"Meet Me Halfway": The Value of Bargaining*

Xu Zhang[†], Puneet Manchanda[‡], and Junhong Chu[§]

[†]London Business School
[‡]Ross School of Business, University of Michigan
[§]NUS Business School, National University of Singapore

November 2016
This version: November 2018

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. We 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. We incorporate the decision to bargain, the bargaining realization, and the purchase decision into a structural model. Using our results, we perform counterfactual analyses to derive the value of allowing bargaining on the platform. We do this by disallowing bargaining such that all sellers on the platform move to a fixed-price mechanism. We find that sellers' profits are higher after the policy change. These benefits are heterogeneous across sellers, with sellers with low reputation levels, high detailed seller ratings, and non-promotion products benefiting more. We also show that buyers' bargaining cost savings are economically significant. Thus, our findings suggest that banning bargaining is beneficial from both buyers' and a social planner's perspective. We provide some reasons for why bargaining still exists on the platform. Finally, we show that our results are robust to our assumptions and can be replicated.

Keywords: Bargaining, Pricing, Platforms, Digital Markets, Structural Models, Alibaba, China

*The authors can be contacted at *xzhang@london.edu*, *pmanchan@umich.edu*, and *bizcj@nus.edu.sg*. This paper is based on one of the essays that comprise the first author's dissertation. The authors are grateful to Alibaba for providing the data, Francine Lafontaine, Dan Ackerberg, Kanishka Misra, Pranav Jindal, Chenyu Yang, Danial Asmat, seminar participants at the Chinese University of Hong Kong, Cornell University, Erasmus School of Economics, Haring Symposium at Indiana University, HEC Montreal, London Business School, McGill University, Nanyang Technological University, National University of Singapore, Peking University (GSM and PHBS), Renmin University of China, Southern Methodist University, Syracuse University, University of Michigan, and UT Dallas for feedback.

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]

Extensive theoretical research has been carried out comparing the bargaining mechanism with the fixed-price mechanism, 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, 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 (we use the terms "buyers" and "consumers" interchangeably throughout the paper). In order to derive these gains (from bargaining), sellers are likely to post higher prices, driving away some buyers as a result. 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 the 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 which pricing mechanism is more beneficial.

In this paper, we 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 on a very large digital marketplace (platform) in China, combined with survey data from Chinese buyers on their propensity to bargain and their expectation of success. We do this by building a structural model of consumer demand and sellers' pricing decisions where we 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.

[1] Amazon introduced "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 had the bargaining mechanism included in the platform by design since its inception in 2003.

Specifically, we estimate a three-stage demand model. 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, 2018). We 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 by the variation in the decision to bargain and the expected bargaining outcome. In the second stage, we model the bargaining realization process between a seller and a buyer. While the Nash bargaining solution has been widely used in the literature (e.g., Beckert et al., 2016; Crawford and Yurukoglu, 2012; Draganska et al., 2010; Ellickson et al., 2017; Gowrisankaran et al., 2015; Grennan, 2013; Jindal and Newberry, 2018), its validity hinges on the correct specification of each party's disagreement payoff and it ties down the gains from bargaining, which is inconsistent with our data pattern. 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, we employ a reduced-form approach by flexibly modeling the bargaining realization using a two-part model to avoid imposing strong assumptions and to better match the data. The first part captures whether bargaining succeeds and the second part captures the discount conditional on success. In the third stage, we use a modified consumer discrete choice approach to model the purchase behavior while taking the bargaining possibility into consideration, including the use of a control function to address any potential endogeneity issues. Finally, on the supply side, we assume that sellers are setting prices to maximize profits while allowing for the possibility of buyer bargaining.

The results from our analysis provide rich insights. First, we 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 which ranges from 11 to 20 yuan. We also find that the variation in the average bargaining costs is related to economic development with the costs being higher in more developed provinces,

[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 our approach, all costs due to bargaining - psychological, time, etc. - result in an anticipated loss of utility by a buyer prior to initiating bargaining.

reflecting the fact that buyers in these provinces place a bigger premium 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), we 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., 2018; 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, we find that a 1% increase in posted price on average leads to a 3.3% 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, we perform counterfactual analysis to derive the value of bargaining in the online marketplace. We do this by disallowing bargaining, i.e., "forcing" sellers to move to a fixed-price mechanism. We find that the average posted price decreases by 1.1% but the average transaction price remains about the same after the move. Interestingly, we 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 turning away consumers, due to the higher posted prices, 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 small. Specifically, 97% of the sellers are better off while the remaining 3% are worse off as a result of the policy change. Not surprisingly, 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 around 160 million yuan (26 million dollars) in the focal category, representing a 0.4% increase. The saved bargaining costs for buyers per day are at 6% of the daily GMV. Thus, the analysis suggests that banning bargaining is modestly beneficial for an average seller, but is greatly beneficial from both buyers' and a social planner's point of view. Note that due to data limitations, we are not able to identify sellers' bargaining costs. Thus, our 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 the platform has not implemented a fixed-price mechanism in the marketplace. While we discuss this in detail later, the short answer lies in the history of the marketplace's evolution, cultural norms, small benefit for an individual seller, and the very large number of heterogeneous sellers. Finally, we show that our results are robust to our model assumptions and structure and not idiosyncratic to one product category (results from one other product category are qualitatively similar).

Our paper differs from previous empirical work on bargaining (Beckert et al., 2016; Grennan, 2013; Keniston, 2011; Huang, 2012; Jindal and Newberry, 2018) in several important ways, and thus makes the following contributions to both marketing and economics literatures on bargaining. First, previous studies have treated transactions with zero bargained discount as equivalent to no-bargain transactions. However, in our paper, we 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 either depends on functional form assumptions or stringent data requirements. In our paper, we instead combine secondary (transaction) data with primary (survey) data, which enables us 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 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 and bargaining gains are not gauranteed. Lastly, we not only examine the effect of pricing policy change on individual players in the market,

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

but also at the platform level. Given the prevalence of platform markets in e-commerce worldwide, our analysis can serve as a template to investigate the impact of different pricing mechanisms.

The rest of the paper is organized as follows. We describe the institutional setting and the data in §2. §3 describes the model. The identification and estimation of the demand side model are detailed in §4 while the supply side is described in §5. The results are discussed in §6 and the counterfactual analysis and the robustness checks in §7. We conclude in §8.

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. Taobao.com's market share of e-commerce in China is over 80%.[4]

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 Aliwangwang - a free Skype-like online chatting service. All buyers and sellers get an Aliwangwang account automatically when they register on Taobao (the tool can also be used within a web browser or via a mobile app). Figure 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

[4] *About Taobao*BBT:2018. 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.

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/his shopping cart (note that this prevents other buyers from seeing this "adjusted" price).

Figure 1: Taobao Item Page

As can be seen from the above, the set up 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, we do not observe the Aliwangwang chat history for transactions. As a result, it is hard for us to impose bargaining theory primitives or equilibrium concepts. We therefore specify a flexible model to capture the bargaining outcome as a function of a rich set of related variables.

2.2 Transaction Data

Our main (secondary) data include a random sample of transactions on cellphones from January 1, 2012 to May 25, 2013 on Taobao.com. This product (cellphone), is particularly suited to study the comparison between the fixed-price and bargaining mechanisms on the e-commerce platform for two

[5]In our Taobao Consumer Survey, described in §2.3, we 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 our sample, 97% have an average online chat waiting time less than 3 minutes.

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, making cellphones the third largest category out of a total of 113 categories on Taobao. Our sample consists of 39,625 transactions made by 24,181 buyers with 8,965 sellers.[6]

For each transaction, we 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), we extract product attributes from each transaction entry to identify a product. Specifically we identify a product using information on brands and models. For certain products with memory capacity options, like iPhones, we also employ memory sizes to facilitate the product identification process. For those cellphone models with less than 200 transactions, we use brands as an analysis unit. Our 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 presents the summary statistics of the transaction sample. From Figure 1, we 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 the 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. Since sellers have full freedom to set both the list price and the promotion depth with zero cost on most occasions, we use only the "Posted Price" (the price consumers should really care about) and the "Promotion Indicator" (the major feature that consumers react to instead of the promotion depth, see Mayhew and Winer (1992)) in the analysis.

A successful bargain occurs if the transaction price is less than the posted price. Based on this

[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, we exclude them from our sample.

[7]To make the product definition clearer, we 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 our results to this product definition where we aggregate less popular models to the brand level, we 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, we believe our results are not sensitive to this product definition.

Table 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 (years)	39,625	2.31	2.00	0	9.9
Site Age (years)	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 (years)	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

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 bargaining 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 our sample, 87% are observed to have offered bargaining discounts in at least one transaction. The high proportion of sellers who have ever offered a bargaining discount and the relatively lower overall proportion of successful bargained transactions together suggest that bargaining is a transaction-specific phenomenon rather than a seller-specific phenomenon, leading us to use transactions as our 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 us to use

[8] According to our conversation 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 our 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. We 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 our results.

these data to compare the value of fixed price and bargaining mechanisms.

About 11% of the transactions are repeat purchases. We therefore create a repeat purchase indicator for each transaction, defined as 1 if the focal transaction is at least the second interaction between the seller and the buyer.[9] Additionally, we 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 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 3) shown on each product page (Figure 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. We therefore include both measures in the analysis. Since the three dimensions in the detailed seller ratings are highly

[9] Given that we do not have information prior to the sample period, this is likely to be a imperfect measure. However, we include as it could have an effect on bargaining outcomes.

Figure 2: Taobao Feedback Page

correlated, we 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 addition to the two reputation measures, we 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. We 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. We see a 70%-30% male-female split in our sample. For those readers who are interested, we report major summary statistics by gender in Appendix A Table A1 and find no statistically significant difference in either the products purchased or the bargaining outcomes across gender. Using the location information, we supplement income data for each buyer using the province average.

Finally, for each seller, though we do not observe individual site visits that do not result in

Figure 3: Taobao Reputation Rule

purchases by a buyer, we do observe the number of unique visitors and unique purchasers on a four-week rolling basis. We use the ratio between the two to define the conversion rate and to get a measure of the proportion of the no-purchase visits at the seller/product combination level.

2.3 Taobao Consumer Survey

One of our 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, we supplement the principal transaction data with a Taobao Consumer Survey. We conducted the survey with Taobao customers via www.sojump.com, the largest online survey website in China. The survey ran for 10 days (December 20-29, 2016) and we obtained 1,009 responses from 31 mainland provinces. We also carried out a (similar) pilot survey in 2015 with 566 respondents.

We asked questions related to general knowledge on the practice of bargaining on Taobao, and 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 yuan?", (2) "How likely would you expect to succeed in bargaining?"

and (3) "How much discount would you expect to get if bargaining succeeds?" The reason that we intentionally hold the product fixed and stay "ignorant" about product and seller characteristics is to make the data collection feasible in practice. Ideally, we would like to design the survey such that the bargaining intention reflects the variation across product characteristics as well. However, the space of all combinations of product and seller characteristics in the transaction is very large. This means that to obtain a precise bargaining intention estimate for each combination, we need to expand our sample size tremendously (e.g., just using products, unique prices and seller reputation levels in our sample gives us over 20,000 combinations, requiring survey samples in the hundreds of thousands). Thus, instead of getting responses on each individual product-seller combination, we picked an abstract product without specifying too many product or seller characteristics except for clearly stating the category of interest and the price at a level that is representative for the transaction sample. Using such a design and under the assumption that consumers' bargaining intention and perceived success rates in the survey are good proxies for those in the transaction sample at the aggregate level, we believe the survey is efficient in terms of getting the necessary information from a reasonable number of respondents, and sufficient in terms of providing the necessary variation to identify bargaining costs. We also test the robustness of our results to this assumption in subsection 7.3.1.

We use the first question to infer the distribution of consumers' bargaining intention. There are three options to choose from: A. Yes; B. No; C. Maybe. We define a variable indicating a consumer chooses "for sure to bargain" as $I(\textit{Certainly Bargain}) = I(\textit{Option} = A)$, and a variable indicating a consumer may bargain as $I(\textit{Certainly + May Bargain}) = I(\textit{Option} = A \textit{ or } C)$. Since we cannot match the survey with the transaction sample at the individual level, we take the perspective of a representative buyer in a province. We assume the survey responses on bargaining intention represent a sample of random draws from the distribution of the bargaining intention of the representative buyer (from the respondent's province). So the average across all survey individuals who responded for $I(\textit{Option} = A)$ denotes the lower bound for the bargaining intention for the representative buyer and the average across all survey individuals who responded for $I(\textit{Option} = A \textit{ or } C)$ denotes the upper bound for the bargaining intention for the representative buyer. The realized bargaining intention for the representative buyer then should lie between these two bounds, i.e., between the probability of "sure to bargain" and "may bargain."

The second question provides a measure of the perceived success rate conditional on bargaining among Taobao consumers. There are ten options to choose from, ranging from 10% to 100%. We 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 gain.

We 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), we believe the survey respondents represent Taobao consumers well and Taobao consumers on average form reasonable expectations on the conditional bargaining discount amounts. Figure 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.

Figure 4: A Comparison of Age Distribution between the Transaction Sample and the Survey

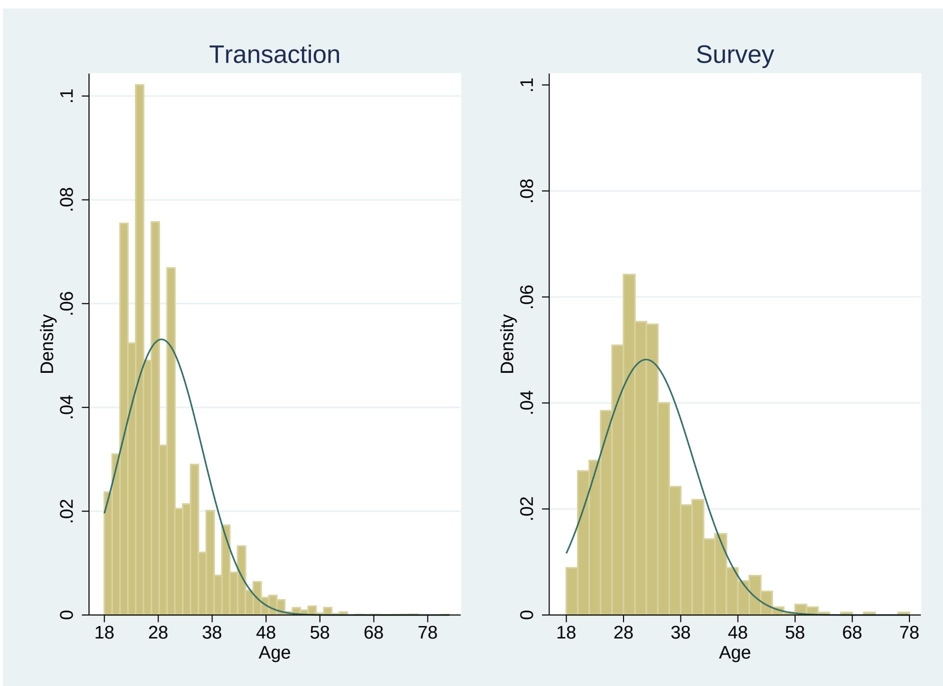

Table 2 summarizes the survey responses. The mean lower bound of bargaining intention is 54.3%, the mean upper bound is 73.6%, and the perceived success rate conditional on bargaining is 49%. In estimation, both bargaining intention and perceived success rate conditional on bargaining are used at the province level. We group Ningxia, Qinghai, and Tibet provinces together given their geographic and socioeconomic proximity to reduce measurement errors. 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%. The variance of the estimated province-level bounds is around 0.006.[10] A detailed summary breakdown by gender is provided in Appendix A Table A2. Even though we 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, consistent with previous findings e.g., Small et al. (2007).

Table 2: 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 “maybe” 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.

3 The Consumer Bargaining and Purchase Model

We first construct an empirical model of consumers’ bargaining and purchasing decisions to describe the demand side of the market. The key objective of our model is to estimate 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 5. In Stage I, a consumer decides whether to bargain by comparing the expected bargaining gain and her/his inherent bargaining cost. If the consumer decides to bargain, s/he proceeds to Stage II (bargaining realization), while if the consumer decides not to bargain, s/he 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, we are not able to model the search behavior explicitly. Instead, we use proxy measures to control for two common search processes employed by

[10] The summary statistics are similar across the main survey and the pilot survey, suggesting that they are stable.

consumers, as detailed later. We describe the model by working backwards from the last stage of the demand model. To facilitate reading, we provide a list of notation definitions in Table 3.

Figure 5: The Bargaining and Purchase Process

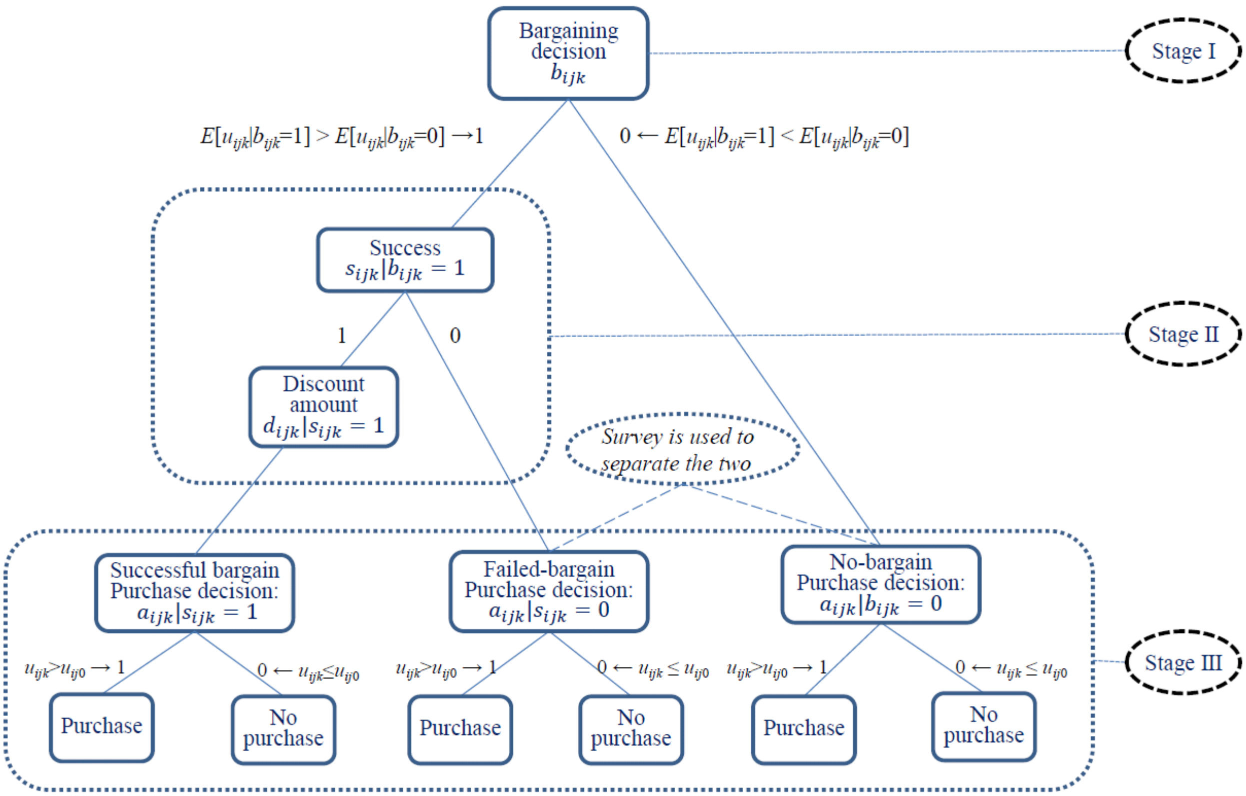

3.1 Stage III: Purchase Decision

While consumers search for products on e-commerce marketplaces, they usually pay great attention to seller characteristics, such as seller reputation and seller average ratings, to infer product quality. Thus, we define a shopping occasion as a product-seller combination, i.e. a consumer considering the product that s/he is interested in being sold by a specific seller. 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

$$u_{ijk} = \delta_k - \alpha_i p_{ijk} + x_{jk}\beta_x + w_{jk}\beta_w + \varepsilon_{ijk} \tag{1}$$

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. α_i captures the consumer's heterogeneous price sensitivity. We formulate α_i as

Table 3: 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
w_{jk}	Two measures to control for search processes and seller competition
x_{ijk}	Represents $\{x_{jk}, w_{jk}, \text{buyer characteristics}, \bar{p}_{ijk}, \delta_k\}$ for brevity
$\alpha_i, \alpha_{inc}, \alpha_{exp}$	Heterogeneous price sensitivity parameters
β_x	Preference for seller/product characteristics
β_w	Impact of competition on outside option valuation
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 success function
ν_d	Unobservable in the discount function
z_{ij}	Instrumental variable - distance to the next reputation level
μ_{jk}	Residual from the control function, which is used to address the price endogeneity
m_{jk}	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 bargaining costs of a representative consumer
mc_{jk}	Marginal cost for seller j on product k
Π_{jk}	Profit for seller j on product k

$(\alpha + \alpha_{inc} log(Income_i) + \alpha_{exp} ShoppingExperience_i)$ such that both the buyer's income level and the shopping experience can have an impact on the price sensitivity. x_{jk} is a vector of seller/product characteristics, including seller reputation level, detailed seller rating, site age, promotion indicator, product age, and the repeat purchase indicator. β_x captures the consumer's 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, we create two measures w_{jk} to capture two common search processes and to control for seller competition. One is the number of sellers who are selling the same focal product in that month and the other is the number of sellers who have the same reputation level as the seller in the focal transaction in that month. Depending on how a buyer values the product *per se* and how the buyer values seller characteristics, these two measures can have different impact on the valuation of the outside option, i.e., different impact on the purchase decision. We added these two variables to u_{ijk}, which is essentially the same as reformulating the outside option u_{ij0}. After the addition of the two variables, a change in the number of competing sellers will result in a change of the utility of the outside option. At the same time, the two variables also reflect the competition at the platform level via two types of buyers' possible search processes. The first is when a buyer searches for a specific product across different sellers and the second is when a buyer searches for products offered by only reputable sellers. Note that these two search strategies mirror the two search paths that the Taobao search engine provides consumers on its landing page (www.taobao.com) – search for product and search for store.

The term ε_{ijk} is the consumer's unobserved utility. An endogeneity problem may arise if ε_{ijk} includes unobserved characteristics that affect utility and price at the same time, such as a seller's word of mouth or the attractiveness of a website design. It is likely that these omitted variables can drive up price and demand at the same time. As a result, without controlling for the omitted variables, the price coefficient would be overestimated (less negative). To solve this problem, we 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 price p_{jk} can be written as

$$p_{jk} = g(x_{jk}, w_{jk}, z_{jk}, \mu_{jk}) \tag{2}$$

where x_{jk} and w_{jk} are defined as above denoting exogenous variables directly entering the utility function; z_{jk} represents the instrumental variables that affect the price but do not affect utility directly (we discuss the specific variables operationalized in §4.4); and μ_{jk} is the unobserved factor that affects the price and is potentially correlated with ε_{ijk}. The existence of μ_{jk} is the source of the dependence between p_{jk} and the error term ε_{ijk}.

The key to the control function approach is that conditional on μ_{jk}, ε_{ijk} is independent of u_{ijk}. Thus, we first estimate a proxy for μ_{jk} using the standard OLS estimator, and then estimate parameters conditional on this proxy variable μ_{jk}. To illustrate, we formulate the utility as follows when there is a single unobserved factor μ_{jk}:

$$u_{ijk} = \delta_k - \alpha_i p_{ijk} + x_{jk}\beta_x + w_{jk}\beta_w + CF(\mu_{jk}; \beta_\mu) + \tilde{\varepsilon}_{ijk} \tag{3}$$

where $CF(\mu_{jk}; \beta_\mu)$ denotes the control function with parameters β_μ. We specify the control function as linear such that $CF(\mu_{jk}; \beta_\mu) = \beta_\mu \mu_{jk}$. Note that, after including the control function, $\tilde{\varepsilon}_{ijk}$ is independent of all the explanatory variables. If we have more than one unobserved factor, the above formulation can be easily extended by using a vector of μ. 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

$$Pr_{ijk} = \int I(u_{ijk} > u_{ij0}) f(\tilde{\varepsilon}_{ijk}) d\tilde{\varepsilon}_{ijk} \tag{4}$$

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

$$m_{jk} = \int Pr_{ijk}(p_{ijk}) f(p_{ijk}) dp_{ijk} \tag{5}$$

where $f(p_{ijk})$ is the realized price distribution 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, we can compute the integral for the conversion rate for each product-seller combination.

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 bargaining solution. However, the validity of Nash bargaining framework depends on the correct specifications of each party's disagreement payoff, which is unavailable in our context, especially the marginal cost at the seller-product level. The second drawback of the Nash bargaining solution is that it ties down the gains from bargaining, which is inconsistent with our data pattern. Also, the prediction power of the Nash bargaining solution in the outcomes of buyer-seller bargaining is questionable (Neslin and Greenhalgh, 1986). 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, we 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, we 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:

$$p(s_{ijk} = 1 | b_{ijk} = 1) = \Phi(\frac{x_{ijk}\gamma}{\sigma_s}) \tag{6}$$

where Φ is the cumulative distribution function of the standard normal. x_{ijk} represents a vector of seller, buyer, and product characteristics. Note that x_{ijk} is different from x_{jk} in equation (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 conversion rate is defined for a seller only. x_{ijk} also includes the posted price, the proxies for seller competition level, and the product fixed effects for brevity, so x_{ijk} essentially represents $[x_{ijk}, w_{jk}, \bar{p}_{ijk}, \delta_k]$. Note that product fixed effects can partially capture the effect of the marginal cost on the bargaining outcomes. 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, we 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, we use a logarithmic transformation. Specifically, we assume a latent variable d^*_{ijk}

following the distribution of:

$$d_{ijk}^* = x_{ijk}\theta + \upsilon_d \qquad \upsilon_d \sim N(0, \sigma_d) \tag{7}$$

$$d_{ijk} = exp(d_{ijk}^*) - 1 \qquad \text{if } d_{ijk}^* > 0 \tag{8}$$

We 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 (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 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.

Note that even though we used a "reduced form" to describe the bargaining realization stage, that does not necessarily mean that the discount from bargaining is exogenous. As can be seen from the model specification, the bargaining success and the bargaining discount both depend on who the seller is and who the buyer is. Incorporating seller and buyer characteristics in the model essentially implies that the bargaining realization process is endogenously co-determined by the seller and the buyer. For example, bargaining with a high reputable seller may yield a smaller discount, while bargaining with a low reputable seller may yield a higher discount. The unobservability of the actual bargaining process limits our ability to specify a particular bargaining protocol, but the goal of this stage - bargaining realization - is not to describe in detail the back-and-forth bargaining process, but instead to accurately predict the outcome of a bargaining process. Given that, a seller can then endogenize the predicted bargaining outcomes when making the pricing decision to maximize profits, as will be shown in Section 5.

3.3 Stage I: bargaining decision

We assume consumers are rational in the market. When a consumer browses the product of interest, s/he first makes a bargaining decision by weighing her/his expected utility if choosing to bargain against that if choosing not to bargain. Thus, a consumer's bargaining decision b_{ijk} follows:

$$b_{ijk} = \begin{cases} 1 & \text{if } E[\text{utility if bargain}] > E[\text{utility if not bargain}] \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

where $E[\text{utility if bargain}] = E[max\{u_{ijk}(\bar{p}_{jk} - d_{ijk} + c_i), u_{ij0} - \alpha_i c_i\}]$, 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 the consumer's 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.

4 Identification and Estimation of the Demand

As explained before, to overcome the data limitations, we 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, we 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 6):

Step one: We first estimate the discount amount conditional on success using the truncated regression model (equation 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, we 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 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, we recover the parameters governing the bargaining success function using the simulated maximum likelihood approach (equation 6).

Step four: For each seller-product combination, we calculate the conversion rate by integrating over potential shopping occasions made by a potential pool of consumers at different transaction prices. The primitives in the utility function (equation 3) are estimated by minimizing the distance between the predicted and the observed conversion rates using the generalized method of moments. We use the control function to deal with the endogeneity problem.

Figure 6: The Flow of the Estimation Process

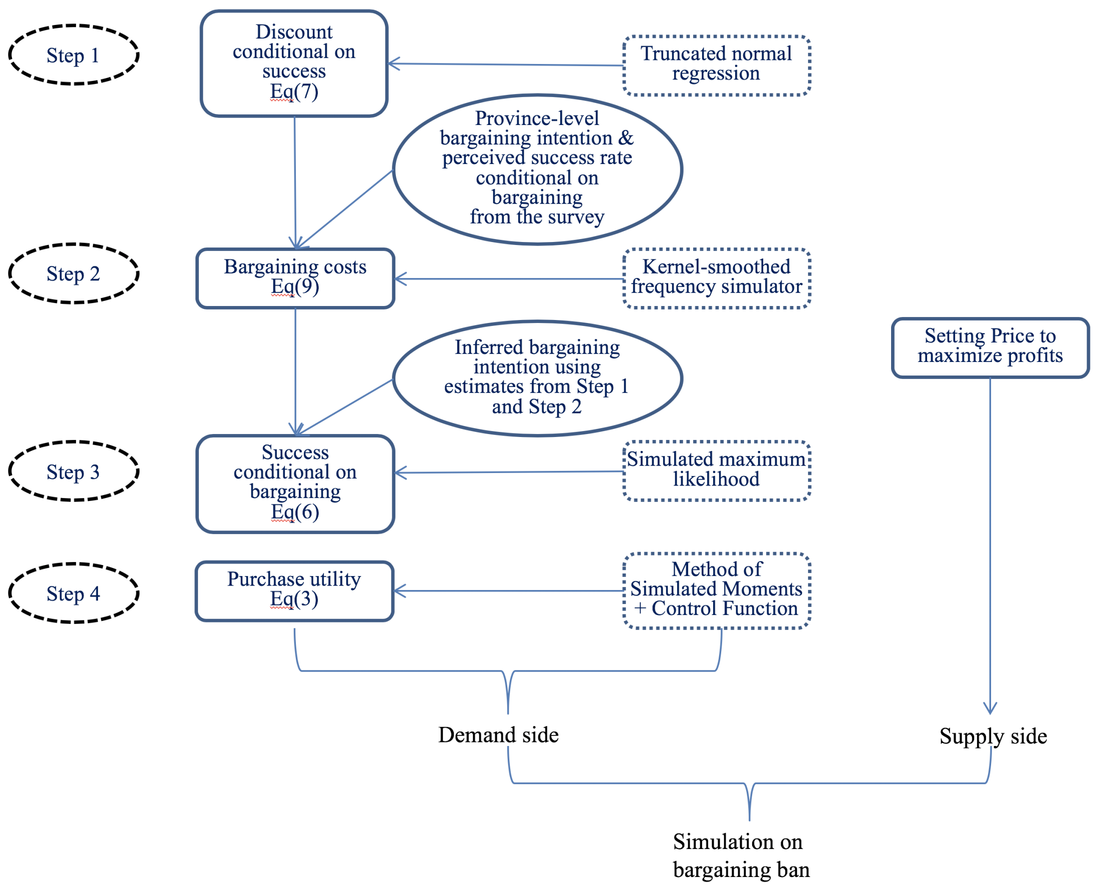

A brief summary of identification is provided in Table 4. We now describe the estimation and identification strategies for each step in detail.

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 from the transaction sample are used in the

Table 4: A Brief Summary of Identification

Parameters	Identified by
δ_k- intrinsic product preference	The fraction of consumers with different income and
$\alpha_i, \alpha_{inc}, \alpha_{exp}$- price preference	shopping experience level who made a purchase
β_x- nonprice preference	among consumers who visited the item page
β_w- impact of competition on outside option	across products, over different prices, and across sellers
c_{zl}, c_{zu}- lower and upper bounds of	(i) the upper and lower bounds of province-level bargaining intention
province-level bargaining costs	(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

regression. The likelihood function for the truncated normal regression is:

$$l_d = \prod_{\{ijk:d^*_{ijk}>0\}} \frac{\phi(\frac{d^*_{ijk}-x_{ijk}\theta}{\sigma_d})}{1-\Phi(\frac{-x_{ijk}\theta}{\sigma_d})} \tag{10}$$

where x_{ijk}, as defined before, includes the logarithm of the posted price, seller, buyer, and product characteristics, product fixed effects, and two proxies for seller competition level. d^*_{ijk} is the latent bargaining discount amount. As we 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, we 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. In §7.3.2, we show that our results are robust to this assumption.

Once equation (10) is estimated, we predict the expected discount amount conditional on success for each shopping occasion:

$$E[d^*_{ijk}|s_{ijk}=1] = x_{ijk}\hat{\theta} + \hat{\sigma}_d\lambda(x_{ijk}\hat{\theta}/\hat{\sigma}_d) \tag{11}$$

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. We 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.

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 we only observe the realized bargaining discounts, not bargaining decisions. That is, we know a consumer chooses to bargain if we see a positive discount, but we do not know if a consumer chooses to bargain when we 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, we use the consumer survey data (described earlier in §2.3). Given how we frame the survey questions, the bargaining decision is simpler in the survey than in the transaction sample. In the transaction sample, the bargaining responses are explained by both differences in bargaining costs as well as the differences in bargain gains (depending on product characteristics). In the survey, the bargaining responses are mainly explained by differences in bargaining costs across respondents, as we keep the product fixed. Thus, the bargaining decision becomes $b_{ijk} = I(E[d_{ijk}|b_{ijk} = 1] > c_i)$ in the survey. As such, bargaining costs are identified by consumers' bargaining decisions across shopping occasions with varying expected bargaining gains. Using the survey, we calculate the lower bounds and upper bounds of bargaining intention of a representative consumer in each province. 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 of the bargaining cost and the upper bound of the bargaining intention $Pr_u(b = 1)$ implies a lower bound of the bargaining cost for the representative consumer in a province. From the perspective of the representative buyer, we think that a uniform distribution between the two bounds is the most natural choice. The interpretation of the uncertainty captured by this distribution is that it represents the uncertainty in the bargaining intention of the representative buyer in a province (as distinct from across heterogeneity which is captured by the collection of all province level distributions). We replicated the analysis with a triangular distribution and a truncated normal distribution and obtained very similar results.

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 for each individual buyer in the transaction sample. This implicitly defines the set of shopping occasions that lead a consumer to choose to bargain. Formally, let this set be

$$B(c_i; x_{ijk}) = \{ijk|E[d_{ijk}|x_{ijk}] > c_i\} \tag{12}$$

Since the individual-level bargaining intention is not available in the transaction sample, we estimate bargaining costs for representative consumers in each province. Of course, if we had the individual buyers' bargaining intention, our model is flexible enough to allow us to estimate the individual-level bargaining cost. In estimating bargaining costs for a representative consumer, we use the province-level perceived success rate conditional on bargaining to proxy for that of the representative consumer 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, we form the following two moment conditions:

$$Pr[b = 1|c_{zl}] = \int I(ijk \in B(c_{zl}; x_{ijk}))f(x_{ijk})dx_{ijk} \tag{13}$$

$$Pr[b = 1|c_{zu}] = \int I(ijk \in B(c_{zu}; x_{ijk}))f(x_{ijk})dx_{ijk} \tag{14}$$

where c_{zl} and c_{zu} are lower and upper bounds of bargaining costs of a representative consumer 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, we use a kernel-smoothed frequency simulator (Mcfadden, 1989) with a univariate survival function as a kernel

$$S(w_{ijk}) = \frac{1}{1 + exp(-hw_{ijk})} \tag{15}$$

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.

4.3 Step three: bargaining success indicator

After recovering the upper and lower bounds of bargaining costs of a representative consumer in step two, we assume that the bargaining cost of consumer i from province z is a random draw from $Uniform(c_{zl}, c_{zu})$. We briefly discuss the identification of the parameters in the bargaining success function. We know that the observed bargaining success equals the bargaining intention times the bargaining success conditional on bargaining. For each transaction, we can calculate the bargaining intention using the predicted bargaining gain and the estimated bargaining cost from the previous steps. And then the variation of the observed bargaining success and various attributes of buyer, seller, and the product help us identify the parameters in the bargaining success function.[11] Thus, the effects of the observed factors on the success rate are identified. The likelihood of observing a bargaining success can be written as

$$\begin{aligned} l_s &= \prod_{ijk} l_{s,ijk} \\ &= \prod_{ijk} Pr(s_{ijk}=1|b_{ijk}=1)Pr(b_{ijk}=1) \\ &= \prod_{ijk} \Phi(x_{ijk}\gamma/\sigma_s)Pr(E[d_{ijk}|x_{ijk}]-c_i>0) \\ &= \prod_{ijk} \Phi(x_{ijk}\gamma/\sigma_s)\int I(E[d_{ijk}|x_{ijk}]-c_i>0)f(c_i)dc_i \end{aligned} \tag{16}$$

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 §4.2, we use a kernel-smoothed frequency simulator to ensure smoothness of the likelihood function. The parameters (γ, σ_s) are estimated using simulated maximum likelihood.

4.4 Step four: purchase utility

In the purchase stage, the unknown parameters are intrinsic product preferences δ_k, heterogeneous price coefficients α_i, α_{inc}, and α_{exp}, seller/product characteristics preference β_x, and the impact of seller competition on the outside option valuation β_w. The parameters are identified by consumers' purchasing patterns between the inside good and the outside good. However, we do not observe

[11] In the estimation we treat consumers as myopic, i.e., they only compare the expected bargaining gain with the bargaining cost but not the expected utilities. We check the robustness of our results to this assumption when we compare the results from the analyses with sophisticated consumers versus naive consumers in §7.3.1. There is no substantive change in the main results.

individual consumers who visited the site but walked away without purchase. Thus, we use seller-level conversion rates to infer the utility preferences. Specifically, we observe the number of unique visitors and the number of unique purchasers for each seller in that month. We assume that one in five unique visitors is a serious shopper, thus the conversion rate is defined as the number of unique purchasers divided by the number of unique visitors, and then 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 we choose two or ten). The average conversion rate is 10.02%. We 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, we think this is a reasonable assumption.

Another challenge here is that we do not observe buyer characteristics of the non-purchasers. Thus, we use several strategies to simulate a potential pool of buyers for each product/seller listing: 1. a pool of buyers who have previously made transactions at the focal seller, 2. a pool of buyers who have made transactions from sellers who have the same reputation level, 3. a pool of buyers who are the same as the buyer in the focal transaction. We find small differences on the estimated price heterogeneity parameters, but the main results are very similar across the three alternative strategies. We use the first strategy to report the main results.

As we use conversion rates 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 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 we fail to control for those unobserved variables, we may underestimate the price coefficient.

As discussed before, the estimation may suffer from endogeneity problem due to omitted variables, such as word of mouth or website design, so we use the control function approach. We first regress the endogenous variable (e.g., 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, 2017). This incentive generates a positive correlation between the instrument and the posted price. For a consumer to obtain the information on a seller's $DistNextLevel$, the consumer 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.

We also include the interaction between price and income, and the interaction between price and shopping experience level to capture the heterogeneity of price sensitivity. In order to instrument for these two interactions, we use the interaction between $DistNextLevel$ and income, the interaction between $DistNextLevel$ and shopping experience level, and a squared term of $DistNextLevel$. The control function approach is detailed in §3.1. The first-stage Stock-Yogo weak identification F tests are 15.01, 18.73, and 245.75 for the price and the two interaction terms, respectively. All pass the weak identification tests to satisfy the relevance condition (Stock and Yogo, 2005).

The probability of making a purchase, i.e., the conversion rate, is:

$$\begin{aligned} m_{jk} &= \int Pr_{ijk}(p_{ijk}) f(p_{ijk}) dp_{ijk} \\ &\approx Pr(u_{ijk} > u_{ij0} | b_{ijk} = 0) Pr(b_{ijk} = 0) \\ &\quad + Pr(u_{ijk} > u_{ij0} | s_{ijk} = 0) Pr(s_{ijk} = 0 | b_{ijk} = 1) Pr(b_{ijk} = 1) \\ &\quad + Pr(u_{ijk} > u_{ij0} | s_{ijk} = 1) Pr(s_{ijk} = 1 | b_{ijk} = 1) Pr(b_{ijk} = 1) \end{aligned} \tag{17}$$

which includes three parts (see the bottom of Figure 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 conversion rates and the observed ones.

Finally, we use the conventional nonparametric bootstrap method to get standard errors for the estimates (Efron and Tibshirani, 1993). Specifically, we draw independent random samples with replacement repeatedly from the sample dataset. We then get all the desired estimates using the whole model with multiple steps corresponding to these bootstrap samples, which form the sampling distribution of each estimate. Last, we calculate the empirical standard deviation of the sampling distribution for each estimate. Note that via the bootstrapping method, the estimates in former steps are used in later steps, thus, any volatility of the estimates in former steps is accounted for in estimating parameters in later steps.

5 Supply Side

In order to perform policy simulations, we need to model sellers' pricing decisions in addition to the demand side model (see Figure 6). We assume that Taobao sellers set prices to maximize profits given the attributes of the seller, the product, the potential pool of buyers, and the competition level.[12] The profit of seller j on product k is given by

$$\begin{aligned}\Pi_{jk} &= (Pr(b_{jk}=0)+Pr(b_{jk}=1)Pr(d_{jk}=0|b_{jk}=1))Pr(a_{jk}(\bar{p}_{jk})=1)(\bar{p}_{jk}-mc_{jk}) \\ &\quad +Pr(b_{jk}=1)Pr(d_{jk}>0|b_{jk}=1)Pr(a_{jk}(\bar{p}_{jk}-\hat{d}_{jk})=1)(\bar{p}_{jk}-\hat{d}_{jk}-mc_{jk}) \\ &= Pr_{jk,zero}(\bar{p}_{jk}-mc_{jk})+Pr_{jk,pos}(\bar{p}_{jk}-\hat{d}_{jk}-mc_{jk}) \end{aligned} \tag{18}$$

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 $a_{jk}(p_{jk}) = 0$. Note that in this price setting process, the competition has been taken into consideration through the purchase probability and the bargaining outcomes where the proxies for seller competition come into play. 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 who have made purchases with this seller before) shopping experience and buyer characteristics of interest. The probabilities and the expected bargaining gains are calculated for this representative consumer. We 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

[12] 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, we 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.

k can be written as

$$\left(\frac{\partial Pr_{jk,zero}}{\partial \bar{p}_{jk}} + \frac{\partial Pr_{jk,pos}}{\partial \bar{p}_{jk}}\right)(\bar{p}_{jk} - mc_{jk}) = \left(\frac{\partial Pr_{jk,pos}}{\partial \bar{p}_{jk}}\hat{d}_{jk} + Pr_{jk,pos}\frac{\partial \hat{d}_{jk}}{\partial \bar{p}_{jk}}\right) - (Pr_{jk,zero} + Pr_{jk,pos}) \quad (19)$$

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

In the model above, we assume a seller maximizes profits by setting the optimal posted price, which can either be a regular or a promotional price. It may seem a little strange that firms are choosing prices that can be either regular or promotional at any given point in time. However, it turns out this is a rather common practice on C2C platforms. For example, on eBay, studies have found substantial price dispersion, both across and within sellers, even with structured price comparison settings (Einav et al., 2015). In our paper, we do not explicitly model the price setting between the regular price and the promotional price due to lack of data availability. However, we believe this should not affect our results materially. The reasons for this come from both the supply side and the demand side. On the supply side, Taobao sellers have full freedom to set both the list price and the promotion depth at any given point in time. Given that, the seller incurs zero cost to "manipulate" their setting simultaneously. As a result, we believe only the promotion-adjusted price is the price that directly impacts the seller's profits. On the demand side, since consumers are aware that sellers can manipulate both the list price and the promotion, the promotion-adjusted price should also be the only price that consumers really care about. However, based on previous research, e.g., Mayhew and Winer (1992), that suggests that consumers react mainly to the promotion indicator rather than the promotion depth, we include a promotion indicator in our specification (equation 1). Since the promotion depth does not directly affect a seller's profit and the promotion depth is not affecting consumers' bargaining or purchase decision, we believe our current price model acts as a reasonable proxy for the behavior of the sellers on this platform.

In conclusion, after detailing the demand and supply side models, we would like to highlight the fact that our model is flexible enough to allow a seller to either benefit from the fixed-price or the bargaining mechanism. This flexibility is achieved through the intercorrelation among the three stages on the demand side and through the sellers' pricing decision after taking bargaining into consideration. First, note that a buyer's characteristics and demographics can influence both the buyer's bargaining outcome and price sensitivity. As a result, a buyer's willingness to pay is naturally correlated with her/his bargaining decision. Second, this correlation introduces a self-selection process into the bargaining decision. Third, due to the existence of bargaining success probability and the expected bargaining discount in the profit function, sellers take this correlation and the selection into the optimal price setting process. As a result, sellers may or may not earn more profits under bargaining compared to the fixed price mechanism. The intuition for this is as follows. Let us consider a situation where a seller has a pool of a combination of high wealth and low wealth customers. If the high wealth customers are less price sensitive and have higher bargaining costs, then the seller can use the bargaining mechanism to effectively price discriminate between the two groups to extract more surplus from consumers, as a result, to earn more profit from bargaining. However, if the high wealth customers are less price sensitive but have higher bargaining intention, then the seller may be better off under the fixed-price mechanism. Note that we formalize this intuition using a stylized analytical model in Appendix B.

6 Results

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

6.1 Bargaining Realization and Bargaining Decision

Table 5 shows the estimates of the two-part model in the bargaining realization stage. We 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, we 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 lower income, and this is a repeat purchase between the seller and the buyer. Also, we find that female buyers on average are more likely to succeed in bargaining, but age does not have a statistically significant effect. To make the results more interpretable, we calculate the average partial effect of each explanatory variable. Specifically, we 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 increases by 0.6% with a 1% decrease in buyers' income, and 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 14%, about 6 yuan, suggesting that higher reputation sellers have more bargaining power. Interestingly, we 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 5.5% decrease, or about 2 yuan, in the bargaining discount amount. Buyer age has a small negative effect on the discount amount, but we 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), we will describe the results from the full sample going forward.

Figure 7 shows the estimated bargaining costs of representative buyers in each province. For each representative buyer, we 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 cost varies across provinces. We attempt to explain this heterogeneity using several factors. Specifically, we 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. We 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)?" We find that income per capita, population density, level of urbanization, and internet penetration are positively correlated with the province-level bargaining

Table 5: Bargaining Realization Estimates

	Bargaining Success Indicator		log(Bargaining Amount)	
	Full Sample	Subsample	Full Sample	Subsample
log(Price)	0.206***	0.198***	1.025***	0.984***
	(0.017)	(0.031)	(0.036)	(0.044)
I(Promotion)	-0.170***	-0.199***	-0.552***	-0.565***
	(0.018)	(0.033)	(0.038)	(0.045)
Seller Reputation Level	-0.022***	-0.013**	-0.140***	-0.139***
	(0.006)	(0.006)	(0.012)	(0.014)
Detailed Seller Rating	0.075	0.093	0.502***	0.511***
	(0.087)	(0.080)	(0.143)	(0.168)
Store Age	-0.001	-0.004	0.013	0.017
	(0.006)	(0.008)	(0.013)	(0.015)
Buyer Shopping Experience	0.055***	0.051***	-0.056***	-0.055***
	(0.004)	(0.010)	(0.012)	(0.017)
log(Buyer Income)	-0.180***	-0.147***	0.181***	0.037
	(0.030)	(0.024)	(0.063)	(0.074)
Buyer Age		-0.001		-0.006**
		(0.002)		(0.003)
Buyer I(Female)		0.065**		0.055
		(0.026)		(0.046)
I(Repeat Purchase)	0.392***	0.437***	0.251***	0.171***
	(0.025)	(0.031)	(0.046)	(0.052)
Product Age	0.006	-0.001	0.143***	0.147***
	(0.011)	(0.015)	(0.018)	(0.023)
Seller Competition at Same Product	-0.016	-0.012	0.051	0.073
	(0.017)	(0.022)	(0.046)	(0.055)
Seller Competition at Same Reputation	-0.011*	-0.017**	0.044***	0.046**
	(0.006)	(0.007)	(0.013)	(0.016)
Product FE	Yes	Yes	Yes	Yes
Number of Observations	39,625	25,201	6,332	4,368

Note: ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels in all tables. Seller competition is measured as the number of sellers who are on the platform in the same month either selling the same product or having the same reputation level with the unit of a hundred.

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 consumers have higher time costs and thus may have a lower propensity to bargain.

In addition to exploring the heterogeneity of bargaining costs across provinces, we also investigate the difference of bargaining costs across gender. We find that on average, female buyers have slightly higher bargaining costs than male buyers, but the difference is not statistically significant. This result is not surprising given the similarity of the products purchased, the bargaining outcomes, and the bargaining intention across gender, as reported in §2. The correlation between the estimated bargaining costs using the main survey data and using the combined main/pilot survey data is 0.95, confirming the robustness of the results.

Using the estimates from the bargaining realization and the estimated bargaining costs, we find that on average bargaining would be initiated in about 78.2% of shopping occasions. Conditional on bargaining, Figure 8 represents the histogram of the estimated success rate across all shopping occasions in the sample. The average success rate conditional on bargaining is 18.3%. Using the estimated bargaining intention and the success rate conditional on bargaining, a back-of-the-envelope calculation of the bargaining success rate shows that it is 14.4% (78.2% $*$ 18.3%) - this is consistent with the 16% success rate observed in the sample.

Figure 7: Estimated Province-level Average Bargaining Costs

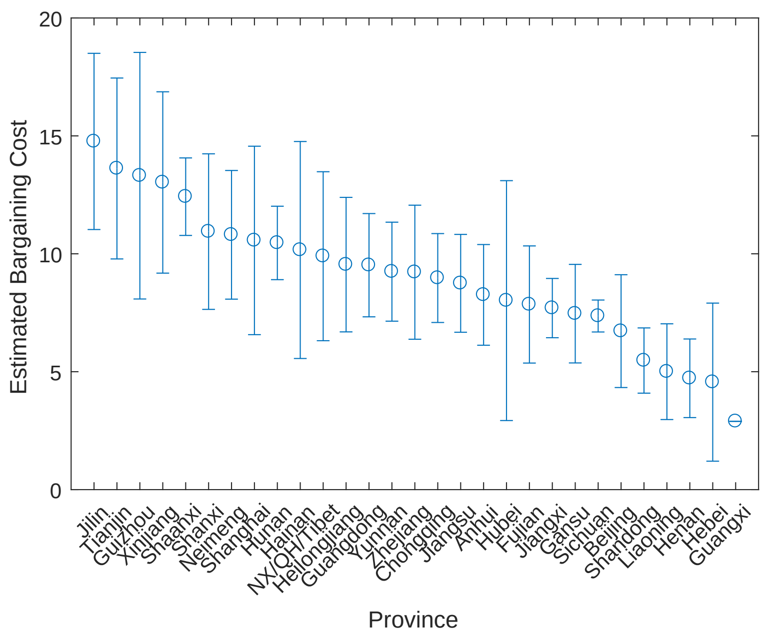

Note: since no respondent in Guangxi province chooses "may bargain," the lower bound and the upper bound for that province overlap.

As highlighted previously, it is important to distinguish between no-bargain and failed-bargain transactions in retail settings, as it can directly affect the estimated bargaining cost and further

Figure 8: Estimated Bargaining Success Probability

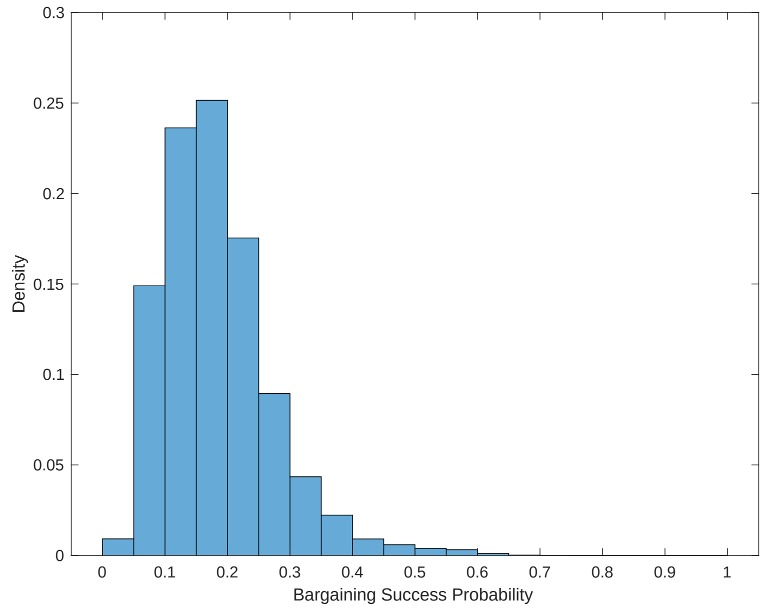

affect the managerial decisions. However, this point has been overlooked in previous literature. To assess the importance of this point, we compare the estimation results with the distinction (between no-bargain and failed-bargain transactions) to that without this distinction. Without the distinction, one would incorrectly conclude that 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%). These relatively large differences suggest that use of these estimates could lead managers to arrive at erroneous conclusions, especially in terms of the overall benefits (in our setting) from switching to a fixed-price mechanism.

6.2 Purchase Decision

Table 6 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 omitted variables, we would underestimate the price coefficient (less negative) and thus after applying the control function approach, the estimated price coefficient becomes more negative. We 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. In terms of the heterogeneity of price sensitivity, not surprisingly, we find that buyers with more shopping experience and higher income are less price sensitive.

Table 6: Purchase Model Estimates

	Without Control Function	With Control Function
log(Price)	-0.505***	-3.967***
	(0.080)	(0.457)
log(Price) * Buyer Shopping Experience	0.001	0.027***
	(0.001)	(0.004)
log(Price) * log(Buyer Income)	0.004	0.013***
	(0.007)	(0.002)
I(Promotion)	0.188***	0.192***
	(0.011)	(0.013)
Seller Reputation Level	0.061***	0.174***
	(0.005)	(0.012)
Detailed Seller Rating	1.125***	4.023***
	(0.187)	(0.325)
Store Age	-0.100***	0.025
	(0.005)	(0.019)
Repeat Purchase Rate	2.409***	2.015***
	(0.429)	(0.353)
Product Age	0.015***	-0.771***
	(0.004)	(0.110)
# of Sellers with Same Product	0.085***	0.330***
	(0.021)	(0.031)
# of Sellers with Same Reputation	-0.042***	-0.033***
	(0.003)	(0.004)
Terms to Correct for Endogeneity		
Control Function	Yes	Yes
Product FE	Yes	Yes
Number of Observations	39,625	39,625

To help interpret the results, we numerically calculate the price elasticity on conversion rate for each product/seller combination (Figure 9). On average, a 1% increase in price leads to a 3.3% decrease in conversion rate, suggesting the average conversion rate would decrease from 10.02% to 9.69%.[13] As there are two types of transactions, i.e., transactions made at the posted price and those made at a bargained price, we separate the price elasticity on conversion rate into two types. The mean implied price elasticities of the above two types are -3.4 (s.e. = 0.28) and -2.7 (s.e. =

[13]Our 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).

0.39) 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.

Figure 9: Implied Price Elasticity on Conversion Rate

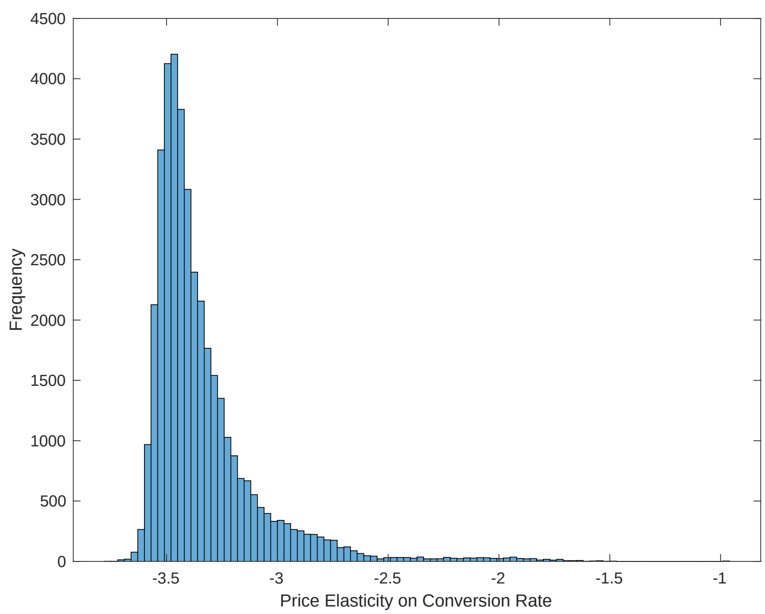

We also calculate the mean price elasticities on conversion rates for the top five brands. Nokia has the lowest price elasticity (-3.2) while Apple has the highest (-3.7), consistent with the belief in the Chinese market (at the time of the sample) that Apple products were considered discretionary products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

7 Counterfactual Analysis: Ban on Bargaining

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 Aliwangwang and became very popular over time. Thus, an interesting question for Taobao is what would happen if it bans bargaining.[14]

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:

$$p^*_{jk} = \underset{p_{jk}}{\operatorname{argmax}}\{Pr(a_{jk}(p_{jk}) = 1)(p_{jk} - mc_{jk})\} \tag{20}$$

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

Table 7 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.[15] 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, we find that when bargaining is banned, the conversion rate becomes higher, increasing from 10.21% to 10.27%, a 0.6% 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

[14]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.

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

Table 7: Counterfactual Analysis: Market Response if Bargaining is Banned

	Before Change	After Change	$\Delta\%$
Average Posted Price (yuan)	1,263	1,249	-1.1%
Average Transaction price (yuan)	1,251	1,249	-0.1%
Conversion Rate (%)	10.21	10.27	0.6%
Number of Transactions per Day	69,622	70,031	0.6%
Total Sales per Day (million yuan)	87.10	87.47	0.4%
Bargaining Costs per Day (million yuan)	4.8	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.

the product at the high posted price before the change, but the consumer 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, we see that the conversion rate is higher with the pure fixed-price mechanism than the mixed-price mechanism.

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 70,031. 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.[16] Using the average transaction price and the number of transactions per day, we impute the GMV per day as 87.10 million yuan before the change and 87.47 million yuan after the ban, which is about a 0.4% increase.

Comparing the average transaction price and the total sales, the difference between the two types of pricing mechanisms is not large. However, the existence of bargaining costs gives rise

[16]We 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 our 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, we would expect the traffic to increase as the majority do not like bargaining, and this would further reinforce our findings and managerial recommendations.

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 estimated bargaining costs incurred by buyers per day in the cellphone category alone are 4.8 million yuan, about 6% of the category total sales - an economically significant number. As we do not observe sellers' marginal costs (for products), we are 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 we document) than the status-quo from a social planner's and sellers' point of view.

7.2 Heterogeneous Effects of the Policy Change

To study whether the effect of the policy change is heterogenous across sellers, we 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 bargaining. In order to compare across sellers, we 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 bargaining mechanism. We 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 are in Table 8, where a negative number indicates bargaining is better for sellers with those attributes, and fixed-price is better for sellers lacking those attributes.

We 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, we find that the fixed-price mechanism is more appealing for sellers with high detailed seller rating. Store age does not seem to matter. 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

Table 8: Determinants of the Heterogenous Effects of the Ban on the Bargaining Mechanism

	(Fixed-price Profit) - (Bargaining Profit)
Seller Reputation Level	-0.012***
	(0.001)
Detailed Seller Rating	0.053***
	(0.009)
Store Age	0.0002
	(0.001)
I(Promotion)	-0.021***
	(0.002)
Product FE	Yes
Number of Observations	39,625

the fixed-price mechanism. This result is only suggestive as we assume the promotion indicator keeps the same before and after the pricing mechanism change. However, the promotion depth is allowed to change as it is implicitly captured in the pricing decision.

While the average profit increase is small, the majority (97%) of sellers are better off under the fixed-price mechanism than under bargaining. A natural question to ask then is why this majority of sellers are not able to take advantage of the price discrimination provided by bargaining. We know that the correlation between the bargaining outcomes and the price sensitivity is a key factor in determining seller profits under the two regimes (see the discussion at the end of §5 and also the detailed analysis in Appendix B). We therefore test this potential underlying mechanism via simulation to answer the question of why sellers are unable to benefit from the price discrimination from bargaining. Recall that in our new specification, we allow buyers' characteristics (income and shopping experience) to impact bargaining outcomes and her/his price sensitivity, consequently affecting the correlation between the bargaining outcomes and price sensitivity. In order to adjust this correlation, we hold the impact of buyers' characteristics on price sensitivity constant and vary the impact of buyers' characteristics on bargaining outcomes (coefficients γ and θ in equations 6 and 7 in the paper). As γ and θ change, the correlation between the bargaining outcomes and the price sensitivity will change, and thus, the proportion of sellers who are better off under bargaining (by examining the difference in profit under the two regimes for each seller) will change.

The results from this simulation are shown in Figure 10. The y-axis shows the percentage of sellers who would favor bargaining, i.e., sellers' profits are higher under bargaining than under fixed-price mechanism. The x-axis shows the parameter that we vary. We manipulate the correlation between bargaining intention and price sensitivity to be more negative by decreasing the coefficient of shopping experience on bargaining success on the left panel and by decreasing the coefficient of income on bargaining discount on the right panel. As can be seen from both panels, as the correlation becomes more negative (moving left on the x-axis), the percentage of sellers who would favor bargaining increase. The highlighted red dots in both panels denote the original results at the estimated model parameters without any adjustment (when the coefficient of shopping experience on bargaining success equals 0.055 and the coefficient of income on bargaining discount equals 0.18, only 3% of sellers prefer bargaining). The location of the highlighted dots implies that more sellers would have preferred bargaining to fixed-price mechanism if the correlation between bargaining intention and price sensitivity were more negative among Taobao buyers (moving left on the x-axis). However, as this correlation is not negative enough (0.17), the majority of sellers on Taobao are not able to benefit from bargaining mechanism.

Figure 10: Simulation on Bargaining Attractiveness

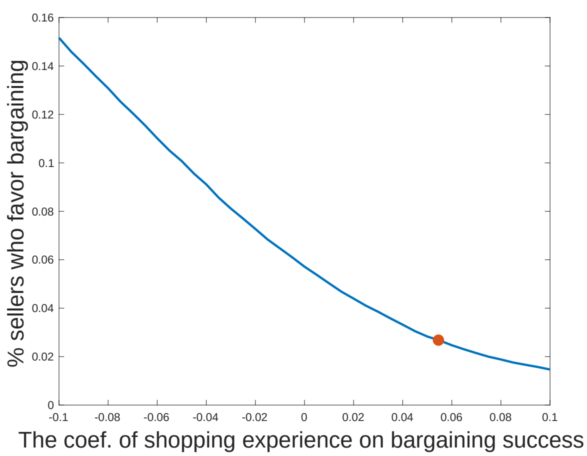

The above results suggest that the majority of sellers would benefit and buyers would save a great amount of bargaining costs if a fixed-price mechanism is implemented. However, we do not see the platform has implemented the fixed-price mechanism. At the platform level, the reason for this is the platform's business model. Specifically, the platform's revenue comes primarily from advertising and not from commissions on transactions - thus the platform does not have any incentive to enact policy changes that impact transactions.

Next, note that our model allows for the possibility that a seller says no to bargaining. In the model, this is operationalized through the bargaining success function. For example, we see that high reputable sellers say no more often than low reputable sellers to the same buyer for similar products. Also, our estimates show that many sellers have bargaining success rates that are close-to-zero, implying that they are trying to say no to bargaining and sticking to fixed-prices. However, a "discrepancy" still remains between the reality that the majority of sellers do not commit to fixed-prices while our counterfactual analysis suggests that they would be better off under fixed-prices. The key to resolving this discrepancy is understanding the nature of the information set under which our counterfactual is carried out. Specifically, in our counterfactual scenario where the platform moves from bargaining to fixed-prices, the "no bargaining" information is public to all buyers and therefore sellers cannot choose to deviate from this. As noted in our results, in this scenario, buyers save bargaining costs and sellers on average benefit from this policy change, even though the benefit is modest at a 0.4% increase in profits.

Consider the contrasting scenario where each individual seller's decision to move to no bargaining is not public information. In such a scenario, (a) many buyers will still bargain (based on our survey data), (b) sellers will run the risk of violating the cultural norm that allows for bargaining, and as a result, may hurt their detailed seller rating (our data/analysis shows that the detailed seller rating is positively correlated with bargaining outcomes), and (c) the cost of implementing a strict no-bargaining rule may offset the (modest) benefit from implementing the fixed-price mechanism. An example of such a cost is that the hassle involved in explaining to customers why no discount is provided to them (especially to repeat customers) via bargaining like what other sellers do on the same platform. Given these three reasons, it is not clear if an individual seller will benefit from a unilateral move to "no bargaining" if it is her/his own i.e., not the platform's, decision. Note that we are unable to carry out the counterfactual analysis where each individual seller gets to decide whether to move (to fixed-price) as we do not have information on sellers' bargaining costs and sellers' cost of implementing the "no bargaining" rule. In addition, even if an individual seller's move is beneficial for the majority of sellers, it is very unlikely that sellers are able to take any collective action given the large (about 7 million in total and about 100,000 in the cell phone category at the time of our data) and heterogeneous seller base on Taobao.

Taking everything together, as long as "no bargaining" is not directly implemented by the platform and is not publicly acknowledged by all the sellers and all the buyers, the major benefit from the saved bargaining costs for buyers and the modest profit increase for average sellers cannot be realized. This is like a prisoner's dilemma and the only way it can get resolved is if the platform acts to ban bargaining and make this public information among all market participants.

In conclusion, while the benefits to individual sellers or the platform may only be minimal after switching to a fixed-price mechanism, this move is very beneficial from the social planner's perspective. Such a platform's policy change would greatly help consumers avoid bargaining costs and at the same time may help sellers avoid operational costs if their bargaining costs are high, and increase the total transaction volume modestly.

7.3 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, we discuss the robustness of our results to several assumptions, particularly those that relate to buyers' bargaining decision and the bargaining realization process. We conclude with a replication of our analysis using data from another product category to judge its generalizability.

7.3.1 Survey Response versus True Belief

In the baseline model, we 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, we 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, we 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.

7.3.2 Assumption for Non-purchasers

In order to distinguish between the three types of non-purchasers at the bottom of Figure 5, s necessary assumption we have to make is that bargaining outcomes conditional on success are not systematically different between purchasers (whose transactions are observed) and non-purchasers (whose transactions are unobserved). This assumption may be seen as strong in the sense that bargaining outcomes for purchasers on average should be better than that for non-purchasers. Unfortunately, without any data from non-purchasers, we are not able to test this assumption directly. However, we will provide three pieces of evidence that the impact of this assumption is not very strong.

The first piece of evidence is that the expected bargaining discount amount revealed in the survey is about the same as the realized bargaining discount amount observed in the transaction sample (170.3 yuan vs. 165.5 yuan). The similarity between the two suggests that buyers' expected bargaining outcomes are likely to be uncorrelated with the purchase decision (as the survey questions do not cover any aspect of the purchase decision).

The second piece of evidence is based on our finding that the correlation between a buyer's bargaining intention and the buyer's price elasticity is low at 0.17 (this is actually the main reason why sellers are not able to effectively price discriminate among buyers). As a result, using the purchasers' bargaining behavior to infer non-purchasers' bargaining behavior is unlikely to create bias.

The third piece of evidence is based on a robustness test of our results to this assumption directly through a simulation. Between purchasers and non-purchasers, the former is likely to have a higher bargaining success rate and a higher bargaining discount amount. Thus, if we infer these two bargaining outcomes for non-purchasers using the parameters estimated from purchasers, we are likely to overestimate them. In order to test the robustness of our results, we consider an extreme case of this overestimation. Specifically, we test how our results would change for a 100% overestimation. In other words, if the model predicts a potential buyer has 40% bargaining success rate and 20 yuan expected bargaining discount, then we will just use the true bargaining success rate of 20% and the true discount amount of 10 yuan for the subsequent estimation steps. Even in this extreme case, we find that our results on total revenues and conversion are robust. This is not surprising given that the baseline change in total revenue and conversion rate is very small - less than 1%. While the total savings in bargaining costs become smaller, they are still not materially different and remain economically significant.

7.3.3 Bargaining Realization Process

One of the critical assumptions underlying our 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 do a reasonable job 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, we have to restrict the analysis to sellers with at least two transactions. Such a restriction introduces a sample selection. Given one of our goals is to estimate the value of bargaining for the platform, we want to keep the sample as representative as possible, and thus we did not include seller fixed effects in the baseline model. Besides the sample selection issue, given the large number of sellers in the sample and the product fixed effects, we found the baseline structural model with seller fixed effects to be computationally intractable.

As a compromise however, we carry out a reduced form analysis to investigate whether the conditional independence assumption could impact our results. Specifically, we use a probit model and a truncated regression model on a subsample (including sellers with at least 20 transactions) and include seller fixed effects in the bargaining realization process. Table 9 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 estimates are similar across columns, suggesting that our results are robust to the conditional independence assumption. Note that the key difference between the reduced form regression results (Table 9) and the structural model results (Table 5) is whether we account for a buyer's bargaining intention. In the structural model, we explicitly estimate the bargaining intention, while the reduced form regressions do not allow us to do so. Nevertheless, the comparison with the reduced-form results increases our confidence in the results.

Table 9: Robustness of Bargaining Realization Process

	Bargaining Success Indicator		log(Bargaining Amount)	
log(price)	0.335***	0.424***	0.776***	0.937***
	(0.035)	(0.054)	(0.090)	(0.109)
I(Promotion)	-0.158***	-0.119**	-0.514***	-0.593***
	(0.037)	(0.045)	(0.081)	(0.075)
Seller Reputation Level	0.052***	-0.045	-0.243***	-0.247**
	(0.017)	(0.050)	(0.035)	(0.097)
Detailed Seller Rating	-0.387***	-0.017	0.634**	-1.033***
	(0.120)	(0.175)	(0.276)	(0.3257)
Store Age	-0.024*	-0.213**	0.067**	0.572***
	(0.011)	(0.074)	(0.027)	(0.141)
Buyer Shopping Experience	0.055***	0.052***	-0.069***	-0.043*
	(0.010)	(0.011)	(0.024)	(0.021)
Buyer log(income)	-0.164***	-0.262***	-0.033	-0.030
	(0.048)	(0.055)	(0.109)	(0.099)
I(Repeat Purchase)	0.425***	0.356***	0.134	0.026
	(0.041)	(0.045)	(0.085)	(0.073)
Product Age	0.016	0.022	0.122*	0.091*
	(0.020)	(0.026)	(0.050)	(0.052)
# of Sellers w/ Same Product	-0.057	-0.059	0.039	-0.137
	(0.041)	(0.046)	(0.097)	(0.087)
# of Sellers w/ Same Reputation	0.015	0.057**	-0.090**	-0.061
	(0.016)	(0.021)	(0.035)	(0.038)
Product FE	Yes	Yes	Yes	Yes
Seller FE	No	Yes	No	Yes
Log likelihood	-4,438	-4,025	-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.

7.4 Generalizability

Our 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, we were able to obtain data on the women's shoes product category. Following the first two steps in §4, we find the the bargaining cost is 8.6 yuan on average in this category.[17] The estimated bargaining cost of 8.6 yuan for women's shoes is very close to our estimate

[17]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 our main analysis to compute the average bargaining success rate conditional on bargaining.

of 9.0 yuan for the cellphone category. Given that the unconditional success rate in women's shoes category is low, we 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 our results are not idiosyncratic to the cellphone category.

7.5 External Validity with Real Chatting Data

One of our data limitations is that we did not have any chatting history between buyers and sellers. This makes the negotiation process a "black box." Even though modeling the process is not the objective of our paper, we think that access to any such data can only increase the external validity of our analysis and results. We were able to obtain three months of chatting data (in text form) from a (small) seller on Taobao. This seller's store sells camera accessories. The data range from April 2013 to June 2013, spanning 306 chatting sessions. We want to state at the outset that we do not consider this seller to be a representative seller for the pool of sellers in our chosen product category of cell phones especially as the mean unit price in the cell phone category is about 1,500 yuan while it is about 250 yuan for this seller. Our objective is to provide some descriptive analysis based on these data to provide more context and hopefully external validity.

The detailed discussion on the analysis of the chatting data is provided in Appendix C. Several key findings from the chatting data are as follows. First, we find that the estimated time cost of the actual time that buyers spend bargaining with a seller is about 10 yuan, which is very similar to the average bargaining costs estimated in the model at 9 yuan. Second, we find that the bargaining intention revealed from the chatting data is comparable to that revealed from the survey, which suggests that our assumption that the bargaining intention in the survey represents the bargaining intention in the transaction sample is reasonable. Third, we find that the bargaining success rate in the chatting data lies between that in the survey and in the transaction sample, implying the numbers that we used in the estimation are consistent with the reality. Lastly, we see support for our use of the repeat purchase indicator in our model.

8 Discussion and Conclusion

This paper contributes to a small but growing body of empirical literature on bargaining. We focus on comparing the value of bargaining versus fixed-price mechanism on an online platform. We 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, we use a consumer discrete choice model augmented with outcomes from the bargaining realization process and use a control function approach to address the potential omitted variable problem. On the supply side, we model sellers making profit-maximizing pricing decisions that are impacted by expected bargaining outcomes and the market competition. Using the estimates from the structural model, we perform a counterfactual analysis to derive the value of bargaining by assuming a counterfactual scenario where bargaining is banned in the marketplace. We find that sellers on average in the marketplace would benefit from this policy change, however, the benefit is small. On the other hand, we find that buyers would greatly benefit from the policy change due to the saved bargaining costs. Further, we investigate the heterogeneity of the benefits from the policy change across sellers and explore the reasons why the majority of sellers benefit from such a policy change. 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. We find that consumers' bargaining costs are comparable to the minimum hourly wage in the geographic setting (China). Further, we 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. We also find that sellers' reputation level and buyers' shopping experience and income 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. Our study thus fills the void.

There are several avenues for future research. First, due to data limitations, we define a market at the seller/product level. This allows us 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. Second, 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, we 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 from the product marginal cost, which would make the sellers' benefit analysis more complete. We hope that future research can carry out these extensions.

References

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 (2018). 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.

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

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.

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.

Efron, B. and R. J. Tibshirani (1993). *An Introduction to the Bootstrap*. Boca Raton, FL: Chapman & Hall.

Einav, L., T. Kuchler, J. Levin, and N. Sundaresan (2015). Assessing sale strategies in online markets using matched listings. *American Economic Journal: Microeconomics 7*(2), 215–247.

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., 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.

Fang, T. and C. Lin (2015). Minimum wages and employment in china. *IZA Journal of Labor Policy 4*(1).

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.

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

Jindal, P. and P. Newberry (2018). To bargain or not to bargain: The role of fixed costs in price negotiations. *Journal of Marketing Research*, forthcoming.

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

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

Mayhew, G. E. and R. S. Winer (1992). An empirical-analysis of internal and external reference prices using scanner data. *Journal of Consumer Research 19*(1), 62–70.

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.

Neslin, S. A. and L. Greenhalgh (1986). The ability of nash's theory of cooperative games to predict the outcomes of buyer-seller negotiations: A dyad-level test. *Management Science 32*(4), 480–498.

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.

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.

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

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

Small, D. A., M. Gelfand, L. Babcock, and H. Gettman (2007). Who goes to the bargaining table? the influence of gender and framing on the initiation of negotiation. *J Pers Soc Psychol 93*(4), 600–13.

Stock, J. and M. Yogo (2005). *Testing for Weak Instruments in Linear IV Regression*, pp. 80–108. New York: Cambridge University Press.

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

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

Press.

Zhong, Z. (2017). Chasing diamonds and crowns: Online reputation systems and seller response. *Working Paper*.

Appendices

A Summary Statistics by Gender

Table A1: 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 A2: 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

B A Stylized Model for Pricing Mechanism Comparison

For simplicity, we assume there are two types of buyers in the market: θ proportion of type A and $(1-\theta)$ proportion of type B. Type A has a demand function of $D_1(p)$ and type B has a demand function of $D_2(P)$. The seller faces a constant marginal cost and without loss of generality, we assume the marginal cost is 0. Under the fixed price mechanism, all the transactions are made at the list price, thus, the seller sets the optimal price (the posted price) by maximizing the profit function

$$\Pi_f = \max_p \{\theta D_1(p)p + (1-\theta)D_2(p)p\} \tag{A1}$$

Under bargaining, we assume the bargaining discount amount only depends on the buyer type. Note that this represents a simplified version of the main empirical model presented in our paper.[A1] Type A gets a discount of d_1 and type B gets a discount of d_2. The seller sets the list price by maximizing the profit function

$$\Pi_b = \max_p \{\theta D_1(p-d_1)(p-d_1) + (1-\theta)D_2(p-d_2)(p-d_2)\} \tag{A2}$$

In order to prove that our model does not predetermine the counterfactual analysis, we need to show that a seller may or may not earn more profits under bargaining compared to the fixed price

[A1] Compared to the stylized analytical model, our main empirical model enriches the description of this bargaining interaction between sellers and buyers in the following aspects: (1) bargaining discounts are influenced by buyer, seller, and product characteristics as well as price and promotion; (2) bargaining success is a continuous measure, which also depends on the above attributes; (3) bargaining decision is endogenously decided via buyers comparing the expected bargaining gain and the bargaining cost; and (4) seller competition is included.

mechanism under one functional form and a given set of parameters. In other words, if we can show this, then it is clear that the conclusion from the counterfactual analysis is not pre-determined.

We set the demand function for type A as $D_1(p) = 1 - p$ and for type B as $D_2(p) = 1 - 2p$. It is clear from the demand functions that type B is more price sensitive and has lower willingness to pay than type A. Now, we consider different correlation patterns between a buyer's willingness to pay and the bargaining discount, and see how these influence the attractiveness of bargaining relative to a fixed price mechanism.

Scenario 1: The correlation between the bargaining discount and the buyer's willingness to pay is negative.

This scenario can arise, for example, if buyers with higher income (type A) prefer not to bargain and have higher willingness to pay. For simplicity, we assume type A does not bargain, i.e., $d_1 = 0$, and type B bargains to get a discount $d_2 = d$, where $d > 0$. Under this scenario, the profit under the fixed price mechanism becomes

$$\Pi_f = \max_p \{\theta(1-p)p + (1-\theta)(1-2p)p\} \tag{A3}$$

By maximizing the profit function, the optimal price can be solved as

$$p_f^* = \frac{1}{4-2\theta} \tag{A4}$$

If we consider the bargaining mechanism, the profit becomes

$$\Pi_{b1} = \max_p \{\theta(1-p)p + (1-\theta)(1-2(p-d))(p-d)\} \tag{A5}$$

By maximizing the profit function, the optimal price can be solved as

$$p_{b1}^* = \frac{4(1-\theta)d+1}{4-2\theta} \tag{A6}$$

Note that when a seller sets the optimal price under bargaining, even though the discount d is exogenously given, the seller endogenizes this amount into the pricing setting process, and the optimal price under bargaining is higher than that under the fixed price mechanism ($p_{b1}^* > p_f^*$). We will compare the profits via simulation later.

Scenario 2: The correlation between the bargaining discount and the buyer's willingness to pay is positive.

This scenario can arise, for example, if buyers with more shopping experience (type A) prefer to bargain and also have higher willingness to pay. For simplicity, we assume that type A bargains to get a discount $d_1 = d$ and type B does not bargain, i.e., $d_2 = 0$. Under this scenario, the profit under the fixed price mechanism is the same as above since the seller does not have the ability to charge different prices between type A and type B. However, with the bargaining mechanism, the profit becomes

$$\Pi_{b2} = \max_p \{\theta(1-p-d)(p-d) + (1-\theta)(1-2p)p\} \tag{A7}$$

By maximizing the profit function, the optimal price can be solved as

$$p^*_{b1} = \frac{1}{4 - 2\theta} \tag{A8}$$

Note that the optimal price under bargaining is the same as that under the fixed price mechanism when the correlation between the bargaining discount and the buyer's willingness to pay is positive ($p^*_{b2} = p^*_f$). In addition, the profits under different pricing mechanisms are different as the transaction prices are different from the optimal list price. As we will show below, the profits under fixed pricing may be higher or lower than that under bargaining, and thus the counterfactual results are not pre-determined.

We now turn to illustrating the above conclusion using a simulation approach. In order to carry out the simulation, We set the discount amount d to an arbitrary number (0.1) and vary the proportion of type A customers. The results from the simulation are presented in Figure A1.

Figure A1: Profits Comparison Under Fixed-Price and Bargaining Mechanisms

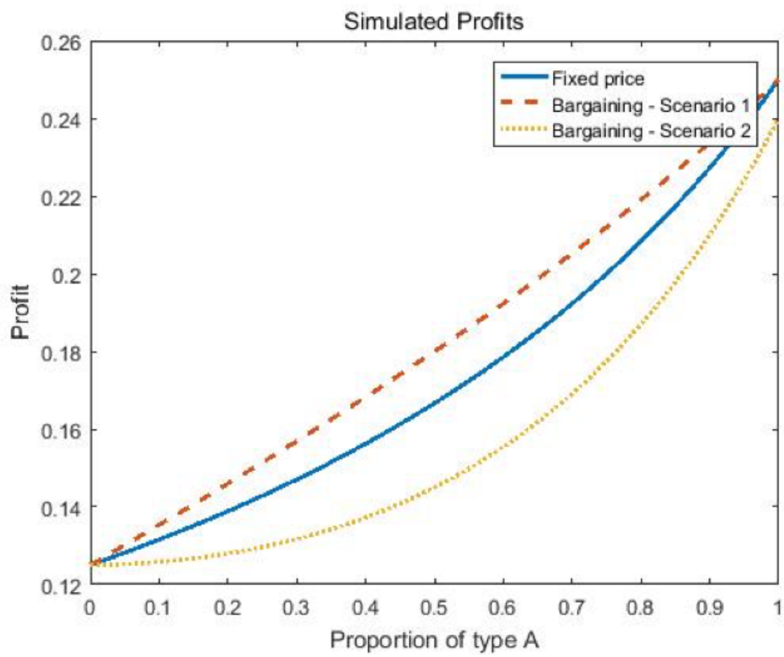

This figure shows the simulated seller profits under three scenarios. The blue solid line represents the fixed-price scenario, the red dashed line represents the bargaining scenario when the correlation between bargaining discount and willingness to pay is negative, and the yellow dotted line represents the bargaining scenario when the correlation is positive. It is clear from the above figure that this simple model does not directly determine the results from the counterfactual analysis. Under some scenarios, like scenario 1, bargaining yields higher profits than the fixed price mechanism. However, under some other scenarios, like scenario 2, bargaining is not as attractive as the fixed price mechanism. Through this exercise, we can see that the counterfactual results will depend on the correlation between buyers' willingness to pay and the bargaining outcomes. As long as we incorporate this correlation into the model, the counterfactual results are not predetermined, even if the bargaining discount is exogenously given.

C Analysis Based on Real Chatting Data

The chatting data come from a small seller on Taobao who mainly sells camera accessories. The data range from April 2013 to June 2013, including 306 chatting sessions. As we noted earlier in §7.5, we do not consider this seller to be a representative seller for the pool of sellers in our chosen product category of cell phones especially as the mean unit price in the cell phone category is about 1,500 yuan while it is about 250 yuan for this seller. Our objective in using these data is to provide more context and hopefully external validity.

First, we use the chatting data to examine whether the bargaining cost that we recover from our structural model is reasonable. We begin by looking at the time buyers spend chatting with the seller. The mean time spent on chatting is 17 minutes (Figure A2 shows the distribution). Unlike the cell phone category, where free shipping is standard, buyers bargain both on price and non-price attributes (e.g., shipping cost) in the camera accessories category. We split the chatting sessions based on whether the buyer bargained over the price, or bargained over non-price attributes, or did not bargain at all - see Table A3. Comparing rows 1 and 3, we see that a buyer spends on average about 10 more minutes chatting with the seller if s/he bargained over the price. Comparing rows 2 and 3, we see that if the subject of bargaining is other than price, then the chatting duration increases by a very modest amount (2 minutes). Using these data, we compute the average time cost of bargaining on average as follows. The median hourly wage in China in the 2012-13 period was reported 60 yuan (Fang and Lin, 2015). Thus the cost of 10 minutes spent on bargaining is 10 yuan. This estimate - 10 yuan - is very close to the average bargaining cost that we backed out from our structural model at 9 yuan. The similarity in these numbers suggests that our estimates are likely to have external validity.

Table A3: Chatting Duration Summary Statistics by Bargaining

	Chatting Duration (min)		
	Number of Obs.	Mean	Std. Dev.
Bargain over price	98	23.6	18.1
Bargain over things other than price	47	15.5	17.8
No bargain at all	161	13.3	11.7
All	306	17.0	15.7

Next, we would like to see whether the bargaining intention revealed from the chatting data is comparable to that revealed from our survey. Among the 306 chatting sessions, we see that 145 buyers tried to bargain over both price and non-price attributes, implying a bargaining intention of 47%. This number is lower than what we got from the survey reports (54% - 74%). We think the main reason is the price difference. For the chatting data, the average price is only about 250 yuan while in the survey we asked buyers about their bargaining intention for a product of 1,500 yuan. As a result, it is not surprising to see that the bargaining intention in the chatting data is somewhat lower than that in the survey. In order to verify this intuition, we used our estimated model parameters

Figure A2: Histogram of Chatting Durations

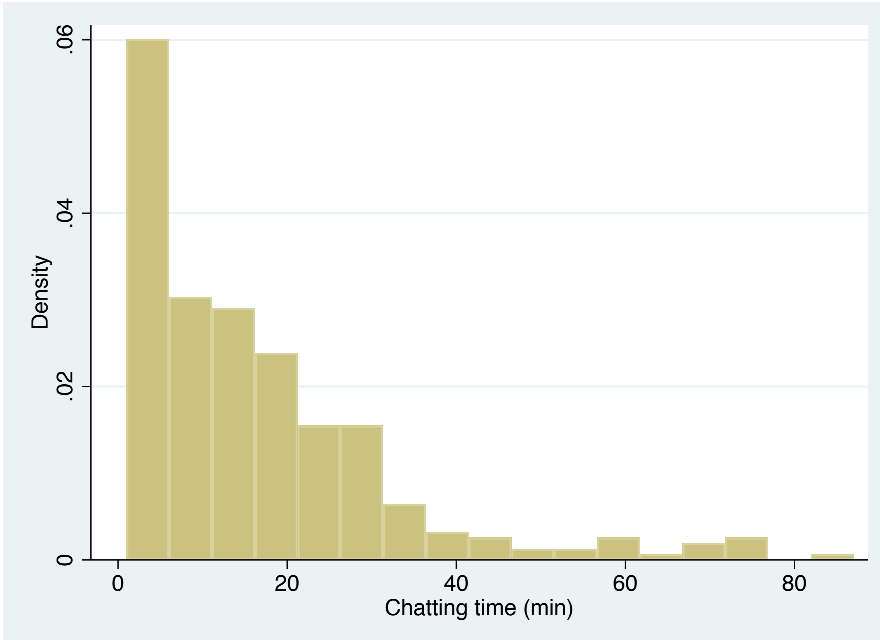

and plugged in a price of 250 yuan to compute buyers' mean bargaining intention. We find that the predicted bargaining intention is 52%, which is quite close to 47% (in the chatting data). This suggests that the use of the survey data to model unobservables (e.g., bargaining intention) is a reasonable approach.

Further, we would like to see whether the bargaining success rate is reasonable based on our estimates compared with that obtained from the chatting data. Among 145 buyers who initiated bargaining, we see that the seller agreed to give 61 buyers a discount. Among these 61 successful bargaining cases, 31 resulted in a transaction. This observation is important as it directly supports our modelling framework in Figure 5, where the bargaining stage is distinct from the purchase stage i.e., a buyer can decide not to purchase even after being successful at bargaining. The above numbers suggest a conditional bargaining success rate of 61/145 = 42%. This lies in the interval between the conditional success rate obtained from the survey (49%) and in the transaction sample (20%). The closeness of the success rate between the chatting data and the survey (42% vs. 49%) suggests that our assumption of using the perceived success rate from the survey in the estimation is reasonable. At the same time, the distance of the success rate between the chatting data and the cell phone transaction data (42% vs. 20%) suggests that it is important to run robustness checks on the above assumption (we do this in Section 7.3.1). Another relevant piece of information from the chatting data is that out of 61 successful bargaining cases, the buyers buy about 50% of the time. This provides an additional piece of evidence that using purchasers' data to infer some of the non-purchasers' behaviour may not be that bad as we see the purchase decision and the bargaining success are not strongly correlated.

Finally, we also see support for our use of the repeat purchase indicator in our model (equations 10 and 16) and the resulting statistically significant coefficient. In 16 of the chatting sessions, buyers cited the fact that they were repeat customers to bargain for a price discount, which is consistent with the statistically significant positive coefficient of the repeat purchase indicator on bargaining outcomes.

Overall, analysis of these data provides us correlationally consistent evidence on bargaining cost, bargaining intention, bargaining success as well as support for our model structure and specification.