

“Meet Me Halfway”: The Costs and Benefits 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

January 2021

Abstract

Bargaining is an important pricing mechanism, prevalent in both online and offline markets. However, there is little empirical work documenting the costs and benefits of bargaining, primarily due to the lack of real-world bargaining data. We leverage rich, transaction-level bargaining data from a major online platform and supplement it with primary data to quantify the costs and benefits of bargaining for sellers, buyers, and the platform. We do this by building a structural model of buyer demand and seller pricing decisions while allowing for the existence of bargaining initiation cost, loss-of-face cost, and price discrimination. Using our results, we perform three policy simulations to quantify the importance of not distinguishing between no-bargain and failed-bargain transactions, ignoring the loss-of-face cost, and not allowing for bargaining. These simulations provide rich details on how the various costs of bargaining impact our understanding of buyer and seller behavior and transaction outcomes. Banning bargaining, in particular, benefits the buyer and the platform greatly, but only has a modest benefit for sellers. Finally, we show that our results are robust to our assumptions and replicate in another product category.

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. They also thank the Senior Editor, the Area Editor, two anonymous reviewers, Francine Lafontaine, Dan Akerberg, 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. The third author gratefully acknowledges research funding from the Singapore Ministry of Education Social Science Research Thematic Grant, MOE2017-SSRTG-023. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of the Singapore Ministry of Education or the Singapore Government.

1 Introduction

Bargaining is an important pricing mechanism all over the world. While it has been prevalent so far in offline settings, it has become quite popular in the online world as well, especially on platforms such as eBay and Alibaba.¹ Extensive theoretical research has been carried out on various aspects of bargaining, 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, to date, there is very little empirical research investigating bargaining, especially on online platforms. This is in no small part due to the difficulty in obtaining granular data on the bargaining process, which is typically carried out verbally in person (even in online settings) over a very short time period. Thus, little is known about the relative costs and benefits of bargaining for all participants in real world settings. Such costs include the bargaining initiation cost and the “loss-of-face” cost (if bargaining fails) for buyers. The benefits of bargaining, e.g., its use as a price discrimination tool, trading off the potential gains from bargaining against the loss of driving away “bargaining-averse” buyers, often appeal to sellers. The existence and the magnitude of these various costs and benefits (for buyers and sellers) have, in turn, business implications for online platforms as well as broader implications for social planners.

In this paper, we quantify the costs and benefits of bargaining in a real-world setting. Specifically, we build a structural model of buyer (consumer) demand and sellers’ pricing decisions where we allow for the existence of bargaining initiation cost, loss-of-face cost, and price discrimination.² 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. We estimate the model using a novel combination of secondary (rich, transaction-level data on bargaining outcomes from a large digital platform in China) and primary data (obtained via online surveys from Chinese consumers). The latter allows us to distinguish between a failed-bargain and a no-bargain transaction as well as to quantify the loss-of-face cost. Using the estimates, we carry out various policy simulations to show how the identified bargaining primitives impact buyers, sellers, and the platform.

¹eBay 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.

²We are grateful to the AE for her/his suggestion of including the loss-of-face cost in our model.

Our modeling strategy on the demand side comprises of three stages, following the data generating process. In the first stage, a consumer decides whether to bargain by comparing the expected utilities with and without bargaining. The bargaining initiation cost is a key factor in determining the attractiveness of the bargaining mechanism. As in Wang (1995), Jindal and Newberry (2018), and Rubin and Brown (1975), this cost represents the cost of time and effort as well as the psychological cost associated with the decision to bargain. Note that this bargaining initiation cost results in an anticipated loss of utility by a buyer *prior* to initiating bargaining.

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., 2018; Gowrisankaran et al., 2015; Grennan, 2013; Jindal and Newberry, 2018), its validity hinges on the correct specification of each party’s disagreement payoff. Also, Nash bargaining’s prediction power is questionable (Neslin and Greenhalgh, 1986), and it may not explain consumers’ behavioral norms in bargaining adequately (Backus et al., 2019, 2020). Furthermore, it does not allow for bargaining breakdown that frequently happens in our data while the total surplus is positive (Loertscher and Marx, 2019). An alternative approach to model the bargaining process is to use an extensive-form bargaining model (e.g., Keniston, 2011; Larsen, 2014). However, that requires detailed alternating-offer data. In sum, the complications that arise in incomplete information settings (see Ausubel et al. (2002) for a detailed review) and the structure of our data (where we do not observe each party’s disagreement payoff and/or alternating offers) guide our approach in this stage. We use a two-part reduced form approach that allows to model the bargaining realization flexibly, avoiding the imposition of strong assumptions while adhering closely to the data generating process (Shelegia and Sherman, 2015). 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 discrete choice methodology to model purchase behavior, incorporating the bargaining realization and the psychological loss-of-face cost (in case bargaining fails).³ In order to address any potential endogeneity issues, we use a control function approach. Finally, on the supply side, we assume that sellers set prices to maximize profits while allowing for

³Previous research, principally in psychology, has shown that failed bargains result in buyers’ experiencing negative emotions (Miles, 2010; O’Connor and Arnold, 2001; Tuncel et al., 2016). This is what the loss-of-face cost in our model captures. Note that, unlike the bargaining initiation cost, this cost becomes relevant only *after* the bargaining process is concluded (and fails).

the possibility of buyer bargaining.

Our results provide rich insights on the bargaining primitives and process. First, we find the mean bargaining initiation 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 initiation costs is related to economic development with the costs being higher in more developed provinces, reflecting the fact that buyers in these provinces place a bigger premium on their time. As for the bargaining success 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, and the transaction is a repeat purchase. Buyer shopping experience drives up success but not discount amount. These findings add to the literature on the determinants of bargaining outcomes (e.g., Ayres and Siegelman, 1995; Backus et al., 2019, 2020; Draganska et al., 2010; Ellickson et al., 2018; 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. Our data also suggest a high level of heterogeneity in the loss-of-face cost, with the majority of customers having a cost of less than one yuan but a substantial proportion having a cost as high as 100 yuan. As a result, the purchase probability of the customers whose loss-of-face cost is at the bottom 10% is almost three times as much as that of the customers whose loss-of-face cost is at the top 10% after an unsuccessful bargain.

Using the results, we carry out three policy simulations. First, we highlight the importance of the bargaining initiation cost by illustrating what happens when the analyst is unable to distinguish between failed-bargain and no-bargain transactions. We show that not doing so leads to a large impact on the estimated bargaining initiation cost, the bargaining intention percentage, and the

value of allowing for bargaining. Second, as the loss-of-face cost can be considered as a market friction, we examine the impact of removing this friction for the platform. We find that it has a big impact on the purchase conversion on the platform. Finally, we pin down the benefit of allowing for bargaining on the platform. We do this by simulating a world where bargaining is not allowed, i.e., by “forcing” sellers to move to a fixed-price mechanism. The (conservative) results show that banning bargaining is only modestly beneficial for the average seller, but is greatly beneficial for both the buyer and the platform (and by implication for the social planner).

The fact that 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, allows us to make the following contributions (to both the marketing and economics literatures). 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 initiation costs (as our policy simulation shows). Second, the identification of bargaining initiation costs in the previous literature depends either on functional form assumptions or on very stringent data requirements. In our paper, we instead combine secondary (transaction) data with primary (survey) data, which enables us to identify bargaining initiation costs. Third, the use of the primary data also allows us to incorporate a cost that has been established in experimental settings, the loss-of-face cost, but never been incorporated into the analysis of market data. Fourth, 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 doing so 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 guaranteed. Finally, we are able to examine the costs and benefits of bargaining for not just buyers and sellers, but also for the platform. 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, e.g., fixed-price or bargaining.

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 policy simulations, robustness checks, generalizability, and external validity 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, allowing individual entrepreneurs and small businesses to do business anywhere with individual buyers. By 2014, Taobao had become a dominant e-commerce player in China (with over 80% market share). At that time, it 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 100 countries. Its online product listings total over 800 million and there were 72 million transactions each day – more than on Amazon and eBay combined. For more detail on Taobao’s business model, growth and success, see Chintagunta and Chu (2021), Chu and Manchanda (2016) and Clark (2016).

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 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 purchases and are well aware of the possibility of using Aliwangwang to carry out bargaining.⁴ 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

⁴In our Taobao Consumer Survey, described in §2.3, we find that the vast majority (81%) of buyers are aware of the possibility of bargaining. Sellers are also very responsive to chat requests, with 97% of sellers in our sample responding to a request in under 3 minutes. In addition, a search on the biggest Chinese search engine – baidu.com – for the term “Taobao Jiangjia” (Taobao Bargaining) returned around 6,710,000 links (carried out on Jan 14, 2021).

this “adjusted” price).

Figure 1: Taobao Item Page



As can be seen from Figure 1, the setup of the bargaining process on Taobao implies that there are no predetermined bargaining rules. Buyers are free to initiate a bargain (or not), carry out alternating-offer bargaining or take-it-or-leave-it bargaining or use any other bargaining heuristic. In addition, 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 category (cellphones) is particularly suited to study the comparison between the fixed-price and bargaining mechanisms on the e-commerce platform for two reasons. First, bargaining is well accepted in most small brick-and-mortar cellphone stores in China. Thus, the bargaining familiarity and habit 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. In 2012, the 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. We restrict the sample to single cellphone purchases as they account for over 95% of all transactions. Our final sample consists of 39,625 transactions made by 24,181 buyers across 8,965 sellers.

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.⁵ 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, as shown in Mayhew and Winer (1992)) in the analysis.

A successful bargain occurs if the transaction price is less than the posted price. Based on this metric, 16% of all transactions are associated with a successful bargaining incidence.⁶ For these transactions, the bargaining discount amount is captured by the difference between the Posted Price

⁵To 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). We test the sensitivity of our results to aggregating the less popular models to the brand level. We found the difference in the predicted conditional discount amounts with and without the aggregation is less than 2%. We therefore believe our results are not sensitive to this product definition.

⁶According 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.

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

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 bargaining. The proportion of successful bargained transactions is large enough to enable us to use 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 (with the caveat that this is an imperfect measure due to left censoring), defined as 1 if the focal transaction is at least the second interaction between the seller and the buyer. 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. 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 Online Appendix §O1 Figure O1 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 Online Appendix §O1 Figure O2) 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 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 the buyer 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 of buyers is about 29 years. We see a 70%-30% male-female split in our sample, but there is no significant difference in the products purchased or the bargaining outcomes across gender (Online Appendix §O2 Table O1). Using the location information, we supplement income data for each buyer using the province level mean.

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

One of our goals is to understand how the costs associated with bargaining affect the buyers, sellers, and the platform. Specifically, we aim to understand two types of bargaining related costs. One is the initiation cost, which is the cost a buyer expects to incur during the bargaining process. As noted earlier, this cost is a combination of a psychological cost and the expected time and effort cost. The other cost is the potential loss-of-face cost, the cost a buyer could incur if s/he decides to purchase the product from a seller when her/his bargaining attempt fails. As both these costs include a psychological component, which cannot be measured directly from the transaction data, we use one of the most common methods in the bargaining literature – primary data collection – to uncover these components (e.g., Ayres and Siegelman, 1995; Neslin and Greenhalgh, 1986; Small et al., 2007). We conducted Taobao consumer surveys to collect primary information on these costs (implementation details are provided in Online Appendix §O3). We then combined the primary data with the transaction data (with appropriate checks and controls – details below and in Online Appendix §O3) to estimate all the relevant costs and benefits of bargaining.

In order to understand the bargaining initiation cost, we asked Taobao consumers the following three questions: (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?” We intentionally hold the product fixed

and stay “ignorant” about product and seller characteristics in the survey in order 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 seller and product characteristics. However, the space of all combinations of product and seller characteristics in the transaction is very large (can be over 20,000 combinations). Thus, 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 (1) getting the necessary information from a reasonable number of respondents, and (2) the necessary variation to identify bargaining initiation costs. We also test the robustness of our results to this assumption in Online Appendix §O4.1.

We use the first question “Would you bargain with a seller if you are going to buy a cellphone priced at 1,500 yuan?” 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 that a consumer chooses “for sure to bargain” as $I(\textit{Certainly Bargain}) = I(\textit{Option} = A)$, and a variable indicating that a consumer may bargain as $I(\textit{Certainly} + \textit{May Bargain}) = I(\textit{Option} = A \textit{ or } C)$. Since we cannot match the survey with the transaction sample at the individual level, we use the survey to infer bargaining intention at the province level. Specifically, the proportion of consumers who answered they *will certainly* bargain serves as the lower bound of bargaining intention and the proportion of consumers who answered they *may or will certainly bargain* serves as the upper bound of bargaining intention in that province. In other words, the bargaining intention at the province level is bounded by the probability of *certainly* bargain and the probability of *may* bargain.

The second question “How likely would you expect to succeed in bargaining?” 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 “How much discount would you expect to get if bargaining succeeds?”

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. In addition, we believe that Taobao consumers on average form reasonable expectations on the conditional bargaining discount amounts. The similarity of the age distribution between the transaction data and the survey once again verifies that the two samples are comparable (see Online Appendix §O3 Figure O3).

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 0.006.⁷ A detailed summary breakdown by gender is provided in Online Appendix §O2 Table O2. Even though we do not see any significant 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: the Bargaining Survey

	N	mean	sd	min	max
I(Certainly Bargain)	1,009	0.543	0.498	0	1
I(May + Certainly Bargain)	1,009	0.736	0.441	0	1
Perceived Success Rate Bargaining	743	0.486	0.250	0.1	1
E[Discount Amount Success] (yuan)	743	165.5	221.5	1	1,000

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.

In order to understand the loss-of-face cost after a failed bargain via survey (see Online Ap-

⁷We compared these statistics to those obtained from a pilot survey carried out prior to this one (see Online Appendix §O3 for details) and found that they were similar, suggesting that they are stable.

pendix §O3 for the survey instrument), we asked the following question: “Assume you plan to buy a cellphone and you are interested in a cellphone model sold by Seller A priced at 1,510 yuan. Though this price is below your willingness to pay, you still started to bargain with Seller A. After quite a bit of bargaining, the seller did not give in. Then, you see that Seller B is selling exactly the same cellphone priced at p^B yuan. Assume all the attributes (reputation, product description, service, logistics, etc.) of Seller A and Seller B are the same, from whom would you like to buy the cellphone, Seller A or Seller B?” Seller B’s price p^B is randomly chosen from the following nine values – 1,511, 1,513, 1,516, 1,520, 1,530, 1,540, 1,550, 1,580, and 1,610 yuan.

Let c_f denote the loss-of-face cost after a failed bargain. Then the full cost of buying the cellphone from Seller A would be $1,510 + c^f$ and the cost of buying from Seller B is p^B . If the respondent selects to buy from Seller A, then we know that $1,510 + c^f \leq p^B$. Using the answers across respondents, we can calculate the cumulative distribution function (CDF) of the loss-of-face cost $Pr(c^f \leq p^B - 1,510)$ at the selected p^B . Essentially, the loss-of-face cost is identified by how much extra a consumer would be willing to pay for the exact same cellphone sold by a different seller after a failed bargain.

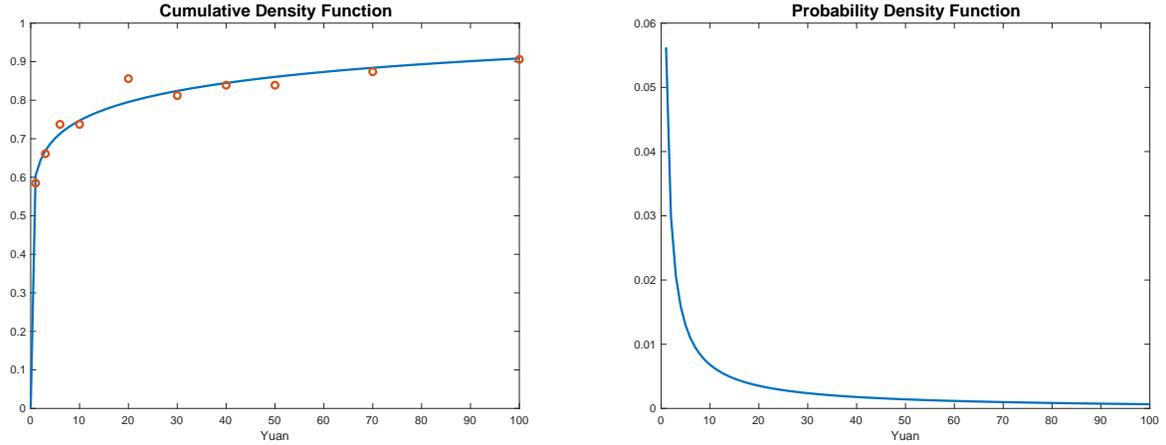
Since no consumer is likely to enjoy a failed bargain, the loss-of-face cost is bounded below at zero. In order to capture the distribution of the loss-of-face cost, we use a Gamma distribution which has the support of $(0, \infty)$ and is very flexible. Using the method of least squares to fit the empirical CDF values, we get the estimated shape parameter to be 0.093 and the estimated scale parameter to be 366.78.⁸ Figure 2 shows the cumulative density function on the left and the probability density function on the right of the estimated distribution of the loss-of-face cost. The red dots represent the empirical CDF values. As we can see, the estimated distribution has a long right tail, suggesting substantial heterogeneity in the population. About 61% consumers have a loss-of-face cost smaller than 1 yuan, while for about 10% of consumers, the loss-of-face costs can be as high as 100 yuan.

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 the costs and benefits of bargaining for all concerned parties. The building block of the model is an individual purchase occasion, which proceeds in three stages, as shown in Figure 3. In Stage I, a consumer decides whether

⁸We check the sensitivity of our results to this distribution by varying these parameters in Online Appendix §O4.4.

Figure 2: The Distribution of the Loss-of-Face Cost



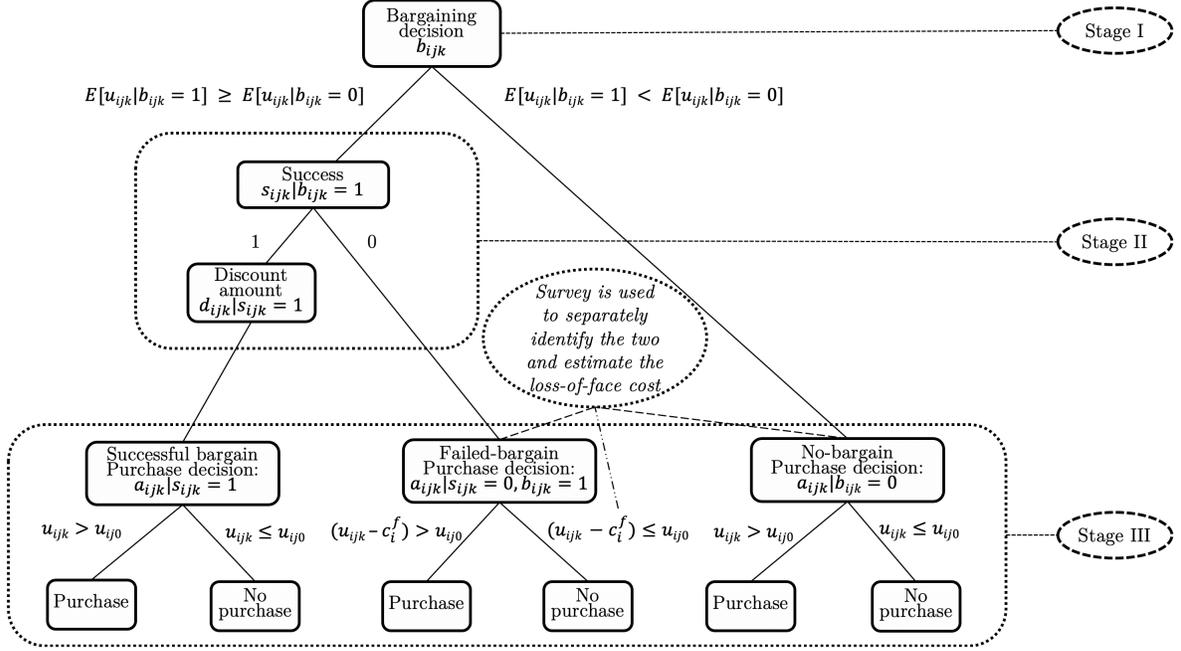
to bargain by comparing the expected bargaining gain and her/his inherent bargaining initiation cost. If the consumer decides to bargain, s/he proceeds to Stage II (bargaining realization), while if the consumer decides not to bargain, s/he skips Stage II and proceeds to Stage III (purchase decision). 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. As can be seen, the model starts with a consumer landing on a product page. We abstract away from modeling the search behavior explicitly due to the lack of the clickstream data. Instead, we use proxy measures to control for the common search processes employed by consumers, as detailed later in §3.1. We describe the model by working backwards from the last stage of the demand model. To make it easier for the reader to follow the model, we list the notation in Table 3.

3.1 Stage III: Purchase Decision

When consumers search for products on e-commerce marketplaces, in addition to product characteristics, they usually also 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 a product sold by a specific seller. A consumer purchases at most one product in a shopping occasion. The indirect utility that consumer i derives from buying product k from seller j is

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

Figure 3: The Bargaining and Purchase Process



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 $(\alpha + \alpha_{inc} \log(\text{Income}_i) + \alpha_{exp} \text{ShoppingExperience}_i)$ 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 and 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. Due to the data limitation, we are not able to model the search process and the seller competition explicitly. Instead, we control for the influence of the search behavior by using two measures w_{jk} to capture two common search processes and thus control for the resulting 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

Table 3: Notation Definition

Notation	Definition
u_{ijk}	Purchase utility
u_{ij0}	Outside option utility
δ_k	Intrinsic preference for product k
\bar{p}_{jk}	Posted price (promotion adjusted)
p_{jk}	Transaction price, which equals \bar{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^b	Individual bargaining initiation cost
c_{zl}^b, c_{zu}^b	Province's distribution of bargaining initiation costs' lower and upper bounds
c_i^f	Individual loss-of-face cost
mc_{jk}	Marginal cost for seller j on product k
Π_{jk}	Profit for seller j on product k

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 add 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 in 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 fixes a product and searches across different sellers and the second is when a buyer fixes a seller reputation level and searches for products. 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 underestimated (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. We discuss the operationalization of the control function in § 4.4.

Consumers make purchase decisions by comparing the utility of the focal product and the outside option. As can be seen in Figure 3, there are three different purchase scenarios, i.e., (1) purchase at a discounted price after a successful bargain, (2) purchase at the posted price after a failed bargain, and (3) purchase at the posted price without bargain. In the first and the third scenarios, consumer i chooses to purchase if and only if $u_{ijk} > u_{ij0}$. The second scenario is different as the consumer may experience a loss-of-face cost if s/he decides to purchase after a failed bargain. As noted earlier, we are able to include this in our model for the first time in the marketing or economics literature. We denote this loss-of-face cost as c_i^f , then consumer i chooses to purchase after a failed bargain if and only if $(u_{ijk} - c_i^f) > u_{ij0}$.

Taking all the three scenarios together, the conversion rate of product k sold by seller j is

$$\begin{aligned}
m_{jk} &= Pr(b_{ijk} = 1)Pr(s_{ijk} = 1|b_{ijk} = 1)Pr(u_{ijk}(p_{ijk}) > u_{ij0}) \\
&\quad + Pr(b_{ijk} = 1)Pr(s_{ijk} = 0|b_{ijk} = 1)Pr(u_{ijk}(\bar{p}_{jk}) - c_i^f > u_{ij0}) \\
&\quad + Pr(b_{ijk} = 0)Pr(u_{ijk}(\bar{p}_{jk}) > u_{ij0})
\end{aligned} \tag{2}$$

where b_{ijk} indicates consumer i 's bargaining decision and s_{ijk} indicates whether the bargaining between consumer i and seller j on product k is successful, i.e., the seller agrees to offer a price discount. \bar{p}_{jk} denotes the posted price and p_{ijk} denotes the discounted price after a successful bargain.

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 the fact that bargaining frequently breaks down (Loertscher and Marx, 2019). Also, the prediction power of the Nash bargaining solution in the outcomes of buyer-seller bargaining is questionable (Neslin and Greenhalgh, 1986) and consumers’ behavioral norms in bargaining may not be explained by existing theories (Backus et al., 2019, 2020). 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:

$$p(s_{ijk} = 1 | b_{ijk} = 1) = \Phi\left(\frac{x_{ijk}\gamma}{\sigma_s}\right) \quad (3)$$

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. 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 v_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:

$$d_{ijk}^* = x_{ijk}\theta + v_d \quad v_d \sim N(0, \sigma_d^2) \quad (4)$$

$$d_{ijk} = \exp(d_{ijk}^*) - 1 \quad \text{if } d_{ijk}^* > 0 \quad (5)$$

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 (3), 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 v_d in the discount amount equation. In the two-part model, v_s and v_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.⁹

Note that even though we use 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. With that, a seller can then endogenize the predicted bargaining outcomes when making the pricing decision to maximize profits, as will be shown in § 5.

⁹We tested the robustness of our findings by (a) adding higher order terms to this parametric specification and (b) using non-parametric methods (kernel density estimation). Details are available from the authors on request.

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 of bargaining against the expected utility of not bargaining. Thus, a consumer's bargaining decision b_{ijk} follows:

$$b_{ijk} = \begin{cases} 1 & \text{if } E[u_{ijk}|b_{ijk} = 1] > E[u_{ijk}|b_{ijk} = 0] \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where $E[u_{ijk}|b_{ijk} = 1]$ is the expected utility if consumer i decides to bargain and $E[u_{ijk}|b_{ijk} = 0]$ is the expected utility if not, specifically

$$\begin{aligned} E[u_{ijk}|b_{ijk} = 1] &= Pr(s_{ijk} = 1|b_{ijk} = 1)max\{u_{ijk}(p_{ijk}), u_{ij0}\} \\ &+ Pr(s_{ijk} = 0|b_{ijk} = 1)max\{u_{ijk}(\bar{p}_{jk}) - c_i^f, u_{ij0}\} - c_i^b \end{aligned} \quad (7)$$

$$E[u_{ijk}|b_{ijk} = 0] = max\{u_{ijk}(\bar{p}_{jk}), u_{ij0}\} \quad (8)$$

Note that bargaining initiation cost c_i^b affects a consumer's bargaining decision (stage I) but does not directly affect the consumer's purchase decision (stage III) as bargaining initiation cost becomes sunk when the bargaining process ends, regardless of the bargaining outcome.

4 Identification and Estimation of the Demand

As explained before, to incorporate the bargaining initiation cost and the loss-of-face cost, we combine the transaction sample with the Taobao Consumer Surveys. With the three-stage demand model explained in § 3, our estimation algorithm follows the following steps (shown in Figure 4):

Step one: We first estimate the discount amount conditional on success using the truncated regression model (equations 4 & 5), and then use the model estimates to predict expected discount amount conditional on success for each shopping occasion, $E[d_{ijk}|s_{ijk} = 1], \forall ijk$.

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 distribution of bargaining initiation cost for each province using the method of moments (equations 6 & 7).

Step three: Once calculated, the estimated expected discount amount and the estimated bargaining initiation cost imply the probability of bargaining. With the bargaining probabilities and the

observed bargaining success, we recover the parameters governing the bargaining success function using the simulated maximum likelihood approach (equation 3).

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, where we also incorporate the loss-of-face costs after a failed bargain. The primitives in the utility function (equation 1) are estimated by minimizing the distance between the estimated and the empirical conversion rates using the generalized method of moments. We use the control function to deal with the endogeneity problem.

Figure 4: The Estimation Process

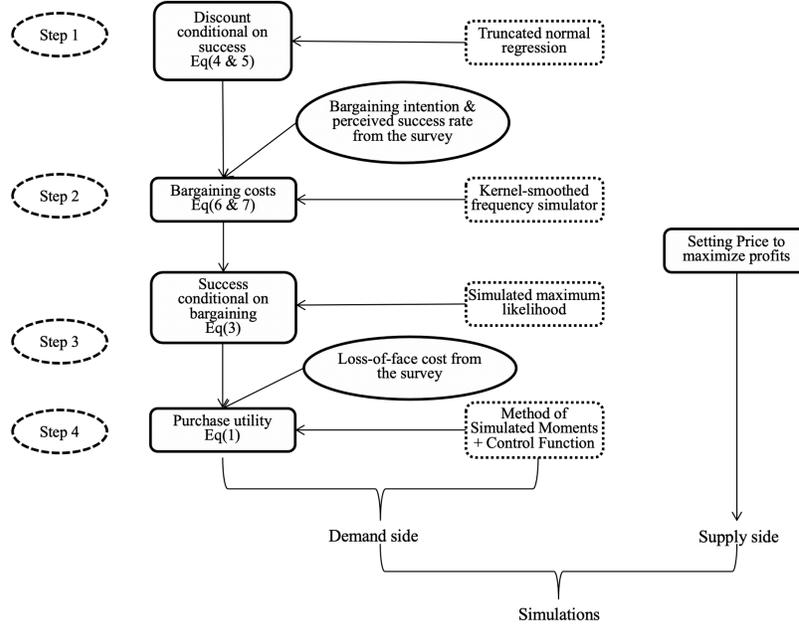


Table 4: A Brief Summary of Identification

Parameters	Identified by
δ_k : intrinsic product preference	The variation in the conversion rate, and
$\alpha, \alpha_{inc}, \alpha_{exp}$: price preference	the variation in consumers' income and experience level, and
β_x : nonprice preference	the variations in seller/product characteristics, and
β_w : impact of competition on outside option	the variation in bargaining decisions and outcomes
c_{zl}^b, c_{zu}^b : lower and upper bounds of province-level bargaining initiation costs	(i) the variation in bargaining intentions (ii) the variation in expected discount amounts across shopping occasions
γ : bargaining success parameters	The variation in bargaining decisions and the observed success rates
θ : bargaining discount parameters	The variation in the discount amounts and seller/buyer/product characteristics
mc_{jk} : marginal cost	(i) the estimated price elasticity (ii) the estimated bargaining outcome (iii) the profit maximizing assumption

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

$$l_d = \prod_{\{ijk:d_{ijk}^* > 0\}} \left\{ \frac{\phi\left(\frac{d_{ijk}^* - x_{ijk}\theta}{\sigma_d}\right)}{1 - \Phi\left(\frac{-x_{ijk}\theta}{\sigma_d}\right)} \right\} \quad (9)$$

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 the competition level. d_{ijk}^* is the latent bargaining discount amount. As the seller does not adjust the posted-price for each individual consumer, the aggregate unobserved shocks are unlikely to cause any systematic correlation between the posted 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 (see Nevo, 2000, p. 544 for more detail). One assumption here is that the discount amount conditional on success is not systematically different across purchasers and non-purchasers. In Online Appendix §O4.2, we show that our results are robust to this assumption.

Once equation (9) is estimated, we predict the expected discount amount conditional on success for each shopping occasion:

$$E[d_{ijk}^* | s_{ijk} = 1] = x_{ijk}\hat{\theta} + \hat{\sigma}_d \lambda(x_{ijk}\hat{\theta}/\hat{\sigma}_d) \quad (10)$$

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 d_{ijk} by transforming the latent variable d_{ijk}^* by using $E[d_{ijk} | s_{ijk} = 1] \approx \exp(E[d_{ijk}^* | s_{ijk} = 1]) - 1$. Given the standard deviation of d_{ijk}^* is small, this approximation is appropriate.

4.2 Step two: bargaining initiation 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 collected information on consumers' bargaining intention from the Taobao Consumer survey (described earlier in §2.3). The bargaining decision is simpler in the survey than in the transaction sample. In the transaction sample, the bargaining responses are explained by differences in both the bargaining initiation costs and the bargaining gains (depending on product characteristics). In the survey, the bargaining responses are mainly driven by differences in bargaining initiation costs as we fix the product and seller characteristics, so the bargaining decision is made by $b_{ijk} = I(E[d_{ijk}|b_{ijk} = 1] > c_i^b)$. As such, bargaining initiation costs are identified by the variation in consumers' bargaining decisions and the variation in the expected bargaining gains across shopping occasions. The survey provides us with the lower bound and upper bound of bargaining intention in each province. As bargaining intention is negatively correlated with the bargaining initiation cost, we are able to derive the upper and lower bounds of the bargaining initiation cost for each province.

Consumers make bargaining decisions by comparing the expected bargaining gain and the bargaining initiation cost. The expected bargaining gains are calculated by multiplying the perceived success rate conditional on bargaining 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

$$B(c_i^b; x_{ijk}) = \{ijk | E[d_{ijk}|x_{ijk}] > c_i^b\} \quad (11)$$

By matching the empirical value of Taobao consumers' bargaining intention and perceived success rate and the estimated values, we form the following two moment conditions:

$$Pr[b = 1|c_{zl}^b] = \int I(ijk \in B(c_{zl}^b; x_{ijk}))f(x_{ijk})dx_{ijk} \quad (12)$$

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

where c_{zl}^b and c_{zu}^b are lower and upper bounds of bargaining initiation cost in province z , respectively.

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

$$S(w_{ijk}) = \frac{1}{1 + \exp(-hw_{ijk})} \quad (14)$$

where $w_{ijk} = E[d_{ijk}|x_{ijk}] - c_i^b$, the difference between the expected bargaining gain and the bargaining initiation cost. h is a tuning parameter in the kernel function, which is set to 0.1. As $w_{ijk} \rightarrow \infty$, $S(w_{ijk}) \rightarrow 1$. 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

With the estimated upper and lower bounds of the bargaining initiation costs for each province, we tried uniform distribution, triangular distribution, and truncated normal distribution for bargaining initiation costs. As the results are very similar, we will report the uniform distribution in what follows and assume that each individual's bargaining initiation cost is a random draw from the province's distribution. The identification of the bargaining success parameters comes from the variation in bargaining decisions and the variation in the observed success rates across different shopping occasions.¹⁰ The likelihood of observing a bargaining success can be written as

$$\begin{aligned} l_s &= \prod_{ijk} l_{s,ijk} \\ &= \prod_{ijk} Pr(s_{ijk} = 1 | b_{ijk} = 1) Pr(b_{ijk} = 1) \\ &= \prod_{ijk} \Phi(x_{ijk}\gamma/\sigma_s) Pr(E[d_{ijk}|x_{ijk}] - c_i^b > 0) \\ &= \prod_{ijk} \Phi(x_{ijk}\gamma/\sigma_s) \int I(E[d_{ijk}|x_{ijk}] - c_i^b > 0) f(c_i^b) dc_i^b \end{aligned} \quad (15)$$

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.

¹⁰In the estimation we treat consumers as myopic, i.e., they only compare the expected bargaining gain with the bargaining initiation 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 Online Appendix §O4.1. There is no substantive change in the main results.

4.4 Step four: purchase utility

In the purchase stage, the unknown parameters are intrinsic product preferences δ_k , heterogeneous price coefficients α , 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, i.e., the conversion rates at the seller level. Conversion rate is defined as the ratio between the number of unique purchasers and the number of unique visitors. We start by assuming that one in five unique visitors is a serious shopper, thus the conversion rate is multiplied by five. Since price elasticities and policy simulations are all relative measures, our results are not sensitive to this multiplier, which we tested using two, five and ten. The average conversion rate is 10.02%. We use the seller-level conversion rate to proxy for the seller-product-level conversion rate as less than 5% sellers sell product across categories.

The potential loss-of-face cost c_i^f of an individual buyer i enters the utility function as a random draw from the distribution estimated in §2.3 (see Figure 2). Note that only those buyers who make a purchase after a failed bargain incur this cost. Given the long tail of this distribution, there is likely to be substantial heterogeneity across buyers. We explicitly investigate the impact of this heterogeneity in §7.2.

In the transaction sample, since we do not observe buyer characteristics of the non-purchasers, 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, and (3) a pool of buyers who are the same as the buyer in the focal transaction. We find small differences in 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.

In the estimation, we include product fixed effects to absorb any correlation between the price and the unobserved product characteristics. However, the potential correlation between the price and the unobserved seller characteristics or 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.

In order to control for the potential price endogeneity problem in the purchase stage, we use the control function approach (Petrin and Train, 2010; Wooldridge, 2015), and take the following three

steps: first, we identify the valid instrumental variables; second, we obtain the separate variations in the residuals (i.e., control functions) by regressing the endogenous variables on the exogenous and instrumental variables, which essentially work as proxies for the omitted variables; and third, we include the control functions (functions of the residuals from the second step) in the model. The endogenous variables include the posted price, the interaction between price and buyer income and the interaction between price and buyer shopping experience level (the latter two are used to capture the heterogeneity of price sensitivity).

In the first step of the control function approach, 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 the information is easily accessible to consumers, *DistNextLevel* is unlikely to affect a consumer’s purchase decision directly. We also use the interaction between *DistNextLevel* and buyer income, the interaction between *DistNextLevel* and buyer shopping experience level, and a squared term of *DistNextLevel* as instruments.

In the second step, we use the posted price, the interaction between price and buyer income, and the interaction between price and buyer shopping experience level as three dependent variables and regress them on all the exogenous variables, including seller, buyer and product characteristics, and the aforementioned instrumental variables. From the three regressions, we get three sets of residuals, i.e., control functions. The Stock-Yogo weak identification F tests are 15.01, 18.73, and 245.75 for the three regressions, respectively, which suggests that they all satisfy the relevance condition (Stock and Yogo, 2005).

In the third step of the control function approach, we include the control functions in the main purchase model, as described in §3.1. After estimating the third step, we conduct the over-identification

test of the instruments to see whether they are uncorrelated with the error term. The Sargan test statistics is 1.59, which is not statistically significant at the 5% test level, which means that our instruments are valid (Sargan, 1958).

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

$$\begin{aligned}
m_{jk} &= \int Pr_{ijk}(p_{ijk})f(p_{ijk})dp_{ijk} \\
&\approx Pr(u_{ijk} > u_{ij0}|b_{ijk} = 0)Pr(b_{ijk} = 0) \\
&\quad + Pr(u_{ijk} - c_i^f > u_{ij0}|s_{ijk} = 0)Pr(s_{ijk} = 0|b_{ijk} = 1)Pr(b_{ijk} = 1) \\
&\quad + Pr(u_{ijk} > u_{ij0}|s_{ijk} = 1)Pr(s_{ijk} = 1|b_{ijk} = 1)Pr(b_{ijk} = 1)
\end{aligned} \tag{16}$$

which includes three parts (as shown at the bottom of Figure 3): 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 empirical conversion rates and the model predicted 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. Through this process, we account for both sources of variance – from the data generating process and from the simulation. Note that via the bootstrapping method, the estimates in the former steps are used in the latter steps, thus, any volatility of the estimates in the former steps is accounted for in estimating parameters in the latter 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 4). 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.¹¹ The profit of seller j on product k is given by

¹¹It 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

$$\begin{aligned}
\Pi_{jk} &= Pr(u_{ijk} > u_{ij0} | b_{ijk} = 0) Pr(b_{ijk} = 0) (\bar{p}_{jk} - mc_{jk}) \\
&+ Pr(u_{ijk} - c_i^f > u_{ij0} | s_{ijk} = 0) Pr(s_{ijk} = 0 | b_{ijk} = 1) Pr(b_{ijk} = 1) (\bar{p}_{jk} - mc_{jk}) \\
&+ Pr(u_{ijk} > u_{ij0} | s_{ijk} = 1) Pr(s_{ijk} = 1 | b_{ijk} = 1) Pr(b_{ijk} = 1) (\bar{p}_{jk} - \hat{d}_{jk} - mc_{jk}) \\
&= Pr_{jk,no}(\bar{p}_{jk} - mc_{jk}) + Pr_{jk,failed}(\bar{p}_{jk} - mc_{jk}) + Pr_{jk,successful}(\bar{p}_{jk} - \hat{d}_{jk} - mc_{jk}) \quad (17)
\end{aligned}$$

where the notations are the same as in §4.4. The profit function consists of three parts: transactions at the posted price \bar{p}_{jk} with no bargain, transactions at the posted price \bar{p}_{jk} after a failed bargain, and transactions at a bargained price $\bar{p}_{jk} - \hat{d}_{jk}$ after a successful bargain, where \hat{d}_{jk} is the seller's expected bargaining discount amount for a representative consumer. 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. In this price setting process, the competition has been implicitly taken into consideration through the purchase probability and the bargaining outcomes where seller competition comes into play. 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 k can be written as

$$\begin{aligned}
&\left(\frac{\partial Pr_{jk,no}}{\partial \bar{p}_{jk}} + \frac{\partial Pr_{jk,failed}}{\partial \bar{p}_{jk}} + \frac{\partial Pr_{jk,successful}}{\partial \bar{p}_{jk}} \right) (\bar{p}_{jk} - mc_{jk}) \\
&= \left(\frac{\partial Pr_{jk,successful}}{\partial \bar{p}_{jk}} \hat{d}_{jk} + Pr_{jk,successful} \frac{\partial \hat{d}_{jk}}{\partial \bar{p}_{jk}} \right) - (Pr_{jk,no} + Pr_{jk,failed} + Pr_{jk,successful}) \quad (18)
\end{aligned}$$

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 present.

On Taobao, sellers have full freedom to set both the list price and the promotion depth at any given point in time, thus the promotion depth decision and the list price decision are combined into one promotion-adjusted posted price decision, as modeled above. As 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, previous research (e.g. Mayhew and Winer, 1992) suggests that consumers react mainly to the promotion indicator rather than the promotion depth, so we include a promotion indicator in consumers' purchase decision (equation 1).

behavior, we conduct a sensitivity test by adding a weighting function, which is formulated as $(1 + \frac{1}{\text{distance to the threshold}})$, into the profit maximizing objective. The results are not sensitive to the added weighting function.

6 Results

We first discuss the estimates for the bargaining realization and bargaining initiation 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 Costs

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.

Table 5: Bargaining Realization Estimates

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

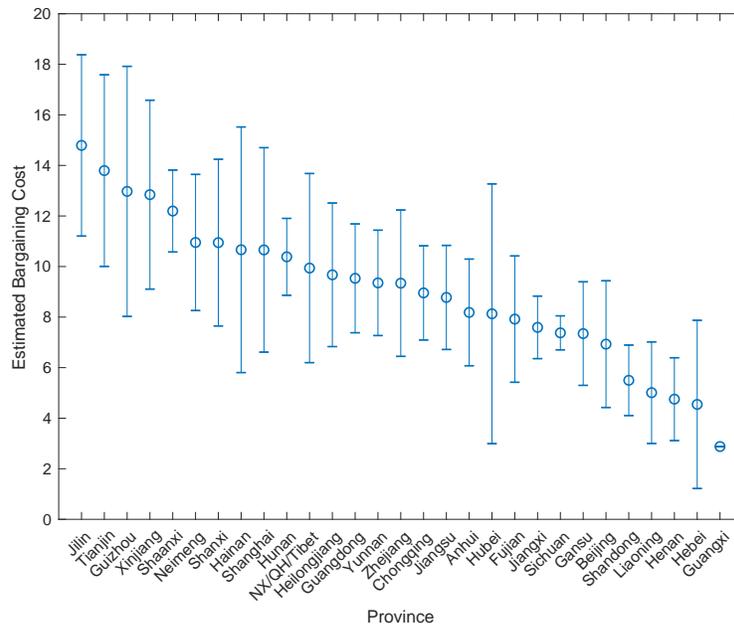
Figure 5 shows the estimated distribution of bargaining initiation cost in each province. On

average, the bargaining initiation 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 initiation 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 initiation 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 initiation 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 initiation 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 initiation costs across provinces, we investigate the difference of bargaining initiation costs across gender. We find that on average, female buyers have slightly higher bargaining initiation 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.

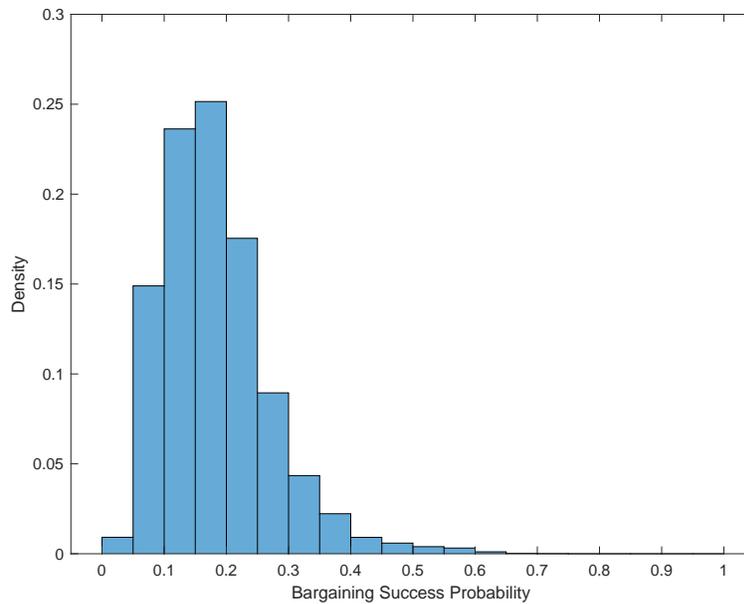
Using the estimates from the bargaining realization and the estimated bargaining initiation costs, we find that on average bargaining would be initiated in about 78.2% of shopping occasions. Conditional on bargaining, Figure 6 represents the histogram of the estimated success rate across all shopping 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 5: Estimated Province-level Average Bargaining Initiation Costs



Note that as no respondent in Guangxi province chooses “may bargain,” the lower bound and the upper bound for that province overlap.

Figure 6: Estimated Bargaining Success Probability



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.

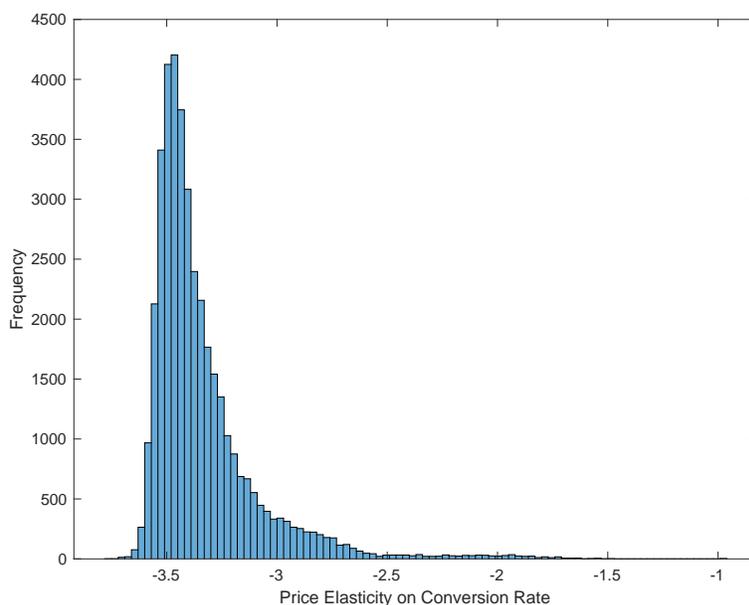
Table 6: Purchase Model Estimates

	Without Control Function	With Control Function
log(Price)	-0.492*** (0.049)	-3.963*** (0.239)
log(Price) * Buyer Shopping Experience	0.001 (0.001)	0.026*** (0.002)
log(Price) * log(Buyer Income)	0.002 (0.004)	0.013*** (0.004)
I(Promotion)	0.188*** (0.011)	0.192*** (0.012)
Seller Reputation Level	0.059*** (0.007)	0.173*** (0.012)
Detailed Seller Rating	1.103*** (0.201)	4.042*** (0.330)
Store Age	-0.101*** (0.005)	0.024 (0.019)
Repeat Purchase Rate	2.353*** (0.455)	2.002*** (0.406)
Product Age	0.010* (0.005)	-0.766*** (0.051)
# of Sellers with Same Product	0.085*** (0.015)	0.328*** (0.018)
# of Sellers with Same Reputation	-0.042*** (0.003)	-0.033*** (0.003)
<i>Terms to Correct for Endogeneity</i>		
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 7). 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%. As there are three types of transactions, we estimate the price elasticity for each of the scenarios. The implied price elasticities are -7.4 (s.d. = 32.6) for transactions made with no bargain, -3.1 (s.d. = 0.4) for transactions made after a failed bargain, and -2.7 (s.d. = 0.4) for transactions made after a successful bargain. The three types of price elasticities are statistically significant different from each other ($p < 0.001$). This implies that conversion of transactions with no bargain is the most elastic, followed by transactions after a failed bargain, and trailed by transactions after a successful bargain. This result makes intuitive sense as an increase in the posted price will not deter consumers who can get price discount from bargaining as much as those who have no choice but to purchase products at the posted price.

Figure 7: Implied Price Elasticity on Conversion Rate



After a failed bargain, consumers would incur a loss-of-face cost if they choose to purchase. Taobao Consumer Survey suggests that the average potential loss-of-face cost is 34 yuan. However, not every consumer would incur this cost after a failed bargain as some of them would walk away rather than making the purchase. The estimated average loss-of-face cost among the transactions after a failed bargain is 2.0 (s.d. = 1.9) yuan. This number is a lot lower than the average among the population as consumers who have a higher loss-of-face cost would be more likely to walk away

and thus do not incur this cost, as suggested by our model in § 3.1.

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 Policy Simulations

In this section, we focus on three policy simulations in order to showcase our results. We first focus on the costs of bargaining by looking at the bargaining initiation cost and the loss-of-face cost. We then shut down bargaining on the platform and examine the impact on the whole system in order to quantify the benefits of bargaining.

7.1 The Importance of Distinguishing Between No-bargain and Failed-bargain Transactions

With only observational data, from online e-commerce markets or offline markets where bargaining is allowed, researchers have typically assumed that transactions at the posted price are due to consumers choosing not to bargain. Given our unique combination of primary and secondary data, we do not have to make this assumption. Specifically, we explore the importance of distinguishing no-bargain transactions and failed-bargain transactions by comparing the results with and without this assumption.

When we do not distinguish no-bargain and failed-bargain transactions, we are essentially assuming that consumers would always get a price discount if they choose to bargain as they would purchase at the posted price only when they choose not to bargain. Under this assumption, we find that the bargaining initiation cost is dramatically overestimated by 744% at 76 yuan (vs. 9 yuan). This estimate (76 yuan) is unlikely to be realistic as it is 4 to 7 times the minimum hourly wage in China (which ranges from 11 to 20 yuan). In addition, the bargaining intention is underestimated by 79% (16% vs. 78%). Given the bargaining culture in China as well as the evidence presented in our primary data, this is also not an accurate depiction of reality. Finally, we show in our third policy simulation (§7.3) that banning bargaining will save buyers a total of 4.6 million yuan in bargaining initiation costs. However, not distinguishing between no-bargain and failed-bargain transactions overestimates this by 72% to 7.9 million yuan.

This policy simulation shows that it is critical to distinguish between no-bargain and failed-bargain transactions in any empirical analysis of bargaining. Not distinguishing between these two situations has a big impact on the estimated bargaining primitives such as the bargaining initiation cost and the propensity to bargain. This can then lead to erroneous decisions, especially with respect to pricing, by sellers which can in turn affect the economics of the platform.

7.2 The Implications of Accounting for Loss-of-face Cost

One of the contributions of this paper is the incorporation of the negative emotions when a bargaining process fails via the use of a loss-of-face cost. The most direct impact of the loss-of-face cost is on deterring consumers from making a purchase after a failed bargain.

As shown in Figure 2, the distribution of this loss-of-face cost is skewed with a long right tail (details on the estimation of the distribution are in §2.3). In this section, we explore the importance of both the existence of this cost and the heterogeneity across consumers.

In order to explore the implications of not incorporating the loss-of-face cost, we run a policy simulation where consumers are assumed to not experience any negative emotion if they make a purchase from a seller after they failed at bargaining. We re-estimate our model by setting this cost to zero for all such transactions.

We find that ignoring the loss-of-face cost results in minor changes to the purchase model estimates, but has a significant impact on assessing the policy implications if the platform were to implement a no-bargaining policy (discussed in detail in §7.3). With and without the existence of the loss-of-face cost, we find that this policy change has the following impacts. First, not accounting for the loss-of-face cost would significantly underestimate the increase in the conversion rate by 14% (0.6% vs. 0.7%). Second, it would underestimate the loss-of-face cost by 0.14 million yuan per day (0 vs. 0.14). Third, it would overestimate the bargaining initiation costs by 0.1 million yuan per day (4.8 vs. 4.7). Finally, it would underestimate the total sales by 0.06 million yuan per day (87.47 vs. 87.53)

In order to show the impact of the heterogeneity in this cost, we compute the purchase probabilities after a failed bargain across consumers with different loss-of-face costs. For consumers who have the median loss-of-face cost, the purchase probability after a failed bargain is about 5.7%. We then look at the purchase probability of consumers in the bottom decile in terms of this cost. Not

surprisingly, given the skewness of the loss-of-face cost distribution, their purchase probability is almost identical to that of a consumer with the median cost. However, if we look at consumers in the top decile of the loss-of-face cost distribution, we find that their purchase probability after a failed bargain drops to 2.2%, about one-third of that of a consumer with the median cost.

In conclusion, this policy simulation not only quantifies the impact of the loss-of-face on transaction outcomes, but also shows for the first time that a cost that has been documented only in the psychology literature is economically significant in real settings.

7.3 The Effect of Banning the Bargaining Mechanism

In this policy simulation, we investigate what would happen if the platform (Taobao) disallowed bargaining completely for all sellers. This exercise has external validity as pricing mechanisms used by the world’s top e-commerce platforms have changed over time. For example, Amazon started as a fixed-price platform and introduced bargaining on certain product categories in 2014, but then discontinued bargaining in 2019. eBay started as an auction site, then introduced a fixed price option followed by bargaining in 2005.¹²

If bargaining is disallowed, sellers will switch from a mixed-price to a fixed-price mechanism, and set their profit-maximizing prices according to:

$$p_{jk}^* = \operatorname{argmax}_{p_{jk}} \{Pr(a_{jk}(p_{jk}) = 1)(p_{jk} - mc_{jk})\} \quad (19)$$

where the primitives are estimated from the model. By setting the first-order condition to zero, we 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 gross merchandise volume (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.¹³ After the change, as no bargaining

¹²Interestingly, though eBay allows bargaining, its setting (as also Amazon’s when it allowed bargaining) is somewhat different from Taobao’s. Specifically, 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 many parts of the world.

¹³Note that the average transaction price reported here is different from the average *observed* transaction price reported in Table 1. Here, the average transaction price is simulated for the 39,625 shopping occasions by an average

Table 7: Policy Simulation: 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 (%)	10.67	10.74	0.7%
Number of Transactions per day	69,622	70,080	0.7%
Total Sales per day (million yuan)	87.10	87.53	0.5%
Bargaining Initiation Costs per day (million yuan)	4.6	0	N/A
Loss-of-face Costs per day (million yuan)	0.14	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 calculated 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 initiation cost is calculated as the number of transactions per day divided by the average conversion rate multiplied by the probability of bargaining and by the average bargaining initiation cost.

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. Due to the additional discounts that consumers could get through bargaining, the final average transaction price between the two scenarios is not very different.

The ban on bargaining increases the conversion rate to 10.74% from 10.67% – a (relative) 0.7% change. This increase in the conversion rate arises from three sources. First, some consumers who chose not to bargain and not purchase at the high posted price due to the high bargaining initiation cost before the policy would purchase at the lower posted price after the policy. Second, some consumers who did not get a price discount after a failed bargain and walked away before the policy would purchase at the lower posted price after the policy. Finally, some consumers who did not purchase after a failed bargain due to the high loss-of-face cost before the policy would purchase at the lower posted price after the policy. Thus, even though bargaining allows sellers to price discriminate, the existence of bargaining initiation costs and loss-of-face costs make the conversion rate higher under a pure fixed-price mechanism than under the bargaining mechanism.

Before the policy change, there are 69,622 transactions per day in the Taobao cellphone category. The estimated number of transactions per day after banning the bargaining mechanism increases to consumer under the current mixed-price mechanism. We do this to facilitate a “fair” comparison between the mixed-price and the fixed-price mechanisms.

70,080. 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.¹⁴ 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 87.53 million yuan after the ban, which is about a 0.5% increase.

In addition to the changes on transaction volume and total sales, the bargaining initiation costs and the loss-of-face costs also have significant welfare implications. Intuitively, these two bargaining related costs function as market frictions, which may lead to market inefficiency. Under the bargaining mechanism, a consumer who initiates bargaining, successfully or not, incurs the bargaining initiation costs. The estimated bargaining initiation costs incurred by buyers in the cellphone category alone are 4.7 million yuan per day. If consumers decide to purchase after a failed bargain, they would also incur a loss-of-face cost. The estimated loss-of-face costs are 0.14 million yuan per day. This implies that banning bargaining could save consumers 4.84 million yuan from bargaining related costs, about 5.6% of the category total sales, a statistically and economically significant proportion. On the sellers' side, it is likely that they would incur some bargaining related costs too. 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, sellers', and buyers' point of view.

We have shown that the average conversion rate and the average sales increase for sellers. Next we investigate how this increase differs across sellers. To study the heterogeneous effect of the banning bargaining policy, we calculate the percentage of the profit change before and after the ban on bargaining and regress the percentage change of profit on a series of seller characteristics. The results are reported 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

¹⁴We 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 anything, we expect the traffic to increase as the majority do not like bargaining. This would further reinforce our findings and managerial recommendations.

Table 8: Heterogenous Effects of the Banning Bargaining Policy

	% change in profit after banning bargaining
Seller Reputation Level	-0.053*** (0.001)
Detailed Seller Rating	0.246*** (0.016)
Store Age	0.010*** (0.001)
I(Promotion)	-0.210*** (0.004)
Product FE	Yes
Number of Observations	39,625

Lippman, 1998). Also, for a similar reason, we find that the fixed-price mechanism is more appealing for sellers with high detailed seller rating and longer store age. Products with promotion indicator benefit more from bargaining. This result has intuitive explanation in that given the posted price being the same, an indicator of promotion leads to a lower bargaining amount under the mixed-price mechanism but it should not have any effect on the transaction price under the fixed-price mechanism. This result is only suggestive as we assume the promotion indicator stays the same before and after the pricing mechanism change (though the promotion depth is allowed to change as it is implicitly captured in the pricing decision).

The goal of this paper is to accurately measure the value of bargaining for sellers, buyers, the e-commerce platform, and the social planner in terms of social welfare. It is important that our results and the policy simulation outcomes are robust to the model assumptions and can be generalized. In Online Appendix §O4, we present the robustness of our results to the following four cases when the key model assumptions are relaxed: (1) consumers' survey responses do not reflect their true belief, (2) bargaining outcomes systematically differ between purchasers and non-purchasers, (3) observed characteristics come into play in the bargaining realization process, and (4) loss-of-face cost follows a different distribution than the one we obtained via survey. In addition, we conduct a generalizability check by replicating our findings in the women's shoes category (see Online Appendix §O5). Lastly, we conduct text and sentiment analysis on some real-world chatting history data to provide external validity for our study (see Online Appendix §O6).

8 Discussion and Conclusion

This paper contributes to a small but growing body of empirical literature on bargaining. Bargaining is an important pricing mechanism all over the world, especially as it is being adopted by e-commerce platforms. In this paper, we quantify the costs and benefits of bargaining in a real-world setting. We do this by building a structural model of buyer (consumer) demand and sellers’ pricing decisions where we allow for the existence of bargaining initiation cost, loss-of-face cost, and price discrimination. 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. We estimate the model using a novel combination of secondary (from a large online platform in China) and primary data (obtained via online surveys) – the latter allows us to distinguish between a failed-bargain and a no-bargain transaction as well as to quantify the loss-of-face cost.

Our results provide rich insights on the bargaining primitives and process. First, we find the mean bargaining initiation cost to be 9 yuan, close to the minimum hourly wage in China. We find that a 1% increase in posted price on average leads to a 3.3% decrease in (browse to buy) conversion rate, with the decrease coming 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. Our data also suggest a high level of heterogeneity in the loss-of-face cost, with the majority of customers having a cost of less than a yuan but a substantial proportion having a cost as high as 100 yuan. We carry out three policy simulations. First, we highlight the importance of the bargaining initiation cost by illustrating what happens when the analyst is unable to distinguish between failed-bargain and no-bargain transactions. We show that not doing so leads to a large impact on the estimated bargaining initiation cost, the bargaining intention percentage, and the value of allowing for bargaining. Second, as the loss-of-face cost can be considered as a market friction, we examine the impact of removing this friction for the platform. We find that it has a big impact on the purchase conversion on the platform. Finally, we pin down the benefit of allowing for bargaining on the platform by comparing the outcomes to a simulated world in which bargaining is not allowed. The results show that banning bargaining is only modestly beneficial for the average seller, but is greatly beneficial for both the buyer and the platform (and

by implication for the social planner). This last policy simulation is likely to be useful for digital platforms in various settings as they explore pricing mechanisms, especially whether to allow for bargaining or not.

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. Third, while we are able to replicate our main results for another product category, we only have data from one platform (albeit a quasi-monopoly) in one country. Fourth, while our data do not support the fact that sellers extract and exploit additional buyer information via chat, it is possible that this may not be the case in other settings. This would then allow for a more in-depth exploration of bargaining strategies. Finally, in this paper, we implicitly include the seller's bargaining initiation 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 initiation cost from the product marginal cost, which would make the analysis of the benefit to the sellers 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, B. Larsen, and S. Tadelis (2020). Sequential bargaining in the field: Evidence from millions of online bargaining interactions. *The Quarterly Journal of Economics* 135(3), 1319–1361.
- Backus, M., T. Blake, and S. Tadelis (2019). On the empirical content of cheap-talk signaling: An

- application to bargaining. *Journal of Political Economy* 127(4), 1599–1628.
- 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.
- Bester, H. (1994). Price commitment in search markets. *Journal of Economic Behavior & Organization* 25(1), 109–120.
- Chintagunta, P. K. and J. Chu (2021). Geography as branding: Descriptive evidence from taobao. *Quantitative Marketing and Economics*. Forthcoming.
- Chu, J. and P. Manchanda (2016). Quantifying cross and direct network effects in online consumer-to-consumer platforms. *Marketing Science* 35(6), 870–893.
- Clark, D. (2016). *Alibaba: The House That Jack Ma Built*. Ecco.
- 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: Hagglng 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.
- Ellickson, P. B., P. Kong, and M. J. Lovett (2018). Private labels and retailer profitability: 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.
- 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* 55(6), 832–851.
- Keniston, D. E. (2011). Bargaining and welfare: A dynamic structural analysis of the autorickshaw market. *Working Paper*.
- Larsen, B. J. (2014). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *The Review of Economic Studies*.
- Loertscher, S. and L. Marx (2019). Bargaining, mergers, and investment in incomplete information io. *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.
- Miles, E. W. (2010). The role of face in the decision not to negotiate. *International Journal of Conflict Management*.
- 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.
- O’Connor, K. M. and J. A. Arnold (2001). Distributive spirals: Negotiation impasses and the moderating role of disputant self-efficacy. *Organizational behavior and human decision processes* 84(1), 148–176.
- 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.
- Sargan, J. D. (1958). The estimation of economic relationships using instrumental variables. *Econometrica: Journal of the Econometric Society*, 393–415.
- 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.
- Tuncel, E., A. Mislin, S. Kesebir, and R. L. Pinkley (2016). Agreement attraction and impasse aversion: Reasons for selecting a poor deal over no deal at all. *Psychological science* 27(3), 312–321.
- 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.
- Wooldridge, J. M. (2015). Control function methods in applied econometrics. *Journal of Human Resources* 50(2), 420–445.
- Zhong, Z. (2017). Chasing diamonds and crowns: Online reputation systems and seller response. *Working Paper*.

Online Appendix to: *“Meet Me Halfway”: The Costs and Benefits of Bargaining*

Xu Zhang
 London Business School, xzhang@london.edu
 Puneet Manchanda
 University of Michigan, pmanchan@umich.edu
 Junhong Chu
 National University of Singapore, bizcj@nus.edu.sg

O1 Taobao Figures

Figure O1: Taobao Feedback Page



O2 Summary Statistics by Gender

Table O1: 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

Figure O2: Taobao Reputation Rule

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

Table O2: Survey Summary Statistics by Gender

	Female		Male	
	(N=526 in Row 1,2; 376 in Row 3,4)		(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

O3 Taobao Consumer Survey

In order to measure the bargaining-related costs among Taobao consumers, we supplement the transaction data with primary information collected via surveys. For robustness and identification purposes, we conducted three rounds of surveys. In 2015, we conducted the first (pilot) survey with 566 respondents. In 2016, we conducted the second survey on bargaining initiation costs with 1,009 respondents. In January 2020, we conducted the third round of survey with 1,417 respondents and added additional questions (Q5 & Q6) about loss-of-face cost after a failed bargain. All surveys were carried out in Chinese – the English translation of the 2020 survey is provided below.

The information from the first and second surveys, combined with the transaction data, allowed us to estimate the bargaining initiation costs. The correlation between the estimated bargaining initiation costs using the main survey data and using the combined main/pilot survey data is 0.95, suggesting the robustness on the bargaining initiation cost estimates. Using the third survey, we derive information on the loss-of-face cost distribution. One concern about the third survey is that consumers in 2020 may not match the consumers in 2015 or 2016. To test to what extent consumers have evolved over time on dimensions related to bargaining, we compared the answers to

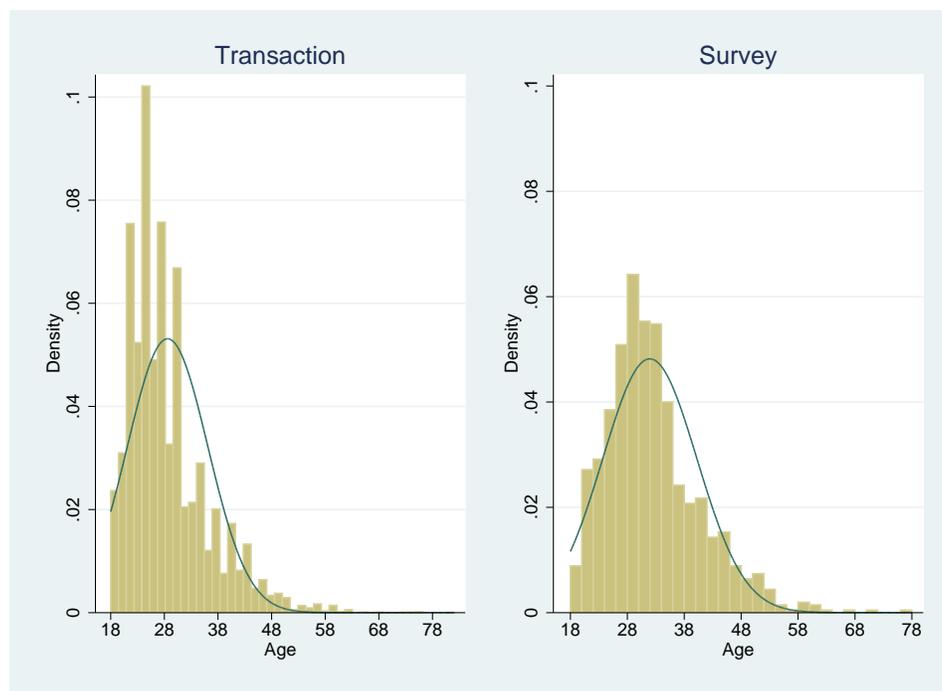
the bargaining intention and expectation questions, which were kept the same across all the surveys. We find that the overall numbers are similar (see Table O3 and Table 2 in the paper), though we do observe that the bargaining intention and the expected bargaining outcome become lower over time. This is expected and consistent with our finding that consumers are less likely to bargain when their opportunity costs become higher as their income increases over the years. Given the fact that the overall bargaining patterns between 2020’s consumers and the 2016’s consumers are similar, we believe consumers’ loss-of-face costs should also be similar. Figure O3 plots a comparison of the age distribution in the transaction data and in the survey. The similarity also verifies that the transaction sample and the survey are comparable.

Table O3: Summary Statistics: 2020 Taobao Consumer Survey

	N	mean	sd	min	max
I(Certainly Bargain)	1,417	0.524	0.499	0	1
I(May + Certainly Bargain)	1,417	0.689	0.462	0	1
Perceived Success Rate Bargaining	977	0.463	0.240	0.1	1
E[Discount Amount Success] (yuan)	977	152.4	143.6	1	1,000

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. The questions related to the above information in the 2020 Taobao Consumer Survey are the same as the ones in the 2016 survey.

Figure O3: A Comparison of Age Distribution between the Transaction Sample and the Survey



Taobao Consumer Survey (translated-version)

The purpose of this survey is to understand Chinese consumers' online shopping behavior. We really appreciate your input.

1. Are you aware that one can bargain on Taobao?
 - (a) Yes
 - (b) No

2. Would you bargain with a seller if you are going to buy a cellphone priced at 1,500 yuan?
 - (a) Yes
 - (b) No
 - (c) Maybe

3. How likely would you expect to succeed in bargaining?
 - (a) About 100%
 - (b) About 90%
 - (c) About 80%
 - (d) ...
 - (e) Less than 10%

4. How much discount would you expect to get if bargaining succeeds?
----- yuan

5. Assume you plan to buy a cellphone and you are interested in a cellphone model sold by Seller A priced at 1,510 yuan. Though this price is below your willingness to pay, you still started to bargain with Seller A. After quite a while of bargaining, the seller did not give in. How would you describe your feelings at that moment?

	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
I do not care.					
I feel embarrassed.					
I feel like a loser.					
I do not feel annoyed.					
I feel upset.					

6. Then, you see that Seller B is selling exactly the same cellphone priced at p_B yuan. Assume all the attributes (reputation, product description, service, logistics, etc.) of Seller A and Seller B are the same, from whom would you like to buy the cellphone, Seller A or Seller B? (Seller B's price p_B is randomly chosen from 1,511, 1,513, 1,516, 1,520, 1,530, 1,540, 1,550, 1,580, and 1,610 yuan.)

- (a) Seller A
 - (b) Seller B
7. What is your age?
----- years old
8. What is your gender?
- (a) Female
 - (b) Male
9. What is your education level?
- (a) Elementary School or below
 - (b) Middle School
 - (c) High School/Vocational School
 - (d) College and above
10. What is your monthly income?
- (a) 2,000 - 3,999 yuan
 - (b) 4,000 - 5,999 yuan
 - (c) 6,000 - 7,999 yuan
 - (d) ...
 - (e) Above 20,000 yuan

O4 Robustness Checks

The goal of this paper is to accurately measure the value of bargaining for sellers, buyers, the e-commerce platform, and the social planner in terms of social welfare. It is important that our results and the policy simulation outcomes are 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.

O4.1 Survey Response versus True Belief

In the baseline model, we assume that buyers' stated bargaining expectations in the survey are consistent with their formed bargaining expectations in reality. This subsection tests this assumption. In this 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 in reality, which means that they are overly optimistic in the survey, or fully naive, i.e., perceive the bargaining success rate as 100% in reality,

which means that they are overly pessimistic in the survey. The sophisticated assumption and the naive assumption provide a lower bound and an upper bound for the bargaining initiation costs, and further test whether our results comparing the fixed-price mechanism and the mixed-price mechanism are robust under the two extreme conditions.

For each assumption considered, we find that the changes in average transaction price and conversion rate are minimal compared with the baseline policy simulation in § 7.3. In contrast, changes in assumptions have a substantial effect on the saved bargaining initiation cost amount. Under the sophisticated and the naive assumptions, the saved bargaining initiation costs per day equal 1.9 and 6.7 million yuan, respectively. However, changes in bargaining assumptions yield qualitatively similar conclusions.

O4.2 Assumption with Respect to Non-purchasers

In order to distinguish between the three types of non-purchasers at the bottom of Figure 3, a 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 her/his 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 revenue and conversion are robust. This is not surprising given that the baseline change in total revenue and conversion rate is small - less than 1%. While

the total savings in bargaining initiation costs become smaller, they are still not materially different and remain economically significant.

O4.3 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 O4 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 O4) 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.

O4.4 Loss-of-Face Cost Distribution

In order to test the robustness of our results to the loss-of-face cost distribution obtained via the survey (see Figure 2), we vary the parameters of this distribution. We use the four following distributions: (a) original shape parameter and 2 times the original scale parameter, (b) original shape parameter and 0.5 times the original scale parameter, (c) 2 times the original shape parameter and 0.5 times the original scale parameter (mean remains the same), and (d) 0.5 times the original shape parameter and 2 times the original scale parameter (mean remains the same). Our results remain materially unchanged across these four different distributions.

Table O4: Robustness of Bargaining Realization Process

	Bargaining Success Indicator		log(Bargaining Amount)	
log(price)	0.335*** (0.035)	0.424*** (0.054)	0.776*** (0.090)	0.937*** (0.109)
I(Promotion)	-0.158*** (0.037)	-0.119** (0.045)	-0.514*** (0.081)	-0.593*** (0.075)
Seller Reputation Level	0.052*** (0.017)	-0.045 (0.050)	-0.243*** (0.035)	-0.247** (0.097)
Detailed Seller Rating	-0.387*** (0.120)	-0.017 (0.175)	0.634** (0.276)	-1.033*** (0.3257)
Store Age	-0.024* (0.011)	-0.213** (0.074)	0.067** (0.027)	0.572*** (0.141)
Buyer Shopping Experience	0.055*** (0.010)	0.052*** (0.011)	-0.069*** (0.024)	-0.043* (0.021)
Buyer log(income)	-0.164*** (0.048)	-0.262*** (0.055)	-0.033 (0.109)	-0.030 (0.099)
I(Repeat Purchase)	0.425*** (0.041)	0.356*** (0.045)	0.134 (0.085)	0.026 (0.073)
Product Age	0.016 (0.020)	0.022 (0.026)	0.122* (0.050)	0.091* (0.052)
# of Sellers w/ Same Product	-0.057 (0.041)	-0.059 (0.046)	0.039 (0.097)	-0.137 (0.087)
# of Sellers w/ Same Reputation	0.015 (0.016)	0.057** (0.021)	-0.090** (0.035)	-0.061 (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.

O5 Generalizability

Our analysis uses data from the cellphone category. However, given that the goal of the paper is to quantify the costs and benefits 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 category. Following the first two steps in §4, we find the bargaining initiation cost is 8.6 yuan on average in this category^{O1}, which is very close to our estimate 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 policy simulation where bargaining is not allowed. Specifically, we expect that major benefits would come from the saved bargaining initiation costs and loss-of-face costs. As women's shoes are smaller

^{O1}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.

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.

O6 External Validity with Real Chat 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 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 cellphones especially as the mean unit price in the cellphone 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.

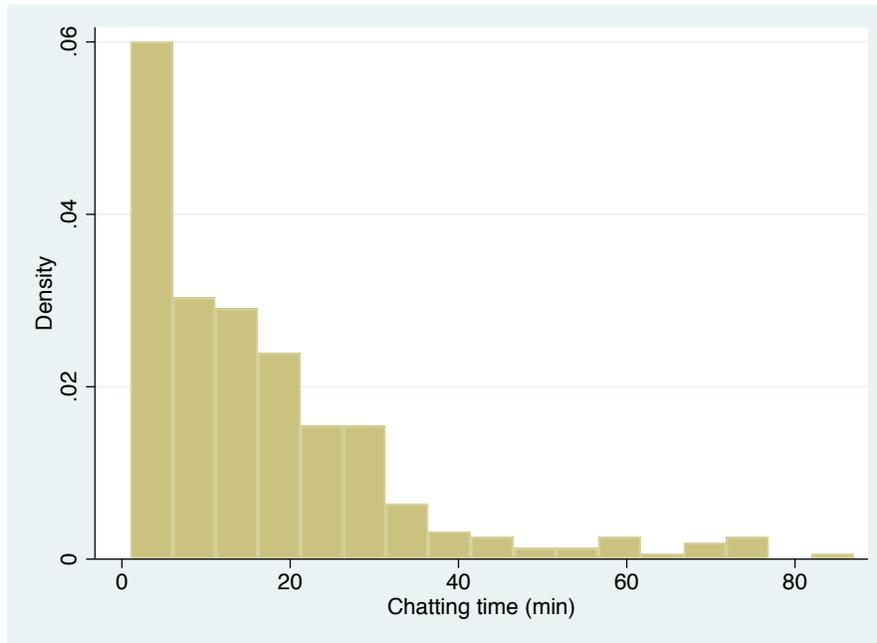
First, we use the chatting data to examine whether the bargaining initiation cost that we recover from our structural model is reasonable. We begin by looking at the time buyers spend chatting with the seller. The mean time spent on chatting is 17 minutes (Figure O4 shows the distribution). Unlike the cellphone 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 O5. 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 initiation 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 O5: 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

Figure O4: Histogram of Chat Duration



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 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 3, 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 cellphone transaction data (42% vs. 20%) suggests that it is important to run robustness checks on the above

assumption (we do this in Online Appendix § O4.1). 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 9 and 15) 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.

In sum, we obtain four key findings from the chatting data analysis. First, 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 initiation costs estimated in the model at 9 yuan. Second, 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, 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. Overall, analysis of these data provides us correlationally consistent evidence on bargaining initiation cost, bargaining intention, bargaining success as well as support for our model structure and specification.

References

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