The Pennsylvania State University
The Graduate School
College of the Liberal Arts

INDUSTRIAL ORGANIZATION OF ONLINE MARKETPLACES

A Dissertation in
Economics
by
Xiaolu Zhou

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

August 2017

The dissertation of Xiaolu Zhou was reviewed and approved* by the following:

Mark Roberts
Professor of Economics
Dissertation Co-Advisor
Co-Chair of Committee

Peter Newberry
Assistant Professor of Economics
Dissertation Co-Advisor
Co-Chair of Committee

Paul L.E. Grieco
Associate Professor of Economics

Charles Murry
Assistant Professor of Economics

Christopher Parker
Assistant Professor of Supply Chain Management

Barry W. Ickes
Chair of the Graduate Program

*Signatures are on file in the Graduate School.

Abstract

The first chapter quantifies the efficiency gains of the rebate contract employed by a dominant e-commerce platform in China. On the platform, sellers receive a rebate of their listing fee if they reach predetermined revenue thresholds. The rebate give sellers incentives to reduce price and produce revenue above the threshold. To quantify the effect of rebates, I construct and estimate a structural model of demand and sellers' dynamic pricing decisions taking into account an unobserved cost heterogeneity. Results show that rebates reduce the average price by $4 (1.7%), leading to 9% gains in total industry revenue. Consumers also benefit from the contract as the lower prices lead to a 9.5% gains in consumer surplus.

The second chapter evaluates the heterogeneous effect of online reputation for sellers on a large Chinese platform that differ in their offline presence. We estimate a demand model that incorporates a structural learning process and allow for the process to vary across sellers who are differentiated by their offline presence. The estimates suggest that the impact of the rating score is larger for local sellers. Using the estimates, we find that the median seller values the reputation system by $4,000 more per month as a local seller than they would have as a national seller, which is 50% of their monthly revenue.

The third chapter explains a methodology for solving the continuous choice dynamic problems. This method reduces the computational burden, and improve the accuracy relative to discretizing the action space. I show that the distribution of optimal actions can be analytically calculated from the first order condition through a change of variables. Therefore, the value function evaluation is able to be approximated by the average of action-specific value functions using the probabilities of optimal actions as weights. Based on the new method, I also formulate a simulated method of moment for the estimation. I provide two examples applicable to the solution method, and use a numerical example to illustrate the computational advantage.

Table of Contents

List of Figures

List of Tables

Acknowledgments

Firstly, I would like to express my sincere gratitude to my advisor Prof. Peter Newberry for the continuous support of my Ph.D study and related research, for his patience, motivation, and immense knowledge. His guidance helped me in all the time of research and writing of this thesis. I could not have imagined having a better advisor and mentor for my Ph.D study.

Besides my advisor, I would like to thank the rest of my thesis committee: Prof. Mark Roberts, Prof. Paul Grieco, Prof. Charles Murry, and Prof. Christ Parker, for their insightful comments and encouragement, but also for the hard question which incentivized me to widen my research from various perspectives.

My sincere thanks also goes to professors in following instutions, the department of economics at the Pennsylvania State University, the department of economics at the University of Finance and Economics, and the Guanghua school of Management at Peking University. Without their teaching, it would be impossible to conduct this research. In particular, I am grateful to Prof. Lixin Ye, Prof. Wanchuan Lin and Prof. Xiaolan Zhou for writing recommendation letters for my application for the graduate program.

I thank the graduate students in the department of economics for the stimulating discussions, and for all the fun we have had in the last four years. In particular, I am grateful to Bo Bian, Jianpeng Deng, Wang Xiang, Chaohai She, Rong Luo, Kai Li, Yingyan Zhao, and Lixiong Li.

Last but not the least, I would like to thank my family: my parents and to my husband, Yifeng Qin, for supporting me spiritually throughout writing this thesis and my life in general.

Dedication

This thesis work is dedicated to my husband, Yifeng Qin, who has been a constant source of support and encouragement during the challenges of graduate school and life. I am truly thankful for having you in my life. This work is also dedicated to my parents and grandparents who have always loved me unconditionally and whose good examples have taught me to work hard for the things that I aspire to achieve.

Chapter 1

The Efficiency and Profitability of Rebate Contracts: A Study of Alibaba's Tmall

1.1 Introduction

Vertical relationships often feature contracts that are designed to align incentives between manufacturers and retailers. A rebate, one form of a vertical contract, is a discontinuous reward to retailers when they exceed a predetermined threshold.[1] Rebates to retailers are used in many industries, such as, computer processors (Intel), tires (Michelin) and aircraft industry (British Airways). Recent literature argues that rebates promote retailers' effort, leading to efficient gains in industry profits.[2] However, little empirical work quantifies how effective rebates are in achieving this. In this paper, I study a rebate-based contract between a dominant e-commerce marketplace and its independent sellers. Specifically, I examine the effect of rebates on sellers' dynamic pricing decisions and, in turn, industry revenue and profits.

The marketplace I focused on is Tmall, the largest business-to-consumer (B2C) platform in China. It is a subsidiary of Alibaba, the worlds' largest retailer by sales volume, and it accounts for 54% of China's B2C market with gross sales of $134 billion in 2015.[3] In order to list their products for sale on the platform, sellers are charged a fixed upfront payment along with a royalty rate on each dollar of revenue. Additionally, sellers have the opportunity to receive a rebate on their fixed payment if their year-end cumulative revenue exceeds a predetermined threshold. I collected data directly from the Tmall platform by scraping price and quantity information for all computer tablets listed between January and December of 2015. The data provided me an unique opportunity to

[1] There are also rebate contracts with a threshold of intermediate goods purchases (see Conlon and Mortimer (2014)).

[2] Taylor (2002), Gérard P. Cachon (2005).

[3] In 2016, Alibaba overtook Wal-Mart as the world's largest retailer by sales volume of merchandise. Information about Tmall's market share and gross sales volume is from Tmall's official website.

quantify the effectiveness of rebates in a large real-world marketplace.

In the reduced form analysis, I find evidence that sellers set prices taking into account the likelihood of reaching the rebate threshold. I regress sellers' prices on the distance of their revenue from the rebate threshold, along with other covariates. The results show that a seller's price decreases as her revenue approaches the threshold. At the end of the year, this effect becomes stronger as sellers attempt to reach the threshold before the year-end deadline. Together, the result means that the rebate contract has effects on sellers' prices over the whole year.

To examine the extent to which the rebate contract encourages lower prices, I formulate a structural model of supply and demand for tablets on Tmall. Sellers dynamically set prices while facing a trade-off between accumulating revenue in anticipation of the rebate and earning profit in the current period. Sellers face uncertainty about future revenue due to, for example, new product characteristics, new consumer tastes, adjustments in retail and wholesale costs. I model the uncertain demand and cost as serially correlated autoregressive processes. Also, sellers' response to the contract can intensify price competition, which may amplify the effect of the rebate. To control the effect of competition on sellers' pricing decisions, I assume that sellers respond to an aggregated competition intensity over all competitors.

I estimate the model in two steps. First, I estimate the demand using standard methods from the industrial organization literature on differentiated products. With the demand estimates, I back out the realizations of market competition intensity. I then estimate the parameters of the serially correlated process in demand and the market competition intensity using the realizations in the data.

Next, I solve for cost shocks from sellers' first order conditions given cost parameters, and I estimate the parameters of the cost function and the cost transitions by matching them with the price data. Solving the cost shocks requires solving ex-ante value function, which is computationally challenging for a continuous choice dynamic problem. I propose a method to resolve the challenge. Specifically, I make use of the inverse mapping from the cost shock to the optimal price. I derive the density of the optimal price from the distribution of the cost shock. Then I integrate the value function over the distribution of optimal prices rather than the distribution of cost shocks. [4]

The estimation results show that rebates induce a 1.06–2.33% reduction in prices within 20% distance to the rebate threshold, with larger price decreases as the year-end contract deadline approaches. The results also indicate that sellers' demand and cost are persistent, with a correlation of 0.95 in demand and a correlation of 0.46 in cost.

To quantify the revenue and profit gains from rebates, I compute a counterfactual in which I eliminate the rebate. Without the rebate, the contract is the same as pure revenue sharing. Under the revenue sharing, sellers' pricing becomes static as they are not induced to be forward-looking. The results indicate that rebates incentivize 84% of sellers to reduce price. Within the whole year, rebates reduce the average price by 1.7% and decrease the standard deviation of prices by 6%. This leads to a 2.4% increase in the number of tablets sold and a 9% increase in total revenue. Additionally, rebates increase industry profits by 12% because rebates better align the platform and sellers' incentives.

[4] Aviv Nevo and Williams (2016) use a similar way to calculate the transition probabilities.

Also, with rebates, consumer surplus increases by 9.5% because of lower prices. The result is robust for various levels of royalty rates. Together, results provide evidence that the rebate effectively reduce price and produce revenue, leading to welfare gains for the industry and consumers.

My study is related to all units discounts/rebate contracts literature in vertical relationships. There are various theoretical explanations for using rebates. Shaffer (1991) and Sreya Kolay and Ordover (2004) argue that a rebate is a second-degree price discrimination, which increase the manufacturer's profits. OB́rien (2014) shows that the rebates help to solve the double moral hazard problem.[5] Taylor (2002) finds that rebates are superior to revenue sharing in terms of promoting downstream efforts. Many other papers point out that rebates reinforce the vertical relationship and thus exclude upstream competitors.[6] Conlon and Mortimer (2014) is the only empirical paper of all units discounts. They provide evidence for both efficiency and anti-competition.

My paper contributes to the above literature in the following ways. First, the accumulated revenue threshold in my study directly induces downstream retailers' dynamic pricing decisions, which is different from the static setting in most of the previous literature. Second, I study the effect of rebates in an environment with a vast of downstream heterogeneity, meaning that the contract is not a perfect one. Then the conclusion that rebates are effective in the single retailer environment studied before may not apply in this environment.[7] Further, my paper adds to platform's pricing design. While previous studies of multi-sided markets have analyzed platform fee policy, they only discuss the effect of the fixed fee or two-part tariff on users' entry decisions.[8] In reality, platforms choose various sophisticated contracts which have not been examined in the existing literature. [9] These fee structures have different implications for seller incentives other than the participation. To the best of my knowledge, this paper is the first empirical study explaining why a platform uses one structure of contract.

This paper is structured as follows. Section 1.2 introduces the institutional background of the market studied and provides a description of the data. Section 1.3 displays the reduced form evidence of sellers' dynamic pricing behavior. Section 3.2 presents the model, sections 3.5 describes the estimation procedure and computational method; section 1.6 and 1.7 displays estimation and counterfactual results; and section 1.8 concludes.

[5]Shaffer (1991) considers downstream firms' private information on the opportunity cost of shelf space. And Sreya Kolay and Ordover (2004) considers the private information on demand. Double moral hazard means that when both upstream and downstream parties have uncontractible actions, leading to the underprovision of the uncontractible actions in both upstream and downstream.

[6] Chao and Tan (2013), Asker and Bar-Isaac (2014) are literatures discussing the anticompetitive effect. Chao and Tan (2013) provide the motive from capacity constraint, and Asker and Bar-Isaac (2014) argue the anticompetitive effect in a situation when upstream entrants needs downstream accommodation. In contrast, Entrigue Ide and Figueroa (2015) cast doubt on the exclusivity effect because all-units discounts induce the exclusivity ex-post once buyers' purchasing decisions are made but not ex-ante when contracts are signed.

[7]Conlon and Mortimer (2014) study the environment with single downstream firm.

[8]Weyl (2010), Rochet and Tirole (2003a), Rochet and Tirole (2003b), Armstrong (2006), Bedre-Defolie and Calvano (2013).

[9]For example, eBay charges a royalty rate plus a fixed fee on each listing, and uses a three-part tariff on the fixed listing fee. Some price comparison sites charge consumers for clicks (**?**).

1.2 Data Description

The reminder of this paper will investigate a rebate-based contract between the online platform, Alibaba's Tmall, and its independent sellers. the goal of this empirical exercise will be to establish whether and to what extent rebates encourage lower price.

Alibaba is a Chinese Internet company that provides consumer-to-consumer, business-to-consumer and business-to-business sales services. In September 2014, Alibaba went public as the largest IPO in history with $25 billion, and its expansion still continues. In fiscal year 2016, Alibaba achieved $455.6 billion gross transaction volume (GMV) with an increase of 39% from the previous year.[10] As a comparison, in 2015, Ebay announced $78.1 billion GMV, and Amazon Marketplace reached over $100 billion in sales.[11] Tmall is the open B2C marketplace of the Alibaba group, which does not sell products by itself. Similar to Amazon Marketplace, it is launched for professional merchants who are either brand owners or authorized distributors. Sellers include some e-commerce platforms (e.g., Amazon and Dangdang) and well-known brands (e.g., Costco and Adidas). Many different products are sold on Tmall, including clothing, shoes, cosmetics, electronics, home appliance, cars and so on.[12]

Tmall's contract is a combination of a royalty rate and an upfront payment, with the opportunity to earn some or all of the upfront payment back in the form of a rebate. The contract requirements vary by the types of products a seller offers. In this paper, I will only focus on one product category, tablets.[13] The fixed upfront payment under tablets is $5,000, and the royalty rate is 2%. The contract terms do not differ with the timing of sellers' entry or the number of listings. Sellers are returned 50% (100%) of the upfront payment at the end of the year if their accumulative revenue exceeds $60,000 ($200,000) by the end of the year and their average rating score is higher than 4.6. The average rating score is the mean of consumers-provided reviews for all of the seller's listings. Similar to Amazon Marketplace, the rating score is scaled from 0 to 5. Sellers are able to sell several products under one primary category, for example, "electronics", which include computers, cameras and tablets. Sellers whose products cover several product categories have to pay the highest fixed fee across these categories. However, the rebate's revenue threshold is based on the product category with the highest sales volume.[14]

The data used in this study include all the listings of tablets on Tmall from January 2015 to December 2015. To collect the data, I scraped the Tmall website every two weeks from September 2014 to May 2015 and every week from June 2015 until February 2016. Importantly, I observed each listing's price and quantity in the last 30 days at the scraped

[10] Fiscal year 2016 starts from April 2015 and ends in March 2016.

[11] Data are from the companies' financial reports.

[12] An example of categories and subcategories is listed in table 1.1 in the Appendix 1.9.

[13] A detailed fee schedule can be found at `http://about.tmall.com/tmall/fee_schedule`. I list a few categories' contract terms in the appendix.

[14] As an example, hats/scarves and women's clothing are two subcategories with the upfront payment as $10k (RMB 60k) and $5k (RMB 30k), respectively, and with the revenue threshold for 50% rebates as $30k (RMB 180k) and $60k (RMB 360k), respectively. If one seller wants to sell both hats and women's jeans, she has to pay $10k (RMB 60k) of women's clothing instead of $5k (RMB 30k) of hats/scarves. If the annual revenue of hats is higher than the revenue of women's jeans, she has to sell more than $30k (RMB 180k) hats in order to have $5k refunded.

date. I used the per-30-day quantity data and created an average monthly price from the per-week scraped prices. [15] Each listing includes product characteristics, such as screen size (inches), hard drive memory (GB), operating system, Internet access and the number of years the tablet has been on the market. The data also contains seller characteristics, including average rating score, the number of ratings, the number of characters and pictures in the product description, a dummy of the variable indicating whether there was a video description, and the percentage of positive consumer comments.

In total, there are 34,864 tablets listed from 686 sellers over the 12 month sample period. To construct my sample, I drop the 14,705 observations with zero sales and the 38 observations with prices lower than \$15.8 (RMB100). Additionally, I drop the 46 sellers who only appear on "singles' day", China's large online shopping festival that occurs on November 11th. As I only have information about the sales of tablets, I want to restrict my sample to sellers who are likely to react to the rebate threshold for this category. Therefore, I restrict the sample to "electronics retailers" and regard all other retailers who sell tablets as the "outside option".[16] However, electronics retailers may also sell computers, laptop, cameras, MP3s/MP4s and electronics accessories. Because I do not have any information about which product category has the highest revenue for the retailers, I assume that all electronic retailers react to the threshold for tablets.

Table 1.2 shows the product characteristics of the electronics sellers on Tmall. The average tablet has an 8.57 inch screen, 24.59 GB of memory and has been on the market for 10 months. Additionally, 24.59% of the tablets have Internet access through a cellular network. The average price is \$234, and the standard deviation is \$222. Sales are concentrated in a few brands, such as Apple, Samsung and some local Chinese brands. The most popular four brands account for over 50% market share, among which, Apple is ranked as 3rd with 13.73% market share, Samsung is ranked as 9th with 2.94% market share and one local Chinese brand, Huawei, is ranked as 5th with 6.63% market share.

Most sellers on Tmall are multi-products sellers, with the average number of different tablets listed per month being six and the percentage of sellers selling more than one product being 71%. However, the revenue threshold is the summation of all the revenue from the sale of tablets, meaning that multi-product sellers have the same incentive from the rebate contract as single product sellers do. Thus, for simplification, I aggregate the seller-tablet-month level data into seller-month level data. To do this, I create a seller-month level product and price by calculating the weighted average of tablet characteristics, where the weights are based on the proportion of sales that month from each tablet. I create a seller-month level quantity by adding up the monthly sales across all tablets for an individual seller. Overall, there are 3,049 seller-month level observations from 470 sellers in 12 months. Table 1.3 displays statistics of the data after this aggregation. As the table shows, there are large variations in the tablet and seller characteristics. This means that it is a heterogeneous product market, where sellers are differentiated in their tablets characteristics and descriptions of their tablets.

The average seller sells 298 tablets with \$59k revenue in one month, but this varies

[15] The average and median number of days between price changes is 23 and 9 days respectively, based on the information of daily listing price from a price tracking website.

[16] Some examples of those sellers' categories in data are cellphones, cars and video games. A total of 3,421 observations are from those sellers.

substantially across sellers. The average seller is above the second rebate revenue threshold, but the median seller is below the first one. Specifically, 12.55% of sellers are above the first revenue threshold but below the second, and 19.79% of sellers above the second revenue threshold. This suggests that sellers are affected by the revenue threshold. However, the average rating score, 4.78, is higher than the rebate's requirement of 4.6, and 99.33% of sellers are above this mark in each month. This suggests that the rebate rating score threshold is not binding.

1.3 Evidence of Dynamic Pricing

Since the rebate contract has discontinuous rewards at the revenue thresholds, sellers may want to accumulate revenue by dynamically setting prices considering the revenue left to the threshold. Specifically, they may decrease price to cross the revenue threshold when their cumulative revenue approaches it. In this section, I test for this dynamic pricing behavior.[17] To do so, I estimate the relationship between sellers' price and their revenue gap to the two thresholds using the following regression specification.

$$\log(p_{st}) = \theta_0 + \theta_t D_{it} + x_{st}\theta_x + \chi_s + \chi_t + \epsilon_{st}, \tag{1.1}$$

where $D_{it} = \bar{r}_i - CR_{st-1}$ is the revenue left to threshold $\bar{r}_i, i = 1, 2$, and $CR = \sum_{\tau=1}^{t-1} R_{s\tau}$ is the cumulative revenue from January to month $t-1$. I run this regression on two different sets of data. In the first, I keep only the sellers who are below the first threshold, and in the second I keep only the sellers who are above the first and below the second. In both regressions, I include time-varying seller/tablet characteristics, x_{st}, a seller level fixed effect, χ_s, and time dummies, χ_t.

The vector of seller and tablet characteristics includes screen size, hard drive memory, Internet access, operating system, product age (in years), 15 brand dummies, the number of products, rating score, the number of ratings, the number of pictures and characters for product description, a dummy of video description and the percentage of positive comments generated by previous consumers, and χ_s and χ_t are seller-level and time-level fixed effect, respectively.[18]

The results of eight different specifications are presented in table 1.4, with (1)-(4) including only sellers who are below the first threshold and (5)-(8) only including sellers who are between the first and the second. When I do not control seasonality, the effects of the threshold are significant and negative in columns (1) and (5), meaning that prices decrease as the cumulative revenue approaches either threshold. In columns (2) and (6), I control for seasonality, and the coefficients become insignificant. This is because the effect of threshold may differ over time. The thresholds possibly have greater effects on sellers' price decisions when the rebate deadline closes. In order to test the heterogeneous effect over time, I look into the coefficient of revenue distance in each period. Specifically, I include interactions between D_{it} and month dummies in columns (3) and (7) and include

[17] As table 1.3 shows, a rating score of 4.6 is not an effective criteria for sellers but the two revenue thresholds are. Therefore, from this section, I only consider the effect of the revenue thresholds of the rebate contract.

[18] Since I use the weighted averaged product, the brand dummies are the brand-level shares among seller s's products at time t, for $\forall s, t$, which range from 0 to 1.

interactions with quarterly dummies in columns (4) and (8). As every seller starts with \$0 revenue in January, I omit the January dummy in the monthly regressions.

The result shows a significant price reduction from August to November when the first threshold is approaching and, the same price pattern in October and November when the second is approaching. This demonstrates that sellers have more incentives to lower price at later months when the deadline closes. The coefficients for interactions with quarterly dummies also verify this effect.

It may be the case that the pricing reductions are not linearly related to the cumulative revenue. For example, an increase of \$100 in revenue surely has a different effect for a seller who is very far from the threshold compared to one who is very close. Because of this, I run additional price regressions, but instead of a continuous measure of the distance to the thresholds, I use four groups indicators, where the groups are defined based on their distance to the threshold. The groups are given by $[0, 0.5\bar{r}_1), [0.5\bar{r}_1, 0.5\bar{r}_2), [0.5\bar{r}_2, \bar{r}_2)$ and $[\bar{r}_2, \infty)$, where the the first group includes sellers who are far from both thresholds, the second group includes sellers who are around the first threshold, the third group includes sellers who are above the first threshold and close to the second threshold, and the fourth group includes sellers who have already passed the second threshold. The regression specification is as follows:

$$\log(p_{st}) = \theta_0 + \theta_{ct} 1\{\underline{c} \leq CR_{st-1} < \bar{c}\} \times 1\{Month\ t\} + x_{st}\theta_x + \chi_s + \chi_t + \epsilon_{st}, \quad (1.2)$$

where $\underline{c}$ and $\bar{c}$ are lower bound and upper bound of the revenue groups, and x_{st}, χ_s and χ_t are the same covariates defined previously. I use the fourth group in each month as a control group because they have passed the second threshold and thus are not incentivized by the rebate.

Since group 2 and 3 sellers are closer to thresholds than group 1 sellers, group 2 and 3 sellers should set lower prices in the same month. Also, the difference should increases over time. The results presented in table 1.5 demonstrate that the group indicators are significant and negative for group 2 and 3 sellers, but not significant for group 1 sellers. The prices of the second and third groups are at least 4% lower than the first group, and this difference intensifies as the deadline approaches, specifically, 7% and 10% in October and November, respectively.

1.4 Model

In this section, I specify a model of demand and supply on Tmall. The primary goal of the model is to capture the downstream sellers' pricing behavior and its relationship to the rebate contract. In the model, I assume that consumers make a one-time decision to purchase a tablet and that sellers dynamically set their prices for time periods $t = 1, \ldots, T$. The platform's contract term is given by $(fc, \gamma, \bar{r}_1, \bar{r}_2, \eta_1, \eta_2)$, where fc is the upfront fixed fee, γ is the royalty rate, $\bar{r}_1$ and $\bar{r}_2$ are two revenue thresholds for rebates, and η_1 and η_2 are the rebate rate on the fixed fee for these two thresholds.

Before go to the model, I provide a summary of the timing of the decision-making process for buyers and sellers. Specifically, I assume that consumers and sellers behave according to the following timeline:

- Within each period $t = 1, \ldots, T$,

    1 N_t potential consumers arrive at the platform, where $\mathcal{S}_t \in \mathcal{S}$ sellers exogenously list their tablets on the platform.[19]

    2 Seller $s \in \mathcal{S}_t$ observes period t's demand and cost forms a belief on competitors' behaviors. With the information, each seller sets a price considering expectations on the potential of receiving a rebate at the end of the contract period.

    3 Consumers observe all sellers' tablet characteristics and choose to purchase a tablet from one of the $\mathcal{S}_t$ sellers or take the outside option.

    4 Sellers s's revenue is realized and she pays γ percent of it to the platform.

- At period $T+1$, seller $s \in \mathcal{S}$ receives the rebate $\eta_1 \times fc$ ($\eta_2 \times fc$) if her cumulative revenue exceeds $\bar{r}_1$ ($\bar{r}_2$).

1.4.1 Demand

In period t, consumers observe $\mathcal{S}_t$ sellers and tablet characteristics including the retail price on the platform. Consumer i's utility of purchasing a tablet from seller s is

$$u_{ist} = \underbrace{x_{st}\alpha_x + \alpha_p \log(p_{st}) + \xi_s + \xi_t + \Delta\xi_{st}}_{\delta_{st}} + \epsilon_{ist}, \tag{1.3}$$

where x_{st} is a vector of seller and tablet characteristics. Because I aggregate the data to the seller level, each seller sells her weighted average tablet, meaning choosing a seller is equivalent to choosing a tablet. The seller's characteristics include the average rating score, the number of ratings, the number of products, the number of pictures and characters for products' description, a dummy variable indicating whether there is a video description, the percentage of positive comments generated by previous purchasers, and 20 brand dummies. The characteristics of the sellers' tablets include the weighted average screen size (inches), hard drive memory (GB), two dummies of Android and Windows operating system, and a dummy of wireless Internet access. The price is given by p_{st}, and ξ_s, ξ_t and $\Delta\xi_{st}$ are seller level, time level and seller-time-level demand shocks, respectively. In the empirical work, ξ_s and ξ_t will be captured by seller-level fixed effect and time dummies, respectively.[20] The idiosyncratic demand shock, ϵ_{ist}, is assumed by the *i.i.d.* extreme value and the mean utility of the outside option is normalized to 0: $u_{i0t} = \epsilon_{i0t}$.

Consumer i chooses the seller which maximizes her utility, meaning the probablity that she chooses seller s is given by:

$$Pr_{ist} \quad = \quad Prob(u_{ist} \geq u_{is't}, \quad \forall s \neq s') \tag{1.4}$$

[19] Since I aggregate data into seller-time level, $\mathcal{S}_t$ is also the product set at period t.

[20] The log assumption of price in utility mimics the Cobb-Douglas utility specification in Berry, Levinsohn and Pakes (1995). What is different is that they assume consumers evaluate price heterogeneously depending on income. However, I assume constant income across consumers, and all consumers' heterogeneity is attributed to idiosyncratic shocks.

Since consumers are homogeneous in their preferences, seller s's market share at period t is

$$sh_{st}(\delta_{st}) = \frac{\exp(\delta_{st})}{1 + \exp(\delta_{st}) + \sum_{s' \neq s} \exp(\delta_{s't})} \tag{1.5}$$

It is important to point out that I assume that consumers are static, meaning they do not anticipate future price changes due to the rebate-based contract. To test for this assumption, I calculate the correlation between the current period's sales with previous periods' price. A high correlation suggests that there exists strategic consumers who wait for future periods' price markdown. Table 1.6 shows the regression of each month's sales on the last three months' price, where all price lag terms' coefficients are below 0.1 and insignificant. Thus, there is no significant evidence for strategic consumers.

1.4.2 Sellers' Information Structure

As the results of the toy model suggest, the expectations on future periods' demand and cost have impacts on sellers' pricing decisions. Therefore, in what follows, I describe how sellers form belief on future periods' demand and cost when making their pricing decisions.

In terms of sellers' demand, there are four variables, $\{\xi_s, \xi_t, x_{st}, \Delta\xi_{st}\}$, where ξ_s is a seller-level demand shock, ξ_t is a time-invariant seller quality, x_{st} is a vector of seller-time level observables and $\Delta\xi_{st}$ is a seller-time level unobservable. Each seller knows ξ_s and $\{\xi_t\}_{t=1}^{T}$ when she first arrives to the market. I assume that sellers care about an utility measure of all characteristics, $X_{st} = x_{st}\alpha_x$, implying that different sets of characteristics may lead to the same utility level. I use the utility measure instead of tracking each characteristic because the dynamic problem is simplified with fewer state variables. I also assume that the seller has only a limited ability to predict future X_{st}. That is, she perceives that the mean utility measure follows a AR(1) process:

$$X_{st} = \theta X_{st-1} + \upsilon_{st} \tag{1.6}$$

$$\upsilon_{st} \sim N(0, \sigma_\upsilon^2), \tag{1.7}$$

where θ is the serial correlation of the tablet and seller observables. I assume that the tablet characteristics are correlated but vary over time because consumers' tastes are consistent over time, but the product set changes over time. For example, new versions of tablets have smaller screen sizes and larger memories, but consumers taste for screen size and memory likely do not change. The process means that sellers cannot perfectly predict a new product's demand and consumer-generated feedback but ensures that the seller is correct on average.

Also, the unobservable $\Delta\xi_{st}$ is independent.

$$\Delta\xi_{st} \sim N(0, \sigma_\xi^2), \tag{1.8}$$

where σ_ξ^2 is a parameter to be estimated. The demand shock includes seller and tablet unpredictable factors which are not included in observables, such as randomness affecting the sellers' service (e.g., shipping delays).

Following Steven T. Berry and Pakes (1995), I assume that the seller has a constant marginal cost that is a log linear function of observed cost characteristics and an unobserved component. I assume that the unobserved component consists of a serially correlated cost process and an independent cost shock.

$$c_{st} = \exp(w_{st}\beta + \zeta_{st} + \omega_{st}), \tag{1.9}$$

where w_{st} is a vector of cost shifters, β is a vector of parameters on cost shifters, ζ_{st} is a serially correlated cost component, and ω_{st} is an *i.i.d.* cost shock. I use the seller-level time-invariant quality ξ_s as the cost shifter. High quality sellers provide better consumer service and faster shipping, which increases their marginal cost for each transaction. A few examples of the correlated cost component are wholesale prices, advertising and shipping cost.

Similar to the demand side, the serially correlated component follows an AR(1) process with correlation ρ and variance σ_{ς}^2. Then, the two cost unobserved components evolve as follows,

$$\zeta_{st} = \rho\zeta_{st-1} + \varsigma_{st} \tag{1.10}$$

$$\omega_{st} \sim N(0, \sigma_{\omega}^2), \tag{1.11}$$

where $\varsigma_{st} \sim N(0, \sigma_{\varsigma}^2)$. When the seller first arrives at the platform, $\zeta_{st_1^s}$ is drawn from an initial distribution, $\zeta_{st_1^s} \sim N(0, \sigma_{\zeta}^2)$, where t_1^s denotes the first period that seller s's tablets appear on the platform. The serially correlated cost process affects all the future period's cost, while the independent cost shock only affects the current period's cost, but does not affect cost in the future periods. Then the parameters ρ, σ_{ς}^2 in the dynamic cost process affects the persistence of price over time, while, σ_{ω}^2 only affects the dispersion of price within one period.

Besides the demand and cost in the future, I also consider the effect of market competition on sellers' price decisions. It is important to capture this effect, as marginal sellers' price reduction intensifies price competition, which may amplify the effect of the rebate. To model this effect, I use the inclusive value of choosing a tablet from seller s's competitors to capture competitors' state and actions:

$$\Gamma_{st} = \log\left\{\sum_{s' \neq s} \exp(\delta_{s't})\right\}, \tag{1.12}$$

where $\delta_{s't}$ is the mean utility of seller s' and her tablet characteristics including price. This assumption means that sellers respond to an aggregation of all competitors instead of responding to each individual competitor, which is reasonable because each seller is small compared to the entire market. With this assumption, the heterogeneous oligopoly game becomes a single-agent problem as in **?** and Gautam Gowrisankaran (2012).

Following Barwick and Pathak (2015), I assume that sellers cannot predict this period's competition intensity, Γ_{st}, but only observe last period's realization. They form a belief of this period's state based on its transition. Similar to the above two papers, I assume sellers' perception about the competition intensity transition $\widetilde{\mathcal{P}}_t(\Gamma_{st+1}|\Gamma_{st})$ is the actual empirical density fitted to a linear autoregressive specification:

$$\Gamma_{st} = \lambda_0\Gamma_{st-1} + \lambda_t 1\{Period\ t\} + \nu_{st} \tag{1.13}$$

where $\nu_{st} \sim N(0, \sigma_\nu^2)$, λ_0, λ_t and σ_ν^2 are parameters to be estimated. This assumption can be thought of as sellers having incomplete information on competitors' shocks. Suppose each seller knows her own demand/cost, but does not know her competitor's, then the seller cannot predict competitors' revenue levels, price decisions and thus the market competition intensity. The equilibrium concept under this assumption is the oblivious equilibrium from Gabriel Weintraub and Roy (2008), where players respond to an industry state which evolves deterministically. Weintraub and Ifrach (2016) also formalize the Markov moment equilibrium of such a model capturing competition by an aggregated statistic.

The inclusive value affects sellers' pricing decisions through the realization of their market share and revenue. When many sellers want to accumulate revenue above the threshold, the decline in their prices will intensify the market competition and thus make sellers earn lower revenue at the same price level. This exacerbates the all sellers' incentive to reduce price.

1.4.3 State Variables

Given the above, the dynamic price decisions are based on five state variables $(CR_{st-1}, \Gamma_{st-1}, \xi_s, \widetilde{\delta}_{st}, \zeta_{st})$. The state which captures the incentive from the rebate threshold is the cumulative revenue, CR_{st-1}, which evolves as revenue is accumulated. The aggregate state Γ_{st-1} captures the market competition intensity, which evolves as equation 1.13. The time-invariant seller quality ξ_s captures sellers' heterogeneity.

The time varying component of demand, $\widetilde{\delta}_{st}$, and the time varying component in marginal cost, ζ_{st}, are included to capture sellers' information on future periods' demand and cost. The dynamic component of demand, δ_{st}, is the combination of the demand observables and unobservable, namely,

$$\widetilde{\delta}_{st} = X_{st} + \Delta\xi_{st} \tag{1.14}$$

Based on the assumption that X_{st} evoles as an AR(1) process (equation 1.6) and $\Delta\xi_{st}$ is independent (equation 1.8), the seller-specific demand variable, $\widetilde{\delta}_{st}$, evolves as an AR(1) process:

$$\widetilde{\delta}_{st} = X_{st} + \Delta\xi_{st} = \theta X_{st-1} + \upsilon_{st} + \Delta\xi_{st} = \theta \underbrace{(X_{st-1} + \xi_{st-1})}_{\widetilde{\delta}_{st-1}} + \underbrace{\upsilon_{st} + \Delta\xi_{st} - \theta\xi_{st-1}}_{\varepsilon_{st}} \tag{1.15}$$

where ε_{st} is an independent shock that follows a normal distribution with the variance of $\sigma^2_{\widetilde{\delta}} = (1+\theta^2)\sigma^2_\xi + \sigma^2_\upsilon$. Parameters θ and $\sigma^2_{\widetilde{\delta}}$ are to be estimated. The dynamic component of cost, ζ_{st}, is the only unobserved state for econometricians, and it evolves as an AR(1) process in equation 1.10.

1.4.4 Seller's Dynamic Pricing

In what follows, I define each seller's dynamic pricing problem. Since the problem becomes a single-agent problem, I omit the seller subscript s below.

The expected per-period profit is

$$E[\pi_t] = E_{\Gamma_t}\left[((1-\gamma)p_t - c_t)sh_t(\widetilde{\delta}_t, p_t, \Gamma_t)\middle|\Gamma_{t-1}\right], \tag{1.16}$$

where γ is the royalty rate and c_t is the period t's marginal cost given by 1.9. The expectation is taken on the current period's market competition, Γ_t. Taking use of the assumption on its distribution, as equation 1.13 shows, the expected value on the per-period profit function has a closed form:

$$E[\pi_t(\Gamma_t)|\Gamma_{t-1}] = ((1-\gamma)p_t - c_t)sh_t^e H_t, \tag{1.17}$$

where

$$sh_t^e = \frac{\exp(\xi + \widetilde{\delta}_t + \alpha_p p_t)}{\exp(\xi + \widetilde{\delta}_t + \alpha_p p_t) + \exp(\lambda_0 \Gamma_{t-1} + \lambda_t + \sigma_\nu^2/2)} \tag{1.18}$$

$$H_t = \left(1 + \left(e^{\widetilde{\sigma}_t^2} - 1\right)(1 - sh_0 - sh_t^e)^2\right) \tag{1.19}$$

and $\lambda_0, \lambda_t, \sigma_\nu$ are parameters in the competition intensity transition defined in equation (1.13).[21]

Based on the state variables $s_t = (\Gamma_{t-1}, \xi, \widetilde{\delta}_t, \zeta_t)$ and the realized value of ω_t, each seller maximizes the total value of current period's profit and the expected rebate. The total value at the final period T is given by

$$\widetilde{V}_T(CR_{T-1}, s_T, \omega) = \max_p E\Big[\pi_T(s_T, p, \omega) + fc \times \eta_1 I_1\{CR_{T-1} + R_T(p)\} + fc \times \eta_2 I_2\{CR_{T-1} + R_T(p)\}\Big], \tag{1.20}$$

where the expectation is taken on Γ_T, $I_1 = 1\{\bar{r}_1 \leq CR_{T-1} + R_T(p, \Gamma_T) < \bar{r}_2\}$ is the index that the seller's cumulative revenue over the whole year exceeds the first threshold but below the second threshold, and $I_2 = 1\{CR_{T-1} + R_T(p, \Gamma_T) \geq \bar{r}_2\}$ is the index that exceeds the second. The two index functions depend on period T's price, state, and the realization of market competition intensity.

For period $t < T$, the seller maximizes the total value of profits from the current period until period T, including the expected rebates, which is given by

$$\widetilde{V}_t(CR_{t-1}, s_t, \omega_t) = \max_{p_t}\ E[\pi_t(s_t, p_t, \omega_t)] + E[V_{t+1}(CR_t + R_t(p_t), s_{t+1})], \tag{1.21}$$

where the first expectation is taken on market competition intensity, and the second expectation is taken on state transition including the competition intensity. The ex-ante value function before the realization of cost shock, ω_t, is

$$V_t(s_t) = E_{\omega_t}\left\{\widetilde{V}_t(s_t, \omega_t, p^*(CR_{t-1}, s_t, \omega_t))\right\}, \tag{1.22}$$

where $p^*(CR_{t-1}, s_t, \omega_t)$ is the optimal price set at state $(CR_{t-1}, s_t, \omega_t)$, and $\widetilde{V}_t$ is defined in equation 1.21 and 1.20.

[21] A detailed calculation is shown in appendix 1.10.

Given all of these pieces, the equilibrium of this game is defined as a set of market shares $\{sh_{st}\}_{\forall s,t}$, optimal prices, $\{p_{st}\}_{\forall s,t}$, and sellers' beliefs on the aggregated state transition, $\left\{\widetilde{\mathcal{P}}_t(\Gamma_{st+1}|\Gamma_{st})\right\}_{\forall t}$, such that

- Given price and state, $sh^*_{st} = sh(\vec{p}_t, \vec{s}_t)$ is defined by equation 1.3 for all s and t.
- Given consumers' preferences and the transition belief at period t, each seller s follows the policy function $p^*_{st} = p(CR_{st-1}, s_{st}, \omega_{st})$ that maximizes the value function in equation 1.22.
- Sellers' expectations are rational such that market competition transition belief corresponds to optimal actions: $\mathcal{P}_t(\Gamma_{st+1}|\Gamma_{st};\sigma_{st}) = \widetilde{\mathcal{P}}_t(\Gamma_{st+1}|\Gamma_{st})$ for all s and t.

1.4.5 Implications of the Model

For sellers at period $t = T$, the optimal price p_T needs to maximize the final period's profit and the expected rebates.

$$\frac{\partial E\pi_T}{\partial p_T} + fc\eta_1 \times \frac{\partial EI_1}{\partial p_T} + fc\eta_2 \times \frac{\partial EI_2}{\partial p_T} = 0, \tag{1.23}$$

where $EI_1 = Prob(\bar{r}_1 \leq R(p_T) + CR_{T-1} < \bar{r}_2)$ and $EI_2 = Prob(R(p_T) + CR_{T-1} \geq \bar{r}_2)$ are the probabilities to reach the first and the second threshold.

For the seller at period $t < T$, the optimal price p_t satisfies the dynamic first order condition:

$$\underbrace{E_{\Gamma_t}\left[(1-\gamma)\frac{\partial R_t}{\partial p_t} - c_t\frac{\partial sh_t}{\partial p_t}\right]}_{\text{per-period marginal profit}} = \underbrace{E_{s_{t+1}}\left[\left(-\frac{\partial V_{t+1}}{\partial CR_t}\right)\left(\frac{\partial R_t}{\partial p_t}\right)\middle| CR_t < \bar{r}_2\right]}_{\text{incentive}} \tag{1.24}$$

The left hand side is the per-period marginal profit where the expectation is taken on Γ_t, and the right hand side is the marginal dis-utility of not getting closer to the threshold, where the expectation is taken on state transition. This equation indicates that sellers' price has to balance the trade-off between per-period marginal profit and revenue accumulation for rebates. If a seller sets a price to maximize per-period profit, she has to face a lower possibility of receiving rebates in the future.

As the right hand side of the equation above shows, the incentive from rebates is affected by the state in all of future periods, including sellers' own demand, cost state and the market competition intensity. When sellers perceive high demand/low cost/non-intensive market competition leading to some possibilities of reaching the threshold in the future, they have incentives to reduce the current period's price. However, when the possibility is too low or too high, the incentive to reduce price is weak. The incentive becomes zero when sellers have passed the second threshold, meaning that they will only maximize per-period profit.

This equation also shows that three effects can affect the incentives to reduce price: the distance to the threshold, the ability to predict future periods' revenue, along with the effect of competition. Sellers are more likely to reduce price and accumulate revenue in order to pass the threshold when the rebate threshold approaches. However, to what extent sellers would like to reduce price is also affected by their expectation on future periods' revenue. With more persistent demand/cost shocks, sellers are more likely to set prices in order to exactly reach the threshold. With less persistence, more sellers are induced to reduce price but may be randomly below or above the threshold. The transition of the competition intensity also affect sellers' pricing decisions. If sellers perceive an intensified competition at the end of year, they prefer to reduce price more and earlier. These factors together determines when and to what extent sellers are willing to reduce price.

To understand how rebates ameliorate double marginalization, I rewrite equation 1.24 as

$$E_{\Gamma_t}\left[\left(sh_t + p_t\frac{\partial sh_t}{\partial p_t}\right)\left(1-\gamma+\frac{\partial E[V_{t+1}|CR_t<\bar{r}_2]}{\partial CR_t}\right)\right] = E_{\Gamma_t}\left[c_t\frac{\partial sh_t}{\partial p_t}\right] \tag{1.25}$$

As the left hand side of equation 1.25 shows, the royalty rate decreases sellers' marginal revenue, making sellers face $1-\gamma$ portion of the total revenue. However, sellers face the total portion of cost, as the equation's right hand side indicates. Sellers who maximize their own profits thus tend to set a higher price than the integrated firm's choice, which is known as double marginalization. Rebates can ameliorate this problem, as rebates subsidize sellers' loss in revenue. As the right hand side of this equation shows, with rebates, sellers face $1-\gamma+E[\partial V_{t+1}/\partial CR_t]$ portion of total revenue.

1.5 Estimation

There are three groups of parameters: (1) demand tastes (α_x, α_p) and demand transition $(\theta, \sigma^2_{\tilde{\delta}})$, (2) aggregate state transition belief $(\lambda_0, \lambda_t, \widetilde{\sigma}^2_\nu)$, and (3) cost parameters β and cost transition $(\rho, \sigma^2_\varsigma, \sigma^2_\omega, \sigma^2_\zeta)$. The first two groups of parameters could be estimated with sales data and realized states, and the cost side parameters are estimated via solving the dynamic pricing problem. I will explain the procedure in two steps.

1.5.1 Demand Preference and State Transition

This subsection explains the estimation of the first two group parameters: demand states, (α_x, α_p), demand transition $(\theta, \sigma^2_{\tilde{\delta}})$ and aggregate state transition $(\lambda_0, \lambda_t, \sigma^2_\nu)$. Remember that demand is characterized by a multinomial discrete choice model in equation 1.3. Invert the market share defined in equation 3.35 to obtain the mean utility as

$$\log(sh_{st}) - \log(sh_0) = \delta_{st} = x_{st}\alpha_x + \alpha_p\log(p_{st}) + \xi_s + \xi_t + \Delta\xi_{st}, \tag{1.26}$$

where x_{st} is a vector of seller s's observed characteristics, p_{st} is the retail price, ξ_t is a time dummy, ξ_s is the seller-level fixed effect, and $\Delta\xi_{st}$ is the seller-time level unobservable. I difference out seller-level fixed effect and estimate taste parameters with the assumption $E[\widetilde{\Delta\xi}_{st}\widetilde{z}_{st}] = 0$ as Berry (1994), where $\widetilde{\Delta\xi}_{st}$ is the after-differenced $\Delta\xi_{st}$, and z_{st} is the instrument for the endogenous price.[22]

[22] I difference out the fixed effect as the fixed effect method used in Stata.

I use the average price for the same product listed by small sellers in the last year as an instrument, where a product is a combination of brand, screen size and memory. This instrument is an approximation of wholesale prices. I choose small sellers' prices because they have little market power, and thus set prices close to wholesale prices.[23] Also, I take the average price over the whole year of 2014, meaning that the instrument does not contain 2014's seasonal demand shock. Therefore, 2014's prices are correlated with 2015's prices through marginal costs but not correlated with 2015's seller-time specific demand shocks. However, there are new products which do not exist in 2014. As they have no instruments, I drop these observations in the demand estimation.[24] I aggregate the instruments into the seller-time level in the same manner as before.[25] Then, the instrument's variation across sellers and time comes from the variations in sellers' product sets. Since I have linear instruments in the model, I use two-stage least square for estimation. That is,

$$(\alpha_x, \alpha_p) = (XZ'(ZZ')^{-1}ZX')^{-1}XZ'(ZZ')^{-1}Y, \tag{1.27}$$

where $X = (\widetilde{x}_{st}, \widetilde{p}_{st}, \widetilde{\xi}_t)$, ξ_t are 11 month dummies, and $Y = \widetilde{\log(sh_{st}) - \log(sh_0)}$.

Having obtained the preference parameters, I am able to recover the seller-level fixed effect ξ_s and the mean utility of characteristics $\widetilde{\delta}_{st}$ and the market competition intensity Γ_{st} as equation 1.14 and equation 1.13, respectively.[26] The next step is to estimate parameters in demand transition and market competition transition. I use the maximum likelihood estimation. Denote $\hat{s}_{st} = \hat{\tilde{\delta}}_{st}$ or $\hat{\Gamma}_{st}$,

$$(\hat{\theta}, \hat{\sigma}^2_{\tilde{\delta}}) \ / \ (\hat{\lambda}_0, \hat{\lambda}_t, \hat{\sigma}^2_{\nu}) = \arg\max \sum_{s,t} -\frac{1}{2}\log(2\pi) - \frac{1}{2}\log(V(\hat{s}_{st})) - \frac{\hat{s}_{st} - E[\hat{s}_{st}]}{2V[\hat{s}_{st}]}, \tag{1.28}$$

where $E[\hat{\tilde{\delta}}_{st}] = \theta\hat{\tilde{\delta}}_{st-1}$, $V[\hat{\tilde{\delta}}_{st}] = \sigma^2_{\tilde{\delta}}$, and $E[\hat{\Gamma}_{st}] = \lambda_0\hat{\Gamma}_{st-1} + \lambda_t$, $V[\hat{\Gamma}_{st}] = \sigma^2_{\nu}$. The normal distribution is the model assumptions described in section 1.4.2.

1.5.2 Cost Parameters

The second step is to identify the parameters in cost shifters and cost transition $(\beta, \rho, \sigma^2_{\varsigma}, \sigma^2_{\omega}, \sigma^2_{\zeta})$ from the seller's single-agent dynamic pricing problem. First, I solve the ex-ante value function backwards from period T to period 1. Then, given price data and the derivative of the next period's value function, I solve the corresponding independent cost shock through the seller's first order conditions for all periods, $\{\omega_t\}_{t=1}^{T}$. Finally, I use the independent assumption between *i.i.d.* cost shock and sellers' observed characteristics.

[23] The small sellers are defined as those with yearly quantities below the 30th percentile among all sellers.

[24] New products account for 0.9% of observations.

[25] That is to say, $z_{st} = \frac{q_{jst}}{\sum_j q_{jst}} z_{jst}$, where j denotes a product and z_{jst} is the 2014's average price for product j.

[26] The seller-level fixed effect is recovered as $\hat{\xi}_s = \overline{Y}_s - \overline{X}_s\hat{\alpha}$, where $\hat{\alpha} = (\hat{\alpha}_x, \hat{\alpha}_p)$, $\overline{Y}_s = \frac{1}{\#t}Y_{st}$ and $\overline{X}_s = \frac{1}{\#t}X_{st}$

The computation of value function will be explained in section 1.5.4. Now suppose I already have the value function over all periods. I will explain how to calculate the independent cost shock from the price data. First, from the dynamic first order condition, marginal cost is a function of price, state, the derivative of next period's value function and parameters. I rewrite 1.24 as

$$c_t = \left(\frac{\partial E[sh_t]}{\partial p_t}\right)^{-1} E_{\Gamma_t}\left[(1-\gamma)\left(sh_t + p_t\frac{\partial sh_t}{\partial p_t}\right) + E\left[\frac{\partial V_{t+1}}{\partial p_t}\Big| CR_t < \bar{r}_2\right]\right], \tag{1.29}$$

where the first expectation is taken on the aggregated state, and the second expectation is taken on the rest of the state variables. The derivative of the next period's value function is calculated numerically, which I will explain later in section 1.5.4. From the cost function defined in 1.9, the independent cost shock conditional on state could be inverted as

$$\omega_t = \log(c_t(p_t, s_t, V_{t+1}; \theta^c)) - w_{st}\beta - \zeta_t, \tag{1.30}$$

where c_t is defined in 1.29, w_{st} is a vector of cost shifters, and ζ_t is the state of this period's serially correlated cost process.

I assume that the independent cost shock is uncorrelated with seller's observed characteristics, $E[\omega_t x_t] = 0$, where x_t is seller and tablet characteristics. However, the independent cost shock is conditional on the state of the correlated cost process. To deal with the unobserved state variable, I integrate the moment condition over the serially correlated cost process from period t to the first period when the seller comes to the market. That is, the moment condition used for estimation is

$$E\left[\int \ldots \int \omega_t x_t P(\zeta_t|\zeta_{t-1}) \ldots P(\zeta_{t_1})\right] = 0, \tag{1.31}$$

where ω_t is defined in 1.30, ζ_{t_1} is the first period when the seller comes to the market.[27] The distribution of correlated shock is defined in section 1.4.2. I estimate the parameters by a generalized method of moment as

$$\max_{\theta} \quad \left(\frac{1}{\#s\#t}\sum_{s,t} m_{st}(\theta^c; p_{st}, s_{st}, V^d_{st+1})\right)' W \left(\frac{1}{\#s\#t}\sum_{s,t} m_{st}(\theta^c; p_{st}, s_{st}, V^d_{st+1})\right), \tag{1.32}$$

where

$$m_{st}(\theta^c; p_{st}, s_{st}, V^d_{st+1}) = \int \ldots \int \omega_{st}(\theta^c; p_{st}, s_{st}, V^d_{st+1}) x_{st} P(\zeta_{st}) \ldots P(\zeta_{st^s_1}) \tag{1.33}$$

and $\omega(.)$ is defined in equation 1.30, and W is a weighting matrix. I use the two-step GMM, where the weighting matrix is an identity matrix at the first step and is replaced by the inverse matrix of the estimated variance at the second step.

[27] In the empirical work, I drop sellers with zero sales in one month. So it is possible that some sellers disappear in one month and show up in the next month; for example, one seller is on the market in January and March, but not in February. For those sellers, I assume the correlated shock transit successively even though the seller does not exist in the data successively. In this example, that is to say the shock transits twice, from January to February and from February to March.

1.5.3 Identificaton

Taste parameters, demand and aggregate state transition are identified from the sales data and the empirical transitions in the data. Here I explain how cost parameters and cost transition are identified from the price data. Cost shifters coefficients, β, is identified by the correlation of price and cost observables. Parameters, $\rho, \sigma_\varsigma^2, \sigma_\omega^2$, are identified by price distribution within one period and the evolution of price distribution over time, where ρ, σ_ς^2 are parameters in the dynamic process in cost and σ_ω^2 is the parameter in the independent cost shock. The realization of the dynamic cost process affects the current period and all future periods' cost distributions. Then, ρ and σ_ς^2 are identified by the evolution of price distribution over time. Specifically, ρ affects the correlation of price and the dispersion of price, and σ_ς^2 only affects the dispersion of price within one period. The variance of the independent cost shock σ_ω^2 is also identified by the dispersion of price data within one period. These two cost variances can be separately identified because of the existence of sellers above the threshold and those below it. Sellers above the threshold set static price independent of the future period's cost, while, sellers who are below the threshold set dynamic price dependent of cost transitions in future periods. Then, taking expectation on ζ_{st}, static-pricing sellers' price variation is totally explained by the independent cost shock, but dynamic-pricing sellers' price variation can be explained by both the independent cost shock and the serially correlated cost state. Then, σ_ω^2 is identified by the variation of static sellers price in one month, ρ is identified by the correlation of price over time, and σ_ς^2 is identified by the variation of dynamic sellers' price relative to static sellers' within one month. Finally, σ_ζ^2 is the variance of the initial distribution of the serially correlated cost process. It is identified by the distribution of dynamic sellers' price variation at the time they first arrive to the marketplace.

1.5.4 Computation

This subsection explains how I computationally solve the model. I discretize all five state variables. I discretize the seller quality, ξ_s, into 15 bins and pick up the mid point of each bin as a grid point. For each bin of seller quality, I take 15 bins of $\widetilde{\delta}_{st}$ and 5 bins for Γ_{st}, respectively. Descriptive statistics for these three state variables are displayed in table 1.8. The fourth state variable, the cumulative revenue state CR_{st-1}, has a wide range, going from 0 to \$200k. Since the rebate incentive intensifies as the revenue thresholds approach, I use 50 grid points with more points around the two thresholds. To estimate the continuous AR(1) process of cost shock, I follow the method from **?**. Specifically, I use five grid points to approximate the continuous process.

With discrete state space, state transitions are approximated by transition matrices. That is,

$$\begin{aligned}\widetilde{V}_t(s_t, CR_{t-1}, \omega) &= E\pi(s_t, \omega) + EV(s_{t+1}, CR_t | s_t, CR_{t-1}, p_t^*) \\ &= E\pi(s_t, \omega) + \sum_{s_{t+1}, CR_t} V(s_{t+1}, CR_t)\Pi(s_{t+1}, CR_t | s_t, CR_{t-1}, p(\omega; s_t, CR_{t-1}, \theta)),\end{aligned} \tag{1.34}$$

where $\Pi(s_{t+1}, CR_t | s_t, CR_{t-1}, p_t^*)$ is the state transition matrix over the discrete state space given the optimal price $p(\omega; s_t, CR_{t-1}, \theta)$. To know the value function at any cumulative revenue level, $V(s_{t+1}, CR_t)$, I use piecewise cubic Hermite interpolation. Specifically, for each combination of the discrete state variables, $(\Gamma_{st}, \widetilde{\delta}_{st}, \xi_s, \zeta_{st})$, I use one-dimensional interpolation over the grid points of the cumulative revenue, CR_{t-1}.[28]

Given the above, the moment condition in 3.29 is approximated as

$$m_{st}(\theta^c) \doteq \sum_{\zeta_t \in \{\zeta^{(i)}\}_{i=1}^5} \cdots \sum_{\zeta_\tau \in \{\zeta^{(i)}\}_{i=1}^5} \{\omega(\theta^c; p_{st}, s_{st}, \zeta_t) x_{st}\} \, \Pi(\zeta_\tau | \zeta_{\tau-1}; \theta^c) \ldots \Pi(\zeta_{t_1^s}; \theta^c), \tag{1.35}$$

where $\omega(.)$ is *i.i.d.* cost shock defined in 1.30, $\{\zeta^{(i)}\}_{i=1}^5$ are five discrete points of the Markov chain, and $\Pi(\zeta_\tau|\zeta_{\tau-1}; \theta^c)$ is the transition matrix.

Next, I explain how I solve the ex-ante value function from period $t = T$ until $t = 2$. Given the state and parameters, I need to compute

$$V_t(CR_{t-1}, s_t; \theta) = \int_\omega \widetilde{V}_t\{CR_{t-1}, s_t, p(\omega; CR_{t-1}, s_t, \theta)\} f_\omega(\omega; \theta) d\omega, \tag{1.36}$$

where $\widetilde{V}_t$ is a known function on the state (CR_{t-1}, s_t), cost shock ω, and optimal price $p_t^* = p(\omega; CR_{t-1}, s_t, \theta)$ at the state and cost shock given parameters, which is shown in equation 1.34, and f_ω is known as a normal distribution by equation 1.30. However, computing the optimal action as a function of state variables and shock is hard because price non-linearly enters the value function. In my case, solving period T's optimal price on 10,000 grids for one shock takes more than one hour. Solving the optimal decision for period $t < T$ consumes even more time because the interpolation of the next period's value function takes time. Then, calculating one objective function for a "one-year problem" will cost at least six hours. Then, one estimation with 100 iterations will cost 25 days.

To improve the speed of estimation and preserve the accuracy, I propose a new method for value function integration taking advantage of the inverse mapping of optimal price. Although solving p as a function of $(\omega, CR_{t-1}, s_t, \theta)$ is hard, computing the inverse c of the function p as a function of ω is relatively straightforward, as I demonstrated above in section 3.5, equation 1.29. More clearly, the method is to use first order conditions to derive the distribution of optimal price from the distribution of cost shock and then change the integration over the distribution of cost shock to the integration over the distribution of optimal price.

I rewrite (1.36) as

$$\int_p \widetilde{V}_t\{CR, s, c(p; CR_{t-1}, s_t, \theta), p\} f_p^*(p; CR_{t-1}, s_t, \theta) dp, \tag{1.37}$$

where f_p^* is the density of optimal prices. It is easy to prove that optimal price is monotone with the independent cost shock. Therefore, given the correlated cost state, the optimal price is monotone with the marginal cost. Then I approximate the *c.d.f.* of the optimal price as

$$F_p^*(p; CR, s, \theta) = F_c(c(p; CR, s, \theta)), \tag{1.38}$$

[28]The reason to use cubic Hermite interpolation is to preserve the monotonicity of the value function with respect to the state of cumulative revenue.

where $c(p; CR, s, \theta)$ is defined in equation 1.29. Note that solving the marginal cost from the first order conditions needs the derivative of value function. To obtain the derivative, I use the first-order difference.[29] I then discretize the action space $\{p^i\}_{i=1}^{n_p}$ and approximate the density of the optimal price over the discrete action space as

$$f_p^*(p^i; CR, s, \theta) = F_c(c(p^i; CR, s, \theta)) - F_c(c(p^{i-1}; CR, s, \theta)), \tag{1.39}$$

where p^i are price grid points, $p^i > p^{i-1}$ for $i = 2, \ldots, n_p$. Then the ex-ante value function is approximated as

$$V_t(CR, s; \theta) = \sum_{p^i} \widetilde{V}_t\{CR, s, c(p^i; CR, s, \theta), p^i\} f_p^*(p^i; CR, s, \theta) \tag{1.40}$$

In summary, I solve the ex-ante value function, $\{V_t(CR, s; \theta)\}_{t=2}^T$, as follows. Start from the final period $t = T$. I solve the marginal cost from this static problem. With the marginal cost, I calculate the distribution of optimal price (as equation 1.39) and the ex-ante value function (as equation 1.40). Next, move to one period before. At period $t = T - 1$, I numerically calculate the derivative of the next period's value function through interpoaltion and finite-difference. With the derivative, I calculate the marginal cost (as equation 1.29), the distribution of optimal price (as equation 1.39), and then ex-ante value function (as equation 1.40). Continue this procedure until period $t = 2$.

1.6 Results

The estimates of the preference parameters are in table 1.7 and statistics summarizing the distribution of price elasticity are in table 1.8. The distribution of price elasticity displays little variation with a mean of -11.85 and a standard deviation of 0.12. The estimated seller quality ranges between -6.16 and 9.63, the mean utility of sellers' characteristics ranges from 68 to 85, and the aggregated state has a mean of 1.35 and a standard deviation of 0.28. The parameters of the cost function and the cost evolution process are displayed in table 1.9, where the correlation ρ is 0.46, the variance σ_ς^2 is 1.04, and the variance of initial distribution is 5.17. This suggests that the distribution of cost shock becomes more condensed over time.

To better understand how these results translate to pricing behavior, I provide an example of an ex-ante policy function with respect to different levels of cumulative revenue. To isolate the effect of the rebate, I assume all the other states are $\widetilde{\delta}_{st} = 73.3084$, $\xi_s = -5.1414$, $\zeta_{st} = 0$.[30] As figure 1.1 shows, the model predicts a V-shape in prices as the seller approaches each threshold, which indicates a greater incentive to lower the price to receive the rebate. Also, this shape becomes deeper as the

[29] That is, with a price point, I can calculate the next period's cumulative revenue and then obtain the next period's value function from interpolation. The derivative of value function can be calculated numerically with a small change in the price. Specifically, to calculate the derivative at the price grid points, I take the difference between two successive grid points. To calculate the derivative at the price data, I use a small change of $1.3 (RMB10) in price.

[30] In the example, the competition intensity varies over time 1.3945, 1.1779, 0.9915 and 0.7061 from September to December.

December deadline approaches, implying that the incentive to lower prices becomes even greater as the deadline approaches.

To give insights for the trade-off between profits and revenue accumulation, I display an example of ex-ante per-period profit loss in figure 1.2. I keep all the other states the same as before. As the figure shows, the willingness to pay for the rebate increases as the rebate threshold approaches, around \$200 close to the second threshold. This indicates a larger incentive to accumulate revenue when the threshold is closer. However, the trend is not significant in the early months. For example, in September, profit loss has a small variation at different revenue level. Also, note that the willingness to pay in November is larger than that in other months because of the online shopping festival on November 11, which leads to more listings with high mean utility and high seller-level fixed effects.

To better understand sellers' willingness to pay for the rebate, I calculate the ex-ante per-period profit loss and the expected future gains relative to a static environment as follows:

$$\Delta E\pi_t = \frac{1}{\#CR_{t-1}, s_t} \sum_{CR_{t-1}}^{s_t} \left(E\pi_t(s_t) - E\pi_t^{stat}(s_t)\right) \tag{1.41}$$

where $E\pi_t(s_t)$ is the ex-ante per-period profits at state s_t with the rebates, and $E\pi_t^{stat}$ is the expected static profits at state s_t without the rebate. The gains from future periods are calculated in the same way:

$$\Delta EV_{t+1} = \frac{1}{\#CR_{t-1}, s_t} \sum_{CR_{t-1}}^{s_t} \left(EV_{t+1}(CR_t, s_{t+1}) - \sum_{\tau=t+1}^{T} E\pi_\tau^{stat}(s_{t+1})\right), \tag{1.42}$$

where the expectation on value function is taken on state transition, and $\sum_{\tau=t+1}^{T} E\pi_\tau^{stat}(s_t)$ represents the expected static profits in future periods. [31]

The results in table 1.10 display the trade-off for two groups of cumulative revenue intervals. The first group includes the sellers whose cumulative revenue level between 80% and 100% of each threshold, and the second group includes that between 50% and 80% level of each threshold. The table shows that the willingness to pay for the rebate enlarges from \$34 to \$402 in August to November but shrinks to \$73 in December. This is because sellers are better to predict the possibility of reaching the threshold in December, thus, those who are still far from the threshold tend to give up accumulating revenue. The expected gains from rebates decrease over time from \$4,293 to \$3,328 because the possibility of reaching the threshold decreases as less time left. Compared to the first group, sellers in the second group are willing to pay more for the rebate from August to November (\$45 to \$439), but pay less in December (\$36). This indicates that sellers look forward to receive the rebate when setting their current period's price. In early months, sellers who are far from the threshold have more incentive to accumulate revenue than those who are close to the threshold, but vice versa in later months when the deadline closes.

[31] I calculate the $E\pi_t(CR_{t-1}, s_t)$ and $EV_{t+1}(CR_{t-1}, s_t)$ in the same way of calculating ex-ante value function in section 1.5.4, for example, $E\pi_t(s_t) = \frac{1}{\#p^i} \sum_{p^i} E\pi_t(p^i, CR_{t-1}, s_t) f_p^*(p^i, CR_{t-1}, s_t)$

1.7 Counterfactual

In this section, I use the estimated parameters and conduct a counterfactual analysis to quantify how effective the rebate contract induce sellers to reduce price and produce revenue for the platform, and to what extent the contract improves industry profits and consumer welfare. I do this by eliminating rebates. The contract thus becomes the same as pure royalty rate contract at the observed level, (γ). Under the pure royalty rate contract, sellers' prices become static as they do not have to look forward rebates gains. I then solve sellers' optimal prices under the pure royalty rate contract, and compare total revenue, profit and consumer surplus with the observed values. Note that I compute these counterfactual outcomes under the royalty rate employed on by Tmall, which is not necessarily the optimal rate without the rebate. Therefore, I compute the counterfactual outcomes under alternative royalty rates for robustness. Before discussing the results, I explain how to solve sellers' optimal price decisions under the adjusted environment.

1.7.1 Methodology

In the counterfactual, I assume sellers observe the revenue sharing contract, γ. With the contract, sellers want to maximize the expected profit over the whole year, i.e.

$$\max_{p_t} E_\Gamma \left[\sum_t ((1-\gamma)p_t - c_t) sh_t^e(p_t, \Gamma_t) \right], \tag{1.43}$$

where sh_t^e is defined in equation 1.18, and the expectation is taken on the market competition intensity with the first period's information. The equation above shows that sellers price decisions are affected by the expected competition intensity. In equilibrium, sellers' expectation on competition intensity should correspond to their optimal price decisions. Thus, to solve sellers' optimal price, I need to obtain the new transition process of Γ_{st}. This state variable is endogenous. It is determined by all competitors' price decisions. In the estimation, I obtain its transition from the data. However, in the counterfactual, I need to solve a new transition that is consistent with the new pay-off function in the adjusted environment.

The methodology that I use to solve for the equilibrium is nested fixed point algorithm. Specifically, I simulate a number of values for the unobserved correlated cost. Given these costs, I solve for sellers' optimal prices and corresponding transition parameters under a guess of the transition parameters. I then iterate over this process until the transition parameters converge. The procedure is as follows:

1 Guess parameters $\theta = (\gamma_0, \gamma_t, \widetilde{\sigma}_t^2)$ for all periods $t = 1, \ldots, T$

2 Solve optimal dynamic pricing problem for each seller, calculate the aggregated state, $\Gamma_t = \sum_s \exp(\delta_{st}(p_{st}(\theta)))$ for all periods $t = 1, \ldots, T$

3 Calculate the new parameters in aggregate state transition belief:

$$\theta' = \arg\max_\theta \sum_{s,t} \log\left(\frac{1}{\widetilde{\sigma}_t} \phi\left(\frac{\Gamma_t - (\gamma_1 \Gamma_{t-1} + \gamma_t)}{\widetilde{\sigma}_t} \right) \right) \tag{1.44}$$

4 Check convergence: $\frac{|\sum_k \theta_k^{'} - \theta_k|}{1e^{-4} + |\sum_k \theta_k^{'} + \theta_k|/2} < \tau$, if not, update parameters and go back to step 1.

1.7.2 The Effect of Rebates

With the procedure proposed above, I am able to quantify the effect of rebates on sellers' pricing, revenue, profit and consumer welfare. Note that this exercise is not intended to understand the optimal contract, as I am ignoring important aspects, such as seller entry. Instead, I aim to quantify the effect of the observed rebate contract on surplus and profit.

1.7.2.1 The Effect on Sellers' Pricing

First, I examine the rebate incentive on sellers' pricing decisions. Figure 1.3 shows the average price path with and without rebates. In the figure, the average price is ascending over time because the mean utility of tablet increases at later months when new tablets are released on the market.[32] The figure shows that there is little difference in the average price with and without rebates at the beginning of the year. However, this difference becomes apparent in later months, as marginal sellers try to reach the revenue thresholds by lowering their price. Competition may also lead non-marginal firms to lower their price as well, further exacerbating this effect. In December, the difference shrinks because nearly all sellers know whether they are able to pass the threshold. Then, sellers who know they will not pass the threshold stop reducing price in December.

Also, to understand the effectiveness of the rebate contract over many sellers, I calculate the distribution of average percentage change in price over sellers' yearly revenue level as follows:

$$\overline{\Delta p_c} = \frac{1}{\#s,t} \sum_{s,t} \frac{(p_{st} - p_{st}^{stat})}{p_{st}^{stat}} \times 1\{\underline{c} \leq CR_{sT} < \bar{c}\}, \tag{1.45}$$

where p_{st} is the price with rebates, p_{st}^{stat} is the price without rebates, $CR_{sT} = \sum_{t=1}^{T} R_{st}$ is seller s's yearly revenue, and $\underline{c}$ and $\bar{c}$ are the lower and upper bound of the histogram revenue bins. To see the rebate outcome, I use the two rebate thresholds as two revenue bins' edges.

The first subfigure 1.4a displays the histogram and the distribution of average percentage change in price. The x-axis is the yearly revenue, the positive part of y-axis is the histogram of sellers' yearly revenue and the negative part of y-axis is the average percentage reduction in price over the interval of yearly revenue. As the figure shows, rebates affect 84% of sellers to reduce price by more than 0.1%. The 3% of sellers with the lowest revenue and the 13% of sellers with the highest revenue hardly respond to the rebate, and the left 84% of sellers whose yearly revenue ranges from 1% of the first threshold to 200% of the second threshold all respond to the rebate contract. Rebates affect a wide range of sellers due to three effects: the incentive to reach the threshold

[32] The mean utility of the first four months is around 76, and the mean utility becomes 78-79 in the last four months.

when sellers are close to the threshold, uncertainty to predict future revenue, and the effect of monopolistic competition. Although many sellers respond to the rebate contract, only 13.5% of them successfully pass the threshold. The left 86.5% of those sellers reduce price but do not pass the threshold because they cannot perfectly predict their future revenue and profits. With the uncertainty in demand, cost and competition intensity, sellers perceive some possibilities of passing the threshold and thus reduce price. But, many sellers unfortunately receive low demand or high cost, allocating them below the threshold. There are also sellers who are above the threshold but also reduce price. Their price reduction can happen before they pass the threshold, which is caused by uncertainty, and also can happen after they pass the threshold, which is caused by the effect of competition.

To see the rebate incentives clearly, I show the distribution of price reduction at different cumulative revenue level in different months. This helps to understand how sellers react to the rebate when they are at different distance to the threshold. I calculate the average percentage change in price across sellers in a similar way:

$$\overline{\Delta p_{ct}} = \frac{1}{\#s} \sum_{s} \frac{(p_{st} - p_{st}^{stat})}{p_{st}^{stat}} \times 1\{\underline{c} \leq CR_{st-1} < \bar{c}\} \tag{1.46}$$

where CR_{st-1} is the cumulative revenue at the beginning of period t, i.e. from January to the period $t-1$. To see the incentive change just below and above the threshold, I use the two rebate thresholds as two revenue bins' edges. The result is in subfigure 1.4b and 1.4c.

The subfigure 1.4b shows the distribution of sellers' price reduction in September, October and November. It shows an intensified price reduction just below both thresholds. This indicates that sellers have more incentives to reduce price and pass the threshold when they are getting closer to the threshold. Also, sellers who are far from the threshold are affected because of the effect of competition and the uncertainty. Sellers who have passed the threshold are only affected by competition. As the figure shows, they hardly decrease their prices, meaning that the competition effect is less than the uncertainty. Most sellers are affected by rebates because the inability to predict future revenue and profits. Subfigure 1.4c shows the distribution of sellers' price reduction in December. In the final period, only sellers who are close to the threshold have incentives to accumulate revenue. Sellers who are far from the threshold perceive little possibilities of reaching the threshold and thus stop reducing price.

1.7.2.2 The Effect on Revenue, Profit and Consumer Welfare

Now I quantify the effect of the rebate-based contract on total revenue, profit and consumer welfare by comparing the market outcomes with versus without rebates. As the results in table 1.11 show, rebates reduce the average price by \$4 (1.7%), and decrease the standard deviation of price by 6%. This leads to a 2.4% increase in total quantity and a 9% increase in total revenue sold on the marketplace. This indicates that rebates effectively reduce price and smooth price, generating consistent revenue and quantities sales over the contract year.

The platform, sellers and consumers all benefit from rebates. The platform's profit increases by \$1.98 million (56%), which consist of a \$0.32 million increase in revenue sharing, a \$0.57 million loss in rebate payment and a \$2.23 million gain in the fixed fee. The platform loss more rebates payment than the increased rents from revenue due to the existence of the large sellers. These large sellers are able to pass the rebate threshold without reducing price, which costs the platform's rebates but not increases revenue. However, since the rebate increases the average seller's profit, the platform is able to charge a fixed fee. The gain from the fixed fee offsets the platform's loss in rebate payment. Thus, the platform still gains with rebates relative to without. Additionally, consumer surplus increases by 9.5%, and an average seller gain \$638 profit with rebates relative to without.

However, the result above does not consider the optimal royalty rate that the platform may choose in the environment with pure revenue sharing. Thus, I consider various royalty rates from 1.5% to 2.5% for the robustness check. The result still shows that rebates effectively reduce price and produce revenue, generating higher surplus for the platform, the average seller and consumers.

1.8 Conclusion

In this paper, I examine how effective are rebates in giving downstream firms' incentives to reduce price. To do this, I focus on a rebate contract between a dominant real-world online marketplace and many heterogeneous downstream independent sellers. I specify and estimate a model of the downstream sellers' dynamic pricing decisions.

Using the estimates, I quantify the profits and welfare gains of the rebate contract by comparing observed outcomes to outcomes under a contract without rebate. The results show that rebates incentivize 84% of sellers to reduce their prices. Within the whole year, rebates lead to a \$4 (1.7%) reduction in the average price and a 6% decrease in the standard deviation of prices. Total quantity sales increases by 2.4%, and total revenue sales increases by 9%. Additionally, consumers, the platform and sellers are all benefit from rebates. Specifically, the consumer surplus increases by 9.5%, the average seller's profit increases by \$638 (1.9%), and the platform gains \$1.98 million (56%) surplus.

There are a few extensions of this exercise. This exercise only considers sellers' pricing decisions but not their entry decisions. The mechanism design of an optimal contract needs both individual rationality and incentive comparability. One extension can be to incorporate sellers' entry decisions. Also, as rebates shift risk to downstream sellers, the effect of rebates will be different if sellers are risk averse. The other extension can be taking sellers' risk aversion into account. Last, I assume sellers have the same ability to predict future demand, cost and competition intensity. It is possible that some experienced sellers have more accurate knowledge than others. It is interesting to estimates sellers' heterogeneous ability to predict their future revenue, and to understand how the heterogeneity affects their success.

1.9 Tables and Figures

Table 1.1: Examples of Categories and Fee policies

Primary Category /Secondary Category	Royalty Rate	Fixed Fee	Rebate Revenue Threshold 50%	100%
Electronics				
Computer	2%	30,000	360,000	1,200,000
Camera	2%	30,000	360,000	1,200,000
MP3/MP4/iPod/Voice Recorder	2%	30,000	360,000	1,200,000
Cellphone	2%	30,000	360,000	1,200,000
PC/Screen/PC Accessories	2%	30,000	360,000	1,200,000
Laptop	2%	30,000	360,000	1,200,000
Electronics Accessories	2%	30,000	360,000	1,200,000
USB/Disk/Flash Card	2%	30,000	360,000	1,200,000
Tablet	2%	30,000	360,000	1,200,000
Office	2%	30,000	360,000	1,200,000
E-Dictionary/E-book	2%	30,000	360,000	1,200,000
Clothings				
Accessories/Belts/Hats/Scarves	5%	30,000	180,000	600,000
Women's Clothing	5%	60,000	360,000	1,200,000
Men's Clothing	5%	60,000	360,000	1,200,000

Fees are measured in Chinese RMB, with the exchange rate of 6.33 to US dollars. Information is from Alibaba's official website, `https://rule.tmall.com/tdetail-3263.htm?spm=0.0.0.0.Tal1Oj&tag=self`.

Table 1.2: Product Characteristics

	Mean	Median	Standard Deviation	Compare w. Ipad Mini 2
Listing price($)	234.22	164.20	222.07	269
Screen size(Inch)	8.56	8.22	1.50	7.9
Memory(GB)	24.59	15.64	32.72	16
Product Age(Years)	0.88	1.00	0.60	1.66
Cellular Internet Access	0.24	0.01	0.35	0
	Market Share(%)		Average Price($)	
Top Brands(Rank)				
Tecalst(Top 1)	16.22		138.43	
Onda(Top 2)	16.13		112.31	
Apple(Top 3)	13.73		573.03	
Cube(Top 4)	11.86		127.67	
Huawei(Top 5)	6.63		246.40	
Samsung(Top 9)	2.94		340.15	
Operating System				
Android	48.53		173.74	
IOS	13.73		573.03	
Windows	47.66		307.28	
Observations	20,075			

Multi-product sellers' products characteristics are weighted averaged into single product. Listed price is converted to US dollars with the exchange rate equal to 6.33, brands are ranked by total sales(in units).

Table 1.3: Sellers Characteristics

	Mean	Median	Standard Deviation
Sales	298.19	25.00	972.39
Revenue($1000)	59.19	4.27	304.46
Listing price($)	234.22	164.20	222.07
Seller Characteristics			
No of Products	6.07	3.00	7.17
Rating Score	4.79	4.80	0.11
$\% \geq 4.6$	0.99	1.00	0.08
Perc. of Positive Reviews	0.58	0.88	0.45
Tablet Characteristics			
Screen size(Inches)	8.57	8.22	1.50
Memory(GB)	24.59	15.65	32.72
Product Age(Years)	0.88	1.00	0.60
Cellular Internet Access	0.25	0.01	0.35
No. of Pictures	79.03	77.67	38.11
No. of Characters/1000	0.67	0.26	0.95
Video Description	0.16	0.00	0.31
Observations	3,049		

One observation is a seller in one month, listed price is converted to US dollars with the exchange rate equal to 6.33.

Table 1.4: Relationship between Price and Revenue Thresholds

	to the 1st threshold				to the 2nd threshold			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cumulative Revenue ($Mil)								
	-1.039***	-0.279			-0.771***	0.0555		
	(0.329)	(0.464)			(0.218)	(0.550)		
Cross Term: Cumulative Revenue ($Mil) ×								
Feb.			1.129	0.319			-0.915	-0.344*
			(0.715)	(0.490)			(1.151)	(0.177)
Mar			0.585				0.161	
			(0.881)				(0.445)	
Apr			1.184**	0.726			-0.330	-0.184
			(0.497)	(0.620)			(0.710)	(0.447)
May			1.554**				-0.328	
			(0.553)				(0.642)	
Jun			-0.0462				0.578	
			(0.668)				(0.565)	
Jul			0.432	-0.687			0.361	0.600
			(0.373)	(0.615)			(0.616)	(0.469)
Aug			-0.750**				1.253*	
			(0.310)				(0.635)	
Sep			-1.593***				1.358*	
			(0.509)				(0.677)	
Oct			-3.588***	-2.240**			-0.579	-1.425**
			(0.549)	(0.773)			(0.739)	(0.634)
Nov			-2.180***				-1.761**	
			(0.580)				(0.751)	
Dec			0.772				-0.974	
			(0.671)				(0.604)	
Constant	3.117***	3.160***	3.165***	3.036***	4.139***	3.588***	3.981***	3.760***
	(0.446)	(0.468)	(0.470)	(0.431)	(0.578)	(0.720)	(0.970)	(0.851)
Time Dummies		Yes	Yes	Yes		Yes	Yes	Yes
Observations	1994	1994	1994	1994	426	426	426	426
R-squared	0.953	0.953	0.953	0.953	0.986	0.986	0.987	0.987

Note: All regressions use log of price in US dollars. Standard errors in parentheses are clustered in month, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.5: Relationship between price and revenue groups

	Cumulative Revenue Group			
	1	2	3	Month
Month	$[0, 0.5\bar{r}_1)$	$[0.5\bar{r}_1, 0.5\bar{r}_2]$	$[0.5\bar{r}_2, \bar{r}_2)$	Dummy
Feb.	-0.0795***	-0.0584***	-0.0418***	0.0723***
	(0.0286)	(0.0104)	(0.0096)	(0.0208)
Mar.	-0.0376*	-0.0418**	-0.0251**	0.0258*
	(0.0288)	(0.0192)	(0.0121)	(0.0197)
Apr.	-0.0296	0.0194	-0.0358***	-0.0224
	(0.0250)	(0.0189)	(0.0147)	(0.0193)
May.	0.0009	0.0433**	-0.0178*	-0.0486***
	(0.0283)	(0.0201)	(0.0134)	(0.0198)
Jun.	0.0208	0.0248	0.0232**	-0.0769***
	(0.0273)	(0.0212)	(0.0135)	(0.0170)
Jul.	0.0119	0.0360**	0.0265*	-0.0920***
	(0.0202)	(0.0201)	(0.0163)	(0.0163)
Aug.	-0.0131	-0.0313**	0.0156	-0.0589***
	(0.0197)	(0.0185)	(0.0163)	(0.0159)
Sep.	0.0005	-0.0414***	-0.0138	-0.0549***
	(0.0242)	(0.0175)	(0.0184)	(0.0165)
Oct.	0.0027	-0.0755***	-0.0667***	-0.0426***
	(0.0164)	(0.0191)	(0.0159)	(0.0141)
Nov.	0.0267*	0.0093	-0.1060***	-0.0384***
	(0.0208)	(0.0173)	(0.0131)	(0.0147)
Dec.	-0.0077	-0.0032	-0.0188	-0.0978***
	(0.0231)	(0.0214)	(0.0209)	(0.0129)

Note: Standard errors in parentheses are clustered in month, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.6: Evidence of Static Consumers

	Log of Sales		
	(1)	(2)	(3)
Log of Price	-0.594***	-0.598***	-0.594***
	(0.0873)	(0.0876)	(0.0880)
Lag Terms of Log of Price			
1st Period Lag	0.103	0.0698	0.0697
	(0.0873)	(0.108)	(0.108)
2nd Period Lag		0.0450	0.0733
		(0.0877)	(0.108)
3rd Period Lag			-0.0402
			(0.0881)
Constant	6.779***	6.732***	6.785***
	(0.385)	(0.401)	(0.413)
Observations	3033	3032	3031
R^2	0.031	0.031	0.032

Standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.7: Demand Estimators

	OLS	2SLS	
		1st Stage	2nd Stage
log of Price	-0.366***	1.26e-5*	-11.89**
	(0.0993)	(5.86e-6)	(5.898)
Product Characteristics			
Screen Size(Inches)	-0.0176	0.228***	2.717**
	(0.0344)	(0.006)	(1.374)
Memory(GB)	0.0077***	0.0021***	0.0299**
	(0.0017)	(0.0003)	(0.013)
Android	-0.637***	-0.0405	-0.598
	(0.161)	(0.0329)	(0.485)
Windows	-0.201	0.171***	2.047*
	(0.171)	(0.0346)	(1.095)
Internet Access	0.115	0.181***	2.181**
	(0.090)	(0.0166)	(1.107)
Product Age(Years)	-0.503***	-0.0354**	-0.894***
	(0.0575)	(0.0118)	(0.222)
Seller Characteristics			
No. of Tablets	0.166***	0.0031*	0.196***
	(0.0069)	(0.0012)	(0.023)
Rating Score	0.506	-0.0331	-0.0149
	(0.305)	(0.0570)	(0.756)
No. of Ratings	0.0272***	-0.0009	0.012
	(0.0077)	(0.0016)	(0.021)
No. of Pictures	0.000639	6.86e-6	0.0005
	(0.0008)	(0.0002)	(0.002)
Vedio Description	0.372**	0.149***	2.120**
	(0.121)	(0.0229)	(0.932)
No. of Characters	0.0903*	0.0053	0.156*
	(0.0371)	(0.0068)	(0.091)
%Pos. Comments	1.509***	-0.0242	1.145***
	(0.0858)	(0.0169)	(0.265)

Table Continues on the Next Page

Continue of Table 1.7

	OLS	2SLS	
		1st Stage	2nd Stage
Month Dummies			
Feb.	-0.198**	0.0064	-0.129
	(0.0761)	(0.0138)	(0.179)
Mar.	-0.216**	-0.0037	-0.261
	(0.0758)	(0.0138)	(0.176)
Apr.	-0.331***	-0.0282*	-0.654***
	(0.0779)	(0.0142)	(0.242)
May.	-0.596***	-0.0379*	-0.990***
	(0.0805)	(0.0149)	(0.290)
Jun.	-0.264**	-0.0528***	-0.870**
	(0.0822)	(0.0150)	(0.363)
Jul.	-0.294***	-0.0679***	-1.103**
	(0.0852)	(0.0158)	(0.450)
Aug.	-0.536***	-0.0613***	-1.264***
	(0.0852)	(0.0159)	(0.416)
Sep.	-0.836***	-0.0558***	-1.544***
	(0.0864)	(0.0162)	(0.397)
Oct.	-1.150***	-0.0618***	-1.960***
	(0.0872)	(0.0166)	(0.433)
Nov.	-1.410***	-0.0628***	-2.247***
	(0.0869)	(0.0166)	(0.432)
Dec.	0.565***	-0.0891***	-0.795
	(0.0982)	(0.0182)	(0.585)
Constant	-7.795***	5.074***	43.488**
	(1.603)	(17.66)	(1.85)
Number of Seller FE	470	354	354
Observations	3,034	2,688	2,688

Table 1.8: Statistics of Estimated Price Elasticities and State Variables

	Mean	SD	p10	p50	p90
Price Elasticity	-11.8565	0.1213	-11.8935	-11.8906	-11.8139
State Variables					
$\widetilde{\delta}_{st}$	77.4237	9.2323	68.2060	77.2483	85.4767
$x_{st}\alpha_x$	77.4159	9.3159	68.7670	77.0479	85.3688
$\Delta\xi_{st}$	0.0078	3.0306	-1.9864	-0.0000	2.0500
ξ_s	0.0531	8.6896	-6.1592	-0.8566	9.6374
Γ_{st}	1.3504	0.2896	0.9942	1.3896	1.5726

Table 1.9: Transition Parameters and Cost Parameters

Transition Parameters				Cost Parameters			
Demand Transition		Aggregate State Transition		Cost Shifters		Cost Transitions	
Constant	3.8031	λ_t	1.3032	β_0	6.8653	σ^2_ω	1.4532
	(0.3297)		(0.0034)				
θ	0.9526	λ_0	0.0450	β_1	0.1599	ρ	0.4661
	(0.0042)		(0.0026)				
$\log(\sigma^2_\delta)$	2.1298	$\log(\sigma^2_\nu)$	-8.6598			σ^2_ς	1.0445
	(0.0146)		(0.0122)			σ^2_ζ	5.1729

The constant term of aggregate state transition is the mean of the coefficients of 12 time dummies.

Table 1.10: Willingness-to-pay for Rebates (US \$)

	Aug.	Sep.	Oct.	Nov.	Dec.
Group 1: $CR_{t-1} \in [0.8\bar{r}_i, \bar{r}_i]$ *for* $i = 1, 2$					
$\Delta E p_t$	-1.81	-2.57	-4.82	-6.08	-5.93
$\% E p_t$	-0.74%	-1.06%	-1.96%	-2.47%	-2.33%
$\Delta E \pi_t$	-34	-90	-171	-402	-73
$\Delta E V_{t+1}$	4,293	4,177	3,972	3,906	3,328
Exp. Rebates	4,233	4,163	4,003	3,824	3,328
Future Profits	60	13	-30	81	0
Group 2: $CR_{t-1} \in [0.5\bar{r}_i, 0.8\bar{r}_i]$ *for* $i = 1, 2$					
$\Delta E p_t$	-0.78	-1.35	-2.36	-2.95	-0.61
$\% E p_t$	-0.32%	-0.56%	-0.96%	-1.20%	-0.24%
$\Delta E \pi_t$	-45	-103	-196	-439	-36
$\Delta E V_{t+1}$	4,053	3,876	3,554	3,410	2,566
Exp. Rebates	3,896	3,788	3,531	3,271	2,566
Future Profits	157	87	23	139	0

This table shows the expected price reduction, profit loss and the expected value gains in future periods. The expected value gains in future periods is the sum of the expected per-period profit gains in future periods and the expected rebate gains. The expectation is taken on the current period's cost shocks.

Table 1.11: Counterfactual: Rebate Contracts' Efficiency and Profitability

	Revenue Sharing(γ)			Rebates($\gamma, fc, \bar{r}_i$)
Royalty Rate(γ)	1.5%	2%	2.5%	2%
Sellers' Price				
Average Price($)	236.64	237.71	239.02	233.16
S.D. of Price($)	236.75	237.19	238.83	222.66
Average Mark-up($)	23.25	23.12	23.20	18.66
Total Sales	958,434	948,361	938,367	971,905
Total Revenue($Mil)	176.90	176.24	175.03	192.17
Surplus Division($Mil)				
Consumer Surplus	0.2924	0.2813	0.2706	0.3080
Platform's Profit				
Total Net Profit	2.65	3.52	4.38	5.50
Royalty Rate	2.65	3.52	4.38	3.84
Fixed Fee	0	0	0	2.23
Rebates	0	0	0	-0.57
Sum of Sellers' Profit				
Total Net Profit	15.72	15.58	15.38	15.88

Price, profit and revenue are converted from RMB to US dollars using the exchange rate of 6.33.

Figure 1.1: Ex-ante Policy Function

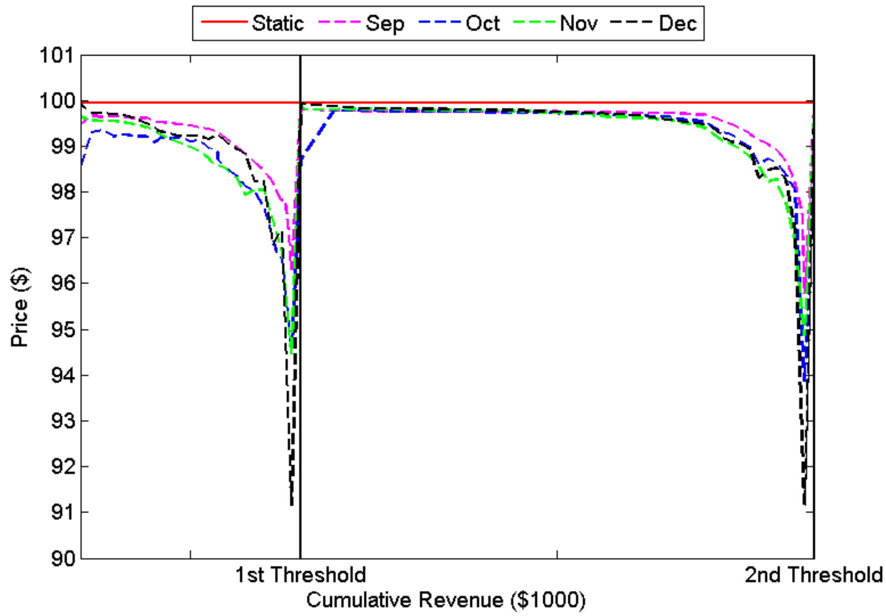

Figure 1.2: Ex-ante Profit Loss

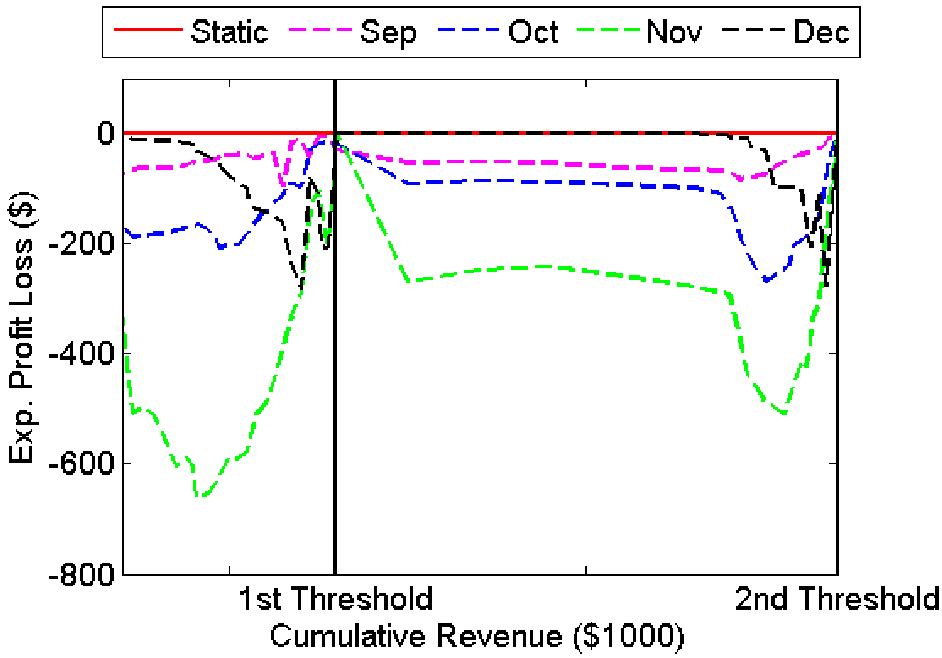

Figure 1.3: Price Path with Rebates versus with Revenue Sharing

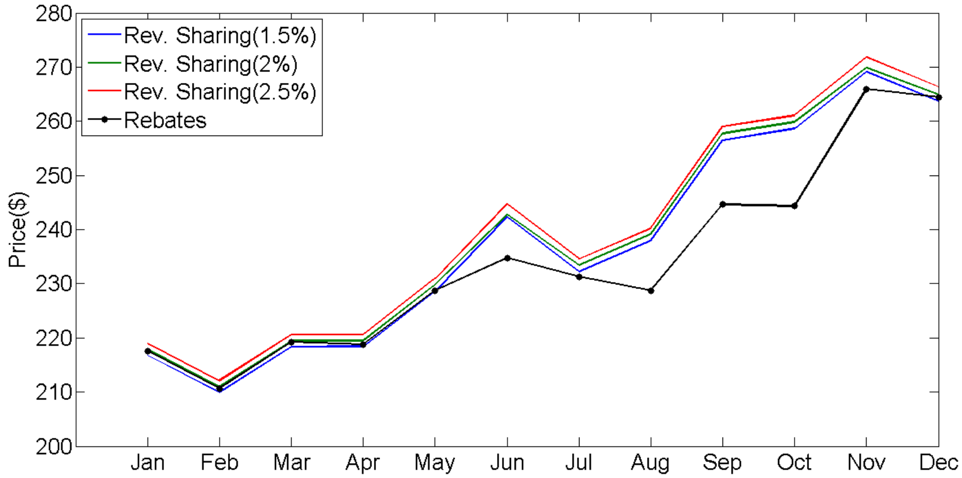

This figure shows the average price path with pure revenue sharing versus with the rebate contract. The average price is ascending over time because the mean utility of tablet increases at later months when new tablets are released on the market.

Figure 1.4: The Distribution of Average Price Reduction

(a) Whole Year Revenue

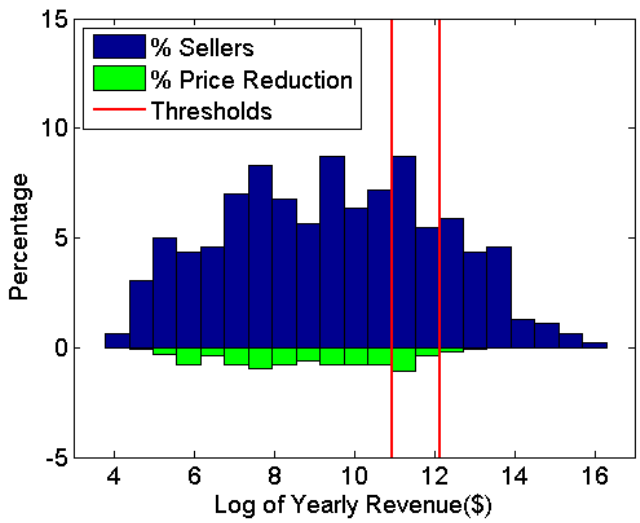

(b) September to November

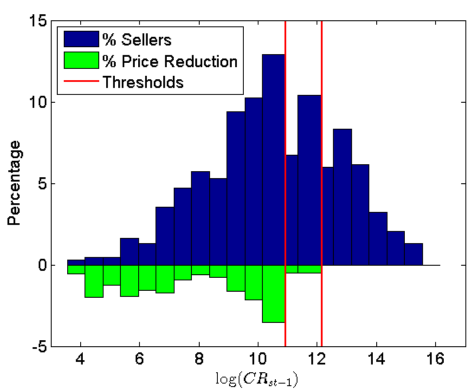

(c) December

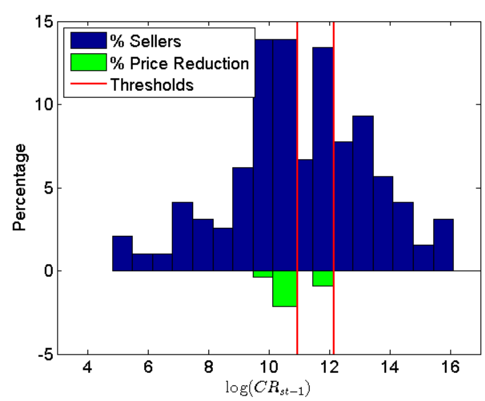

This figure shows the histogram of sellers' revenue and the distribution of percentage change in price with rebates relative to without. In the subfigure 1.4a, the x-axis is log transformation of yearly revenue in US dollars, the positive part of y-axis is the percentage of sellers at each revenue bin, and the negative part of y-axis is the average percentage change in price. The subfigure 1.4b and 1.4c displays the histogram of the cumulative revenue in different months along with the distribution of percentage change in price with rebates relative to without over the cumulative revenue level. The x-axis is the cumulative revenue at the beginning of each month, i.e. CR_{st-1} is the sum of revenue generated from January to the previous month.

1.10 The Functional Form of the Expected Per-period Profit

Denote $a = \sum_j ((1-\gamma)p_{jst} - c_{jst}) \exp(\delta_{jst})$, $b = 1 + \sum_k \exp(\delta_{kst})$, where $\delta_j = x_{jst}\alpha_x + \alpha_p p_{jst} + \Delta\xi_{st} + \xi_{jst}$, then per-period profit function could be written as

$$\pi_{st} = \frac{a}{b + \exp(\Gamma_t)} \tag{1.47}$$

Since $\Gamma_t \sim N(\mu_t, \sigma_t^2)$, where $\mu_t = \gamma_{0t} + \gamma_{1t}\Gamma_{t-1}$, $\exp(\Gamma_t) \sim logN(\mu_t, \sigma_t^2)$. Then $b + \exp(\Gamma_t)$ follows a shifted log-normal distribution with support $\exp(\Gamma_t) \in (-b, +\inf)$, $b + \exp(\Gamma_t) \sim logN(\widetilde{\mu}_t, \widetilde{\sigma}_t^2)$, where

$$\begin{aligned} E[\exp(\Gamma_t) + b] &= E[\exp(\Gamma_t)] + b & (1.48) \\ Var[\exp(\Gamma_t) + b] &= Var[\exp(\Gamma_t)] & (1.49) \end{aligned}$$

that is,

$$\begin{aligned} e^{\mu_t + \sigma_t^2/2} &= e^{\widetilde{\mu}_t + \widetilde{\sigma}_t^2/2} + b & (1.50) \\ (e^{\sigma_t^2} - 1)e^{2\mu_t + \sigma_t^2} &= (e^{\widetilde{\sigma}_t^2} - 1)e^{2\widetilde{\mu}_t + \widetilde{\sigma}_t^2} & (1.51) \end{aligned}$$

then

$$e^{\widetilde{\sigma}_t^2} = \frac{(e^{\sigma^2} - 1)e^{2\mu_t + \sigma_t^2}}{(e^{\mu_t + \sigma_t^2/2} + b)^2} + 1 \qquad e^{\widetilde{\mu}_t} = \frac{e^{\mu_t + \sigma_t^2} + b}{e^{\widetilde{\sigma}_t^2/2}} \tag{1.52}$$

Also, $\frac{a}{b+\exp(\Gamma_t)} \sim logN(-\widetilde{\mu}_t + \ln(a), \widetilde{\sigma}_t^2)$, then its expectation is

$$\begin{aligned} E\left[\frac{a}{b + \exp(\Gamma_t)}\right] &= \frac{a}{b + \exp(\Gamma_t)} \left(\frac{(e^{\sigma_t^2} - 1)e^{2\mu_t + \sigma_t^2}}{(e^{\mu_t + \sigma_t^2/2} + b)} + 1 \right) & (1.53) \\ &= \bar{\pi}_t \left((e^{\sigma_t^2} - 1)(1 - s_0 - \bar{s})^2 + 1 \right) = \bar{\pi}_t H_t & (1.54) \end{aligned}$$

where $\bar{\pi}_t = \sum_j (1-\gamma)p_j - c_j)\bar{s}_j$, $\bar{s} = \sum_j \bar{s}_j$, and $\bar{s}_j = \frac{\exp(\delta_j)}{1 + \sum_k \exp(\delta_k) + \exp(\mu_t + \sigma_t^2/2)}$

Chapter 2

Heterogeneous Effects of Online Reputation for Local and National Retailers

2.1 Introduction

The last decade has seen a rapid increase in e-retail as more and more individuals are choosing to make purchases through online platforms rather than brick-and-mortar stores. In the US alone, the percentage of total retail sales made online has risen from 2.2% to 6.7% since 2005.[1] However, this growth is accompanied by potential problems associated with information asymmetries due to the impersonal nature of the internet. Prior to their purchase decision, online shoppers may be unsure about the quality of a particular retailer in terms of the accuracy of the product descriptions, the seller's after-sale service, the speed of processing and shipping, etc. To lessen the impact of these asymmetries, most online markets have introduced peer review systems which rely on user supplied ratings and/or recommendations to provide consumers a signal of a seller's quality (i.e., reputation).[2] Survey and empirical evidence suggest that these rating systems are an important part of a consumer's decision-making process.[3]

At the same time, there exists heterogeneity in online sellers in terms of their prevalence in offline markets. For example, many third party sellers on Amazon Marketplace do not have offline stores, whereas Best Buy, Target, and Walmart have both an online store and a large offline presence. The goal of this paper is to study the heterogenous effect of online reputation across sellers who differ along this dimension. Specifically, we ask: does having a large offline presence make measures of online reputation less impor-

[1] See `http://ycharts.com/indicators/ecommerce_sales_as_percent_retail_sales`.

[2] For example, Amazon.com displays the distribution of seller and product ratings (1 through 5) given by previous shoppers, while ebay.com has a feedback system in which users indicate whether their experience was positive or negative.

[3] For survey evidence see `http://marketingland.com/survey-customers-more-frustrated-by-how-long-it-takes-to-resolve-a-customer-service-issue-than-the-resolution-38756`. Both Chevalier and Mayzlin (2006) and Dellarocas (2003) provide empirical evidence.

tant and if so, by how much? On the one hand, a national offline presence implies that consumers are familiar with a particular retailer from shopping at its brick-and-mortar store, which may lower the impact of measures of online reputation. On the other hand, consumers may see the online shopping experience as separate from the offline experience, and use reputation in their decision-making process.

By answering this, we inform platforms and policy makers as to the efficiency of such systems in various contexts. For example, having a rating system for well-established national retailers may be wasteful in terms of the retailer's costs and/or the customer's time. Additionally, identifying the impact of reputation for different types of sellers may allow businesses to tailor their efforts towards or away from improving their online reputation. Finally, as retailers may adjust their pricing strategies in order to manage their online reputation, our results can generally speak to how differential effects of this reputation may impact competition.

To accomplish our goal, we quantify the importance of the online rating system for sellers on Tmall, a branch of China's largest e-commerce company, Alibaba. Tmall is a business-to-consumer (B2C) platform which features thousands of professional sellers offering a wide variety of different products. The platform is a leader in China's online B2C market, as it had 54 percent market share and total transactions reached $39 billion in Q3 2015.[4] Tmall features a rating system that is similar to that of Amazon, where each customer purchasing a product from a given seller rates the quality of the transaction on a scale of 1 to 5. Customers who arrive thereafter can observe both the average rating and the total number of ratings for a given seller.[5]

Our data include monthly prices and total monthly quantities for tablets sold on Tmall between September, 2014 and April, 2015. We observe both product and seller characteristics, where the latter include the rating score (i.e., average rating) and the complete distribution of ratings. Additionally, we observe a classification of sellers, as defined by Tmall, which is based on a seller's offline presence. The classification allows us to separate retailers into two groups which we called 'types': those who are likely to have a large offline presence and be well-known nationally and those who are local or online only retailers. Our aim is to estimate the impact of the rating system across these two types.

We do this by specifying a discrete choice model of demand in which a consumer chooses to buy a tablet from one of the sellers on Tmall. Before making her purchase decision, the consumer forms a belief about seller quality using the average rating score and the number of ratings. In order to provide evidence that consumers are using ratings in their purchase decision, we first assume that this belief formation can be represented by a linear function of the average rating, the total number of ratings, and an interaction between the two. Importantly, we allow for the effect of these variables to vary across local and national sellers. When accounting for product characteristics and utilizing the covariation between ratings and sales within a given seller, we find evidence that ratings significantly impact the sales of both types, but the impact is the higher for local sellers.

Next, we enrich the demand model by replacing the linear function of ratings with a structural learning model, allowing us to estimate the learning parameters that determine

[4]See http://www.chinainternetwatch.com/15957/chinas-b2c-sales-q3-2015/.

[5]See Figures 2.1 and 2.2 for screen shots of a search page and a seller's home page.

the relationship between sales and ratings. Specifically, we assume that consumers use the average rating and the number of ratings to infer the expected quality of a seller via a Bayesian learning process and we allow the updating process to be functions of the seller's type and other seller characteristics.

As is common in learning models, consumers form their belief through knowledge of how the learning signals are generated. We assume that ratings are generated from an independent private draw of a 'shopping experience' for each transaction, where the distribution of experiences is centered around the true seller quality, and the variance in these experiences depends on the seller's type. This allows for the informativeness of the rating score to depend on whether the seller is a national or a local seller. The independence of signals implies that, while consumers can learn about the seller's true quality through ratings, each rating is not a function of the previous ratings (i.e., it is not state dependent). This simplifies the estimation procedure substantially, as it implies that consumers (and the econometrician) do not have to integrate overall the possible combinations of ratings that would have led to the observed average as in Newberry (2016). Instead, the consumer sees the rating score as an average of independent draws from the signal distribution and calculates the updated belief based on these draws. Along with simplifying estimation, we believe that it is reasonable to believe that consumers behave in this manner.

We estimate the demand parameters and the parameters of the learning process via a two-step procedure that combines a standard model of demand for differentiated products and the structural learning process. The first step identifies the preferences over observable product characteristics, the price sensitivity, and the expected quality of the seller (i.e., a seller/month fixed effect). In the second step, we use a Bayesian learning process in order to specify expected seller quality as a function of a seller-level prior, the variance in signals, the average rating, and the number of ratings. We then estimate the primitives of the learning model by relating the estimated seller quality to the observed ratings.

The primary learning parameters of interest are the impact of seller-type on the prior mean and the weight placed on learning signals (i.e., the ratio of the prior variance and signal variance). The results indicate that the mean of the prior belief varies across the two types of sellers, but the difference is not large. However, we find that the weight placed on the rating score is significantly higher for local sellers. This result comes from two possible sources. First, there could be a larger variance in the ex ante beliefs about the quality of a local seller as compared to a national seller. Second, the variation of shopping experiences (i.e., signals) and, hence, ratings could be larger for local sellers. While our model cannot separate these stories directly, we show using the raw data that the latter source is more likely, suggesting that consumers are more certain about the ex ante quality of national sellers.

Using the estimates of the model, we quantify the heterogeneous effect of online reputation through two simulation exercises. In the first exercise, we estimate the comparative 'value' of the rating system between local and national sellers. We do this by calculating the amount of revenue which is lost (or gained) when the ratings are removed for a seller assuming it is a local seller (i.e., takes on the 'local' value of the learning parameter), and then comparing it to the same calculation assuming the seller is a national seller.

Results suggest that local sellers value the ratings system substantially more than they would if they were national sellers. Specifically, the median difference in value for sellers who benefit from ratings is over $4,000 per month and 75th percentiles are over $28,000 and $500, respectively. Much of this variation is coming from the fact that there are big sellers and small sellers. Therefore, to put this in perspective, we calculate the difference in value as a percentage of a seller's monthly sales and find the average to be 45% and the median to be 48%.

For sellers who are hurt by ratings, which is nearly a quarter of them, the median difference in value is just above -$4,000 implying that local sellers are willing to pay $4,000 more to have their ratings hidden than national sellers. The average difference in the value as a percentage of monthly revenue is -215% and the median is -77%.

In the second simulation exercise, we measure and compare the marginal impact of the rating score across seller types. To do this, we artificially increase or decrease a seller's average rating by setting it to the average score observed in the data. We then calculate the associated change in monthly revenue assuming a seller is a local seller and compare it to the change in revenue assuming the seller is a national seller. We find that the median seller who has rating scores above the average loses over $2,000 per month more as a local seller than it would have as a national seller. The results are similar for sellers who have a rating score below the average, as the median seller gains just under $2,000 more per month. Finally, the average difference in the value as a percentage of monthly revenue is 68% for sellers who have a ratings score above the mean and -87% for sellers who have a ratings score below the mean.

The source for the results in both exercises is the fact that consumers place more weight on the prior mean for national sellers then they do for local sellers. Because of this, removing or changing the rating score has a comparatively lower impact on the ex ante quality belief for national sellers and, hence, a lower impact on revenue.

The results of these exercises suggest that, while both types of sellers value the rating system, the local sellers value it by more. Therefore, depending on the costs to implement such a policy, it may be beneficial for the platform and consumers to have a type-specific reputation system. For example, they may want to highlight ratings for local sellers, while highlighting other information such as prices and product characteristics for national sellers. Additionally, the results imply that local sellers may have more of an incentive to manage their online reputation through pricing policies and/or improvements in customer service, compared to national sellers. Finally, the analysis indicates that ratings play more of a role for local sellers in determining degree of differentiation between sellers, possibly resulting in different levels of competition on this platform.

This paper is related to several different strands of literature. First, there are a number of papers which estimate the impact of online reviews and reputation. Chevalier and Mayzlin (2006) estimate the effect of reviews at the product level using data from book sales on Amazon, while Dellarocas (2003), Cabral and Hortacsu (2010), Nosko and Tadelis (2015) and Saeedi and Sundaresan (2016) study the impact of seller reputation on eBay. Zhu and Zhang (2006) and Su et al. (2016) do the same but focus on Taobao. Zhu and Zhang (2010) estimate the impact of online reviews in the video game industry and find that the effect of reviews varies across games that differ in their characteristics.

Elfenbein et al. (2015), Hollenbeck (2016), and Luca (2011) have a similar goal to this paper as they examine the effects of reviews across sellers who differ by their experience or chain affiliation. Elfenbein et al. (2015) finds that reputation effects are stronger for eBay sellers who are relatively inexperienced, and Hollenbeck (2016) and Luca (2011) find that online ratings are more impactful for independent restaurants and hotels relative to chains.

Our paper extends the above by studying the role of an offline presence in determining online sales, rather than focusing on the impact of online reputation on offline sales. As well known national retailers such as Target and Wal-Mart continue to focus efforts on increasing their online presence, our paper can generally speak to the possible role of their offline presence in determining the outcomes of this effort. Additionally, we connect ratings and reputation to demand through a structural learning model, allowing us to estimate a rich set of learning parameters and, therefore, quantify the differential impact of ratings. Similarly, Zhao et al. (2013) estimate a structural learning model where expected quality of a book's genre is based on user ratings, but do so for a single seller. We study how differences in ratings affect substitution across heterogenous sellers.

Finally, our paper is related to the large amount of industrial organization literature studying demand for differentiated products that follow Berry (1994) and Steven T. Berry and Pakes (1995). Specifically, we add structure to the product quality term, allowing it to be formed based on a Bayesian learning process where consumers update their belief using the online reputation system. Grennan and Town (2015) model learning in the medical device industry in a similar manner, but do not estimate the full set of learning parameters.

The rest of the paper is organized as follows. Section 2 discusses the relevant institutional details and introduces the data, while Section 3 specifies the model demand with learning. The estimation of the model is in Section 4, the results are presented in Section 5, and Section 6 concludes.

2.2 Data

Alibaba Group Holding Limited, launched in 1999, is the leading online commerce provider in China. While the company offers a broad spectrum of e-commerce services, the two primary e-commerce platforms are taobao.com and tmall.com. Taobao is similar to eBay, as it is an open platform connecting individual sellers and buyers, also known as a consumer-to-consumer (C2C) market. On the other hand, Tmall specializes in providing a marketplace for professional sellers and enterprisers, meaning it resembles Amazon Marketplace, a business-to-consumer (B2C) market. These two platforms sell a wide range of products, including apparel, shoes, books, electronics, smart phones and televisions. Their sales account for nearly 80% of China's total online shopping market, and their gross merchandise volume is the largest in the world.[6] Total sales reached 2950 billion RMB (around $466 billion) in 2015, with $285.78 billion on Taobao and $180.25

[6] Statistics can be found at `http://www.alibabagroup.com/en/ir/glance`.

billion on Tmall.[7] For comparison purposes, eBay's sales from 2015 were \$81.7 billion.[8]

Much like other e-commerce platforms, Tmall has a system that allows consumers to rate their experience with a given seller/product. After each transaction, the buyer is required to post a rating score for three aspects of the experience: (1) the seller's descriptions on goods, (2) the customer service, and (3) the shipping service. Scores range from 1 to 5, with 1 being the lowest score in terms of satisfaction. The average rating score is calculated based on the posted ratings for all transactions that occurred in the last six months.

In order to observe the average rating along with the total number of ratings, an individual must either display the search results 'by seller' by clicking a link on the default search page, navigate to a page that displays all sellers selling a particular product, or go to an individual seller's page. In the former two scenarios, the consumer would observe the information for many sellers at once, while in the latter, she would only observe the information for one seller. An example of a search result sorted by seller is shown in Figure
refChX-figure:search. The top listed seller, Kindle Official Store, is rated 4.9, 4.8, and 4.8 for their description, customer service, and shipping services, respectively, and the first Kindle listed has a total of 8,253 ratings.

If a customer clicks to a specific seller's page, she can observe more detailed records about a seller's reputation. Specifically, the distribution of ratings given to the seller in last six months and the total number of ratings for that store are displayed. Figure 2.2 shows that the Kindle Official Store's customer service has been rated by 22,887 consumers and that 94.11% (21,539) of the ratings are the highest and 0.52% of the ratings are the lowest.

Another important feature of Tmall is that it employs a classification system of sellers. Sellers are separated into three different groups that are based on the authorization contract with the product's manufacturer and the scale of the business. Group-one sellers are the legal representatives of a registered trademark, which could be a product brand such as Apple (iPad), Amazon (Kindle), or Samsung (Galaxy Pad). Additionally, some large retailers with a registered store brand (e.g., Best Buy) could be invited to be a group-one seller.[9] Group-two sellers are enterprises which are authorized by the brand owners to sell their products, meaning they typically sell only one brand. Finally, for the most part, group-three sellers are authorized retailers who sell multiple brands that they purchase directly from suppliers.

Table 2.1 displays information on the number of offline stores, the number of employees, and the existence of a website separate from Tmall for the top five sellers of each group.[10] Note that some of this information was not available for every group. The

[7]Information is from Alibaba's financial report, downloaded from `http://www.alibabagroup.com/en/ir/financial`.

[8]Information is from eBay's annual financial report downloaded from `https://investors.ebayinc.com/annuals.cfm`.

[9]We use retailers popular in the United States as examples to fix ideas. These retailers do not necessarily operate on Tmall.

[10]This information was gathered by searching for each seller's official website where, often times, the number of offline retail outlets are displayed. When the seller did not have an official website, we search for the store in a business directory that lists the number of employees.

primary difference between group-one sellers and the other two groups is in the store's brand recognition. Group-one sellers generally have a large offline presence and have their own website, meaning they are likely well-known to customers throughout China. Thus, we call them national sellers. These measures imply that group-two and three sellers are similar in size and usually lack a website. Because of this, we join group-two and three into one group, which we call local sellers. These sellers generally have either a small local presence or are retailers who only sell online.

2.2.1 Descriptive Statistics

The data include all sales of tablets which occurred between September, 2014 to April, 2015. We collect the data by scraping the Tmall website on the 11th or 12th of each month, depending on how many days occur in that month. For each seller who sells tablets on Tmall, we observe the quantity sold in the last month and the price at the time of the scrape for all the products which it sells. We formulate the price by taking the average of the two prices: the one observed last month and the one observed this month. If the product was not sold by the seller last month, we use only this month's price. We also observe the value of the three different rating scores and the number of ratings at the time of the scrape for each seller. Finally, we observe the following product characteristics for each product: brand, operating system, screen size, memory, cellular internet capability, and the number of years since the product's introduction (i.e., product age).

The market for tablets on Tmall is described in Table 2.2. The first two columns display statistics for all the sellers offering tablets, while the last two columns show only sellers who are identified as electronics retailers. The non-electronic retailers are general merchandise stores selling goods in many different product categories (i.e., a retailer like Wal-Mart). Moving forward, we focus our analysis on the electronic retailers and assume that buying from a general merchandiser is the 'outside good'. This allows us to limit the amount of heterogeneity across sellers to a certain extent.

For the four month sample period, there were nearly 671 different sellers, with a majority of those (562) being electronics retailers. The average number of electronics sellers per month is about 391. There were 94 total local sellers and 62 national sellers within the electronics category, making up 14% and 11% of sellers. The table also shows that there are many different brands and models sold on Tmall. Finally, Tmall is a very large seller of tablets with nearly 84 thousand units sold and $21 million in revenue per month for electronic retailers alone.

Next, we present seller-level descriptive statistics in Table 2.3. The average quantity sold and revenue over the eight month time period are 1,430 and $342,970, but these distributions are heavily skewed with the medians being only 79 and $15,894. The median seller sells only one brand and four different models. The average price of tablets on Tmall is about $270. Overall, the market is extremely competitive with the average market share being less than 1% and an HHI of 0.02. The rating scores on Tmall are all quite high, as the means of all three scores are very close to 4.75. There is little variation in scores across stores, with the standard deviation of the rating being around 0.10 for all three scores. Not displayed in the table is the correlation across the three ratings, which is 0.92 for the product description and the customer service rating, 0.94 for the product description and the delivery rating, and 0.97 for the customer service and the delivery rating.

Table 2.4 shows the same statistics, but breaks them down by the two different seller-types. For the most part, the sellers appear to be similar along most dimensions. One exception is that national sellers are larger in terms of the number of tablets sold and revenue, which also leads to more ratings. The average and variation in rating scores do not differ much across seller-type.

Table 2.5 presents the average characteristics for the tablets sold by the electronics retailers on Tmall. The last column shows the information across all seller-types, while the others display these statistics within a seller-type. While some brands familiar in the United States are also popular on Tmall (e.g., Apple is #4 and Samsung is #9), the top three brands are Teclast, Onda, and Cube. Android and Windows operating systems appear on nearly 61% and 33% of tablets, respectively, while iOS is on about 10% of the tablets sold. Note that there are some tablets which have multiple operating systems available, resulting in the total across the three major ones being over 100%. Also, there are a few lesser-known operating systems which appear in the data, but account for a very small percentage of the market. About 26% of the tablets have internet access through cellular networks, the average screen size is 8.5 inches, and the average memory is just under 7 GB. While this seems like a low level of memory, many of the tablets have the ability to add storage. Finally, the average tablet is sold about 7 months after its initial release date.

The first two columns look at the heterogeneity in the tablets sold across the seller-types. Again, sellers do not seem to differ much in the characteristics of the tablets they offer, but do seem to differ somewhat in the brands that they sell. Specifically, local sellers sell more Apple products than national sellers. Despite this, the descriptive statistics mostly suggest that the primary heterogeneity between national and local sellers is in their offline presence. Therefore, we assume that the seller-type reflects the degree of the national offline presence for the majority of our analysis, but do allow for other possible dimensions of heterogeneity to affect consumer demand.

2.3 Model

We specify a model of demand that incorporates consumer learning from seller ratings. Before making her purchase decision, we assume that the consumer observes both the rating score and the total number of ratings for each seller. While there are three rating scores displayed for each seller, we assume that the consumer uses the product description score to form her belief about seller quality. As mentioned in Section 2.2.1, the three scores are highly correlated, meaning forming a belief based on one score is approximately equivalent to forming a belief based on all three. However, the model could be extended to include all three scores without much difficulty. While the rating scores don't appear on the default search page, they do appear when an individual sorts by seller within the search results. Additionally, if a consumer is looking for a particular tablet,they will navigate to a page displaying all the sellers selling that tablet, along with the rating score information. Finally, once a consumer navigates to a seller's page before purchasing the tablet, the information on the entire distribution of ratings is visible.

While the total number of ratings is only visible once the consumer clicks to the seller's home page, the number of ratings for a seller/tablet pair are visible on the search page sorted by seller. To the extent that we see the learning process as an approximation of the observed behavior, it is reasonable to assume that consumers can infer the accuracy of the average rating from the total number of ratings for the product(s) which appear on the search page. However, taken literally, the assumption implies that consumers are made aware of the total number of ratings by clicking to each of the seller's home pages. Finally, we assume that consumers know the ratings information for all sellers. One may worry that any variation in the effect of ratings we estimate may be due to the fact that potential customers do not observe the information for all sellers. However, as long as the observable information doesn't systematically vary across seller-type, then we will still be able to capture the heterogeneity we are interested in.

2.3.1 Demand Model

We assume that M_t consumers arrive in month t and choose to purchase a tablet from one of the electronics retailers on Tmall or to purchase from a seller who is not an electronics retailer (option 0). Consumer i's utility of purchasing product j from store s in month t is:

$$U_{ijst} = \beta X_{jt} + \xi_s + \nu_{jst} + \epsilon_{ijst} \tag{2.1}$$

where X_{jt} is a vector of product characteristics including the price of the tablet. The quality of seller s is ξ_s, ν_{jst} is a product-level demand shock, and ϵ_{ijst} is an idiosyncratic taste shock assumed to be iid extreme value. We assume that ξ_s is fixed, meaning stores can not adjust their quality over our sample period. As it is only 8 months, it is reasonable to assume that it is difficult to adjust quality in the short run. Finaly, we normalize the mean utility of the outside option to 0:

$$U_{i0t} = \epsilon_{i0t} \tag{2.2}$$

Consumers face uncertainty about the quality of a given seller prior to purchase, implying that they make purchase decisions based on expected utility:

$$E[U_{ijst}] = \beta X_{jt} + \underbrace{E[\xi_s | r_{st}, n_{st}]}_{\zeta_{st}} + \nu_{jst} + \epsilon_{ijst} = \beta X_{jt} + \zeta_{st} + \nu_{jst} + \epsilon_{ijst} \tag{2.3}$$

where ζ_{st} is the expected store quality given the average rating, r_{st}, the number of ratings, n_{st}, at the beginning of month t. The demand shock for product j, ν_{jst}, is known to the consumers prior to their purchase decision, and it captures the variation in demand across products within the given seller. This term may represent variation in the expected quality across products for a given seller/month, but we assume that it is not a function of the seller-level ratings.

2.3.2 Evidence of Learning

To provide evidence of learning, we first assume a reduced form measure of consumer beliefs. That is, beliefs are formed through a linear function of the number of ratings, the average rating and an interactions between the two along with a seller-level fixed effect. Specifically, we assume that:

$$\zeta_{st} = \sigma_1^{c_s} n_{st} + \sigma_2^{c_s} r_{st} + \sigma_3^{c_j} n_{st} \times r_{st} + \tilde{\xi}_s$$

where $\tilde{\xi}_s$ is considered the ex ante expected quality of seller s when there are no ratings, or in other words, the 'prior' belief about the quality of the seller. Additionally, we allow the demand shock to include a month fixed effect:

$$\nu_{jst} = \tilde{\xi}_t + \tilde{\xi}_{jst}$$

The function by which consumers form beliefs varies by a seller's type, $c_s \in \{L, N\}$ where L and N indicate local and national. Given these assumptions, we have the familiar logit estimating equation:

$$log(s_{jst}) - log(s_{0t}) = \beta X_{jt} + \sigma_1^{c_s} n_{st} + \sigma_2^{c_s} r_{st} + \sigma_3^{c_s} n_{st} \times r_{st} + \tilde{\xi}_s + \tilde{\xi}_t + \tilde{\xi}_{jst}$$

We present results for four different specifications in Table 2.6, where each specification varies by the restrictions we place on the way consumers form their beliefs.[11] Note that in the presented results, we assume that prices are exogenous conditional on the covariates and the fixed effects.[12]

Specifications (1) and (2) do not include the interaction effect between the number of ratings (i.e., $\sigma_3^{c_s} = 0$), while specifications (1) and (3) restrict the parameters to be the same across seller-type. While we focus on the effect of ratings, we note that all specifications have negative and significant price effects. Additionally, we omit the estimates of preferences over product characteristics for these regressions, but will analyze them in results of model presented in Section 2.5. Finally, the number of ratings is measured in thousands.

In specifications (1) and (2), we find that the average rating has a positive and significant impact on sales and that this effect exists for both national and local sellers. However, interestingly, we find that the effect of the rating score is larger for national sellers, which goes against the hypothesis that consumers use ratings more in forming their beliefs about the quality of local sellers. Specifications (3) and (4) account for the fact that as the number of ratings increase, the average rating becomes more informative, something that is more in line with how consumer learning models are specified. The results of these specifications again show that there is a positive effect of the average rating and that the effect is larger for national sellers compared to local sellers. However, the coefficient on the interaction term is positive and significant for local sellers and

[11]Products/sellers who have 0 sales in a month are dropped (i.e., included in the outside option). We have run the same regressions correcting for market shares as in Gandhi et al (2013), with similar results.

[12] We have also explored the use of instrumental variables with little change in the results. Specifically, we have used both the stock of tablets for a given seller and the product characteristics for competing sellers as instruments.

insignificant for national sellers. Using the median number of ratings for local sellers (around 840), the marginal effect of the average rating is significantly higher for local sellers (0.43) compared to national sellers. Further, for an n low as 250, the marginal effect for local sellers is higher than the marginal effect for national sellers.

Overall, these results provide evidence that consumers are learning from ratings. However, the results are mixed on whether or not this effect is significantly different between national and local sellers. There are a few weakness of this exercise which we address by imbedding a structural model of consumer learning into the demand model. First, the effect of the 'prior' is fixed in that it is not allowed to vary as the number of ratings changes. It is common in learning models to assume that the effect of the prior (or the weight on the prior) may become smaller as consumers receive more signals. Second, the linearity of the learning function is restrictive in that it doesn't allow for interpretation of coefficients to be weights placed on the signal(s) and the prior. Finally, and similar to the above points, these specifications don't allow for an estimation of what are thought of as 'structural' learning parameters or the deep parameters of the learning process such as the prior beliefs of seller quality and the distribution of learning signals.

2.3.3 Baysian Learning Model

In order to specify the learning model, we first discuss how ratings are generated. After purchasing a tablet, each customer is required to rate the transaction.[13] We assume that a customer posts a rating based on her shopping 'experience', which is generated from a distribution centered around the true seller quality. It is important to emphasize the fact that the reviewer does not intend to reveal her *belief* about the quality of the seller, but rather she wishes to share the quality of her *individual shopping experience*. Under this assumption, the distribution of reviews is not state-dependent.

Not only does this assumption make the model more tractable, but we also believe it is realistic in online markets which feature review systems. That is, suppose that a customer has a very bad experience with a seller. It is reasonable to believe that she will give a very bad rating irrespective of the previous rating score. To further justify this assumption, a reviewer cannot observe the ratings of her peers or the average rating while she is submitting her own, making it less likely that the number she gives is state-dependent.

Specifically, we assume that shopping experiences arrive as a random signal measured in utility that is distributed normally around the true store-level quality:

$$e_{\tilde{i}jst} \sim N(\xi_s, \widetilde{\sigma}_s^2) \tag{2.4}$$

where the subscript $\tilde{i}$ identifies consumers who purchased a good from seller j before the arrival of consumer i. Upon receiving this signal, the consumer posts a rating, which is a transformation of the experience into a number which is in line with the rating system:

$$r_{\tilde{i}jst} = \rho e_{\tilde{i}jst} + y$$

where the constant shifter y and the scale term ρ play the role to transform the distribution of the signals, which is measured in utility, into the distribution of ratings.

[13] If a rating is not posted, a 5 is automatically given. We account for this in a robustness check.

The platform displays the average of these ratings across the products for seller s as the rating score, r_{st}, along with the total number of ratings for that seller, n_{st}.

Here, we assume that the posted ratings and, hence, the rating score are both continuous and have infinite support, something that is at odds with the actual setting. While we make these assumptions in order to simplify the updating procedure, we believe that this model serves as a good approximation of the process in this application. We provide further discussion of this model as an approximation after presenting the updating procedure.

Upon arrival, a consumer observes the current rating score, the number of ratings, and the seller's identity. Based the sellers identify, the consumer believes that the seller's quality is drawn from a normal distribution:

$$\xi_s \sim N(\mu_s, \sigma_s^2) \tag{2.5}$$

which is referred to as the prior belief. The size of the variance in this distribution represents the level of uncertainty that the consumer has about the store's quality based purely on that store's identity. For example, a consumer may observe that the seller is walmart.com and be very certain about the seller's quality (i.e., a low value of σ), or she may observe a local seller and be uncertain about the quality (i.e., a high value of σ).

Then, using the average rating and number of ratings, the consumer updates her beliefs about the expected quality of the seller using Bayes' rule:

$$E[\xi_s | r_{st}, n_{st}] = \frac{\widetilde{\sigma}_s^2}{\widetilde{\sigma}_s^2 + n_{st}\sigma_s^2}\mu_s + \frac{n_{st}\sigma_s^2}{\widetilde{\sigma}_s^2 + n_{st}\sigma_s^2}(\frac{r_{st} - y}{\rho}) \tag{2.6}$$

This is a weighted average of the shifted rating score and the prior mean. The weights are based on the variance in the prior belief, the variance in the distribution of signals and the total number of signals which generated the rating score. Intuitively, the posterior is heavily weighted towards the prior when there are no ratings and converges to the shifted rating score as the number of ratings gets very large. Also note that because we assume that the experience signals are distributed around the true quality of seller s, the average rating, and thus, the belief of store quality, converges to the true store quality as the number of ratings gets large.

Importantly, we allow for observable heterogeneity in the learning process across sellers. That is, we specify the learning parameters as a function of the seller's type, c_s, and other seller covariates, Z_s. The mean of the prior and the ratio of the prior variance to the signal variance are written as function of these observables:

$$\mu_s = \sum_{k=L,N} \mu_k \mathbf{1}(c_s = k) + \mu_z Z_s \tag{2.7}$$

$$\frac{\sigma_s^2}{\tilde{\sigma}_2^2} = \bar{\sigma}_s^2 = exp\left(\sum_{k=L,N} \sigma_k \mathbf{1}(c_s = k) + \sigma_z Z_s\right) \tag{2.8}$$

where the vectors μ and σ are parameters to be estimated. The reason we parameterize the ratio of the variances rather than each of them separately is because we are not able to identify the ex ante variance of quality separately from the distribution of signals using only the demand model. As we discuss further in Section 2.4, we would need to use either the sellers with zero sales or the distribution of the individual ratings within a seller in order to separately identify these objects. However, the ratio is a sufficient measure of the learning process for our purposes, as it is a measure of the relative weight individuals place on the prior versus the rating score. A relatively large weight on the prior is either because of a relatively small prior variance or a relatively large signal variance. In either case, this large weight implies that the learning signals are not very informative, as either the consumers are confident in their ex ante belief on seller quality and/or the experience signals do not provide much information.

Because of the linearity in consumer preferences, the expected utility from buying product j from seller s is:

$$\begin{aligned} E[U_{ijst}] &= \beta X_{jt} + \zeta_{st} + \nu_{jst} + \epsilon_{ijst} \\ &= \beta X_{jt} + \frac{1}{1+n_{st}\bar{\sigma}_s^2}\mu_s + \frac{n_{st}\bar{\sigma}_s^2}{1+n_{st}\bar{\sigma}_s^2}(\frac{r_{st}-y}{\rho}) + \nu_{jst} + \epsilon_{ijst} \end{aligned} \tag{2.9}$$

Given this information, consumer i chooses to purchase the product/seller which gives her the highest expected utility.

2.3.4 Rating Model

The primary reason for assuming that ratings come from a continuous and infinitive support is that it simplifies the updating procedure substantially. That is, suppose we assumed a model of how ratings are generated in which a consumer receives the experience signal, $e_{\tilde{i}jst}$, and then posts a rating based on a series of thresholds:

$$r_{\tilde{i}jst} = \begin{cases} 1, & \text{if } \rho e_{\tilde{i}jst} + y \in [-\infty, G_1) \\ 2, & \text{if } \rho e_{\tilde{i}jst} + y \in [G_1, G_2) \\ 3, & \text{if } \rho e_{\tilde{i}jst} + y \in [G_2, G_3) \\ 4, & \text{if } \rho e_{\tilde{i}jst} + y \in [G_3, G_4) \\ 5, & \text{if } \rho e_{\tilde{i}jst} + y \in [G_4, \infty) \end{cases} \tag{2.10}$$

After observing the average rating and the number of ratings, consumer i calculates the expectation of the $e_{\tilde{i}jst}$ which replaces $(\frac{r_{st}-y}{\rho})$ in equation 2.6. If the consumer only observes the total number of ratings, then this calculation becomes computationally intractable. For example, suppose the consumer observes $n_{st} = 100$ and $r_{st} = 4.5$. She would have to integrate over all the possible combinations of the 100 ratings that would lead to an average rating score of 4.5 in order to calculate the expected value of the signals. Considering the fact that most sellers have n_{st} in the thousands, calculating beliefs under this of assumption is infeasible. Under our assumption of a continuous and infinite distribution of ratings, this calculation is not necessary.[14]

[14]Another option is to use the discrete model and assume that consumer i observes the complete distribution of ratings, something that is observable on the seller's page, which simplifies the problem

Additionally, we believe that our model serves as a reasonable approximation of belief formation in this market for two reasons. First, the scale parameter, which will be estimated, serves to adjust the variance of ratings such that a realization of a rating outside of the range of 1 to 5 occurs with low probability. Second, with only a small number of ratings (e.g., 100), the average rating from the discrete model above approximates the average rating from the continuous model quite closely. The median seller in our sample has over 800 ratings.

2.4 Estimation

The parameters of interest are separated into the preference parameters, β, and the learning parameters, $\theta = [\mu_c, \mu_z, \sigma_c, \sigma_z]$. We estimate these parameters via a two-step procedure. In the first step, we obtain consistent estimates of β and the expected quality, ζ_{st}. We then map this expected quality into our learning model and estimate θ via non-linear least squares.

The estimation of β is standard. Given the assumptions of the model, the probability that consumer i purchases product j from seller s is:

$$P_{ijst} = \frac{exp(\beta X_{jt} + \zeta_{st} + \nu_{jst})}{1 + \sum_{j,s} exp(\beta X_{jt} + \zeta_{st} + \nu_{jst})}$$

We do not assume any heterogeneity in preferences, meaning this maps directly into the market share of each product:

$$s_{jst} = \frac{exp(\beta X_{jt} + \zeta_{st} + \nu_{jst})}{1 + \sum_{j,s} exp(\beta X_{jt} + \zeta_{st} + \nu_{jst})}$$

Dividing by the outside share and taking logs results in the linear estimating equation:

$$log(s_{jst}) - log(s_{0t}) = \beta X_{jt} + \zeta_{st} + \nu_{jst}$$

We estimate β via OLS with seller/month fixed effects. Importantly, the fixed effects allow us to consistently estimate the preference parameters as we can account for any unobserved seller-level quality which may affect prices and/or ratings. In practice, we estimate the model based on first differences between products sold by the same seller in order to difference out the seller/month effect. We then calculate the value of the structural error term:

$$\tilde{\xi}_{jst} = log(s_{jst}) - log(s_{0t}) - \hat{\beta} X_{jt}$$

---

substantially. That is, the expectation of the signal becomes:

$$E[e_{\tilde{i}jst} | \{n_{cst}\}, g \in [1,2,3,4,5]] = \sum_{g \in [1,2,3,4,5]} [e_{\tilde{i}jst} | \rho e_{\tilde{i}jst} + y \in [G_{m-1}, G_m)] \times \frac{n_{cst}}{n_{st}}$$

where n_{cst} indicates the number of ratings which are equal to g. This is a straightforward calculation because the assumed distribution of signals is normal. We have estimated the model under this structure, with little qualitative change in results.

and take the mean across the J products sold by seller s:

$$\hat{\zeta}_{st} = \frac{1}{J} \sum_{j} \tilde{\xi}_{jst}$$

to obtain the estimated expected quality. We assume that this is composed of the updated belief of seller quality based on the ratings information and a mean zero random shock:

$$\hat{\zeta}_{st} = \zeta_{st} + \varepsilon_{st}$$

where ε_{st} is iid across sellers and time. This assumption implies that any variation in the estimated expected quality which is not captured by the observables in the learning process is exogenous.

For a vector of parameters, we can solve for the quality belief at every (r_{st}, n_{st}):

$$\zeta_{st}(r_{st}, n_{st}|\theta) = \frac{1}{1 + n_{st}\bar{\sigma}_s^2}\mu_s + \frac{n_{st}\bar{\sigma}_s^2}{1 + n_{st}\bar{\sigma}_s^2}(\frac{r_{st} - y}{\rho}) \tag{2.11}$$

which implies that:

$$\hat{\zeta}_{st} = \zeta_{st}(\theta) + \varepsilon_{st}$$

at the true values of the parameters. Therefore, we estimate the learning parameters via non-linear least squares. That is, we find the parameters that minimize the sum of squared errors:

$$\hat{\theta} = \arg\min_{\theta} \sum_{s,t} \left(\hat{\zeta}_{st} - \zeta_{st}(\theta)\right)^2 \tag{2.12}$$

2.4.1 Identification

The taste parameters are identified by variation in market shares based on variation in product characteristics and/or price across products within the same seller and month. The identifying assumption in estimating these parameters is that the seller/product/month level shock is not correlated with prices and/or ratings. As we have included a seller/month fixed effect and a complete set of product characteristics, we believe this is a reasonable assumption. However, as in the linear model presented earlier, we have also completed the analysis taking an instrumental variables approach with little change in the results.

The learning parameters are then identified by variation in the estimated $\hat{\xi}$ across sellers and time. The prior mean, μ, is a seller-level constant, as it shifts ξ up or down conditional on the average rating and number of ratings. Due to computational burden, we do not estimate a prior for each seller. Instead, we allow it to depend on seller-type and the seller observables contained in Z. Therefore, covariation in demand and these observables holding the number of ratings and the rating score fixed identifies the μ_c and μ_z. The weighting parameters, σ, are identified by the variation in demand as a result of a change in the rating score and the number of ratings. The key identifying assumption to estimate the learning parameters is that the shock, ε_{st}, is independent of ratings and other seller-level observables. This assumption is valid if ε_{st} represents measurement error and/or pure randomness in the estimated $\hat{\zeta}_{st}$. Finally, the transformation terms,

y and ρ, are identified by the relationship between the distribution of rating scores and the distribution of $\hat{\xi}$.

As previously mentioned, we do not identify the signal variances, $\tilde{\sigma}$, separately from the prior variance, σ. These parameters would be separately identified by utilizing information for the sellers who have zero ratings and/or using the distribution of ratings within an individual seller. We choose not to utilize these pieces of information because there are very few sellers in our data with zero ratings and because utilizing the within-seller distribution of ratings would require richer model of how ratings are generated. Although a decomposition of $\bar{\sigma}$ is an interesting exercise, it is not necessary to accomplish the goals of this research. Specifically, $\bar{\sigma}$ is informative of the weight individuals place on the prior versus the signals, meaning it is a measure of how much people use the ratings in their decision-making. This is exactly the behavior we are trying to identify. However, we use the raw data to provide some evidence as to which of the factors is more important in driving the results.

2.5 Results

We present the results of the estimates in Tables 2.7 and 2.8.[15] The preference estimates are all significant and have the predicted sign. People value screen size, memory, and access to the internet, in addition to having stronger preferences for new tablets. iPad's iOS is the most valued, followed by Windows and then Android. The price coefficient implies an average elasticity of about -6.7 and a median of -4.25.

We now turn to the learning parameters in Table 2.8. We present two sets of estimates, with standard errors in parentheses. For both sets of estimates we specify the the learning parameters as a function of observables. These observables are seller-type, seller experience, which is defined as the number of months they have sold goods on Tmall, a multi-product dummy variable indicating whether or not the seller offers more than one brand of tablet, and dummy variables indicating the seller's location.[16]

The set of estimates in column 1 is estimated under the assumption that consumers believe that the average rating was generated from the full set of ratings according to the model discussed in Section 2.3.3. A potential problem is that some of the rating scores which are equal to 5 could be what we refer to as 'default fives'. Recall that a consumer is required to rate every transaction and if he or she fails to do so, the rating is recorded as a 5. We account for this in column 2 by assuming 47.9% of the ratings are generated by this process. This rate comes from Dellarocas (2013), which says that the rate at which buyers take time to leave feedback ratings on eBay is 52.1%.[17] This assumption implies that every customer knows that 47.9% of the fives are due to the default system and adjust their beliefs accordingly. Because the results do not qualitatively vary to a large degree between these two assumptions, we focus on the parameters estimated

[15] As before, the seller/product combinations which had 0 sales are dropped from the regression (i.e., included the outside option). We explored correcting the market shares as in Gandhi et al (2013) and obtained similar results.

[16] We include dummies for each of Shenzhen, Shanghai, Guangzhou, and Beijing.

[17] We have estimated the model under various rates at which consumer submit ratings, and the results do not change significantly.

under the assumption that all ratings are informative.[18]

Focusing on column one, the estimates of the prior mean are about -11.99 and -12.69 for national and local sellers, respectively, suggesting that the former are perceived to be of slightly higher quality, ex ante. Additionally, younger sellers and those who only sell one brand are perceived as being higher quality ex ante, although the multi-brand dummy is not significant.

The σ parameters are the ones which determine the effect of ratings, as they pin down the weight which is placed on the prior relative to the average rating. As is shown in equation 2.8, the function of the parameters is transformed by the exponential function in order to restrict the variances to be positive. A smaller ratio, or a large signal variance relative to the prior variance, implies a larger weight on the prior and a smaller weight on the average rating. Intuitively, a smaller weight on the average rating means that the ratings don't affect the belief about seller quality as much and therefore, the ratings don't affect consumer behavior as much.

Before moving to the effect of seller-type on the learning, we discuss the estimates of the effect of other seller observables. Interestingly, seller experience, measured by how long the seller has been on Tmall, implies an increased reliance on ratings. While this contrasts with the findings of Elfenbein et al (2014), it may be due to the fact that the rating score is more impactful to consumers when the store has been around longer. That is, similar to the effect of the number of ratings, consumers may see the rating score for a new seller as a consequence of randomness, whereas the score for an experienced seller is a measure of the true quality. Results also indicate that if a seller offers multiple brands, the weight on the prior is significantly lower. The reason for this could be that rating score for single-brand sellers is coming from sales of a similar product, whereas the score for a multi-brand seller is coming from many different products, and therefore, they are less informative for any one product.

We now turn our attention to the effect of a seller's offline presence on the learning parameters. The results who that national sellers have a significantly lower value of σ, implying that ratings are more influential for local sellers compared to national sellers. The test at the bottom of the table indicates we can reject a test of equal parameter values at the 0.1% level.

The estimates imply that demand for sellers with a large national offline presence is not as sensitive to online reputation as demand for local or online-only sellers. This result could be because the distribution of shopping experiences (or ratings) for national sellers is relatively wide and/or the ex ante belief about seller quality (i.e., the prior variance) is

[18] A potential problem with the estimated model is that sellers differ in unobservable ways that are correlated with the rating score and/or the number of ratings, which will lead to bias estimates of the parameters. In the reduced form model, we are able to account for this through a seller fixed effect, but the non-linearity of the structural model makes this approach infeasible. However, we have estimated the model adding a linear seller-level effect to equation 2.11 and estimate the model based on changes in ratings over time for an individual seller. The results do not vary substantially from the results presented and, importantly, the estimates of the seller-type learning variables do no change significantly. However, we do not include these estimates in the primary results because it is difficult to interpret the model under this assumption. The best interpretation of this model is that the linear term is the 'known' quality of the seller that remains fixed over time, and the consumer uses the ratings to update her belief about the 'unknown', or uncertain, portion of seller quality.

relatively tight. In order to provide some evidence for the source of this result, we perform an exercise with the raw data. For each seller, we calculate the standard deviation of their ratings, which serves as a measure of the dispersion of signals. We then compare the average of the dispersion across seller groups. For local sellers it is 0.62 and for national sellers it is 0.67, indicating that the signals are not as informative for national sellers. However, we cannot reject the hypothesis that these means are equal at the 5% level (t value of 1.61). So while we cannot rule out that consumers receive more informative signals and have a wider prior belief for local sellers, this provides suggestive evidence that the latter is a more important factor in determining the differences in the learning process.

Given that, one could speculate that consumer's ex ante beliefs about a local seller is wider, or less certain, due to the fact that they do not have many prior experiences or interactions with these sellers, and therefore rely more on the online reputation system to guide them in their purchase decisions. For a national seller, on the other hand, consumers are fairly certain as to their quality (i.e., have a small prior variance), likely due to the fact that it is a seller or brand they know and have dealt with in the past. For example, a consumer may choose to put less emphasis on the reputation for the online stores of Target.com, BestBuy.com, and Wal-Mart.com while at the same time using the reputation system to help them determine whether or not a local seller is a reputable dealer.

2.5.1 Simulations

We use the results of the model to quantify the effect of online reputation across the different seller-types. To do this, we perform two simulation exercises. First, for each seller in each month, we set $\sigma_k = \sigma_L$ for all sellers and then calculate the change in revenue that results from removing the ability for consumers to observe the average rating score and the number of ratings, keeping prices and all other observables constant. Removing the rating score acts to restrict the consumer to make her purchase decisions based on the ex ante expected quality, or the prior mean. In other words, she does not update her belief about seller quality using ratings and places all the weight on the prior mean for seller s.

Specifically, we calculate the predicted revenue for store s in month t using the estimates of the model and assuming $\sigma_k = \sigma_L$:

$$R_{st}(\mu_s, \sigma_s | \sigma_k = \sigma_L) = \sum_{j \in J_s} \hat{s}_{jst}(L) * M_t * p_{jst}$$

where J_s are the products which seller s offers, and M_t is the total market size in month t. The expected quality for seller s is given by:

$$\hat{\zeta}_{st}(k = L) = \frac{1}{1 + n_{st}\hat{\bar{\sigma}}(k = L)_s^2}\mu_s + \frac{n_{st}\hat{\bar{\sigma}}(k = L)_s^2}{1 + n_{st}\hat{\bar{\sigma}}(k = L)_s^2}(r_{st} - \hat{y})$$

where $\hat{\bar{\sigma}}(k = L)$ is the estimated value of $\bar{\sigma}$ when $\sigma_k = \sigma_L$. The market share is then:

$$\hat{s}_{jst}(k = L) = \frac{exp(\hat{\beta}X_{jt} + \hat{\zeta}(k = L)_{st} + \hat{\nu}_{jst})}{1 + exp(\hat{\beta}X_{jt} + \hat{\zeta}(k = L)_{st} + \hat{\nu}_{jst}) + \sum_{j' \neq j, s} exp(\hat{\beta}X_{jt} + \hat{\zeta}_{st} + \hat{\nu}_{jst})}$$

Here, we use the estimates for $\hat{\beta}$ and $\hat{\nu}$ from specification in column 1 of Table 2.8.

We then calculate revenue for seller s in month t when ratings are removed and assuming:

$$\tilde{R}_{st}(\mu_s, \sigma_s) = \sum_{j \in J_s} \tilde{s}_{jst} * M_t * p_{jst}$$

where $\tilde{s}_{jst}$ is the market share calculated assuming that the expected seller quality for seller s is equal to the prior for that seller:

$$\tilde{s}_{jst} = \frac{exp(\hat{\beta} X_{jt} + \mu_s + \hat{\nu}_{jst})}{1 + \sum_{j,s} exp(\hat{\beta} X_{jt} + \mu_s + \hat{\nu}_{jst})}$$

Note that the consumers are not updating, meaning that σ_k does not play a role, which is the reason we do not condition on $\sigma_k = \sigma_L$. We define the difference in revenue between the observed environment and the counterfactual environment as:

$$\Delta^1_{st}(\sigma_k = \sigma_L) = R_{st}(\sigma_k = \sigma_L) - \tilde{R}_{st} \tag{2.13}$$

where the superscript indicates the first simulation exercise. For local sellers, this exercise provides an estimate of how much revenue she would lose or gain if ratings were removed in month t, which can be thought of the 'value' of the rating system to these sellers, or the amount the seller would pay to have its reputation observable. The value of reputation can be positive or negative depending on whether or not the ratings result in a belief that is higher or lower than the prior mean. Because the average rating reveals information about the true quality of the seller, we can say that, in general, the lower quality sellers are hurt by ratings and the higher quality sellers are helped by ratings. For national sellers, on the other hand, this exercise provides an estimate of value of reputation if consumers update their belief as if the seller was a local seller.

We then preform the same procedure with $\sigma_k = \sigma_N$ and calculate the "difference in the differences" in the effect of reputation:

$$\tilde{\Delta}^1_{st} = \Delta^1_{st}(\sigma_k = \sigma_L) - \Delta^1_{st}(\sigma_k = \sigma_H)$$

That is, we estimate the differences in the changes in revenue that result from removing the ability of consumers to update their belief about seller quality. The second term in equation 2.13 cancels out, meaning this amounts to a difference between $R_{st}(\sigma_k = \sigma_L)$ and $R_{st}(\sigma_k = \sigma_N)$. This measures the relative value of ratings between local and national sellers. By taking the differences in differences for each seller and month, we are able to isolate the effect of the rating score apart from the product offering, other seller observables, and the number of ratings. Additionally, because an individual seller is small compared to the market, doing it in this way allows us to reasonably assume that sellers do not adjust their pricing due to the removal of ratings.

We present the mean and percentiles of the distribution of $\tilde{\Delta}^1_{st}$ for sellers who benefit from ratings in the first row of Table 2.9, which we call the 'high quality' sellers. The mean value is nearly \$55,000 per month, but the distribution is quite skewed with the median being \$4,166 per month. This implies that the median seller is willing to pay over \$4,000 *more* per month to have their reputation visible as a local seller than they would as a national seller. In order to provide a measure of how large this value is compared to monthly revenue, we also present the distribution $\frac{\tilde{\Delta}^1_{st}}{R_{st}(\mu_s,\sigma_s|\sigma_k=\sigma_L)}$, or the relative monthly value of ratings as a percentage of monthly revenue, in the second row. For the average seller, the relative value of the rating system is about 45% of their monthly revenue as a local seller.

Mechanically, the reason for these results is the fact that the weight on the average rating for national sellers is smaller than the weight for local sellers. This combined with the fact that the sellers who benefit from ratings have a relatively "high" rating score, implies that sellers will have higher revenue as a local seller than they would have as a national seller. Hence, the value of reputation for local sellers is greater than the value for national sellers.

In the first row of Table 2.10, we present the same statistics of the distribution of $\tilde{\Delta}^1_{st}$, but do so for sellers who are hurt by ratings, or the 'low quality' sellers. The distribution is similar to the distribution in the first panel in terms of magnitudes as the average seller loses about \$40,000 more per month as a local than they would as a national seller. The median value is \$4,138, implying that the median seller would be willing to pay over \$4,000 more per month to have their reputation removed from site as a local seller than they would if they were a national seller. The second row shows that the average seller who is hurt by ratings has a relative value which is about -215% of their monthly revenue as a local seller.[19] Overall, the results of this exercise imply that that the ability of consumers to observe a signal of reputation affects local sellers to a substantially greater degree than it affects national sellers.

The second simulation exercise is similar, but instead of removing the ability of consumers to observe the rating score and number of ratings, we change the rating score from the observed score to the average score in the data, 4.82.[20]. Specifically, we calculate:

$$R_{st}(\mu_s,\sigma_s|\sigma_k=\sigma_L,\bar{r})=\sum_{j\in J_s}\hat{s}_{jst}(L,\bar{r})*M_t*p_{jst}$$

where $\hat{s}_{jst}(L,\bar{r})$ is the estimated share when $\sigma_k=\sigma_L$ and the rating score is set to the average, $\bar{r}$. We calculate the difference between this and the revenue with the observed rating score:

$$\Delta^2_{st}(\sigma_k=\sigma_L)=R_{st}(\sigma_k=\sigma_L,\bar{r})-\tilde{R}_{st}(\sigma_k=\sigma_L)$$

We then calculate the same under the assumption that consumers update as if the sellers where all national sellers:

$$\Delta^2_{st}(\sigma_k=\sigma_N)=R_{st}(\sigma_k=\sigma_N,\bar{r})-\tilde{R}_{st}(\sigma_k=\sigma_N)$$

Finally, we form the difference in differences:

$$\tilde{\Delta}^2_{st}=\Delta^2_{st}(\sigma_k=\sigma_L)-\Delta^2_{st}(\sigma_k=\sigma_H)$$

[19]Note that the reason this can be above 100% is because we are measuring the revenue gains from removing ratings, which can more than the revenue itself.

[20]This average is different than the one presented in Table 2.2 because it is calculated using the full distribution of ratings, whereas the average in Table 2.2 is calculated using the two digit rating observed on the sellers' home page.

The results of this exercise represent the relative 'value' of the seller being able to differentiate itself from the average seller. That is, it estimates how much extra revenue does a local seller lose (gain) when their rating score is mechanically decreased (increased), providing a measure of the relative marginal effect of changing the rating score.

We present the mean and percentiles of the distribution of $\tilde{\Delta}^2_{st}$ for sellers whose observed rating score is above the average in second row of the first row of the second panel in table Table 2.9. Again, we see that sellers are significantly more affected by this change as local sellers than they are as national sellers, with the median seller losing over \$2,000 more a month when its rating score is reduced. The average relative value of changing the rating score is around 68% of the monthly revenue.

The results in terms of absolute changes are similar for the sellers who have a rating score below the average, as the median seller gains just under \$2,000 more as a local seller. Finally, the average low quality seller's relative value of changing the rating score is about 88% of their revenue. The revenue effect of changing ratings scores is larger for local sellers because consumers put more weight on their rating score for these sellers, and therefore, beliefs change to a greater degree. That is, the effect of increasing or decreasing the ratings score has more impact on a consumer's belief about the seller's quality.

These results suggest that while both types of sellers value the reputation system, it is significantly more valuable for the local sellers. Because of this difference, it may be beneficial for the platform and for consumers to have a type-specific reputation system. For example, on the seller's page, the platform could highlight ratings for local sellers while highlighting other aspects, such as price and product characteristics, for national sellers. Of course, this also depends on the costs of such a policy. Further, it could even be the case that the cost of implementing and maintaining a reputation system, in terms of the platform's costs and the consumer's time, outweighs the benefits for national sellers, implying that the current system is inefficient. Additionally, if sellers are able to impact their long run average rating through pricing policies and/or improvements in quality, then these results suggest that local sellers have more of an incentive to engage in such activities. Finally, the fact that ratings play a bigger role in determining expected quality and, hence, the degree of differentiation for local sellers implies that online reputation has a larger impact on the level of competition for these sellers. Therefore, researchers and policy makers need to consider such effects when analyzing market structure in this environment.

2.6 Conclusion

We studied the heterogenous affect of online reputation across sellers on Tmall, a large business-to-consumer platform in China. We focus on how sellers differ in their national presence. Results indicate demand for retailers who have a large network of offline stores is not as responsive to online reputation as retailers who do not and that the differential value of the rating system across these seller-types can be substantial.

The fact that ratings affect local and online-only sellers to a greater degree is an intuitive result, as it is likely that individuals already have a very good idea about a retailer's quality if they have shopping experiences offline. This suggests that online reputation systems for large offline retailers may be unnecessary and possibly inefficient if there are significant costs of maintaining the system and/or costs for individuals to complete a rating. Additionally, sellers with a well-established offline presence may be able to shirk on their online business without hurting their sales to a large degree. However, this would be somewhat counteracted by the extent that shirking would hurt repeat online or offline business.

One area we did not focus on is how sellers manage their reputation. Our analysis implies that sellers with a large network of brick-and-mortar stores may spend little time and energy managing their online reputation. Results in Fan et al. (2016) suggest that sellers on Taobao adjust their pricing strategy in order to build their reputation. Connecting this to our study implies that there may be competitive effects associated with online reputation systems, as offline retailers may have the advantage of not having to set prices in order to build their online reputation. Studying this aspect of online markets would be an interesting way to push this research forward.

2.7 Tables and Figures

2.7.1 Tables

Table 2.1: Information for Top 5 Sellers of Each Group

Group	Rank	# of offline stores	# of employees	Website
1	1	22	.	Yes
1	2	177	.	Yes
1	3	838	.	Yes
1	4	33	.	Yes
1	5	900	.	Yes
2	1	.	200-499	No
2	2	.	1-50	No
2	3	.	50-99	No
2	4	.	20-99	No
2	5	.	1-49	No
3	1	.	1-49	No
3	2	.	1-49	No
3	3	.	20-50	No
3	4	.	.	No
3	5	.	10	No

Notes: Displayed are the number and employees for the top 5 seller of each type in terms of sales. This information was gathered by searching for each seller's official website where, often times, the number of offline retail outlets are displayed. When the seller did not have an official website, we search for the store in a business directory that lists the number of employees. Note that a '.' indicates that the information was not available for that seller.

Table 2.2: Description of the Market

	All		Electronics Only	
	Total	Average per month	Total	Average per month
No of sellers	671	461.50	562	391.63
No of brands	87	73.63	85	69.13
No of models	1,242	696.63	1,208	672.38
Sales	771,450	96,431.25	670,919	83,864.88
Revenue($)	206,964,736	25,870,590	174,521,504	21,815,188

Notes: Displayed are the aggregate statistics for sellers of tablets on Tmall both over the entire 8 month sample period and the average per month. Prices are converted to US dollars using an exchange rate equal to 6.33.

Table 2.3: Description of Electronics Retailers

	Mean	Median	Sd
Sales	1,430.39	78.00	5214.69
Revenue($)	342,970.48	15,894.38	2,025,822.77
Price($)	268.01	197.31	232.73
No. of Brands	1.27	1.00	0.72
No. of Models	9.47	4.00	12.85
Market share(%)	0.19	0.01	0.68
No. of Ratings	5,125.95	860.00	18,737.15
Rating Score 1 (Description)	4.78	4.80	0.09
Rating Score 2 (Service)	4.76	4.76	0.09
Rating Score 3 (Shipping)	4.76	4.76	0.10
HHI	0.02	-	-

Notes: Displayed are the seller-level statistics for electronics retailers on Tmall. All statistics are averages across stores for the entire 8 month sample except for price and the rating information. Ratings are averages across stores and months, while prices are averages across products, stores, and months. Prices are converted to US dollars using an exchange rate equal to 6.33.

Table 2.4: Description of Electronics Retailers by Type

	Mean	Median	Sd
National Retailers			
Sales	2,324.74	170.50	7,181.57
Revenue($)	435,806.22	30,699.43	1,308,086.39
Price($)	241.63	174.76	193.92
No. of Brands	1.20	1.00	0.81
No. of Models	8.91	5.00	10.52
Market share(%)	0.30	0.02	0.93
No. of Ratings	8,036.51	1,404.72	24,621.20
Rating Score 1	4.76	4.76	0.08
Rating Score 2	4.75	4.75	0.09
Rating Score 3	4.73	4.71	0.10
Local Retailers			
Sales	1,283.92	63.00	4,813.05
Revenue($)	327,766.62	14,219.41	2,121,185.17
Price($)	272.33	203.44	238.41
No. of Brands	1.28	1.00	0.70
No. of Models	9.56	4.00	13.20
Market share(%)	0.17	0.01	0.62
No. of Ratings	4,649.28	840.88	17,580.79
Rating Score 1	4.78	4.80	0.09
Rating Score 2	4.76	4.77	0.09
Rating Score 3	4.76	4.77	0.10

Notes: Displayed are seller-level statistics separated by seller-type. All statistics are averages across stores for the entire 8 month time period except for price and the rating information. Ratings are averages across stores and months, while prices are averages across products, stores, and months. Prices are converted to US dollars using an exchange rate equal to 6.33.

Table 2.5: Average Characteristics of Tablets

	National	Local	All
Top Brands			
Teclast(#1)	0.3034	0.1609	0.1935
Onda(#2)	0.2217	0.1005	0.1282
Cube(#3)	0.0002	0.1647	0.1271
Apple(#4)	0.0060	0.1265	0.0989
Sumsung(#9)	0.0000	0.0503	0.0388
Operating System			
Android	0.6019	0.6075	0.6062
IOS	0.0060	0.1265	0.0989
Windows	0.4643	0.3002	0.3377
Cellular Internet Access	0.2725	0.2599	0.2628
Screen Size(Inches)	8.6373	8.4608	8.5012
Memory(GB)	6.8325	6.8231	6.8252
Product Age(year)	0.4295	0.6462	0.5967
Observations	2,051	12,241	14,292

Notes: Displayed are the average characteristics for tablets sold by electronics retailers. Tablets may feature more than one operating system resulting. Many tablets offered have very little storage, but have the ability to expand through external sources.

Table 2.6: Learning from Ratings

	(1)	(2)	(3)	(4)
$\log(p_{jst})$	-0.436***	-0.435***	-0.437***	-0.435***
	(0.0916)	(0.0916)	(0.0915)	(0.0915)
n_{st}	0.0451		-0.546	
	(0.0295)		(0.674)	
r_{st}	0.175***		0.172***	
	(0.0371)		(0.0369)	
$r_{st} \times n_{st}$			0.126	
			(0.144)	
$n_{st} \times \mathbb{1}\{c_s = N\}$		0.108		0.177
		(0.0630)		(0.778)
$n_{st} \times \mathbb{1}\{c_s = L\}$		0.0255		-1.431*
		(0.0237)		(0.626)
$r_{st} \times \mathbb{1}\{c_s = N\}$		0.247***		0.250***
		(0.0542)		(0.0526)
$r_{st} \times \mathbb{1}\{c_s = L\}$		0.173***		0.167***
		(0.0392)		(0.0388)
$r_{st} \times n_{st} \times \mathbb{1}\{c_s = N\}$				-0.0147
				(0.164)
$r_{st} \times n_{st} \times \mathbb{1}\{c_s = L\}$				0.309*
				(0.133)
Constant	-5.733***	-5.771***	-5.679***	-5.665***
	(0.589)	(0.581)	(0.595)	(0.582)
Adjusted R^2	0.409	0.409	0.409	0.410
Observations	12167	12167	12167	12167

Notes: Standard errors in parentheses, $^{*}(p < 0.05),^{**} (p < 0.01),^{***} (p < 0.001)$. Displayed are the results of regressions of the difference in log shares on different functions of the average rating and the number of ratings. All regressions include product characteristics, dummies for the top 20 brands, a seller fixed effect and month fixed effect. Sellers/products with zero sales are dropped in all specifications.

Table 2.7: First Stage Estimation Results

Variable	Coefficient	Variable	Coefficient
Price	-0.004***	IOS	10.633***
	(0.000)		(0.462)
Screen Size(Inches)	1.200***	Windows	2.137***
	(0.001)		(0.001)
Memory(GB)	0.083***	Cellular Internet Access	0.485***
	(0.000)		(0.001)
Android	0.897***	Product Age(Years)	-1.085***
	(0.003)		(0.001)
20 Brand Dummies	Yes	Obs.	12178
Price Elasticity	**Mean**		**Median**
Type I	-5.947		-4.225
Type II	-6.923		-4.403

Notes: Standard errors in parentheses, $^{*}(p < 0.05)$,** $(p < 0.01)$,*** $(p < 0.001)$
Displayed are the results of the first-stage estimation of preference parameters.

Table 2.8: Second Stage Estimation Results

		No Default Fives	47.9% Default Fives
μ_{st}	National	-11.989***	-11.992***
		(0.68)	(0.68)
	Local	-12.689***	-12.7***
		(0.778)	(0.778)
	Experience	-3.221***	-3.226***
		(0.624)	(0.623)
	Multi-Brand	-0.118	-0.115
		(0.269)	(0.269)
$\bar{\sigma}_{st}$	National	-7.941***	-7.286***
		(0.629)	(0.629)
	Local	-3.736***	-3.071***
		(0.573)	(0.573)
	Experience	1.958**	1.963**
		(0.732)	(0.734)
	Multi-Brand	-1.39***	-1.395***
		(0.343)	(0.343)
y		5.406***	5.779***
		(0.092)	(0.177)
ρ		0.054***	0.103***
		(0.008)	(0.015)
Location Dummies		Y	Y
Observations		1760	1760
Hypothesis Test		$H_0 : \bar{\sigma}_L - \bar{\sigma}_N \leq 0$	
		$H_1 : \bar{\sigma}_L - \bar{\sigma}_N > 0$	
T-value		4.28***	4.28***

Notes: Standard errors in parentheses, $^{*}(p < 0.05),^{**}(p < 0.01),^{***}(p < 0.001)$ Displayed are the results of the second-stage estimation of the learning parameters. The first specification includes all ratings while the second assumes that 47.9% of the ratings are 'default fives'.

Table 2.9: Impact of Ratings for High Quality Sellers

	25th	50th	75th	Mean	N
			Remove Ratings		
Absolute Change	$500	$4,166	$28,074	$54,921	1,480
% Change	23.7	47.5	66.2	44.8	1,480
			Replace r_{st} with Mean		
Absolute Change	$346	$2,184	$14,474	$35,564	1,524
% Change	25.4	50.3	70.2	67.5	1,524

Notes: The first panel displays results of a simulation where we remove ratings for a single seller/month assuming the seller is a local seller and then compare the change in revenue to the change if the seller was a national seller. Row 1 presents the distribution of the difference in levels, while row 2 presents the distribution of the difference as a percentage of monthly revenue. The second panel displays the same statistics, but instead of removing ratings, we change the observe ratings score to the average rating score. Both panels limit the sample to seller/months that lose revenue under the counterfactual scenario.

Table 2.10: Impact of Ratings for Low Quality Sellers

	25th	50th	75th	Mean	N
			Remove Ratings		
Absolute Change	-$21,783	-$4,128	-$518	-$40,379	474
% Change	-258.3	-76.8	-22.9	-215.4	474
			Replace r_{st} with Mean		
Absolute Change	-$11,579	-$1,963	-$209	-$30,198	430
% Change	-88.8	-34.4	-6.2	-87.7	430

Notes: The first panel displays results of a simulation where we remove ratings for a single seller/month assuming the seller is a local seller and then compare the change in revenue to the change if the seller was a national seller. Row 1 presents the distribution of the difference in levels, while row 2 presents the distribution of the difference as a percentage of monthly revenue. The second panel displays the same statistics, but instead of removing ratings, we change the observe ratings score to the average rating score. Both panels limit the sample to seller/months that gain revenue under the counterfactual scenario.

2.7.2 Figures

Figure 2.1: Rating Information on the Search Page

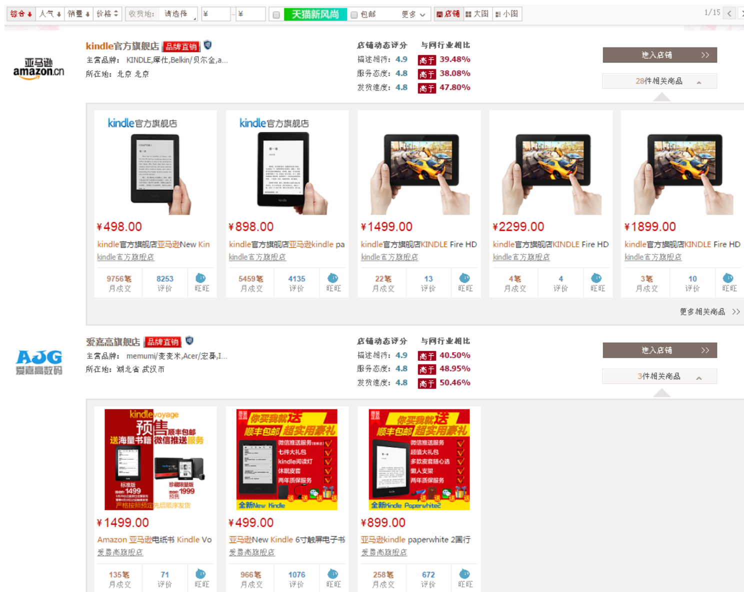

Notes: Displayed is a screen shot of a search page after a consumer sorts the default search by seller. A seller's average ratings are displayed at the top, while the number of ratings for each product are displayed below the products price and description.

Figure 2.2: Ratings Information on the Seller Page

Notes: Displayed is a screen shot of seller's home page. The average and the distribution of ratings are prominently displayed and the number of ratings appear above the distribution.

Chapter 3

A Computational Method to Solve Continuous-choice Dynamic Problems

3.1 Introduction

Economists care about stochastic continuous-choice dynamic problems, such as, capital accumulation, investment and dynamic pricing. Firms want to choose investments in order to maximize the long-run profit. Airline and theatre want to intertemporally price discriminate consumers considering a limit number of seats. These models are usually solved from the Bellman equation, a mapping from the next period's value function to current period's value function, through value function iteration for infinite-horizon problems, or backward induction for finite-horizon problems.

However, it is hard to iterative over the ex-ante Bellman equation in continuous-choice problems, because the value function evaluation requires simulations and the optimal action at each state and simulated action-specific shock.[1] If the continuous

[1] In the discrete choice dynamic problems, the Bellman equation has a functional form under certain assumptions of the action-specific shocks (Rust (1987)).

action nonlinearly enters the dynamic optimization problem, finding optimal actions at each state and shock takes a long time and slow down the estimation substantially.

This paper proposes a new computational method to solve the continuous-choice dynamic programming. Under some assumptions for the action-specific shock, the value function iteration can be computed without solving optimal actions at each state and each shock. Instead, it can be computed via taking average of action-specific value functions at some points in the action space using the probability that each point is optimal conditional on states as a weight.

So my method need some points from the action space. Within each iteration, my method requires to compute probabilities that each point is the optimal action conditional on states. The point from the action space can be random pseudo Monte-Carlo points, or evenly distributed points. These points need to cover the mass of optimal actions, which is usually known from the data, but not required to be grid points with a small step. The probabilities are easy to be computed via a change of variables. The simplicity comes from some assumptions for the action-specific shock, which leads to an analytical form of the inverse mapping from optimal actions to shocks conditional on states.

My paper makes two contributions. First, my method approximates the continuous policy function better than the canonical method calculating optimal actions from the discretized action space. In the discretization method, when the "real" optimal action locate within the intervals between the grid points, we are only able to obtain the nearest point as the optimal. The resulting policy function has the computational error increasing with the step size. However, my method is able to approximate a continuous function because the the distribution of optimal actions smooth the discrete points in the action space.

Second, my method avoids the curse of dimensionalities for multi-dimensional actions. In the discretization method, the number of grid points increases exponentially with the dimension of the action space. For a firm with k products, who want to set optimal prices p for each product, we need n^k grid points for the k-dimensional action space with n points for each dimension. However, my method use the idea of multi-dimensional integral. It can avoid the curse of dimensionality by (quasi) Monte Carlo integration, where the number of Monte-Carlo points are much less than the discretization, and the weight for each Monte-Carlo point only needs a k-by-k Jacobian matrix.

My method has some restrictions. First, the action space should be bounded. Otherwise, the numerical integration over the distribution of optimal actions may not be accurate. Second, the dynamic optimization problem is differentiable and has interior global solutions, meaning that the first order condition is the sufficient and necessary condition for optimal actions. To demonstrate the applicability of my method, I provide two examples: a dynamic pricing problem with a limited inventory, and a dynamic investment decision of plants. The distribution of optimal actions are derived for each model.

Using my computational method, I formulate a simulated method of moment for estimation for finite and infinite horizon problems. The moment is the cumulative probability of observing a certain action conditional on a certain state. To show the computational advantage, I solve an numerical example of a multi-products retailers' dynamic pricing problem given parameters. My method saves 48% computational time for one solution, and approximate a more smoother ex-ante policy function relative to the discretization method.

The rest of this paper is organized as follows. Section 3.2 describes a framework for continuous-choice dynamic problems. Section 3.3 explains the computational method, and section 3.4 gives two examples. Section 3.5 illustrates the estimation procedure. Section 3.6 demonstrates the performance of my method with an numerical example.

3.2 A Framework

Consider a dynamic programming problem in which a firm makes a series of continuous choices a_t over his lifetime $t \in \{1, \ldots, T\}$. A firm can be described by a vector of characteristics (z_t, ϵ_t), where z_t is the state variable, and $\epsilon_t(a_t)$ is an action-specific shock. The state variable z_t is discrete, or is continuous but discretized. The shock ϵ_t is independent and identically distributed over time with continuous support and distribution function $G_\epsilon(\epsilon_t)$, and the vector z_t evolves as a Markov process, depending stochastically on the choices of the firm. The probability of z_{t+1} conditional on being in z_t and making choice a_t at time t is given by $f(z_{t+1}|z_t, a_t)$. At the beginning of each period t, the firm observes $(z_t, \epsilon_t(a_t))$. The firm makes a continuous choice a_t to sequentially maximize the expected discounted sum of profits

$$\max_{a_t} \sum_{t=1}^{T} \beta^{t-1} \pi(z_t, a_t, \epsilon_t) \tag{3.1}$$

where $\pi(z_t, a_t, \epsilon_t)$ denotes the current profit of a firm with characteristics z_t from choosing a_t. The discount factor is denoted by $\beta \in (0, 1)$, and the state z_t is updated at the end of each period.

Let $V(z_t)$ be the expected value of lifetime profit at date t as a function of the current state z_t:

$$V(z_t) = E_{\substack{z_{t+1},\ldots,z_T|z_t \\ \epsilon_t,\ldots,\epsilon_T}} \left\{ \max_{a_t} \sum_{\tau=t}^{T} \beta^{\tau-t} \pi(z_\tau, a_\tau, \epsilon_\tau)|z_t\} \right\} \tag{3.2}$$

where the expectation is taken on all of future states $(z_{t+1}, \ldots, z_T)$ conditional on current state z_t, and the shock from the current period until the final period $(\epsilon_t, \ldots, \epsilon_T)$. The final period T goes to infinity for infinite-horizon problems, and T is a finite number for finite-horizon problems. Write above problem with the recursive form:

$$V(z_t) = \int_{\epsilon_t} \left\{ \max_{a_t} \pi(z_t, a_t, \epsilon_t) + \beta \sum_{z_{t+1}} V(z_{t+1}) f(z_{t+1}|z_t, a_t) \right\} dG_\epsilon(\epsilon_t) \tag{3.3}$$

The value functions are given by current period profit plus the expected value of future utility. Let $\widetilde{a}^* = \widetilde{a}(z_t, \epsilon_t, V(z_{t+1}))$ be the optimal continuous choice the firm would like to choose conditional on state z_t after observing ϵ_t given the next period's value function $V(z_{t+1})$.

$$\widetilde{a}^* = \arg\max_{a \in A} \left\{ \pi(z_t, a, \epsilon_t) + \beta \sum_{z_{t+1}} V(z_{t+1}) f(z_{t+1}|z_t, a) \right\} \tag{3.4}$$

where A is the action space. Write the Bellman equation as a contraction mapping from next period's value function to current period's value function conditional on the state z_t, and the optimal action $\widetilde{a}^* = \widetilde{a}(z_t, \epsilon_t, V(z_{t+1}))$:

$$T(V_{t+1}; z_t, \widetilde{a}^*) = \int_\epsilon \left\{ \pi(z_t, \widetilde{a}^*, \epsilon_t) + \beta \sum_{z_{t+1}} V(z_{t+1}) f(z_{t+1}|z_t, \widetilde{a}^*) \right\} dG_\epsilon(\epsilon_t). \tag{3.5}$$

where $V_{t+1} = \{V(z_{t+1})\}_{z_{t+1}}$ is a vector of the value function at $t+1$.

The solution of the continuous-choice dynamic problem $V(z_t)$ can be computed by iterating over 3.5 until the value function converges $V = TV$ in infinite-horizon problems, or until the time period goes to the initial period in finite-horizon problems $V_t = T(V_{t+1})$ from $t = T-1$ to $t = 1$. In the dynamic discrete choice model, the Bellman equations in 3.5 have functional forms if the action-specific shock linearly enter the value function and follows the generalized extreme value distributions. However, in continuous-choice problems, each iteration over 3.5 has to be computed through simulations. The process is to simulate shocks ϵ_t from its distribution $G_\epsilon(\epsilon_t)$, calculate the optimal action $\widetilde{a}(z_t, \epsilon_t, V(z_{t+1}))$ at each state z_t and each shock ϵ_t given next period's value function, and integrate current value function over the simulated shocks.

The above Bellman equation iteration needs the policy function $\widetilde{a}(z_t, \epsilon_t, V(z_{t+1}))$ at each state and shock given next period's value function. However, if the action nonlinearly enter the current period profit function or the transition probabilities, directly solving the optimal action for all states and shocks at each iteration is time-consuming. This slows down the estimation in the outer loop as well.

One approximation method for the policy function is to discretize the action space, calculate the action-specific value function at each grid point, and choose the point generating the maximum value as the optimal action. This method avoids solving optimal actions from maximization problems or nonlinear first order conditions. However, it encounters two problems. First, the optimal action chosen from the grid points may not be the real optimal one because optimal actions may locate at the intervals between the points. The computational error increases with the step of the action space grid points. Second, the number of grid points increases with the dimension of the action space exponentially. For example, a firm with 5 products want to set a 5-dimensional action to maximize the long-run profits. If there are 10 number of grid points for each dimension, the total number of points is 10^5. This is computational intractable.

3.3 The Solution Method

Under some conditions, the continuous-choice dynamic programming can be solved without directly solving optimal actions. This approach comes from the assumptions:

Assumption 1. *The action space is bounded.*

Assumption 2. *Current period's profit function $\pi(z_t, a, \epsilon_t)$ and transition probabilities $f(z_{t+1}|z_t, a)$ are both differentiable with respect to actions a, and the maximizers a^* of the dynamic optimization problem are interior solutions.*

Assumption 3. *The derivative of current profit function with respect to actions is a simple function of the action specific shock. Let $\pi_a(a_t, z_t, \epsilon_t)$ be the derivative of current period's profit function with respect to the action a_t. For any number of x that*

$$\pi_a(a_t, z_t, \epsilon_t) = x, \tag{3.6}$$

there is an analytical form of its inverse function:

$$\epsilon_t = \pi_a^{-1}(a_t, z_t, x) \tag{3.7}$$

The third assumption can be satisfied if, for example, the derivative of current period's profit function is a linear function of shocks. Such functional form assumption is applicable to a variety of problems as the examples shown in section 3.4.

With the above assumption, we can solve an analytical form of the inverse policy function from actions to independent shocks conditional on states and next period's value functions from first order conditions. The first order condition at optimal action a_t^* and state (z_t, ϵ_t) given next period's value function $V(z_{t+1})$ is

$$\pi_a(a_t^*, z_t, \epsilon_t) + \beta \sum_{z_{t+1}} V(z_{t+1}) f_a(z_{t+1}|z_t, a_t^*) = 0 \tag{3.8}$$

where $f_a(a_{t+1}|a_t, z_t)$ is the derivative of transition probabilities with respect to the action a_t. The inverse policy function from actions to shocks conditional on states can be written as

$$\epsilon(a_t^*, z_t, V_{t+1}) = \pi_a^{-1}\left(a_t^*, z_t, -\beta \sum_{z_{t+1}} V_{t+1} f_a(z_{t+1}|z_t, a_t^*)\right) \tag{3.9}$$

where $V_{t+1} = V(z_{t+1})$ is a vector of value functions at next period's states.

Remember that shocks distributes with the function of $G_\epsilon(\epsilon_t)$. The distribution of optimal actions can be derived through a change of variables:

$$g_a^*(a_t^*; z_t, V_{t+1}) = g_\epsilon\left(\epsilon(a_t^*, z_t, V_{t+1})\right) \left| det \left\{ \frac{\partial \epsilon}{\partial a}(a_t^*, z_t, V_{t+1}) \right\} \right|. \tag{3.10}$$

With the distribution of optimal actions, the ex-ante value function can be integrated over the mass of optimal actions rather than shocks. Through a change of variable, the ex-ante value function in equation 3.5 can be written as

$$
\begin{aligned}
& V(z_t) && (3.11) \\
= & \int_{a^*} \left\{ \pi(z_t, \epsilon(a^*, z_t, V_{t+1}), a^*) + \beta \sum_{z_{t+1}} V(z_{t+1}) f(z_{t+1}|z_t, a_t^*) \right\} \; g_a^*(a_t^*; z_t, V_{t+1}) \; da_t^*,
\end{aligned}
$$

where $g_a^*(a_t^*; z_t, V_{t+1})$ is derived in 3.10.

With some points $\{a^i\}_{i=1}^n$ from the action space, current period's value function given next period's ones can be calculated by taking the weighted average of values at these points. The weight is the probability that each point a^i is the optimal action satisfying the first order condition. I will illustrate the choice of the action points at the end of this section. With the action points $\{a^i\}_{i=1}^n$, the mapping from next period's value function to current period's value function in equation 3.5 can be written as

$$
V(z_t) = \sum_{i=1}^{n} \left\{ \pi(a^i, z_t, \epsilon(a^i, z_t, V_{t+1})) + \beta \sum_{z_{t+1}} V(z_{t+1}) f(z_{t+1}|z_t, a^i) \right\} w(a^i; z_t, V_{t+1}), \tag{3.12}
$$

where

$$
w(a^i; z_t, V_{t+1}) = \frac{g_a^* \left(\epsilon(a^i, z_t, V_{t+1}) \right)}{\sum_{r=1}^{n} g_a^* \left(\epsilon(a^r, z_t, V_{t+1}) \right)}, \tag{3.13}
$$

and $g_a^*(a_t^* = a^i; z_t, V_{t+1})$ is in 3.10.[2]

3.3.1 The Procedure

The above method is applicable to both finite-horizon and infinite-horizon problems with little differences.

Finite Horizon

With a finite horizon problem, the ex-ante value functions are solved backwards. At the

[2]The probabilities of optimal actions contains the Jacobian matrix. This can be approximated by the one-side difference. The procedure is to shift the action a with a small number $a + h$, calculate the action-specific shock at these two points, $\epsilon(a^i + h, z)$ and $\epsilon(a^i, z)$, and approximate the derivative with the change rate, $\partial\epsilon/\partial a = (\epsilon(a^i + h, z) - \epsilon(a^i, z))h$. If the action space is J-dimensional, the derivative is approximated by $\partial\epsilon/\partial a_j = (\epsilon((a_1^i, \ldots, a_j^i + h, \ldots, a_J^i; z) - \epsilon((a_1^i, \ldots, a_j^i, \ldots, a_J^i; z))/h$ for $\forall j = 1, \ldots, J$.

final period T, the maximization problem becomes static. It is easy to solve the inverse policy function $\epsilon(a_T, z_T)$ exempt from the next period's value function. Then put the inverse policy function into the distribution of shocks, and calculate the final period's value function through

$$V(z_T) \quad = \quad \sum_{i=1}^{n} \left\{\pi(a^i, z_T, \epsilon(a^i, z_T))\right\} w_i(a^i; z_T). \tag{3.14}$$

Then goes to one period ahead. With the final period's value function, we can solve the the inverse policy function $\epsilon(a_t, z_t, V_T)$. Similarly, derive the distribution of optimal actions from the inverse policy function, and then calculate the ex-ante value function and policy function as equation 3.12 described. The process goes through

1 Generate n points $\{a^i\}_{i=1}^n$ from the bounded action space.

2 Start from the final period $t = T$, where $V_{t+1} = V(z_{T+1})$ has terminal values.

  2.1 calculate $\epsilon(a_t^* = a^i, z_t, V_{t+1})$ in 3.9 for all $\{a^i\}_{i=1}^n$ and z_t;
  2.2 calculate $g_a^*(a^i, z_t, V_{t+1})$ in 3.10 for all $\{a^i\}_{i=1}^n$ and z_t;
  2.3 calculate $V(z_t)$ in 3.12 for all z_t.

3 Stop if $t = 1$. Otherwise, go to $t' = t - 1$, and update $V_{t'+1} = V(z_t)$, and continue step 2.1-2.3.

Infinite Horizon

For an infinite-horizon problem, the ex-ante value function can be solved iteratively from the Bellman equation. The procedure goes through

1 Generate n points $\{a^i\}_{i=1}^n$ from the bounded action space, and guess an initial value function $V^0(z)$.

2 Given $\{a^i\}_{i=1}^n$ and $V^r(z) = V^0(z)$,

  2.1 calculate $\epsilon(a^i; z, V^r(z))$ in 3.9 for all a^i and state z;
  2.2 calculate $g_a^*(a_t^* = a^i, z, V^r(z))$ in 3.10 for all a^i and state z;
  2.3 solve $V^{r+1}(z)$ in 3.12 for all z.

3 Check convergence: $||V^{r+1}(z) - V^r(z)|| < \tau$, where τ is a small number. If not, update $V^r(z) = V^{r+1}(z)$ and continue step 2.1-2.3.

3.3.2 The Choice of Action Space Points

The action points $\{a^i\}_{i=1}^n$ can be pseudo random Monte-Carlo points, evenly distributed grid points, or random points sampled from the action data. The action points need to cover the range of optimal actions, which are usually known from the action data, but not have to be grid points with a small step size. The weighted average of value function evaluation is actually the Monte-Carlo integral. Thus, to guarantees the accuracy, the action space needs to be bounded. The boundedness can be satisfied in various cases. Given model parameters, the boundary of the action space is usually known directly from the model. For example, in dynamic pricing, optimal prices needs to be higher than zero and lower than the price that demand approaches to zero. In estimation, where parameters are unknown, the observed action data provides information for the boundary. Suppose in a dynamic investment problem, data contains investment measures. The action points can be chosen from a larger box than the range of investment measures in the data.

3.4 Examples

3.4.1 Example 1: Dynamic Pricing with Limited Inventory

A retailer has z_t inventory for a product at period t. She want to set prices to maximize the expected discounted profits over a finite or infinite horizon. This problem is prevalent in airline, theater, hotels and retailing industries. For example, airline have limited number of seats and want to sell out its tickets before the departure date. A theatre has limited number of seats and want to sell out the tickets before the concert. Retailers want to sell out the limited inventory of seasonal products before the next season. Also, hotels have limited number of rooms, and want to set prices contingent on the number of rooms left.

The demand for the product at the price p_t is $Q(p_t)$, where demand decreases with prices $\partial Q(p_t)/\partial p_t < 0$ for $\forall z_t$. Suppose the retailer's current profit has a constant marginal cost, and the marginal cost function contains an action-specific shock.

$$\pi(z_t, p_t, \epsilon_t) = p_t \min\{Q(p_t), z_t\} - c(\epsilon_t) \min\{Q(p_t), z_t\} \tag{3.15}$$

where $c(\epsilon_t)$ is the marginal cost at the action-specific shock ϵ_t. Based on the distribution function of the action-specific shock $G_\epsilon(\epsilon)$, the distribution of marginal cost is known. Let $G_c(c)$ be the cumulative distribution, and $g_c(c)$ is the probability of the marginal cost is c.

The cumulative probability that the next period's inventory is less than $z \geq 0$ conditional on current period's inventory z_t and price p_t is

$$\begin{aligned} F(z_{t+1} = z | z_t, p_t) &= Pr(z_{t+1} = \max\{z_t - Q(p_t), 0\} \leq z | z_t, p_t) \qquad (3.16) \\ &= 1 - \exp(-\lambda(z_t - Q(p_t) - z)) \qquad (3.17) \end{aligned}$$

This functional form of the state transition means that $\max\{z_t - Q(p_t), 0\} - z_{t+1}$ is a random variable and is distributed as an exponential function with the parameter of λ. This assumption tells that the retailer's inventory depreciate over time.

After observing the inventory z_t and shocks ϵ_t, the retailer want to sequentially choose prices p_t to maximize the expected discounted profit over an infinite time horizon. Omit the time notation t, the retailer's problem becomes

$$V(z) = \int_c \max_{p \in \{\widetilde{p} | Q(\widetilde{p}) \leq z\}} (p - c)Q(p) + \beta \sum_{z'} V(z') f(z'|z, p) \;\; dG_c(c). \tag{3.18}$$

The model maximize the expected discounted profits under the constraint $Q(p) \leq z$. To solve the optimal prices, let us consider the case with no constraint first. For the price $p^0 = p^0(z, c, V')$ satisfying the first order condition without the constraint, the corresponding marginal cost is

$$c(p^0, z, V') = \left(\frac{\partial Q(p^0)}{\partial p}\right)^{-1} \left(Q(p^0) + \frac{\partial Q(p^0)}{\partial p} p^0 + \beta \sum_{z'} V(z') \frac{\partial f(z'|z, p^0)}{\partial p}\right) \tag{3.19}$$

The distribution of the optimal price without the condition $Q(p) \leq z$ can be derived from the distribution of costs as equation 3.10 describes. Let $g_p^0(p; z, V')$ be the distribution of optimal prices without the constraint $Q(p) \leq z$. From $c(p^0, z, V')$ in equation 3.19, $f_p^0(p, z, V')$ can be computed through a change of variables as 3.10 shows.

Let $p^* = p(z, c, V')$ be the optimal price at the state z and the marginal cost shock c. If $Q(p^0) \leq z$, the optimal price is to charge $p^* = p^0$. Otherwise, the optimal price is to charge the price at which the quantities sold equal to the inventory, $Q(p^*) = z$. The optimal price is thus $p^* = \max\{p^0, Q^{-1}(z)\}$, where the maximum comes from the assumption that $Q(p)$ decreases with p. The distribution of optimal prices in this example is therefore

$$g_p^*(p, z, V') = \frac{g_p^0(p, z, V')}{1 - G_p^0(Q^{-1}(z), z, V')} \tag{3.20}$$

for $p \geq Q^{-1}(z)$, where $F_p^0(p, z, V')$ is the cumulative distribution of $f_p^0(p, z, V')$.

3.4.2 Example 2: Plant Production

This example is borrowed from Arcidiacono and Miller (2010). A plant has the condition of $z_t \in \{1, \ldots, Z\}$, where higher levels of z_t indicate that the plant is in worse condition. The plant has the current profit from state z_t as $\pi(z_t)$. At each period t, a plant manager chooses an input $c_t \in (0, \infty)$. The cost of choosing the factor c_t is a quadrature function in the logarithm of c_t:

$$C(c_t, z_t, \epsilon_t) = (\epsilon_t + \alpha_1 z_t) \log(c_t) + \alpha_2 z_t (\log(c_t))^2, \tag{3.21}$$

where ϵ_t is an independent and identically distributed action-specific shock, and (α_1, α_2) are model parameters. The density function and the cumulative distribution function of

ϵ_t is $g_\epsilon(\epsilon_t)$ and $G_\epsilon(\epsilon_t)$ specifically. Also, increasing inputs c_t raises the probability that the machinery is in bad condition next period $t+1$,

$$f(z_{t+1}|z_t, c_t) = \begin{cases} \gamma_0/(\gamma_0 + c_t^{\gamma_1}) & \text{, if } \ z_{t+1} = z_t + 1 \\ 1 - \gamma_0/(\gamma_0 + c_t^{\gamma_1}) & \text{, if } \ z_{t+1} = z_t \\ 0 & \text{, otherwise} \end{cases} \tag{3.22}$$

where $\gamma_0, \gamma_1 > 0$ are parameters.

After observing the state z_t and an action-specific shock ϵ_t, the plant manager sequentially decides an input level to maximize the discounted expected profit.

$$V(z_t) = \int_\epsilon \max_{c_t} \pi(z_t) - C(c_t, z_t, \epsilon_t) + \beta \sum_{z_{t+1}} V(z_{t+1}) f(z_{t+1}|z_t, c_t) \ \ dG_\epsilon(\epsilon_t) \tag{3.23}$$

Let $c_t^* = c(z_t, \epsilon_t, V_{t+1})$ be the optimal input at state z_t and shock ϵ_t. From the first order condition, the corresponding shock is

$$\epsilon(c_t^*, z_t, V_{t+1}) = -(\alpha_1 z_t + 2\alpha_2 z_t \log(c_t^*)) + \beta c_t^* \left\{V(z_t) - V(z_t+1)\right\} \frac{\gamma_0 \gamma_1 (c_t^*)^{\gamma_1 - 1}}{(\gamma_0 + (c_t^*)^{\gamma_1})^2} \tag{3.24}$$

The distribution of optimal inputs is thus computed as equation 3.10.

3.5 Estimation

Model parameters can be estimated via simulated method of moments, where the moments are the probabilities of observing a action at a state in the data. This requires that the actions need to be observed in the data. However, the model can contain unobserved states, which will be explained at the end of this section.

3.5.1 Finite Horizon

Suppose the data contains the action and state observations $\{\{a_{it}, z_{it}\}_{i=1}^m\}_{t=1}^T$, where m is the total number of observations, T is the total number of periods. For any two numbers of a and z, the frequency of observing actions less than a at states less than z at period t is

$$m(a,z,t) = \frac{1}{m}\sum_{i=1}^{m} 1\{a_{it} \le a\}1\{z_{it} \le z\} \tag{3.25}$$

Using the distribution of optimal actions described in section 3.3, the moment can be calculated from the model. The cumulative probability of states can be calculated by aggregating the probability of paths arriving to state $z_t = z$ at period t. The unconditional distribution of state transition can be calculated as

$$f(z_t = z|z_{t-1};\theta) = \sum_{i=1}^{n} f(z_t = z|z_{t-1}, a_{t-1} = a^i;\theta)g_a^*(a_{t-1} = a^i; z_{t-1},\theta) \tag{3.26}$$

where $\{a^i\}_{i=1}^n$ are points from the action space as section 3.3 described, $F(z_t|z_{t-1}, a_{t-1})$ is known by the model assumption, and $g_a^*(a_{t-1} = a^i; z_{t-1})$ is calculated in section 3.3. Suppose all states start from an initial state z_1, i.e. $f(z_1)$ is known by the model assumption. The probability of observing the state $z_t = z$ at period t is

$$f(z_t = z;\theta) = \sum_{z_1}\sum_{z_2}\dots\sum_{z_{t-1}} f(z_t = z|z_{t-1};\theta)f(z_{t-1}|z_{t-2};\theta)\dots f(z_2|z_1;\theta)f(z_1). \tag{3.27}$$

With the unconditional distribution of states, the probability of observing actions less than a at period t is

$$\widetilde{m}(\theta; a, z, t) = \sum_{z_t} G_a^*(a_t^* = a; z_t,\theta)f(z_t;\theta)1\{z_t \le z\} \tag{3.28}$$

Calculate the above moments in equation 3.25 and 3.28 for a vector of numbers $\{a, z\}$ at each period t. Parameters are thus estimated through matching the model-generated moments with the moment in the data:

$$\min_\theta \left(\frac{1}{T}\sum_{t=1}^{T}\big(m(a,z,t) - \widetilde{m}(\theta;a,z,t)\big)\right)' W \left(\frac{1}{T}\sum_{t=1}^{T}\big(m(a,z,t) - \widetilde{m}(\theta;a,z,t)\big)\right) \tag{3.29}$$

where W is the weighting matrix.

3.5.2 Infinite Horizon

Suppose the data contains the action and state observations $\{a_i, z_i\}_{i=1}^{m}$, where m is the total number of observations. The frequency of observing actions less than a at states less than z is

$$m(\theta; a, z) = \sum_{i=1}^{m} 1\{a_i \leq a\}1\{z_i \leq z\} \tag{3.30}$$

Using the solution method described in section 3.3, the unconditional state transition can be calculated through aggregating over the distribution of optimal actions:

$$f(z'|z) = \sum_{i=1}^{n} f(z'|z, a = a^i) g_a^*(a = a^i; z) \tag{3.31}$$

where $\{a^i\}_{i=1}^{n}$ are points from the action space as section 3.3 described, $F(z'|z, a)$ is known from the model assumption, and $g_a^*(a = a^i; z)$ is calculated in section 3.3. Let $f^0(z)$ be the stationary distribution of states. From the transition probabilities, the stationary distribution of states can be calculated from

$$f^0(z) = f(z'|z) f^0(z) \tag{3.32}$$

In practice, the stationary distribution of states can be calculated by multiply the unconditional state transition with an initial guess for many times.[3] With the stationary distribution of states and the distribution of optimal actions conditional on states, the model-generated moment is

$$\widetilde{m}(\theta; a, z) = \sum_{z} G_a^*(a; z) f^0(z) \tag{3.33}$$

The model estimation is through simulated method of moments, where the objective function is similar to equation 3.29.

[3]For example, for an arbitrary guess of the distribution f^r, iterate the calculation $f^{r+1} = f(z'|z)f^r$ for $r = 1, \ldots, R$ times until the distribution converges $|f^{r+1} - f^r| < \tau$, where τ is a small number.

3.6 Application

The above computational method can approximate smoother value/policy function than the method of discretizing the action space. Also, the method proposed in this paper is able to solve a multi-dimensional action. To demonstrate the advantage, I consider an numerical example of a multi-product firm's dynamic pricing problem as example 1 in section 3.4.1.

Suppose the retailer has J number of products. Each product has z_j number of inventory. Suppose consumer i's utility of purchasing product $j \in \widetilde{J}$ is

$$u_{ij} = \delta_j - \alpha p_j + \epsilon_{ij} \tag{3.34}$$

where δ_j is the mean utility of product j, α is the price coefficient, and ϵ_{ij} is the idiosyncratic shock. Also, normalize consumer i's utility from choosing the outside option as zero: $u_{i0} = \epsilon_{i0}$. With the assumption that idiosyncratic shocks are independent and follow the type I extreme value distribution, product j's market share as a function of price p and inventory z is

$$Q_j(p, z) = \begin{cases} \frac{N \exp(\delta_j - \alpha p_j)}{1 + \sum_{k \in \widetilde{J}} \exp(\delta_k - \alpha p_k)}, & \text{if } \ z_j > 0 \\ 0 & \text{otherwise} \end{cases} \tag{3.35}$$

where N is the total number of potential consumers.

Discretize the state space $z_j \in \{s_j^1, \ldots, s_j^k, \ldots, s_j^K\}$, where $s^k \leq s^{k+1}$ for $\forall k = 1, \ldots, K-1$. Let $f_{mk}(j)$ be the probability that the inventory of product j transit from $z_{jt} = s^m$, to the inventory z_{jt+1} where $s_j^k < z_{jt+1} \leq s_j^{k+1}$, is

$$f_{mk}(j) = \begin{cases} H(z_j - Q_j(p) - s_j^{k-1}) - H(z_j - Q_j(p) - s_j^k) & ,\text{if } s^k > 0 \\ 1 - F_\epsilon(z_j - Q_j(p) - s_j^k) & ,\text{otherwise} \end{cases}, \tag{3.36}$$

where $H()$ is the cumulative distribution function of exponential distribution. The probability that the vector of inventory $z_t = s^m = (s_1^m, \dots, s_J^m)$ to the inventory $z_{t+1} = s^k = (s_1^k, \dots, s_J^k)$ is

$$f_{mk} = \Pi_j f_{mk}(j) \tag{3.37}$$

With the above functional form assumptions, I can solve costs corresponding optimal prices conditional on states and next period's value function from the first order condition as equation 3.19. Since the retailer has multiple products, the derivative of demand with respect to prices is a J-by-J matrix, and the derivative of state transition with respect to prices is a J-by-1 vector. From the demand function in equation 3.35, the element at the jth row and lth column of the market share derivative $\nabla Q(s,p)$ is

$$(\nabla Q(s,p))_{jl} = \frac{\partial Q_l(p)}{\partial p_j} = \begin{cases} -\alpha sh_j(p)(1 - sh_j(p))N & \text{if } j = l \\ \alpha sh_j(p) sh_l(p) N & \text{if } j \neq l \end{cases} \tag{3.38}$$

From the inventory transition in equation 3.37, the jth element of the state transition derivatives is

$$(\nabla f_{mk})_j = \frac{\partial f_{mk}}{\partial p_j} = \sum_l \frac{f_{mk}}{f_{mk}(l)} \frac{\partial f_{mk}(l)}{\partial p_j} \tag{3.39}$$

The numerical example used in this paper has parameters: $\alpha = -0.08, \beta = 0.99$. The marginal cost of product j has the function $c_j = \bar{c}_j + \epsilon_j$, where ϵ_j is independent and identically distributed as $N(0,9)$. Thus, the distribution of marginal cost c is $G_c = N(\bar{c}, \Sigma)$, where $\Sigma_{ij} = 9I$, and I is an identity matrix. The action space is $[20,40]^J$, where J is the total number of available products. In the method of discretizing the action space, I draw 5-25 number of quasi-newton points from the distribution of cost shocks $G_c(c)$ depending on the number of products.

Table 3.1 shows the performance for these two solution methods, where Panel A, B, C display the result for a retailer with one, two and three products respectively. In the case of single-product retailer, I use 1001 state points $z_t \in \{0, 1, \dots, 1000\}$. To discretize the action space, I use 1000 grid points from the action space with the step of 0.05, and 5 quasi-newton points from the distribution function of marginal costs. To use the new method, I use 20 price points $\{a^i\}_{i=1}^{20}$, where the price points are evenly distributed along the action space $[20, 40]$. The result shows that my method saves 37 seconds (48%) of the computational time for one solution. The value/policy function computed with my method is within two digits, i.e. -4.7×10^{-3} and -1.5×10^{-3}, relative to the discretization method.

Figure 3.1 displays the difference of policy functions under these two solution methods at the parameters $\lambda = 0.5$ and $\lambda = 0.1$ respectively. With the parameter of $\lambda = 0.5$, these two methods compute similar values at different inventory levels. However, when the parameter is $\lambda = 0.1$, my method computes a smooth policy function, but the discretization method computes a step function. This is because, using discrete action space, the optimal action may allocate within the intervals between grid points, but the computed optimal action is chosen from the grid points. In the two-products case, my method still saves computational time to solve the model, and the ex-ante policy function shown in figure 3.2 is smoother in my method relative to the other method.

3.7 Tables and Figures

Table 3.1: Performance of Two Solution Methods

	Discretization	New Method	Differences
Panel A: $J = 1$ ($\lambda = 0.5$)			
$\#z$	1001	1001	-
$\#a$ (Δa)	401 (0.05)	20	-
$\#\epsilon$	5	-	-
Computational Time	76s	39s	-37s (-48%)
Accuracy			
$\overline{V(z)}$	214.1455	214.1407	-4.7$\times 10^{-3}$
$\overline{p(z)}$	28.2649	28.2634	-1.5$\times 10^{-3}$
Panel B: $J = 2$ ($\lambda = 3$)			
$\#z$	961	961	-
$\#a$ (Δa) (step size)	441 (1)	75	-
$\#\epsilon$	25	-	-
Computational Time	164s	82s	-81s (-49%)
Accuracy			
$\overline{V(z)}$	336.8751	336.5058	-0.3693
$\overline{p_j(z)}$	30.3181	30.3338	15.7$\times 10^{-3}$
Panel C: $J = 3$ ($\lambda = 0.5$)			
$\#z$	961	961	-
$\#a$ (Δa) (step size)	1331 (2)	256	-
$\#\epsilon$	9	-	-
Computational Time	609s	174s	-435s (-71%)
Accuracy			
$\overline{V(z)}$	30.9359	30.5946	-0.3413
$\overline{p_j(z)}$	29.6069	27.9979	-1.6091

The action space is $[20, 40]^J$. The starting point of value function iterations is $(0, \ldots, 0)$. The convergence criteria for the ex-ante value function is $\sum_z (V^r(z) - V^{r-1}(z))^2 / \sum_z (V^{r-1}(z))^2 < 10^{-6}$.

Figure 3.1: Policy Function for $J = 1$

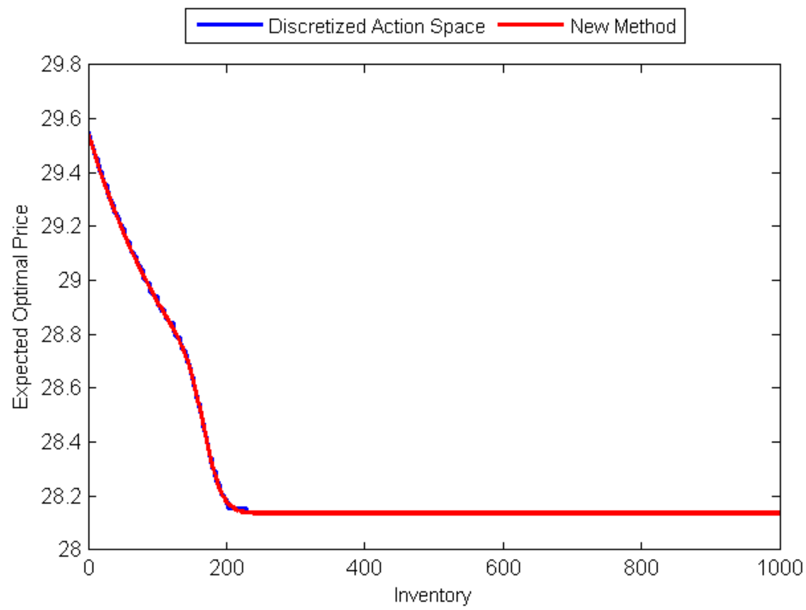

(a) $\lambda = 0.5$

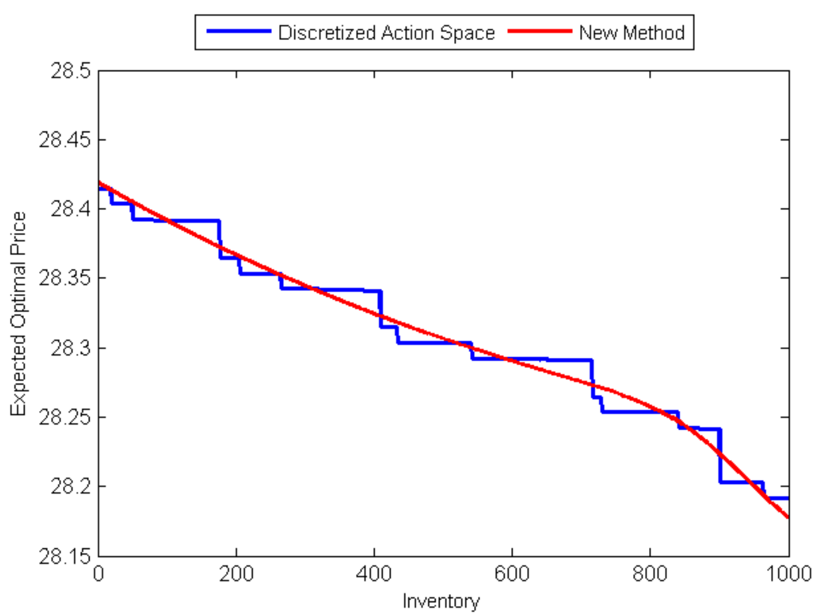

(b) $\lambda = 0.1$

Figure 3.2: Policy Function for $J = 2$

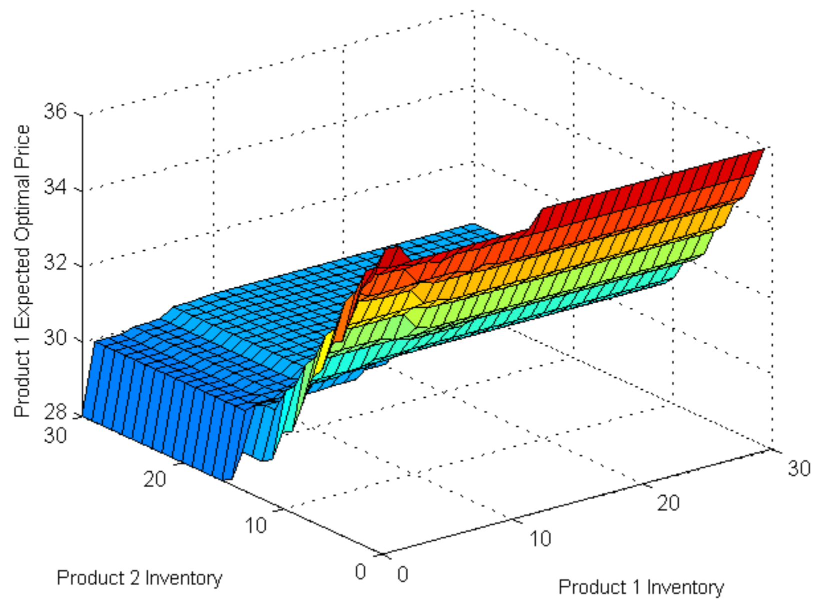

(a) Discretization

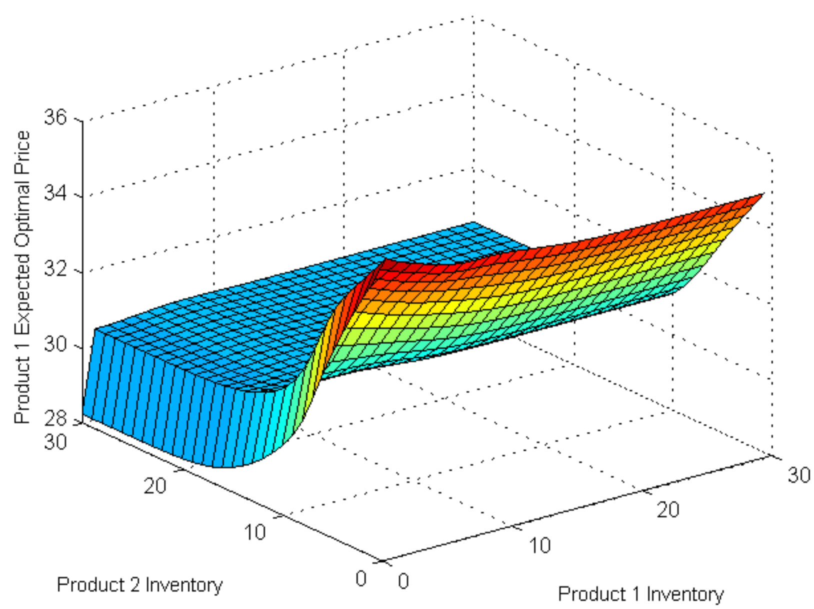

(b) New Method

Bibliography

Berry, S. (1994, Summer). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics 25*(2), 242–262.

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

Cabral, L. and A. Hortacsu (2010). The dynamics of seller reputation: Evidence from ebay. *The Journal of Industrial Economics 58*(1), 54–78.

Chevalier, J. A. and D. Mayzlin (2006, August). The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research 43*(3), 345–354.

Dellarocas, C. (2003). The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management science 49*(10), 1407–1424.

Elfenbein, D. W., R. Fisman, and B. McManus (2015). Market structure, reputation, and the value of quality certification. *American Economic Journal: Microeconomics 7*(4), 83–108.

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.

Grennan, M. and R. Town (2015). Regulating innovation with uncertain quality: Information, risk, and access in medical devices. NBER working paper 20981.

Hollenbeck, B. (2016). Online reputation mechanisms and the decreasing value of brands. Working Paper.

Luca, M. (2011). Reviews, reputation, and revenue: The case of yelp.com. Harvard Business School Working Paper, No. 12-016.

Newberry, P. W. (2016). An empirical study of observational learning. *The RAND Journal of Economics 47*(2), 394–432.

Nosko, C. and S. Tadelis (2015). The limits of reputation in platform markets: An empirical analysis and field experiment. NBER working paper 20830.

Saeedi, M. and N. Sundaresan (2016, March). The value of feedback: An analysis of the reputation system. Working Paper.

Su, L., P. Xu, and H. Ju (2016). Common threshold in quantile regressions with an application to pricing for reputation. *Econometric Reviews*. Forthcoming.

Zhao, Y., S. Yang, V. Narayan, and Y. Zhao (2013). Modeling consumer learning from online product reviews. *Marketing Science 32*(1), 153–169.

Zhu, F. and X. Zhang (2006). The influence of online consumer reviews on the demand for experience goods: The case of video games. *ICIS 2006 Proceedings*, 25.

Zhu, F. and X. Zhang (2010). Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics. *Journal of marketing 74*(2), 133–148.

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

Asker, J. and H. Bar-Isaac (2014). Raising retailers profits: On vertical practices and the exclusion of rival. *American Economic Review 104*(2), 672–686.

Aviv Nevo, J. L. T. and J. W. Williams (2016). Usage-based pricing and demand for residential broadband. *Econometrica 84*(2), 411–443.

Barwick, P. J. and P. Pathak (2015). The costs of free entry: An empirical study of real estate agents in greater boston. *The RAND Journal of Economics 46*(1), 103–145.

Bedre-Defolie, O. and E. Calvano (2013). Pricing payment cards. *The Amercian Economics Journal: Microeconomics 5*(3), 206–31.

Chao, Y. and G. Tan (2013). All-units discount, quantity forcing, and capacity constraint. *Working Paper*.

Conlon, C. T. and J. H. Mortimer (2014). Efficiency and foreclosure effects of all-units discounts: Empirical evidence. *Working Paper*.

Entrigue Ide, J.-P. M. and N. Figueroa (2015). Discounts as a barrier to entry. *Working Paper*.

Gabriel Weintraub, L. B. and B. V. Roy (2008). Markov perfect industry dynamics with many firms. *Econometrica 76*(6), 13751411.

Gautam Gowrisankaran, M. R. (2012). Dynamics of consumer demand for new durable goods. *Journal of Political Economy 120*, 1173–1219.

Gérard P. Cachon, M. A. L. (2005). Supply chain coordination with revenue-sharing contracts: Strengths and limitations. *Management Science 51*(1), 30–44.

OB́rien, D. P. (2014). All-unit discounts and double moral hazard. *Working Paper*.

Rochet, J.-C. and J. Tirole (2003a). Platform competition in two-sided markets. *Journal of the European Economic Association 1*(4), 990–1029.

Rochet, J.-C. and J. Tirole (2003b). Two-sided markets: A progress report. *The RAND Journal of Economics 37*(3), 645–667.

Shaffer, G. (1991). Capturing strategic rent: Full-line forcing, brand discounts, aggregate rebates, and maximum resale price maintenance. *The Journal of Industrial Economics 39*(5), 557–575.

Sreya Kolay, G. S. and J. A. Ordover (Fall 2004). *Journal of Economics and Management Strategy 13*(3), 429459.

Steven T. Berry, J. L. and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica 63*(4), 841–890.

Taylor, T. A. (2002). Supply chain coordination under channel rebates with sales effort effects. *Management Science 48*(8), 992–1007.

Weintraub, G. and B. Ifrach (2016). A framework for dynamic oligopoly in concentrated industries. *The Review of Economic Studies (forthcoming) 9*(16).

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

Arcidiacono, P. and R. A. Miller (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica 79*(6), 1823–1867.

Jerome Adda, R. W. C. (January 2003). *Dynamic Economics: Quantitative Methods and Applications* (1st Edition ed.). MIT Press.

Judd, K. (1999). *Numerical methods in economics* (1st Edition ed.). MIT Press.

Rust, J. (September 1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica 55*(5), 999–1033.

Vita

Xiaolu Zhou

I was born in Hubei, China, on November 1, 1987. Between 2006 and 2010, I studied economics at Shanghai University of Finance and Economics in China. I received an M.A. in economics from Peking University in June 2012, and a Ph.D. in economics from the Pennsylvania State University in August 2017.