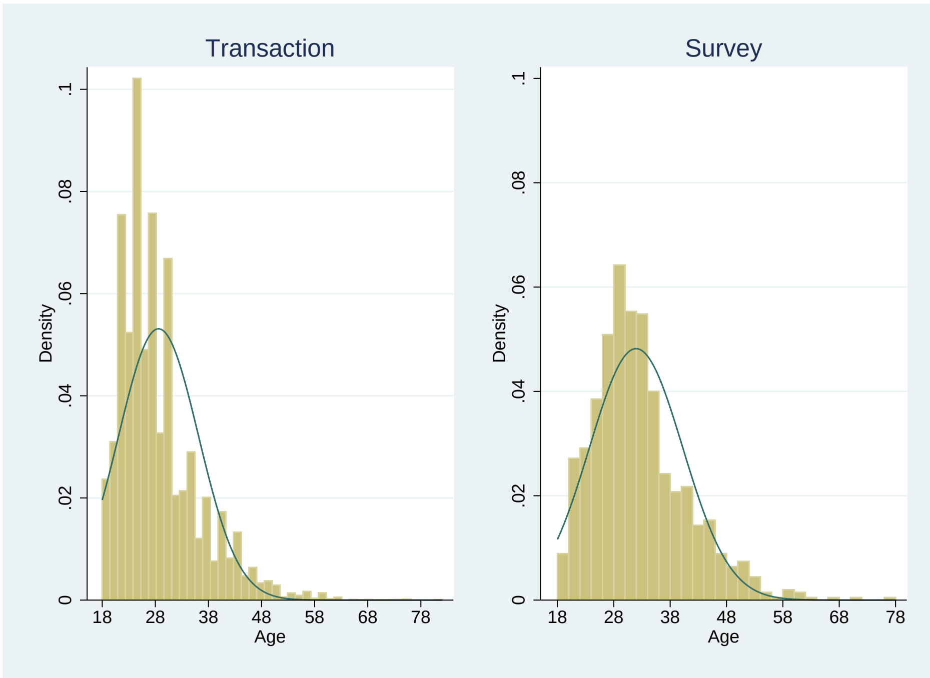

Figure 1.5: The Bargaining and Purchase Process

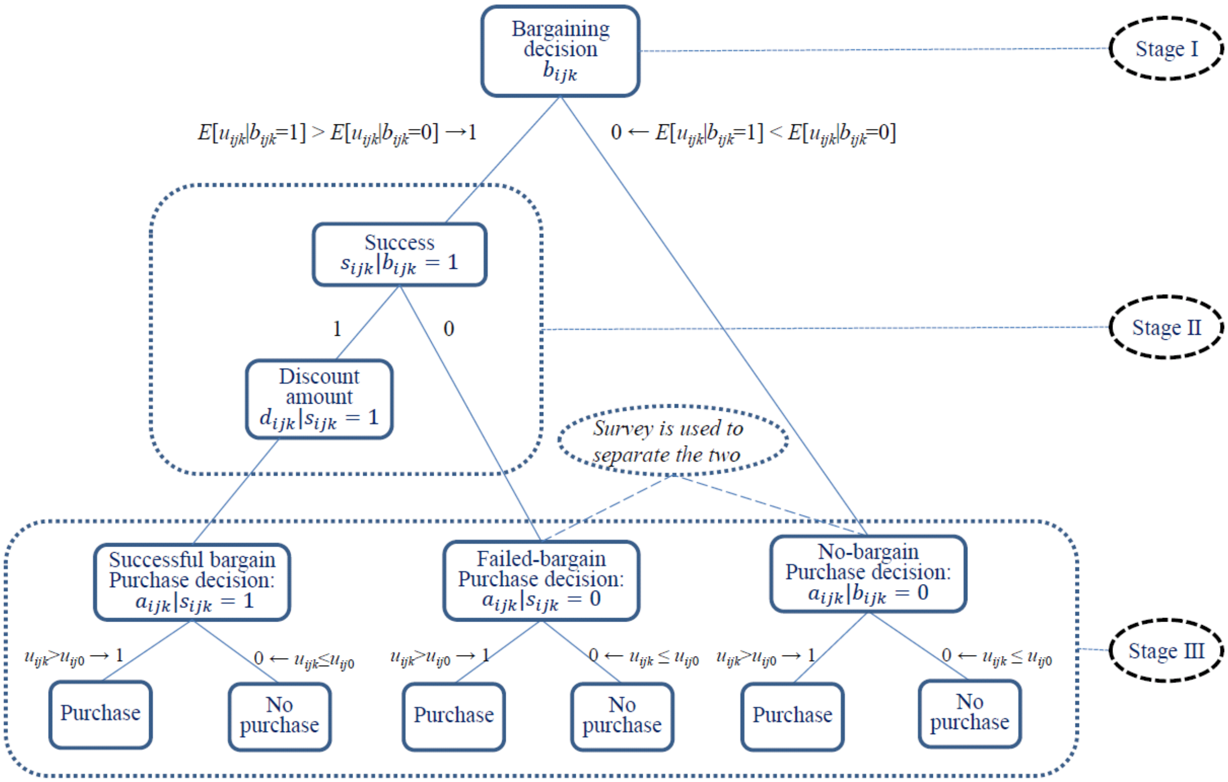

Figure 1.6: The Flow of the Estimation Process

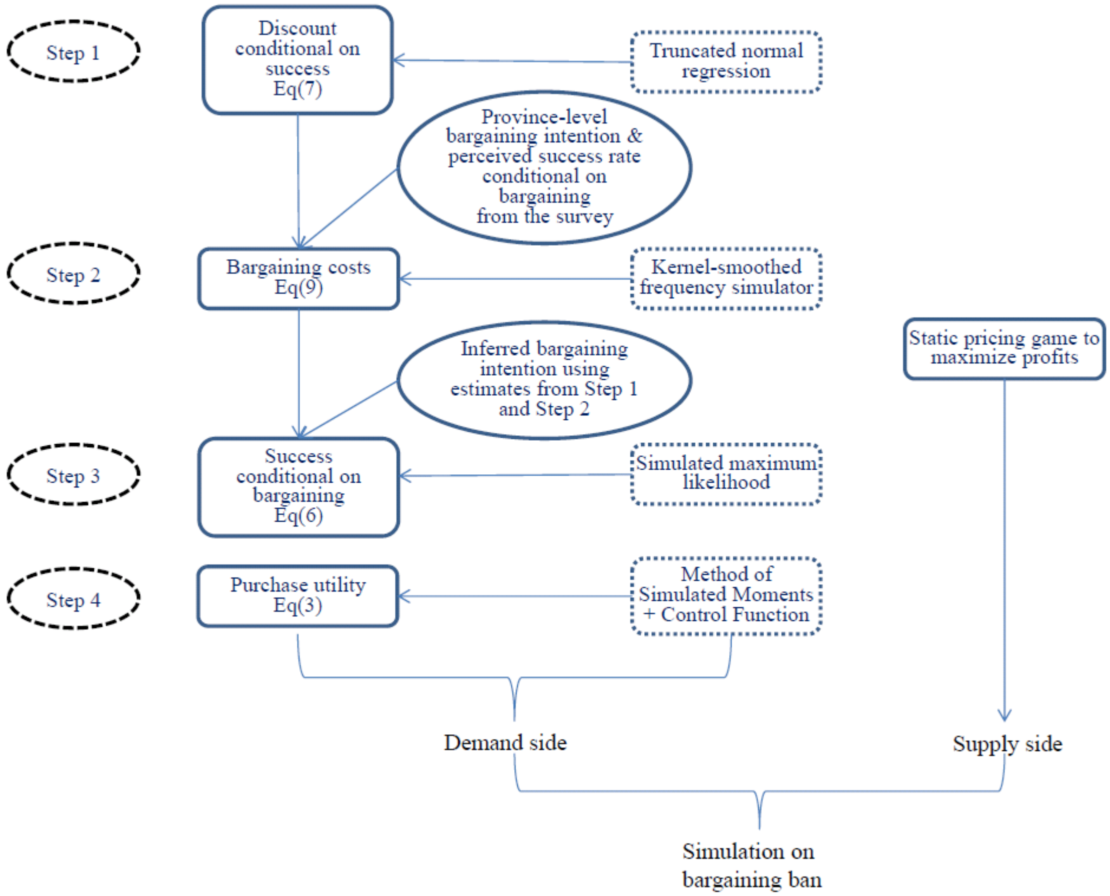

Figure 1.7: Estimated Province-level Average Bargaining Costs

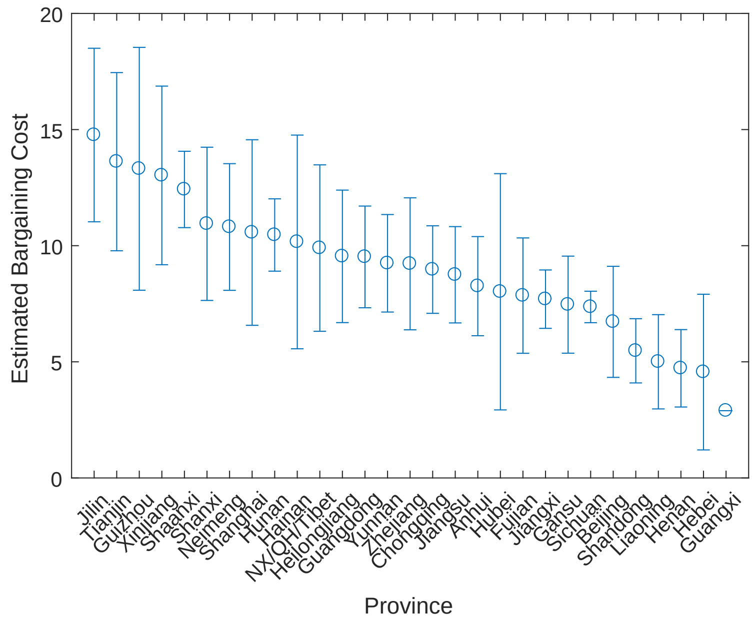

Note: since no respondent in Guangxi province chooses "may bargain," the lower bound and the upper bound for that province overlap.

Figure 1.8: Estimated Bargaining Successful Probability

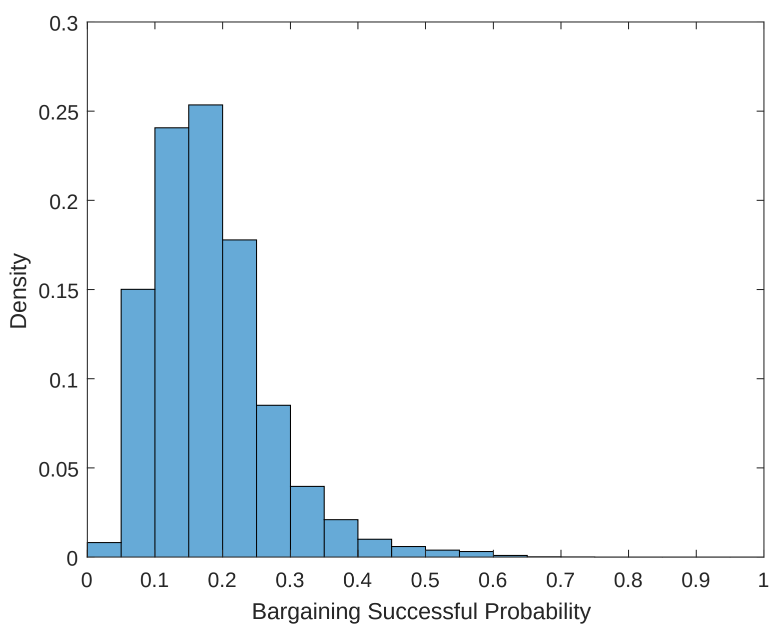

Figure 1.9: Implied Price Elasticity on Conversion Rate

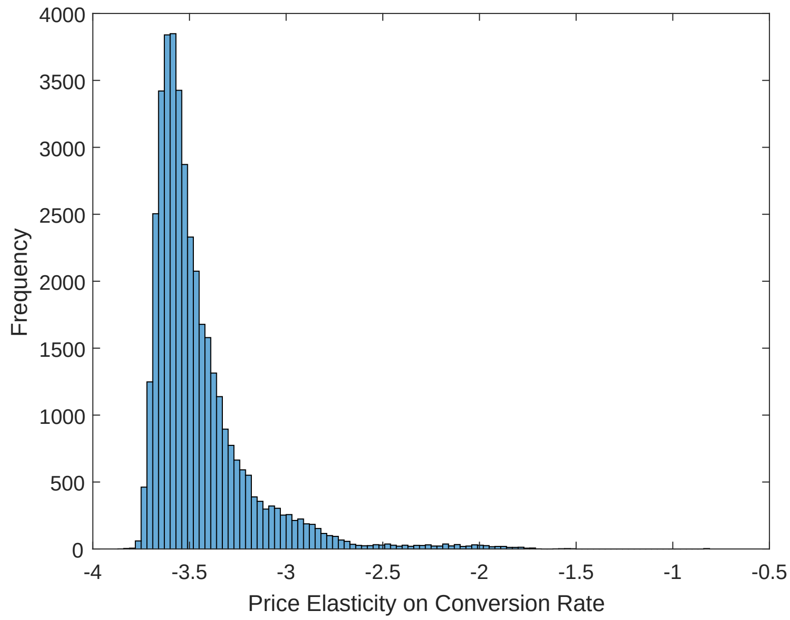

1.10 Tables

Table 1.1: Summary Statistics: Transaction Sample

	N	mean	sd	min	max
Posted Price (yuan)	39,625	1,263	1,133	45	8,650
I(Promotion)	39,625	0.645	0.479	0	1
I(Bargaining success)	39,625	0.160	0.366	0	1
Bargaining Discount Amount (yuan)	39,625	27.21	133.5	0	1,515
Bargaining Discount Amount | Success (yuan)	6,332	170.3	295.3	0.01	1,515
I(Repeat Purchase)	39,625	0.11	0.31	0	1
Product Age (yrs.)	39,625	2.31	2.00	0	9.9
Site Age (yrs.)	39,625	3.027	1.946	0	8.764
Seller Reputation Level	39,625	8.423	2.225	0	14.35
Detailed Seller Rating	39,625	4.807	0.128	1	5
Seller Repeat Purchase Rate	39,625	0.077	0.053	0	1
Buyer Shopping Experience Level	39,625	4.174	1.639	0	9.862
Buyer Age	25,201	28.6	7.5	18	82
Female	25,201	0.29	0.45	0	1
Number of Unique Visitors (past 4 weeks)	39,625	4,721	11,291	0.2	145,338
Number of Unique Purchasers (past 4 weeks)	39,625	58	113	0	1,158

Table 1.2: Transaction Summary Statistics by Gender

	Female (N = 7,357)		Male (N = 17,844)	
	Mean	Std. Dev.	Mean	Std. Dev.
Posted Price	1,339	1,190	1,361	1,179
Promotion Indicator	0.65	0.48	0.64	0.48
Bargaining Success Indicator	0.18	0.39	0.17	0.38
Bargaining Discount (Yuan)	28.6	136	29.5	140
Buyer Shopping Experience	5.1	1.3	4.6	1.4

Table 1.3: Summary Statistics: Survey

	N	mean	sd	min	max
I(Certainly Bargain)	1,009	54.3%	0.498	0	1
I(May + Certainly Bargain)	1,009	73.6%	0.441	0	1
Perceived Success Rate|Bargaining	743	48.6%	0.250	10%	100%
E[Discount Amount|Success] (yuan)	743	165.5	221.5	1	1000

Note: Only those respondents who answered "yes" or "not sure" to the bargaining intention question are required to provide information for the perceived success rate conditional on bargaining and the expected discount amount conditional on success. Thus, we see a decrease in the sample size in the last two rows.

Table 1.4: Survey Summary Statistics by Gender

	Female (N=526 in Row 1,2; 376 in Row 3,4)		Male (N=492 in Row 1,2; 373 in Row 3,4)	
	Mean	Std. Dev.	Mean	Std. Dev.
I(Certainly Bargain)	0.50	0.50	0.58	0.49
I(May + Certainly Bargain)	0.71	0.45	0.76	0.43
Perceived Success Rate | Bargaining	0.46	0.25	0.51	0.25
E[Discounted Amount | Success] (yuan)	154	209	178	236

Table 1.5: Notation Definition

Notation	Definition
u_{ijk}	Purchase utility
u_{ij0}	Outside option utility
δ_k	Intrinsic preference of product k
$\overline{p}_{jk}$	Posted price (promotion adjusted)
p_{jk}	Transaction price, which equals $\overline{p}_{jk}$when the bargaining discount amount is zero
x_{jk}	Seller and product characteristics, and proxy measures for search processes
x_{ijk}	x_{jk} plus buyer characteristics and the posted price
γ	Effects of various factors on the bargaining success rate
θ	Effects of various factors on the bargaining discount amount
ν_s	Unobservable in the conditional bargaining success function
ν_d	Unobservable in the conditional bargaining discount function
z_{ij}	Instrumental variable - distance to the next reputation level
μ_j	Residual from the control function, which is used to address the price endogeneity
m_{jk}	Market share / Conversion rate for seller j on product k
s_{ijk}	Bargaining success indicator
b_{ijk}	Bargaining decision
c_i	Individual bargaining cost
c_{zl}, c_{zu}	Lower and upper bounds of province-level bargaining costs
mc_{jk}	Marginal cost for seller j on product k
Π_{jk}	Profit for seller j on product k

Table 1.6: A Brief Summary of Identification

Parameters	Identified by
δ_k- intrinsic product preference β_p- price preference β_x- nonprice preference	The fraction of consumers who made a purchase among consumers who visited the item page across products, over different prices, and across sellers
c_{zl}, c_{zu}- lower and upper bounds of province-level bargaining costs	(i) the upper and lower bounds of province-level bargaining intention (ii) the expected discount amount across shopping occasions
γ- bargaining success parameters	The estimated bargaining decision and the observed success rate
θ- bargaining discount parameters	The realized discount amounts conditional on success
mc_{jk}- marginal cost	(i) the estimated price elasticity (ii) the estimated bargaining outcome (iii) the profit maximizing assumption

Table 1.7: Relevance Condition Check on IV

	log(price)
log(DistNextLevel)	0.053***
	(0.008)
log(DistNextLevel)-squared	-0.004***
	(0.001)
x_{ij}	Yes
Product FE	Yes
R-squared	0.675
Cragg-Donald Wald F Statistics (weak IV test)	18.47
Number of Observations	39,625

Table 1.8: Bargaining Realization Estimates

	Bargaining Success Indicator		log(Bargaining Amount)	
	Full Sample	Subsample	Full Sample	Subsample
log(posted price)	0.203***	0.175***	1.028***	0.985***
	(0.014)	(0.015)	(0.036)	(0.044)
I(Promotion)	-0.169***	-0.184***	-0.553***	-0.565***
	(0.012)	(0.033)	(0.038)	(0.045)
Seller Reputation Level	-0.021***	-0.021***	-0.141***	-0.139***
	(0.003)	(0.005)	(0.012)	(0.014)
Detailed Seller Rating	0.045	0.047	0.532***	0.525**
	(0.061)	(0.152)	(0.142)	(0.168)
Store Age	-0.003	0.004	0.014	0.016
	(0.006)	(0.004)	(0.013)	(0.015)
Buyer Shopping Experience	0.055***	0.052***	-0.054***	-0.055***
	(0.005)	(0.005)	(0.012)	(0.017)
Buyer Age		-0.001		-0.006*
		(0.001)		(0.003)
Buyer I(Female)		0.063**		0.055
		(0.023)		(0.046)
I(Repeat Purchase)	0.402***	0.425***	0.260***	0.192***
	(0.020)	(0.021)	(0.046)	(0.052)
Product Age	0.004	-0.004	0.143***	0.147***
	(0.007)	(0.006)	(0.018)	(0.022)
# of Sellers w/ Same Product	-0.021	-0.018	0.052	0.074
	(0.023)	(0.033)	(0.046)	(0.055)
# of Sellers w/ Same Reputation	-0.011	-0.017*	0.044**	0.046**
	(0.007)	(0.006)	(0.013)	(0.015)
Product FE	Yes	Yes	Yes	Yes
Number of Observations	39,625	25,201	6,332	4,368

Note: ***, **, and * indicate statistical significance at the 0.1%, 1%, and 5% levels in all tables

Table 1.9: Purchase Model Estimates

	Without Control Function	With Control Function
log(price)	-0.464***	-3.852***
	(0.010)	(0.370)
I(Promotion)	0.185***	0.196***
	(0.014)	(0.011)
Seller Reputation Level	0.057***	0.181***
	(0.006)	(0.013)
Detailed Seller Rating	1.108***	4.206***
	(0.114)	(0.302)
Store Age	-0.100***	0.035
	(0.005)	(0.017)
Repeat Purchase Rate	2.372***	2.016***
	(0.294)	(0.313)
Product Age	0.012**	-0.806***
	(0.004)	(0.089)
# of Sellers w/ Same Product	0.088***	0.334***
	(0.016)	(0.027)
# of Sellers w/ Same Reputation	-0.041***	-0.027***
	(0.004)	(0.004)
Terms to Correct for Endogeneity		
Control Function for log(price)		3.393***
		(0.371)
Product FE	Yes	Yes
Number of Observations	39,625	39,625

Table 1.10: Counterfactual Analysis: Market Response if Bargaining is Banned

	Before Change	After Change	Δ%
Average Posted Price (yuan)	1,263	1,249	-1.1%
Average Transaction price (yuan)	1,251	1,249	-0.1%
Conversion Rate (%)	10.06	10.13	0.7%
Number of Transactions per Day	69,622	70,095	0.7%
Total Sales per Day (million yuan)	87.0	87.5	0.6%
Bargaining Costs per Day (million yuan)	4.9	0	N/A

Note: The posted price, the transaction price, and the conversion rates are calculated based on the sample. The other three measures are back-of-the-envelope calculations with supplemental information on several aggregate metrics of Taobao cellphone category in 2012. The increase in the number of transactions per day is proportional to the increase in the conversion rate given that the market size for each seller/product is assumed to be unchanged. GMV is calculated by multiplying average transaction price and the number of transactions. The bargaining cost incurred per day is calculated as the number of transactions per day divided by the average conversion rate times the probability of bargaining times the average bargaining cost.

Table 1.11: Determinants of the Heterogenous Effects of the Ban on the Bargaining Mechanism

	(Post-profit) - (Pre-profit)
Seller Reputation Level	-0.008***
	(0.0001)
Detailed Seller Rating	0.098***
	(0.002)
Site Age	0.0003*
	(0.0001)
I(Promotion)	-0.037***
	(0.0006)
Product FE	Yes
Number of Observations	39,625

Table 1.12: Robustness of Bargaining Realization Process

	Bargaining Success Indicator		log(Bargaining Amount)	
log(price)	0.329***	0.404***	0.782***	0.943***
	(0.035)	(0.054)	(0.090)	(0.109)
I(Promotion)	-0.159***	-0.119**	-0.510***	-0.591***
	(0.036)	(0.045)	(0.080)	(0.075)
Seller Reputation Level	0.043**	-0.073	-0.243***	-0.244***
	(0.016)	(0.050)	(0.035)	(0.097)
Detailed Seller Rating	-0.691***	-0.442**	0.554***	-1.055**
	(0.079)	(0.149)	(0.185)	(0.323)
Store Age	-0.021	-0.195**	0.065**	0.554***
	(0.011)	(0.074)	(0.027)	(0.141)
Buyer Shopping Experience	0.053***	0.050***	-0.072**	-0.046*
	(0.010)	(0.011)	(0.024)	(0.021)
I(Repeat Purchase)	0.375***	0.290***	0.223**	0.178*
	(0.041)	(0.046)	(0.085)	(0.075)
Product Age	0.010	0.022	0.122*	0.093
	(0.020)	(0.026)	(0.050)	(0.052)
# of Sellers w/ Same Product	-0.062	-0.063	0.035	-0.138
	(0.041)	(0.046)	(0.096)	(0.087)
# of Sellers w/ Same Reputation	0.014	0.058**	-0.090*	-0.059
	(0.016)	(0.021)	(0.035)	(0.038)
Product FE	Yes	Yes	Yes	Yes
Seller FE	No	Yes	No	Yes
Log likelihood	-4,455	-4,047	-2,810	-2,388
Number of Observations	12,190	12,190	1,695	1,695

Note: Columns (1) and (2) are Probit regressions and columns (3) and (4) are truncated regressions. The sample is restricted to the sellers who have at least 20 observed transactions.

CHAPTER II

Negotiation Pricing on a Health Platform Market: Bringing Hospitals and Patients Together

2.1 Introduction

Many top internet companies around the world host platform-based businesses, such as Amazon and eBay in the US, Alibaba and JD.com in China. One of the most important strategies that a platform firm can choose is pricing (Rysman, 2009). On business-to-consumer (B2C) platforms, when the number of firms on the business side is large, platforms often set standard terms and conditions for firms to join, like Amazon's and Alibaba's category-based standard terms. This type of pricing has been extensively studied both theoretically and empirically (see Rysman, 2009; Sriram et al., 2014, for an overview). But when the number of firms on the business side is not large,[25] platforms often negotiate or bargain with individual firms on the terms and conditions for them to join the platform (as with Expedia with Marriott, and managed care organizations with hospitals). Although a few theoretical studies have explored the impact of price discrimination within one side of the platform markets on the platform's performance,[26] to the best of my knowledge, no empirical study

[25] This can either happen when the platform is at the startup stage or when the platform is facing a limited number of critical companies that they want to collaborate with.

[26] Even though Caillaud and Jullien (2003) and Weyl (2010) allow for price discrimination based on users' identity in their theoretical models, the price discrimination they model is a one-sided decision - completely set by the platform, which is different from the negotiation pricing strategy - a mutual agreement between two parties.

has investigated the value of a negotiation pricing strategy in B2C online platform markets.

In this paper, I take a first step towards understanding the determinants and the value of negotiation pricing between an intermediary platform (as a business entity) and the business side of a B2C platform market. I look into the following questions: First, how does network size, measured by the number of participants on both sides of the market, affect the businesses' willingness to participate with the platform? Second, how does the network size affect the negotiation outcomes? I focus on three key contract terms - price, payment method, and clearing cycle. Price is important as it directly influences the platform's and businesses' profitability. Payment method and clearing cycle are important as they affect the platform's Gross Merchandise Value (GMV) and cash holdings, which are key determinants of a startup's success (Brown et al., 2009; Holtz-Eakin et al., 1994; Oliveira and Fortunato, 2006). Third, since the platform has two network size measures, one on the B-side and the other on the C-side, does the B-side network size have the same impact on the negotiation outcomes as the C-side? If not, how different are they? The answer helps to evaluate whether the platform is exerting effort on the right side of the market. Fourth, when the businesses on the B-side are very different, can I detect heterogeneous network size effect on negotiation outcomes from different types of businesses? The answer to this question can shed light on the value of negotiation compared with a uniform pricing for online platform markets, especially at the startup and growth stage. Lastly, I explore mechanisms behind some of the findings by linking the findings with some business characteristics, such as profit margin, product variety, and flexibility of the business.

The empirical context I focus on is one of the earliest online physical examination (PE) platform markets in China, established on August 1st, 2015. The platform brings hospitals and consumers together. Hospitals provide physical exam packages, and consumers make purchases and make appointments on the platform, and then go to hospitals for checkups. The data span from the first day of the platform's launch date to the end of 2017. By the end of the sample period, the platform had around 700,000 registered users and more than 700 hospital contracts in about 200 cities. One unique feature of my data is that it includes detailed information on the contracts and on the updates of contract terms after each hospital visit or a business call done by the platform's sales teams. The fact that the platform is in the startup and growth phase provides with the necessary variations on the network sizes and the frequent hospital visits and contract updates yield variations of negotiation outcomes. Thus, this is an ideal setting to study the effect of the network size on negotiation outcomes in a B2C platform market.

Although a random field experiment is an ideal way to get a casual relationship, it is nearly impossible to conduct a random experiment in a business-to-business setting given the high implementation costs and the high opportunity costs. As a result, I rely on the unique observational data to assess the effect of the network size on negotiation outcomes and to test the heterogeneity among businesses.

I find that both the consumer network size and the business (here I refer to the hospital) network size, measured as the number of participants, have a positive impact on hospitals' willingness to participate, and this impact is similar between public and private hospital network sizes. In terms of the key contract terms, the overall pattern suggests that an increase in the consumer network size and an increase in the hospital size both lead hospitals to agree on more preferable contract terms toward the platform, suggesting an increase in the platform's bargaining power.

The contribution of the paper is twofold. First, the paper adds to the literature on two-sided markets by being the first empirical study looking at questions on negotiation pricing, one of the two most common practices in B2C platform markets. Second, I add to the literature on negotiation and bargaining by focusing on the effect of network size in two-sided markets. The aspect that the total market size is changing over time in the focal two-sided market gives a distinguishing angle from the previous studies.

The rest of the paper is organized as follows. I review the relevant literature in §2.2, introduce the institutional setting in §2.3, and describe the data in §2.4. In §2.5, I set up the research design. The results and discussions are presented in §2.6. Finally, §2.7 concludes.

2.2 Literature Review

Platform markets are markets that facilitate economic interactions between at least two types of end users wherein the decisions of one set of users are likely to have an effect on the other via network effects / externalities (Sriram et al., 2014). Given the fact that (1) a platform needs to attract different sides of the market, (2) the network effects of each side vary, and (3) the incentives of each side to join and to use the platform differ, platforms need to find the right balance to "get the two sides on board" (Rochet and Tirole, 2006). Such a balance is typically operationalized through a pricing structure, i.e., an allocation of charges across the two sides. Armstrong (2006), Caillaud and Jullien (2003), Rochet and Tirole (2006), and Weyl (2010) each provide a theoretical framework to analyze pricing strategies that allow platforms to protect their business and to gain new business. Empirically, rigorous research has also shown that the pricing decisions of the platform depend on the demand elasticities and cross network elasticities on each side of the market (Ackerberg and Gowrisankaran, 2006; Clements and Ohashi, 2005; Rysman, 2004; Fan, 2013; Kaiser and Wright, 2006; Pattabhiramaiah et al., 2018; Song, 2013; Wilbur, 2008). However, all the previous studies have treated the pricing decision as purely a platform's decision; in other words, the users or businesses have no say in the pricing structure and they can only choose to take it or leave it at the preset price. My study is the first that looks at a two-sided market allowing the pricing structure to be determined by both the platform and businesses through negotiation pricing. In particular, I investigate how the network size influences the pricing structure and the underlying mechanism that drives this effect. The flexibility in pricing introduced by negotiation can be especially valuable to a growing platform when the number of businesses is not large. I also quantify the heterogenous network effects and highlight how the heterogeneous network effects can play a role in the pricing structure, which has not been emphasized in the previous literature.

This paper also contributes to the literature on negotiation and bargaining. Negotiation and bargaining are widely used in B2B settings and understanding their impact is currently a very active area of research (Lee, 2015). Draganska et al. (2010) and Meza and Sudhir (2010) study price negotiations between wholesalers and supermarket retailers, highlighting the role of firm size, store-brand, and service level in determining the bargaining power. Crawford and Yurukoglu (2012) focuses on bargaining in multichannel television markets to study the value of bundling. Gowrisankaran et al. (2015), Ho (2009), and Lewis and Pflum (2017) focus on insurer-hospital markets, and find that hospital size, capacity, merger, and cross-market dependencies can affect the negotiated price with insurers. Grennan (2013, 2014) focus on medical devices markets to quantify bargaining ability. My paper adds to this stream of literature by focusing on a B2C online platform market and highlighting the role of network sizes on negotiation outcomes. Studying such a topic has been proved very challenging due to the inaccessibility and incomplete observability of the contracts as contracts are of significant strategic importance for companies (Lee, 2015).[27]

2.3 Institutional Setting

The market for physical exams is enormous in China. Figure 2.1 plots the total number of health checkups in China over the recent years. In 2017, 501 million health checkups have been done, with the value worth hundreds of billions RMB, i.e., about a hundred billion dollars. Even though the total number is huge, it only covers about 35% of the total population in China, which is much smaller compared with some developed countries, such as 74% in the United States and 97% in Germany. This industry in China is expected to grow

[27] A few exceptions that have the detailed contract data include Dafny (2010), Grennan (2013), and Lafontaine and Slade (2008).

rapidly in the coming years. I focus on one of the first physical exam online platforms in China, established in August 2015. The platform consists of a hospital side and a consumer side. Hospitals provide various physical exam packages to be listed on the platform. These physical exams are generally initiated by consumers instead of a doctor, and thus, are not covered by insurance companies. There is a wide variety of physical exam packages, designed for different age groups (teenager, middle-aged, elderly, etc.), for different concerns (cardiovascular, diabetes, women's health, etc.), for different purposes (preventive care, driver's license application, job application, marriage certificate application, etc.), and performed by different hospitals (public vs. private with different ratings). Accordingly, this online platform is very much like other types of e-commerce markets where differentiated products are available for consumers to voluntarily choose from, but with a focus on healthcare.

On the consumer side, consumers pay nothing to register on this platform. The only information required for registration is a cellphone number for verification purpose. The cellphone number contains location information up to the city level.[28] On the hospital side, hospitals do not pay any direct fee to the platform for participation. The hospitals incur costs, such as administrative costs, including preparing hospital information to be uploaded, creating physical exam packages to be uploaded, integrating the hospital information system with the platform's, and bargaining over the terms and conditions for participation.

In order to grow and expand the market, the platform puts a lot of effort into seeking participation from hospitals. The platform adopts geographic exclusivity for its sales teams, who go and meet with their respective hospital administrators to negotiate. If a contract is agreed upon, it often includes a one-year automatic renewal clause and requires one-month advance notice for termination by either party. The contracts are kind of informal in the sense that even though the typical length of a contract is one year, the contract terms can be and often are renegotiated while the current one is still in effect. In particular, there are

[28]The cellphone number in general can accurately predict a user's location in China due to the household registration system and the roaming charges by telecommunications companies.

three key contract terms to be negotiated on. First is the depth of price discount, second is the acceptance of online payment, and third is the clearing cycle. I next explain why the above three contract terms are important for the hospital-platform collaboration.

The first key contract term is the depth of price discount or the wholesale price the platform can get. The way that the platform makes money is selling PE packages at a price higher than the wholesale price. The profit depends on two things: the depth of price discount in the contract, and the online price decided by the platform. There are three relevant prices here - original price, contract price, and the online price. An example can help illustrate the differences among the three prices and how the profit is calculated. A hospital sets the original price (walk-in price which cannot be negotiated) for a physical exam package at \$100 and agrees with the platform on a 20% discount. The platform will be allowed to set the online price anywhere between \$80 and \$100, where \$80 is like resale price maintenance.[29] If the platform sets the price at \$95, then the \$15 price difference between the online price and the contract price is the profit margin earned by the platform. As this example illustrates, the depth of the discount can directly influence the platform's online price setting, and thus further influence the profit margin and the gross merchandise volume (GMV).

The second key contract term is the hospitals' acceptance of online payments. When a hospital agrees to accept online payments, a consumer will pay the platform first and then the platform pays the hospital later at a pre-specified payment cycle (see next paragraph for details). If a hospital does not agree to such payment, a consumer can use the platform to make an appointment but then pays at the hospital reception desk. As a result, the transaction bypasses the platform, even though the profit for the platform will be paid by the hospital later. Hospitals' acceptance of online payment benefits the platform in two ways. First, though the acceptance of online payment does not literally affect the profit of

[29] See Israeli et al. (2016) for an overview of the literature on resale price maintenance.

the platform,[30] the online payment greatly contributes to the GMV of the platform, one of the most important metrics investors look for when valuing an online platform, especially at the startup stage before the network effects and profits kick in.[31] Second, the acceptance of online payment increases the platform's cash holdings and thus alleviates its potential liquidity constraints, a benefit similar to that obtained from a longer clearing cycle, as detailed below.

The third key contract term is the clearing cycle, at which the platform transfers the collected payments to the specific hospitals at the contract price, which can be per transaction, weekly, monthly, or quarterly. For hospitals that agree to online payments, all the sales go to the platform first and then the platform pays the hospital at the end of the clearing cycle. Thus, a longer clearing cycle means the platform gets to hold the sales revenue for a longer period, which increases the platform's cash holdings and alleviates liquidity constraints. The liquidity constraint is one of the key factors determining a startup's survival prospect and business operational decisions (Brown et al., 2009; Holtz-Eakin et al., 1994; Oliveira and Fortunato, 2006), and thus a longer clearing cycle is always preferred by the platform. For the hospitals that do not accept online payment, the clearing cycle is often discussed in the contract as well, as it may be used in the future if any change on the acceptance of the online payment takes place.

2.4 Data and Summary Statistics

The data span almost two-and-a-half years from the platform launch date August 1^{st}, 2015 to December 31^{st}, 2017. The platform has been growing rapidly and is still in the market expansion phase. Figure 2.2 shows the new and cumulative membership registrations. By the end of 2017, about 700,000 consumers have registered on the platform, with an average

[30] As the informal contracts are not legally binding, the hospital can easily dispute. As a result, the platform may not be able to get its fair share back.

[31] https://www.forbes.com/sites/ryancaldbeck/2015/11/10/how-to-value-online-marketplaces

monthly growth rate of 70%. Consumers are twice more likely to register on weekdays compared to weekends.

As I mentioned before, the platform seeks collaboration opportunities by visiting hospitals.[32] Table 2.1 shows the number of hospitals visits by the platform. By the end of 2017, the platform has paid visits to 1,950 hospitals, with 4 visits on average for each hospital. Among them, over 700 hospitals agreed to sign contracts. The average number of visits to contracted hospitals is larger than that to non-contracted hospitals, 6.9 vs. 2.5. The average number of visits to hospitals after a contract being signed is 4.5, suggesting that the platform keeps renegotiating with hospitals on contracts. The trend of new and cumulative number of hospital contracts can be found in Figure 2.3. Figure 2.4 presents a comparison of the 700 contracted hospitals and the rest 1,250 non-contracted hospitals on the hospital type distributions. Hospitals in China are broadly separated into two groups - public hospitals and private hospitals. Public hospitals are rated by the government into different levels based on the following measures: number of beds, number of doctors, number of departments, specialty areas, patient safety, medical equipments, management, and other care-related indicators. Public AAA is the highest level, followed by AA, A, BBB, BB, and public B is the lowest level among the six. Public comprehensive hospitals are the public hospitals that do not provide level information to the platform. Private hospitals on average are much smaller than public hospitals. Given the short history of private hospital existence in China, they are trying hard to gain trust from the public to compete with public hospitals over the years. By the end of October 2017, there are 12,200 public hospitals and 17,771 private hospitals. Very few hospitals have multiple branches and the contracts are all at the branch level. Even though the number of private hospitals is larger than that of public hospitals, there are many more admissions in public hospitals than in private hospitals, 2.40

[32] The general process goes like this. After a sales lead is generated, a sales person first calls the hospital to further check its interest and schedule a visit. If a prospect is identified, the sales person visits the hospital, which can be in the format of formal meetings, informal talks, and lunches. For important hospitals, regional managers and sometimes even the platform's senior managers may get involved in the visits.

billion vs. 0.37 billion.[33] From the histogram, we can clearly see that public hospitals are less likely to collaborate with the platform than private hospitals. In terms of location, the 1,950 hospitals cover 304 cities and the contracted 700 hospitals cover 190 cities.

In addition to the aggregate platform performance information, I am fortunate to also have information on the platform's sales teams' hospital visits and negotiated results on the key contract terms. The platform has paid in total over 8,000 visits to the 1,950 hospitals, so about 4 visits per hospital. Contracted hospitals receive on average 6.9 visits by the platform, which is a lot more than that of non-contracted hospitals (2.5 visits on average). Given the fact that 48% of hospitals signed a contract upon the first visit by the platform, and contracted hospitals on average receive 2.3 visits prior to signing a contract, the average 6.9 total visits towards contracted hospitals suggest that the platform constantly tries to re-negotiate contract terms with hospitals while a contract is still in effect.

The first key contract term is the depth of discount level, i.e. one minus the ratio between the contract price and the original price. Table 2.2 presents the summary statistics of the discount level. The average discount is 11.1% in the original contract between the platform and any contracted hospital. As the platform keeps trying to renegotiate contracts by visiting hospitals, the average discount increases to 14.5% by the end of 2017. Over time, about 6% of hospitals (43 hospitals) terminated contracts with the platform, so I see a decrease in the number of observations in the first two rows in Table 2.2. Comparing the discount level in the original contracts and the end-of-2017 contracts, we see the platform is able to get better discounts over time. The next two rows provide a breakdown between public and private hospitals and show that private hospitals are willing to give a discount level almost three times higher than that of public hospitals. It is worth mentioning that 53% of the discount is set to zero in the original contracts. In other words, the platform earns zero profit from almost 400 hospitals when they first signed the contract. This kind of penetration pricing

[33]Source: National Health and Family Planning Commission of the People's Republic of China http://www.nhfpc.gov.cn/mohwsbwstjxxzx/s7967/201712/cb133528a46d4e98af757e6a657bdb9c.shtml, http://www.nhfpc.gov.cn/mohwsbwstjxxzx/s7967/201712/9e82bc727b8d4c6ba3c718bd04bf7629.shtml

is prevalent in platform markets, especially during the introduction and expansion phase, as the platform relies heavily on the network effects induced by hospitals to attract consumers rapidly.

The second key contract term is the payment method. Figure 2.5 plots a histogram of this attribute in both the original contracts and in the updated contracts at the end of 2017. Since the platform benefits from larger cash holdings, it prefers online payment to a combination of online and offline payment to offline payment. In about 10% of all the contracts, the payment method is not specified, denoted as "none". As can be seen from the plot, there is a significant move from "none" or "offline" to "online" over time, suggesting that the platform is getting better terms via updating the contracts. If one separates public and private hospitals, in the original contracts, there is a distinct difference between the two types on "online" (55% vs. 51%) and "none" (31% vs. 22%). However, by the end of 2017, the difference almost disappears with "online" (77% vs. 77%) and "none" (13% vs. 10%). The changes imply that the platform's renegotiation on the original contracts is working and may have different effects on public and private hospitals.

The third key contract term is the clearing cycle. Figure 2.6 plots a similar histogram as the one above for the payment method. Over time, the platform is able to renegotiate for longer clearing cycles, resulting in better prospects in cash holdings for the platform. Similar to the discount and the payment method, in the data we see private hospitals are more likely to agree to the platform's preferred terms than public hospitals. But unlike the payment method, the gap between the public and private hospitals on the clearing cycle persists as the contracts are being re-negotiated and updated.

In addition to the internal contract information provided by the company, I also used web scraping to collect information on the physical exam packages. Since August 2015, the platform launch date, over 40,000 physical exam packages have been created (see Figure 2.7). The hospitals and the platform keep updating the packages over time. At the end of 2017, about 11,000 packages are available on the platform. Table 2.3 provides the summary statistics for those packages. The average original price and online price are 1,608 Yuan and 1,249 Yuan respectively, suggesting an average 22.3% online discount. The average number of cumulative orders per package is 190.

By analyzing the text data from the website, I extract key features for each physical exam package on multiple dimensions. Regarding the gender, 41% of packages target male, 56% of packages target female, and a small proportion of packages do not specify the gender. In terms of the age, 48% are suitable for young adults, 47% for adults, and 24% for seniors. In terms of health concerns, the top five most frequently covered tests include HIDA scan (77%), lung scan (68%), cancer screening (58%), GYN exam (38%), and thyroid test (35%), note that a package may target more than one type of users and always covers multiple tests, on average 23 different tests. 8.4% of packages target married people, 5.5% target unmarried people, 2% target pregnant women, and the rest do not specify. White-collar workers are a special target group as 22% of packages use words such as "white collar", "stress", "overtime", "government officials", etc. in the package descriptions. All the packages sold on the platform come with a free test result interpretation and a one-year doctor consultation on the test results. 56% of packages also come with instant online result checking services so that a consumer does not need to go back to or wait in the hospital for a print-out result.

2.5 Research Design

2.5.1 The Effect of Network Sizes on Contract Negotiations

Two unique features of the platform enable me to study the effect of network sizes on contract negotiations. First, similar to Expedia and Priceline.com, the focal platform negotiates with different business entities (hospitals in this case) to achieve customized contracts. The customized contracts provide variations on the negotiation outcomes, which will be the variables of interest. Second, unlike Expedia or Priceline.com which are in the mature operation stage, the changes in the network sizes since the platform launch date are necessary variations to derive insights for online platforms that are still in the startup and growth stage.

I focus on four negotiation outcomes. First is whether the hospital is willing to collaborate with the platform (i.e. to sign a contract), second is the depth of discount level, third is the payment method, and the fourth is the clearing cycle. As the platform has two sides, the network sizes are going to be defined on both sides as well - the consumer network size and the hospital network size. To study the main effect of network sizes on negotiation outcomes, let i index hospital, j index city, k index province, and t index time. Let N^h_{jt} indicate the total number of contracted hospitals in city j at time t and N^c_{jt} indicate the total number of registered users in city j at time t. As consumers have to be physically present in order to use the physical exam package, I define the market at the city level. In other words, when deciding to sign a contract, a hospital only cares about the number of registered users in the same city and the spillover effect or competition effect from other hospitals in the same city. The service nature of physical exams dictates the local network size in this industry. Thus, the effect of an increase in the network size on negotiations is estimated with regressions of the form

$$Y_{ijkt} = f(\gamma_h ln(N^h_{jt-\kappa}) + \gamma_c ln(N^c_{jt-\tau}) + X_j\alpha_X + \alpha_V V_{it} + \eta_t + \eta_k + \epsilon_{ijkt}) \tag{2.5.1}$$

where X_j is a vector of hospital characteristics, including the hospital type indicators and the number of branches, V_{it} is the number of visits that the platform has paid to hospital i until time t, η_t are month fixed effects, and η_k are city fixed effects. Since there may exist temporal lags between a contract signing date and the time when consumers are able to see this hospital's packages online, and between a consumer's registration date and the time when a hospital is able to get this piece of information, I set the lag τ and κ to 1 in the main specification, and test the robustness to other choices of 4 and 7 as well. Using the lagged terms can also alleviates some simultaneity concerns. The functional form f

depends on the focal negotiation outcome. When I consider the probability of accepting a contract, Y_{ijkt} is a binary variable denoting whether a contract is signed between hospital i and the platform at time t, which is modeled as a probit function. When I consider the discount level, I use a linear regression model. When it comes to the payment method (none, offline, offline/online, online) and the clearing cycle (per transaction, weekly, monthly+), I use ordered logit models. The network size effects of the consumer side and the hospitals side on negotiation outcomes are captured by the coefficients of interests γ_c and γ_h, respectively. To allow for heteroskedasticity, I use robust standard errors.

As I mentioned in the previous section, public hospitals and private hospitals are very different in terms of size, reputation, trust level, pricing, and other hospital attributes, thus they may have different spillover or competition effects on other hospitals after signing a contract with the platform. I further investigate how the hospital network size effects on negotiation outcomes differ between public hospitals and private hospitals by running regressions of the following form:

$$Y_{ijkt} = f(\gamma_{hpub} ln({N^{hpub}}_{jt-\kappa}) + \gamma_{hpri} ln({N^{hpri}}_{jt-\kappa}) + \gamma_c ln(N^c_{jt-\tau}) + X_j \alpha_X + \alpha_V V_{it} + \eta_t + \eta_k + \epsilon_{ijkt}) \tag{2.5.2}$$

where ${N^{hpub}}_{jt}$ and ${N^{hpri}}_{jt}$ indicate the total number of contracted public and private hospitals in city j up to time t, respectively. And the rest variables are the same as before. The difference between the estimated coefficients γ_{hpub} and γ_{hpri} suggest how big the heterogeneous effects are between public and private hospitals on the network effects.

2.6 Results

2.6.1 Network Size Effects on Collaboration

I report the results of estimating the effects of the network sizes on the collaboration probability in Table 2.4. Columns (1) and (3) report the overall effects of the hospital network size and the consumer network size. Columns (2) and (4) report the heterogeneous effects after splitting the hospital side into public and private ones. In addition to the network sizes, the covariates include hospital characteristics(number of branches, and the hospital type fixed effects), month fixed effects, and city fixed effects in columns (1) and (2), and further include the number of visits paid to the hospital by the platform in columns (3) and (4). Instead of reporting the index function coefficient estimates in the probit model, I report the marginal average effects, calculated as the average of $\phi(X'\beta)\beta_j$ across all the observations, where $\phi(\cdot)$ is the density function of the standard normal distribution and β_j is the coefficient on the corresponding variable.

Across the four different specifications, I find that both the hospital network size and the consumer network size have positive effects on a same-city hospital's probability to collaborate. Columns (1) and (3) show that one percent increase in the hospital network size would yield a 12% - 18% increase in the next hospital's collaboration probability. One percent increase in the consumer network size would yield a 3% - 6% increase in the next hospital's collaboration probability. In specifications (2) and (4), after splitting the hospital side into public hospitals and private hospitals, I expect that the two have different effects. However, despite the fact that the two types of hospitals differ greatly in various characteristics, I find that the effects are similar between the two, with one percent increase in the network size leading to about 6% - 10% increase in the next hospital's collaboration probability.

In columns (3) and (4), I further control for the number of visits paid to the hospital by the platform. Not surprisingly, I find that the number of visits is positively correlated with the collaboration probability. The hospitals with one more visit from the platform's sales team have about 3.8 percent higher chance to sign a contract. Of course, given that the number of hospital visits is a decision for the platform, this result could mean that additional visit pays off through higher collaboration probabilities, or that the sales team select the hospitals that are more likely to collaborate to visit. Still, the observed correlation between the number of visits and the collaboration probability is worth noting. The estimated effects of the hospital network size on collaboration is weaker, but is still significant in magnitude.

In terms of control variables, the estimated coefficients are very similar across the four columns. The number of branches that a hospital has is positively associated with the collaboration probability. Compared with private hospitals (the default category), public hospitals are clearly less likely to sign a contract on average except for public BB - the lowest level for the public hospitals in the analysis.[34] And specialty hospitals, whether they are chain or non-chain, are more likely to collaborate with the platform compared with private hospitals.

2.6.2 Network Size Effects on Contract Terms

Table 2.5 presents the estimates of the effects of network sizes on the discount level. The layout of the table is similar to Table 2.4, where columns (1) and (3) present the effect of the consumer network size and the effect of the overall hospital network size, and columns (2) and (4) explore the heterogeneous effects between public and private hospitals on the discount level. Columns (1) and (2) control for hospital characteristics, month fixed effects and city fixed effects, and the next two columns further control for the number of visits to a hospital by the platform. The sample include all the hospitals that have signed contracts with the platform.

With the linear regression model for the discount level, it is straightforward to inter-

[34]As there are less than five public B hospitals, I have grouped them with the public BB category in the analysis.

pret the estimated coefficients as the average effects. The results are generally stable across various specifications. Columns (1) and (3) show that one percent increase on the hospital network size leads to a 4 percentage point increase on the discount level in the future contracts. After separating public and private hospitals, interestingly, columns (2) and (4) suggest that this effect mainly comes from public hospitals, but not private hospitals. This result is not surprising given the significant differences between the two types of hospitals, which will be discussed at length in the following section. On the other side of the market, one percent increase on the consumer network size leads to 0.4 - 0.5 percentage points increase on the discount level.

Table 2.6 reports the estimates of the effect of network sizes on the payment method. I use an ordered logit model as four types of payment methods are observed in the contracts - "none", "offline", "online/offline", and "online" - in the order of least preferred to most preferred by the platform. This table follows the same layout as the previous two tables. I find that the hospital network size does not have a significant impact on the choice of the payment method in the future contracts. However, the consumer network size has a significant positive impact. One percent increase in the consumer network size would yield 0.4 - 0.5 increase in the ordered log-odds of being in a more preferred payment method category, i.e., moving from no term being specified to the offline payment to the online payment, while the other variables in the model are held constant. This makes economic sense as moving towards online payment can decrease transaction costs, especially when the number of potential transactions (positively correlated with the consumer network size) is large. Given the fact that the "none" category (no term being specified on payment method) in the contract is essentially the same as accepting "offline" payment only, I put "none" and "offline" categories together and re-estimate the ordered logit model. The results show that the estimated effects are robust to this variation on category re-definition.

Table 2.7 presents the results of the network size effects on the clearing cycle, same layout and the same ordered logit model as in the previous table. There are five different categories in the baseline analysis - none, per transaction, weekly, monthly, and quarterly - in the order of least preferred to the most preferred by the platform. Columns (1) and (3) indicate that the hospital network size has a positive effect on the clearing cycle term. If the hospital network size were to increase by one percent in a city, the ordered log-odds of being in a more preferred clearing cycle category would increase by 1.3. Columns (2) and (4) present the heterogeneous results showing that one percent increase on the public hospital network size leads to 0.9 increase in the ordered log-odds of moving up to a better clearing cycle category in the future contracts for the platform, while one percent increase on the private hospital network size leads to a smaller 0.5 increase. The consumer network size also has a positive effect on the clearing cycle. One percent increase on the consumer network size would yield a 0.2 increase in the ordered log-odds of moving up to a more favorable clearing cycle category for the platform. I also test the robustness of the results to different category definitions. As no term being specified on clearing cycle in the contract (the "none" category) is essentially the same as agreeing to clearing after each transaction, and as the number of hospitals that agree to "quarterly" clearing cycle is very small, less than five, I redefine the clearing cycle into three categories - per transaction, weekly, and monthly+. The estimated results are similar as above, suggesting the robustness of the conclusion.

Across the three key contract terms, the sign of the control variables is as expected in most cases. Repeated visits to hospitals by sales people on average result in better contract terms, including discount, payment method, and clearing cycle, for the platform. Once again, however, this correlation may either suggest that repeated visits pay off or suggest that the selection of which hospital to visit is an endogenous decision by sales people. The agreed-upon contract terms with public hospitals are tougher for the platform than with private hospitals on discount level and payment method. There is no clear pattern on the clearing cycle between public and private hospitals. Chain specialty hospitals are willing to agree on more preferable terms with the platform than private hospitals on all dimensions. As hospitals in the "others" category may include a mixture of different types, I refrain from reading much into it.

Differences Between Public and Private Hospitals

The results in previous subsections reveal that the public and private hospitals vary in their negotiation outcomes with the platform and their respective network sizes have different impacts on negotiation in the future contracts. In this subsection, I explore the factors that may drive those hospitals' different behaviors and outcomes.

I first focus on explaining the difference of bargaining outcomes between the private and public hospitals. In subsection 2.6.2, I see a consistent pattern that it is easier for the platform to bargain with private hospitals than public hospitals in terms of participation and contract terms. For example, private hospitals are more willing to accept a deeper discount in the contract. One explanation is that public hospitals have bigger bargaining power as they enjoy better reputation among the public and they are facing a higher demand in China. But in addition to higher bargaining power enjoyed by public hospitals, is there any other major drivers that can lead to this result? As the negotiated discount level cannot be isolated from hospitals' pricing decisions, another possible explanation/hypothesis is that if private hospitals on average set the original price higher than public hospitals, then they would have a higher profit margin and thus they are more willing to give a deeper discount to the platform. To test this hypothesis, I analyze the text information that I collected from the website. I extract more than twenty features for each physical exam package, including the type of the operating hospital, number of total tests, the targeted gender, the targeted age, and the targeted health concerns. I regress the original price (walk-in price) on these attributes and also control for the hospital location. The results are presented in Table 2.8. Compared with the private hospitals (the baseline category), public A-level hospitals and public comprehensive hospitals are clearly charging a higher price, while chain specialty hospital, which is a specific type of private hospitals, is charging even less than the general private hospitals. Thus, the estimated results reject my hypothesis that the deeper discount that the private hospitals are willing to offer, which we see in the data, is a result of a higher original price.

The next factor I consider is the product variety. The idea here is that the product variety and its change over time, can partially capture how well a hospital adjusts to the market demand and how well a hospital could gain a niche market. If there is a systematic difference on the product variety offering and its over-time variation between public hospitals and private hospitals, this may provide insights on why the effects of network sizes are different between public and private hospitals. Table 2.9 presents a comparison of product variety between the two groups. The first row shows that at the end of 2017, public hospitals on average have 13.4 physical exam packages available compared with 11.6 of private hospitals. Since hospitals may update their package offerings from time to time, I also provide a measure on the all-time product variety - the number of distinct physical exam packages offered by a hospital since it joined the platform. The second row shows that when it comes to all-time product variety, public hospitals still have a larger number. The differences between public hospitals and private hospitals are statistically significant at 5% level for both measures on product variety. Brynjolfsson et al. (2003) have shown that increased product variety can enhance consumer surplus and market efficiency. Sun et al. (2016) demonstrates the importance of software product variety in hardware adoptions. A similar idea applies here that the increased product variety makes the platform more attractive to consumers, which as a result can increase the platform's bargaining power and yield a higher agreed-upon discount level. To further strengthen this idea, I check the number of orders placed on each package offered by public hospitals versus that by private hospitals.[35] The third row shows that on average public hospitals have 50 more orders being placed on a package compared with private hospitals despite the higher price set by the public hospitals. Even though I cannot ascertain that the higher number of orders is a result of the increased product variety and better catering to the markets, as this observation may also be driven by the

[35]The numbers are the ones displayed on the platform, which are collected by web scraping.

better reputation of the public hospitals, I think the data pattern of the number of orders is consistent with the story of the product variety.

The third factor I consider is the reputation. It is a general belief in China that public hospitals have better doctors, better service, and better technology. Even though private hospitals are trying hard over the years to catch up, this rooted belief is hard to change,[36] which can be another source why public hospital network size has different effect on negotiations compared with private hospital network size. Without having measures on the public's belief of the reputation between public hospitals and private hospitals, I cannot test this factor directly. However, I collected data on hospital ratings from the platform, which reflect a post-service evaluation on hospitals. The ratings have three dimensions - expertise, service, and environment. Table 2.10 shows that, contrary to the public belief, private hospitals are operating as good as public hospitals, if not better. The ratings on expertise is statistically significantly higher for private hospitals and the ratings on service and environment are similar between the two. If the public's evaluation on private and public hospitals prior to purchase is consistent with that of post-purchase, this finding strengthens my aforementioned story that the product variety is a major factor that drives up the effect of network size on future negotiations for the platform. However, if the public's evaluation is different from the post-purchase ratings, then the ratings implies that the public are gradually changing the rooted belief between the public hospitals and private hospitals and further implies a brighter future for the private hospitals in China is on the way.

The last factor I consider is the appointment flexility between public and private hospitals. I collected data on hours of operation and plotted a bar graph in Figure 2.8. Almost 100% of the public hospitals are open to physical exams from Mondays to Fridays. Private hospitals are similar except that about 15% are closed on Mondays. The big difference between the two groups happens during weekends. The proportion of hospitals that are open for private and

[36] See a report written by The Economist Intelligence Unit on "Understanding China's Emerging Private Healthcare Market". https://perspectives.eiu.com/healthcare/understanding-chinas-emerging-private-healthcare-market

public hospitals are 78% vs. 38% on Saturdays and 46% vs. 10% on Sundays, respectively. Even though, as we have seen, the additional flexibility doesn't offset the fact that demand for public hospitals is still higher, this additional flexibility of the private hospitals is a good way to differentiate themselves from public hospitals.

2.7 Conclusion

This paper adds to the literature on platforms by focusing on negotiation pricing, which has been a prevalent practice in business settings but an understudied topic in academic research. This paper also adds to the literature on negotiation of bargaining by highlighting the role of network sizes on negotiation outcomes in online platform markets. To study the relationship between the network size and the negotiated contract terms essentially brings two hitherto disparate strands of literatures together.

I use novel data on detailed contract terms from an online physical exam platform, which brings hospitals and consumers together. In particular, I quantify how the consumer network size and the hospital network size influence hospitals' willingness to collaborate, and the key negotiation outcomes between the platform and the hospital side, including the discount level, the payment method, and the clearing cycle. I find positive significant effects of both network sizes on collaboration probability and on contract terms that are more favorable to the platform. I separate the effects between public hospitals and private hospitals and highlight the heterogeneous effects between the two. To understand the mechanisms behind these findings, I analyze hospital characteristics, such as the profit margin, the product variety, the hospital reputation, and operating hour flexibility. My findings provide managerial insights on the nature of negotiation pricing for B2C platform markets, especially in the early growth phase.

There are several avenues to extend the current study. First, I would like to build a richer characterization of the heterogeneous effects between public and private hospitals once I have access to detailed hospital characteristics. Second, I would like to build a structural model to describe this platform market so that I can quantify the value of negotiation pricing by comparing the current situation to a counterfactual scenario where the platform uses a fixed fee instead of customized contracts.

2.8 Figures

Figure 2.1: The Market of Health Checkups in China

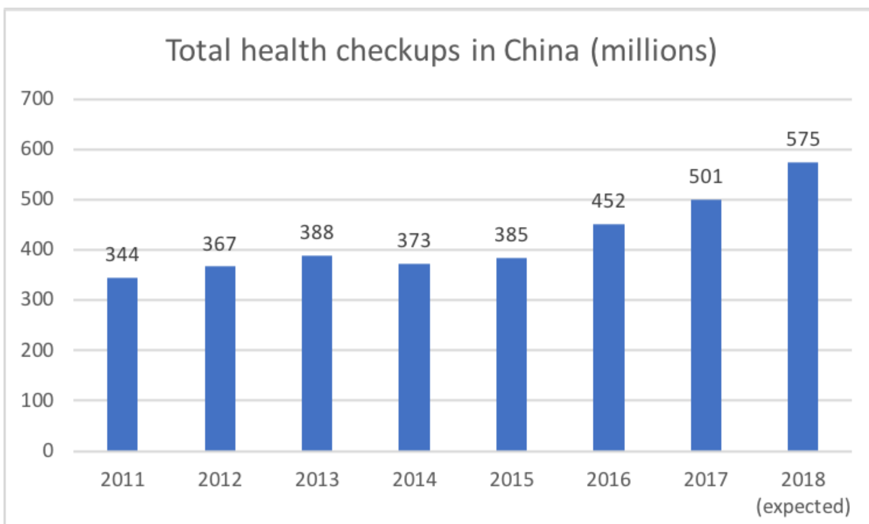

Source: China Statistical Yearbook of Public Health.

Figure 2.2: Membership Registration

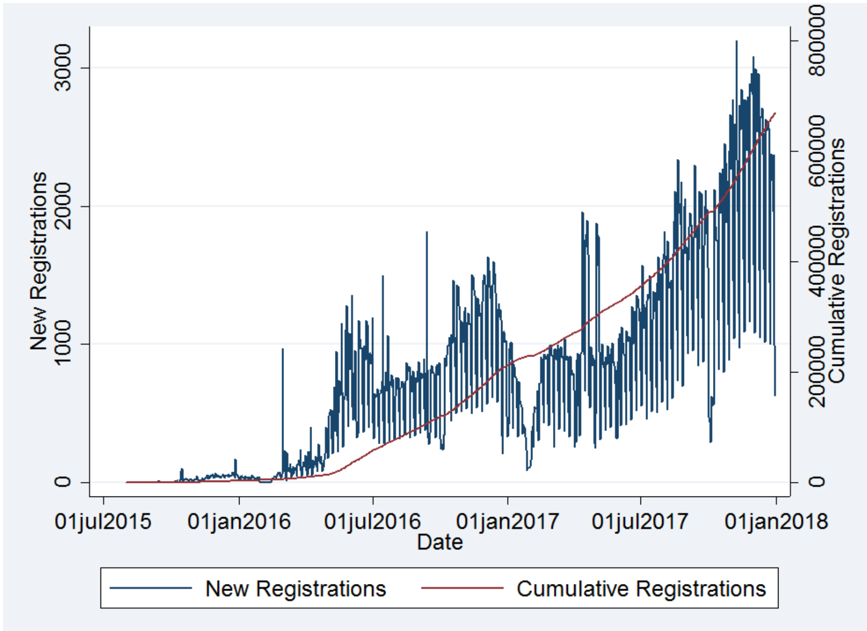

Figure 2.3: Hospital Contracts

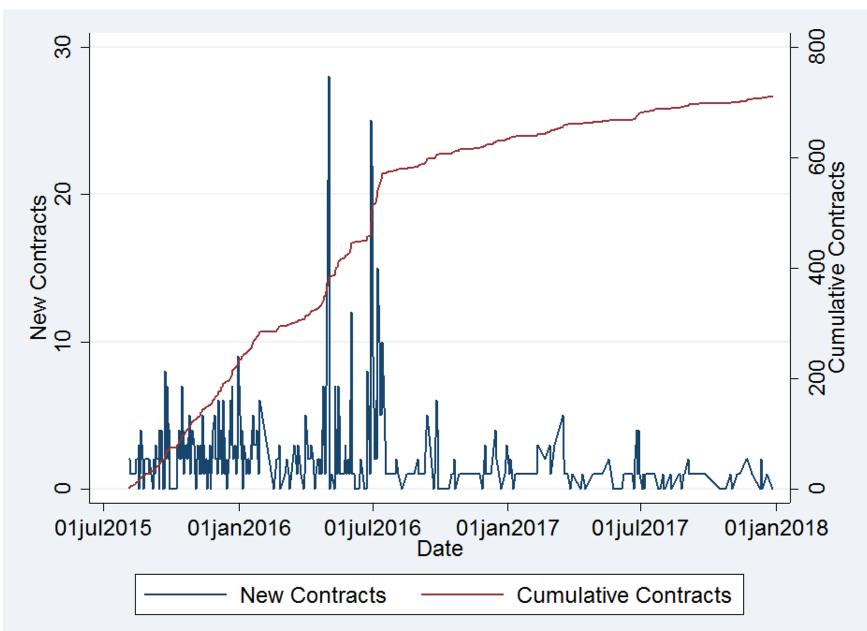

Figure 2.4: Hospital Type Distributions of contracted and non-contracted hospitals

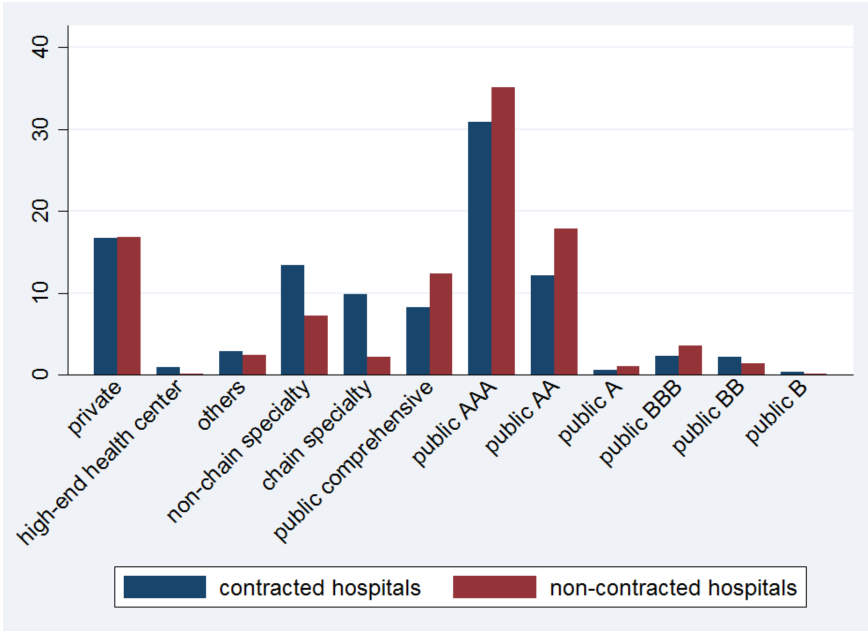

Figure 2.5: Payment Methods

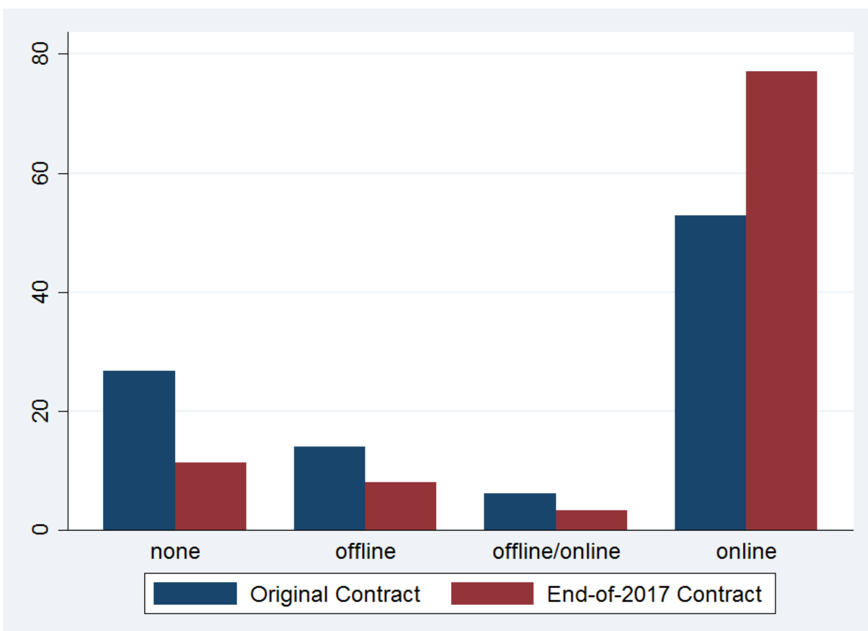

Figure 2.6: Clearing Cycles

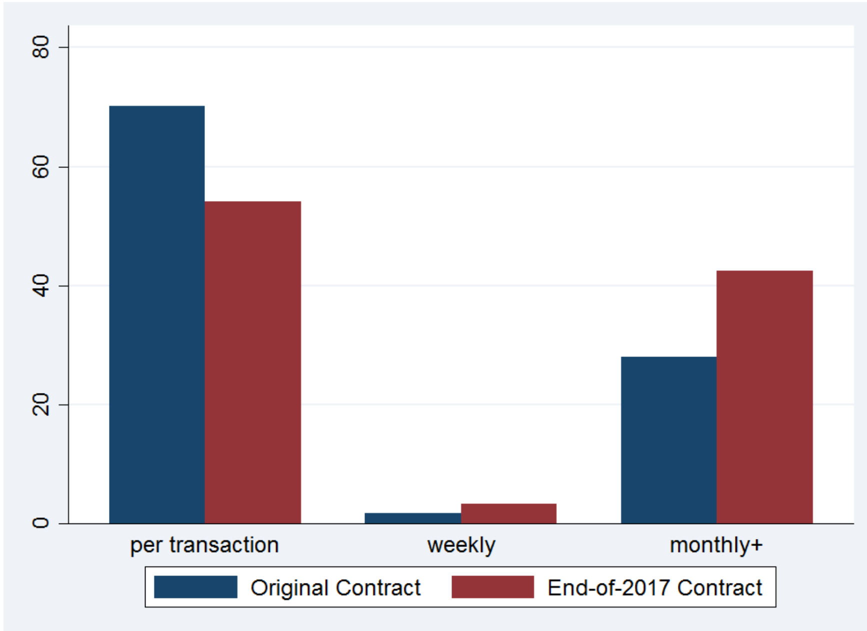

Note: No term specified on the clearing cycle is essentially the same as clearing after each transaction, so the first bar represents both the "clearing per transaction" and the "none specified" categories. Only two hospitals agree to quarterly clearing cycles and thus I group them with the monthly clearing cycle, which is plotted in the third bar.

Figure 2.7: Cumulative Physical Exam Packages Trend

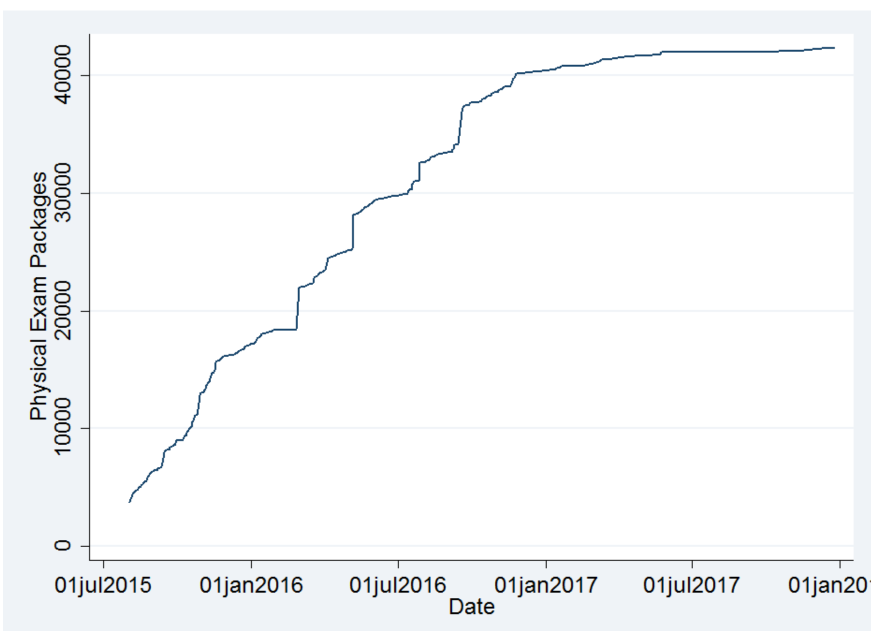

Figure 2.8: Hospital Hours

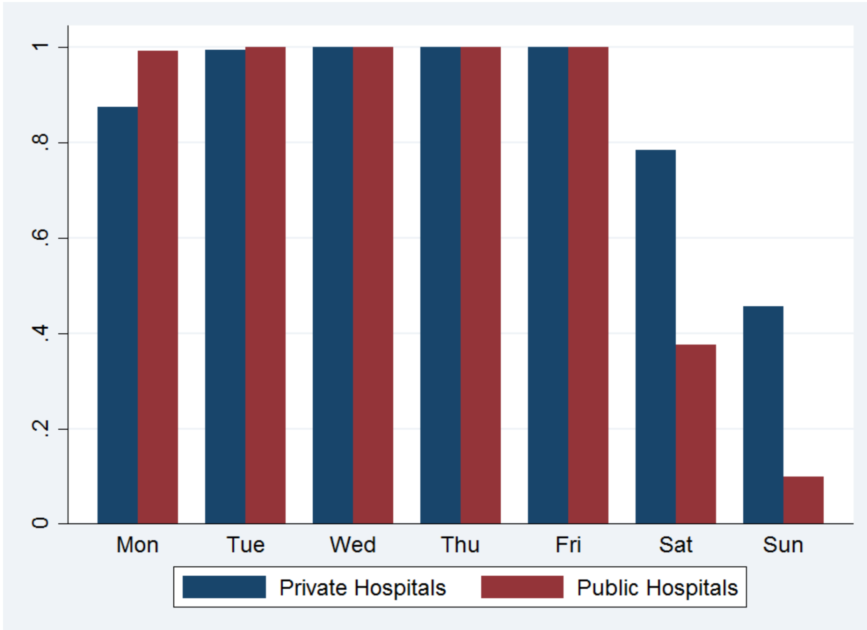

2.9 Tables

Table 2.1: Summary Statistics of Number of Hospital Visits

	Number of Obs.	Mean	Std. Dev.	Min	Max
All Hospitals	1,950	4.13	4.19	1	29
Contracted Hospitals	736	6.87	5.30	1	29
Non-contracted Hospitals	1,214	2.47	1.96	1	20

Table 2.2: Summary Statistics of the Depth of Discount Level

	Number of Obs.	Mean	Std. Dev.	Min	Max
First Contract	734	11.1%	18.2%	0	90%
The End-of-2017 Contract	690	14.6%	19.5%	0	90%
First Contract for Public Hospitals	417	6.5%	12.5%	0	90%
First Contract for Private Hospitals	317	17.1%	22.2%	0	90%
End-of-2017 Contract for Public Hospitals	394	9.2%	14.0%	0	90%
End-of-2017 Contract for Private Hospitals	296	21.7%	23.1%	0	90%

Table 2.3: Summary Statistics of Physical Exam Packages

	Number of Obs.	Mean	Std. Dev.	Min	Max
Original Price (Yuan)	11,544	1,608	1,857	0	40,500
Online Price (Yuan)	11,544	1,249	1,607	0	38,700
Number of Orders per Package	11,544	190	239	0	9,076

Table 2.4: Network Size Effects on Collaboration

VARIABLES	(1)	(2)	(3)	(4)
ln(Hospital Network Size)	0.180***		0.119***	
	(0.025)		(0.021)	
ln(Public Hospital Network Size)		0.097***		0.065***
		(0.023)		(0.020)
ln(Private Hospital Network Size)		0.116***		0.061***
		(0.023)		(0.021)
ln(Consumer Network Size)	0.060***	0.061***	0.026***	0.029***
	(0.005)	(0.005)	(0.005)	(0.005)
Branches	0.080***	0.081***	0.068***	0.068***
	(0.012)	(0.011)	(0.010)	(0.010)
Public AAA	-0.091***	-0.091***	-0.115***	-0.115***
	(0.015)	(0.016)	(0.015)	(0.015)
Public AA	-0.118***	-0.117***	-0.126***	-0.127***
	(0.019)	(0.019)	(0.018)	(0.018)
Public A	-0.222***	-0.223***	-0.152***	-0.153***
	(0.050)	(0.049)	(0.048)	(0.048)
Public Comprehensive	-0.130***	-0.129***	-0.110***	-0.110***
	(0.022)	(0.022)	(0.020)	(0.020)
Public BBB	-0.207***	-0.209***	-0.204***	-0.206***
	(0.034)	(0.034)	(0.031)	(0.031)
Public BB	0.086**	0.087**	0.090***	0.089**
	(0.038)	(0.038)	(0.035)	(0.035)
Others	-0.003	-0.003	-0.016	-0.016
	(0.035)	(0.035)	(0.032)	(0.032)
Chain Specialty	0.031	0.030	0.022	0.022
	(0.025)	(0.025)	(0.026)	(0.026)
Non-chain Specialty	0.046**	0.045**	0.045**	0.044**
	(0.022)	(0.022)	(0.020)	(0.020)
Number of Visits			0.038***	0.038***
			(0.002)	(0.002)
City FE	X	X	X	X
Month FE	X	X	X	X
Observations	8,036	8,036	8,036	8,036

Robust standard errors in parentheses. The default category of hospital type is "private", so all the estimated coefficients on hospital types are relative to private hospitals. When the number of hospital visits by the platform is less than 20 in a city, I use province fixed effects instead in the final four columns. The results are robust to different thresholds - 2, 5, and 10, in constructing city fixed effects. With the cutoff of 20, the city FE can still perfectly predict some outcomes, so the number of observations is smaller than that in the summary table.

Table 2.5: Network Size Effects on Discount

VARIABLES	(1)	(2)	(3)	(4)
ln(Hospital Network Size)	4.034***		3.931***	
	(1.216)		(1.215)	
ln(Public Hospital Network Size)		3.929***		3.862***
		(1.122)		(1.122)
ln(Private Hospital Network Size)		-0.276		-0.406
		(1.072)		(1.072)
ln(Consumer Network Size)	0.461*	0.568**	0.349	0.464*
	(0.264)	(0.270)	(0.273)	(0.278)
Branches	0.083	0.051	0.064	0.032
	(0.350)	(0.350)	(0.351)	(0.350)
Public AAA	-7.148***	-7.232***	-7.247***	-7.332***
	(0.788)	(0.791)	(0.784)	(0.787)
Public AA	-8.497***	-8.670***	-8.551***	-8.727***
	(0.925)	(0.933)	(0.924)	(0.932)
Public A	-6.558***	-7.002***	-6.152***	-6.597***
	(2.059)	(2.059)	(2.058)	(2.058)
Public Comprehensive	-2.057*	-2.231*	-1.969*	-2.147*
	(1.166)	(1.175)	(1.169)	(1.178)
Public BBB	-3.666	-4.296*	-3.654	-4.292*
	(2.235)	(2.260)	(2.231)	(2.257)
Public BB	-0.266	-0.484	-0.122	-0.339
	(1.658)	(1.664)	(1.659)	(1.665)
Others	-9.580***	-9.424***	-9.527***	-9.373***
	(1.489)	(1.494)	(1.500)	(1.504)
Chain Specialty	20.928***	21.051***	21.097***	21.223***
	(1.552)	(1.550)	(1.561)	(1.558)
Non-chain Specialty	4.955***	5.114***	5.006***	5.167***
	(1.210)	(1.214)	(1.213)	(1.216)
Number of Visits			0.091*	0.092*
			(0.050)	(0.050)
City FE	X	X	X	X
Month FE	X	X	X	X
Observations	3,794	3,794	3,794	3,794
R-squared	0.372	0.372	0.372	0.372

Robust standard errors in parentheses. The default category of hospital type is “private”, so all the estimated coefficients on hospital types are relative to private hospitals.

Table 2.6: Network Size Effects on Payment Method

VARIABLES	(1)	(2)	(3)	(4)
ln(Hospital Network Size)	-0.021		-0.144	
	(0.241)		(0.216)	
ln(Public Hospital Network Size)		0.183		0.114
		(0.204)		(0.188)
ln(Private Hospital Network Size)		-0.072		-0.186
		(0.209)		(0.196)
ln(Consumer Network Size)	0.535***	0.509***	0.446***	0.425***
	(0.052)	(0.050)	(0.050)	(0.049)
Branches	0.232***	0.232***	0.213***	0.212***
	(0.083)	(0.083)	(0.080)	(0.080)
Public AAA	-0.632***	-0.641***	-0.700***	-0.710***
	(0.146)	(0.146)	(0.143)	(0.143)
Public AA	-0.828***	-0.840***	-0.842***	-0.857***
	(0.178)	(0.179)	(0.181)	(0.182)
Public A	0.358	0.362	0.605	0.623
	(0.909)	(0.922)	(0.866)	(0.878)
Public Comprehensive	-0.423*	-0.434**	-0.350	-0.367*
	(0.220)	(0.221)	(0.217)	(0.218)
Public BBB	-0.135	-0.156	-0.183	-0.210
	(0.413)	(0.410)	(0.407)	(0.405)
Public BB	-0.807**	-0.821**	-0.673**	-0.686**
	(0.346)	(0.346)	(0.339)	(0.337)
Others	-0.009	0.004	-0.026	-0.014
	(0.262)	(0.265)	(0.255)	(0.258)
Chain Specialty	0.363*	0.372*	0.481**	0.491**
	(0.201)	(0.202)	(0.203)	(0.204)
Non-chain Specialty	-0.528***	-0.518***	-0.509***	-0.499***
	(0.167)	(0.168)	(0.164)	(0.164)
Number of Visits			0.092***	0.092***
			(0.014)	(0.014)
City FE	X	X	X	X
Month FE	X	X	X	X
Observations	3,794	3,794	3,794	3,794

Robust standard errors in parentheses. The default category of hospital type is “private”, so all the estimated coefficients on hospital types are relative to private hospitals.

Table 2.7: Network Size Effects on Clearing Cycle

VARIABLES	(1)	(2)	(3)	(4)
ln(Hospital Network Size)	1.309***		1.293***	
	(0.233)		(0.232)	
ln(Public Hospital Network Size)		0.954***		0.944***
		(0.200)		(0.199)
ln(Private Hospital Network Size)		0.492***		0.481**
		(0.189)		(0.189)
ln(Consumer Network Size)	0.222***	0.225***	0.212***	0.217***
	(0.053)	(0.053)	(0.055)	(0.054)
Branches	0.051	0.050	0.050	0.049
	(0.045)	(0.045)	(0.045)	(0.045)
Public AAA	-0.116	-0.132	-0.121	-0.136
	(0.120)	(0.120)	(0.121)	(0.120)
Public AA	-0.450***	-0.477***	-0.453***	-0.480***
	(0.151)	(0.151)	(0.151)	(0.151)
Public A	0.259	0.210	0.297	0.242
	(0.380)	(0.383)	(0.383)	(0.385)
Public Comprehensive	0.625***	0.612***	0.635***	0.620***
	(0.203)	(0.203)	(0.203)	(0.203)
Public BBB	1.266***	1.119***	1.271***	1.123***
	(0.407)	(0.398)	(0.406)	(0.397)
Public BB	-0.093	-0.133	-0.078	-0.120
	(0.225)	(0.224)	(0.226)	(0.225)
Others	-1.315***	-1.312***	-1.308***	-1.305***
	(0.245)	(0.247)	(0.247)	(0.248)
Chain Specialty	0.858***	0.873***	0.874***	0.886***
	(0.181)	(0.182)	(0.182)	(0.184)
Non-chain Specialty	0.103	0.117	0.109	0.122
	(0.162)	(0.161)	(0.162)	(0.162)
Number of Visits			0.007	0.006
			(0.007)	(0.007)
City FE	X	X	X	X
Month FE	X	X	X	X
Observations	3,794	3,794	3,794	3,794

Robust standard errors in parentheses. The default category of hospital type is "private", so all the estimated coefficients on hospital types are relative to private hospitals.

Table 2.8: Regression of Original Price on Package Attributes

	Original Price		
VARIABLES	Coef. (Std. Err.)	VARIABLES	Coef. (Std. Err.)
public AAA	249.960*** (74.566)	Cancer Screening Tests	-98.346* (58.681)
public AA	9.884 (84.174)	Gynecological Disorders	-118.443 (72.797)
public A	396.738** (189.749)	Thyroid Disease	402.993*** (60.834)
public comprehensive	198.101** (95.806)	Cardiovascular Disease	285.043*** (59.953)
public BB	-173.678* (105.292)	Gastroenterology Disease	427.727*** (62.556)
others	-247.709 (184.268)	Breast Cancer	-113.661 (71.725)
chain specialty	-267.199** (124.536)	Cervical Cancer	177.192** (73.616)
non-chain specialty	80.166 (87.364)	Osteoporosis	216.078*** (64.424)
Number of Total Tests	75.491*** (2.732)	Prostate Exam	-108.410 (71.523)
Instant Report Service	-26.276 (54.381)	Cervical Spine Exam	-236.575*** (80.664)
Male	-158.899*** (56.267)	Hepatitis B Test	21.857 (64.492)
Female	-95.553* (53.428)	Arthritis Test	197.759** (88.179)
Senior	167.327*** (63.530)	Nutrition Metabolism Test	317.822*** (83.713)
Adults	-115.661** (47.964)	Lumbar Disease	12.811 (110.022)
Young Adults	-138.622*** (46.823)	Blood Type Test	-226.744** (107.751)
Married	228.806*** (82.840)	STD test	385.317*** (118.032)
Unmarried	316.089*** (103.351)	Diabetes Diagnosis	1,643.192*** (198.336)
Expectant Mother	-297.162* (162.123)	Prenatal Postnatal Care	109.037 (193.716)
General Public	133.133*** (47.225)	Immune Disorder	-503.038* (258.185)
White-Collar Group	102.561* (60.542)	Genetic Disorder	2,343.852*** (251.889)
Hepatobiliary Disease	-321.910*** (58.230)	Allergy Test	2,971.747*** (280.369)
Lung Disease	-408.930*** (58.427)		
Province FE		X	
Observations		5,735	
R-squared		0.484	
*** p<0.01, ** p<0.05, * p<0.1			

The default category of hospital type is "private", so all the estimated coefficients on hospital types are relative to private hospitals.

Table 2.9: Product Variety and Number of Orders Between Public and Private Hospitals

	Public Hospitals			**Private Hospitals**			**t-tests**
	Obs.	Mean	Std. Err.	Obs.	Mean	Std. Err.	t-stat
Dec 2017 Product Variety	291	13.41	0.56	203	11.58	0.59	2.19
All-time Product Variety	291	18.25	0.88	203	15.41	0.92	2.17
Number of Orders per Package	2,453	238.9	8.5	1,472	185.5	4.3	4.66

The two-sample t test is run with the equal variance assumption. I find statistically significant differences in all outcomes at 0.05 level.

Table 2.10: Hospital Ratings Between Public and Private Hospitals

	Public Hospitals			**Private Hospitals**			**t-tests**
	Obs.	Mean	Std. Err.	Obs.	Mean	Std. Err.	t-stat
Rating - Expertise	223	4.71	0.015	119	4.78	0.02	2.64
Rating - Service	223	4.49	0.014	119	4.52	0.021	1.34
Rating - Environment	223	4.49	0.014	119	4.51	0.021	0.65
Number of Ratings	223	17.9	1.76	119	10.1	0.7	3.16

The two-sample t test is run with the equal variance assumption. I find statistically significant differences in Rating - Expertise and Number of Ratings at the 5% level.

Bibliography

Ackerberg, D. A. and G. Gowrisankaran (2006). Quantifying equilibrium network externalities in the ach banking industry. The RAND Journal of Economics 37(3), 738–761.

Armstrong, M. (2006). Competition in two-sided markets. The RAND Journal of Economics 37(3), 668–691.

Arnold, M. A. and S. A. Lippman (1998). Posted prices versus bargaining in markets with asymmetric information. Economic Inquiry 36(3), 450–457.

Ausubel, L. M., P. Cramton, and R. J. Deneckere (2002). Bargaining with incomplete information. In Handbook of Game Theory, Volume 3, Chapter 50, pp. 1897–1945. Elsevier.

Ayres, I. and P. Siegelman (1995). Race and gender discrimination in bargaining for a new car. The American Economic Review 85(3), 304–321.

Backus, M., T. Blake, and S. Tadelis (2016). On the empirical content of cheap-talking signaling: An application to bargaining. Working Paper.

Beckert, W., H. Smith, and Y. Takahashi (2016). Competition and price discrimination in a spatially differentiated product market with negotiated prices. Working Paper.

Belotti, F., P. Deb, W. G. Manning, and E. C. Norton (2015). twopm: Two-part models. Stata Journal 15(1), 3–20.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. Econometrica 63(4), 841–890.

Bester, H. (1994). Price commitment in search markets. Journal of Economic Behavior & Organization 25(1), 109–120.

Brown, J. R., S. M. Fazzari, and B. C. Petersen (2009). Financing innovation and growth: Cash flow, external equity, and the 1990s r&d boom. The Journal of Finance 64(1), 151–185.

Brynjolfsson, E., Y. Hu, and M. D. Smith (2003). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. Management Science 49(11), 1580–1596.

Caillaud, B. and B. Jullien (2003). Chicken & egg: Competition among intermediation service providers. The RAND Journal of Economics 34(2), 309.

Chu, J. and P. Manchanda (2016). Quantifying cross and direct network effects in online consumer-to-consumer platforms. Marketing Science 35(6), 870–893.

Clements, M. T. and H. Ohashi (2005). Indirect network effects and the product cycle: Video games in the u.s., 1994-2002*. Journal of Industrial Economics 53(4), 515–542.

Cragg, J. G. (1971). Some statistical models for limited dependent variables with application to the demand for durable goods. Econometrica 39(5), 829.

Crawford, G. S. and A. Yurukoglu (2012). The welfare effects of bundling in multichannel television markets. American Economic Review 102(2), 643–685.

Dafny, L. S. (2010). Are health insurance markets competitive? American Economic Review 100(4), 1399–1431.

Desai, P. S. and D. Purohit (2004). "let me talk to my manager": Haggling in a competitive environment. Marketing Science 23(2), 219–233.

Draganska, M., D. Klapper, and S. B. Villas-Boas (2010). A larger slice or a larger pie? an empirical investigation of bargaining power in the distribution channel. Marketing Science 29(1), 57–74.

Ellickson, P. B., P. Kong, and M. J. Lovett (2017). Private labels and retailer protability: Bilateral bargaining in the grocery channel. Working Paper.

Fan, Y. (2013). Ownership consolidation and product characteristics: A study of the us daily newspaper market. American Economic Review 103(5), 1598–1628.

Fan, Y., J. Ju, and M. Xiao (2016). Reputation premium and reputation management: Evidence from the largest e-commerce platform in China. International Journal of Industrial Organization 46, 63–76.

Gowrisankaran, G., A. Nevo, and R. Town (2015). Mergers when prices are negotiated: Evidence from the hospital industry. American Economic Review 105(1), 172–203.

Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. American Economic Review 103(1), 145–177.

Grennan, M. (2014). Bargaining ability and competitive advantage: Empirical evidence from medical devices. Management Science 60(12), 3011–3025.

Ho, K. (2009). Insurer-provider networks in the medical care market. American Economic Review 99(1), 393–430.

Holtz-Eakin, D., D. Joulfaian, and H. S. Rosen (1994). Sticking it out: Entrepreneurial survival and liquidity constraints. Journal of Political Economy 102(1), 53–75.

Huang, G. (2012). Posted price and haggling in the used car market. Working Paper.

Israeli, A., E. T. Anderson, and A. T. Coughlan (2016). Minimum advertised pricing: Patterns of violation in competitive retail markets. Marketing Science 35(4), 539–564.

Jindal, P. and P. Newberry (2016). To bargain or not to bargain - the role of fixed costs in price negotiations. Working Paper.

Kaiser, U. and J. Wright (2006). Price structure in two-sided markets: Evidence from the magazine industry. International Journal of Industrial Organization 24(1), 1–28.

Keniston, D. E. (2011). Bargaining and welfare: A dynamic structural analysis of the autorickshaw market. Working Paper.

Lafontaine, F. and M. E. Slade (2008). Exclusive contracts and vertical restraints : empirical evidence and public policy, pp. 391–414. Cambridge, MA: MIT Press.

Larsen, B. (2015). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. NBER Working Paper.

Lee, R. S. (2015). Empirical models of bilateral contracting. pp. 1–9.

Lewis, M. S. and K. E. Pflum (2017). Hospital systems and bargaining power: evidence from out-of-market acquisitions. The RAND Journal of Economics 48(3), 579–610.

Mcfadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical-integration. Econometrica 57(5), 995–1026.

Meza, S. and K. Sudhir (2010). Do private labels increase retailer bargaining power? Quantitative Marketing and Economics 8(3), 333–363.

Morton, F. S., J. Silva-Risso, and F. Zettelmeyer (2011). What matters in a price negotiation: Evidence from the U.S. auto retailing industry. Quantitative Marketing and Economics 9(4), 365–402.

Nevo, A. (2000). A practitioner's guide to estimation of random-coefficients logit models of demand. Journal of Economics & Management Strategy 9(4), 513–548.

Oliveira, B. and A. Fortunato (2006). Firm growth and liquidity constraints: A dynamic analysis. Small Business Economics 27(2-3), 139–156.

Pattabhiramaiah, A., S. Sriram, and S. Sridhar (2018). Rising prices under declining preferences: The case of the u.s. print newspaper industry. Marketing Science 37(1), 97–122.

Petrin, A. and K. Train (2010). A control function approach to endogeneity in consumer choice models. Journal of Marketing Research 47(1), 3–13.

Riley, J. and R. Zeckhauser (1983). Optimal selling strategies - when to haggle, when to hold firm. Quarterly Journal of Economics 98(2), 267–289.

Rochet, J.-C. and J. Tirole (2006). Two-sided markets: a progress report. The RAND Journal of Economics 37(3), 645–667.

Rubin, J. Z. and B. R. Brown (1975). The Social Psychology of Bargaining and Negotiation. Elsevier.

Rysman, M. (2004). Competition between networks: A study of the market for yellow pages. The Review of Economic Studies 71(2), 483–512.

Rysman, M. (2009). The economics of two-sided markets. Journal of Economic Perspectives 23(3), 125–143.

Shelegia, S. and J. Sherman (2015). Bargaining at retail stores: Evidence from Vienna. Working Paper.

Song, M. (2013). Estimating platform market power in two-sided markets with an application to magazine advertising. Working Paper.

Sriram, S., P. Manchanda, M. E. Bravo, J. Chu, L. Ma, M. Song, S. Shriver, and U. Subramanian (2014). Platforms: a multiplicity of research opportunities. Marketing Letters 26(2), 141–152.

Sun, L., S. Rajiv, and J. Chu (2016). Beyond the more the merrier: The variety effect and consumer heterogeneity in system markets. International Journal of Research in Marketing 33(2), 261–275.

Wang, R. Q. (1995). Bargaining versus posted-price selling. European Economic Review 39(9), 1747–1764.

Weyl, E. G. (2010). A price theory of multi-sided platforms. American Economic Review 100(4), 1642–1672.

Wilbur, K. C. (2008). A two-sided, empirical model of television advertising and viewing markets. Marketing Science 27(3), 356–378.

Wooldridge, J. M. (2010). Econometric analysis of cross section and panel data (2nd ed.). MIT Press.

Zhong, Z. (2016). Chasing diamonds and crowns: Online reputation systems and seller response. Working Paper.