Essays on Digital Economy

by

Mengxi Wu

B.A., Wellesley College (2011)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2017

Author . . . Signature redacted . . . .
Department of Economics
May 15, 2017

Certified by . . Signature redacted . . . .
Sara Fisher Ellison
Senior Lecturer of Economics
Thesis Supervisor

Certified by . . . Signature redacted . . . .
Glenn Ellison
Gregory K. Palm Professor of Economics
Thesis Supervisor

Accepted by . . . . . Signature redacted . . . .
Ricardo Caballero
Ford International Professor of Economics
Chairman, Department Committee on Graduate Theses

Essays on Digital Economy

by

Mengxi Wu

Submitted to the Department of Economics
on May 15, 2017, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This dissertation is a collection of three empirical essays studying consumer behavior and firm strategy in the digital economy. The first chapter examines how consumers learn from their market experience in an online marketplace. Using consumers' six-month purchase history data in a unique empirical setting in one of China's largest e-commerce platforms, I find that consumers buy from cheaper sellers as they gain market experience. To investigate how market experience improves the outcomes of price search, I incorporate price learning into a flexible structural search model, in which consumers have Dirichlet priors and update their price beliefs based on their past purchase prices in a Bayesian fashion. The results suggest that consumers have an upward bias in their prior price beliefs and are increasingly more price sensitive as they gain market experience. Early in a consumer's purchase history, the memory of sellers and prices from previous purchases accounts for a large portion of the price improvements, whereas an increasing price sensitivity plays a larger role in explaining the price advantage later on.

The second chapter investigates whether the new form of quality disclosure in the digital age – online reviews – incentives restaurants to improve quality. With little local information, tourists rely more on online reviews for restaurant recommendation than locals. Exploiting this source of variation in the impact of online reviews on restaurants, I study the trend of Yelp ratings for chain-affiliated restaurants in Las Vegas between 2005 and 2015. After controlling for common trends of restaurant chains and zip-code areas, I find that for chains with a moderate size, the customer reviews of their units closer to the Strip – the center of Las Vegas tourist activity – improve significantly more during the eleven-year data period when online reviews gain popularity, while the Strip units initially had worse ratings than the off-the-Strip units in the early days of online reviews. No such difference is found for very small, regional chains or large, multinational chains.

While market transparency is expected to increase as a result of the digital economy, in the third chapter I document the obfuscation strategies that merchants implement on an e-commerce platform with a price-comparison feature. Furthermore, I present evidence that sellers intentionally engage in the price obfuscation strategy to be more profitable and find a systematic relation between seller experience and their choice of obfuscation strategies: experienced sellers are more likely to use the bait-pricing strategy by advertising a low price, while new sellers tend to combine similar products into one listing to appear more popular and also offer the lowest price. In addition, consistent with results found in the first chapter, consumers with less market experience are more prone to being exploited by price obfuscation.

Thesis Supervisor: Sara Fisher Ellison
Title: Senior Lecturer of Economics

Thesis Supervisor: Glenn Ellison
Title: Gregory K. Palm Professor of Economics

Acknowledgments

First and foremost, words cannot express my sincere gratitude to my advisor Sara Fisher Ellison. Throughout my time in graduate school, Sara has given me not only tremendous guidance and encouragement on research, but also invaluable advice on career and beyond. Warm-hearted and lively, she is the mother figure to me and the best mentor I could ever ask for.

I am also indebted to my advisors Glenn Ellison and Nancy Rose. Glenn went through every draft of the thesis with me meticulously, even when I asked at the last minute. His immense knowledge and continuous support made the thesis tremendously better. By meeting with me on a weekly basis, Nancy guided me through the early stage of my research – it later turned into the first chapter of the thesis – when I needed the most help.

Fortunate enough, I got to know many of my amazing peers at MIT and form a strong bond with them during graduate school. I thank my study buddy and workout partner, though we occasionally took a year-long break between two workout sessions, Tianjiao Dai, for always being there for me. My thanks also go to Hongkai Zhang, Mark Hou, Yufei Wu, David Colino, Enrico Cantoni, and Ben Brookins for their tremendous help all along, especially during the job market process; the outcome would have been very different without the many informative conversations with them. I thank German Gieczewski, Jie Bai, and Sophie Sun for the precious friendships. I am grateful to Lori Breslow and Dale Deletis, whose expertise helped my job market presentation immensely.

I would also like to take the opportunity to express my appreciation to people who have been a constant, encouraging force in my life. I cannot express enough my gratitude to my undergraduate advisor, Philip Hirschhorn, who has sparked my interest in math, inspired me to embark on the graduate school journey, and brightens up my day with his witty sense of humor whenever we meet. I am thankful to my undergraduate economics advisor, Olga Shurchkov, for her continued encouragement and support in college and graduate school. I thank my dear friends from college, Kathy Chen and Maria Victoria Abrenica, for staying close to me even we were thousands of miles apart, for having more confidence in me even than I do, and for feeding my sweet tooth with cookies, chocolates, and bubble tea. Thank you to Shijie Zhao, for his love and support.

Last but by no means least, I would like to thank my parents. They have raised me in a loving and nurturing environment with their unconditional love and care, guide me in this world with their wisdom and strength, and support me to pursue my dreams in every possible way they could. I feel more than grateful to be their child. I am extremely lucky.

Contents

List of Figures

List of Tables

Chapter 1

The Mechanisms of Consumer Learning and Price Search in a Homogeneous Goods Market

1.1 Introduction

Consumers face an overwhelming amount of information when making a purchase decision. They care about quality, price, and match-specific attributes of a product, but the process of learning this information can be time-consuming. Even in homogeneous goods markets, gathering price and retailer information is not easy, especially today when there are not only the traditional local retailers but also numerous online merchants. Because of this vast amount of information, market experience is believed to play an important role in consumers' purchase behaviors. Indeed, many researchers in economics and marketing have studied how consumers learn about their matches with product attributes in differentiated product markets (e.g., Erdem et al., 2005; Crawford and Shum, 2005). This paper focuses on the role of market experience on price search in a homogeneous goods setting.

Two main challenges arise from studying consumer learning in homogeneous goods markets. One is unobservable differences in product quality, services, and sellers that are known to consumers but not researchers. If not all controlled for, these subtle differences would be reflected in the different prices consumers are willing to pay for different sellers, making the market essentially not homogeneous and more importantly, leading to biased estimates. To avoid this problem, I study a very special product market – the cell phone refill card market in one of the largest online marketplaces in China, described in more detail in Section 1.3. Cell phone refill cards are a type of stored-value cards used for prepaid cell phone plans, commonly used in countries like China, India and western European countries.[1] They are essentially gift cards for wireless services and, therefore, have no unobserved quality difference across sellers. Another special feature of the Chinese online refill card market in this study is that the purchased credit is automatically transferred to a designated cell phone account by software within moments after a transaction. Furthermore, merchants in the e-commerce platform are small verified firms with no name recognition and operate under the same policies and regulations set by the platform. With many potentially unobservable heterogeneities being stripped away from this market, it is difficult to imagine an empirical setting to be more homogeneous than this cell phone refill card market: searches occur over a single attribute, namely, price.

The other challenge in investigating the impact of consumer experience on purchase decisions is that experience is often difficult to measure. Individual purchase history is needed to accurately measure a consumer's experience, because if age or other general time measures are used as proxies for experience, it may well be the case that experience gained in other markets has a spillover effect. However, data on individual purchase history is difficult to collect. Fortunately, my data includes every purchase that consumers made with all the vendors in the market over a period of several months; among them, many are repeat purchases, which allows me to trace the purchase history of every consumer during the data period. This unusually rich data, along with the special features of this market, provide an almost ideal setting for an in-depth study of consumer learning and price

[1] http://en.wikipedia.org/wiki/Prepay_mobile_phone

search in a homogeneous goods market. Section 1.4 presents data and summary statistics.

Section 1.5 presents reduced-form evidence that consumers purchase from cheaper sellers as they become more experienced. The prices paid throughout a consumer's purchase history decrease significantly over time. With six or more previous purchases, the probability of buying from the cheapest seller increases by 10% to 30% relative to that in a consumer's first purchase. Moreover, I show that learning occurs in a passive way: consumers absorb information and become savvy progressively, instead of actively searching for a good deal early in their purchase history. These patterns motivate me to incorporate learning into a structural search model to investigate the mechanisms of consumer learning in a homogeneous goods market. Section 1.6 discusses the model in detail.

A potentially key aspect of learning is consumers' knowledge of prices and sellers from previous purchases. For instance, experienced consumers may have a better idea of the available market prices and therefore make better judgments about whether a price is low or high in the market. To incorporate this type of price learning, I follow Koulayev (2013) and De los Santos et al. (2015) by assuming that consumers' price beliefs have a Dirichlet distribution and are updated in a Bayesian fashion based on previous transaction prices. To accommodate a flexible prior belief, my model extends earlier works in two ways. First, instead of assuming that consumers know the empirical distribution prior to the search, I estimate prior beliefs in a flexible and computationally tractable way. The parameter estimates show that consumers generally believe prices to be higher than the actual prices. Second, as opposed to assuming the same price belief for all sellers, the model allows consumers to infer prices based on the observable seller-specific information before price search – the number of refill cards a seller has sold. The estimation results indicate that consumers indeed expect popularity to signal low prices. In addition to the price information, the model also accommodates the possibility that consumers have a precise memory of the sellers and prices from their previous transactions. In other words, it investigates whether consumers are able to identify the expensive and cheap sellers from their previous purchases.

Another type of learning considered in the structural model is that consumers improve their search and purchase behaviors. For example, with some level of market experience, navigating the websites and finding useful information on this marketplace may become easier, which translate into a lower search cost. Moreover, experienced consumers could potentially care more about prices after they realize that given the number of purchases they have made, they could have accumulated a great deal by making some small savings every time. To incorporate these potential changes in search and purchase behaviors, search cost and price sensitivity are modeled as functions of past market experience in the structural search model.

Section 1.7 presents the estimation results. Experienced consumers are indeed found to be better informed about market information in a way that they can precisely recall the prices and sellers from their previous purchases, which suggest that experienced consumers can search more effectively by avoiding the expensive sellers they discovered earlier. Another key driver of price improvement is consumers' increasing price sensitivity. One additional purchase increases price sensitivity by 8%. Search cost, on the other hand, changes little with experience. The counterfactual analysis suggests that the price improvement, occurring earlier in a consumer's purchase history, is largely due to the market information acquired in previous purchases, whereas price sensitivity kicks in and plays an increasingly important role in explaining the price advantages later on.

1.2 Literature Review

This study is closely related to the consumer learning literature that investigates how consumers draw inferences about product attributes from their previous market participation.[2] Based on the Bayesian learning framework, recent works show that consumers learn about product characteristics in various ways, such as from the price signal (Erdem et al., 2008), by actively gathering product information (Erdem et al., 2005), and with strategic trials (Che et al., 2015). While building upon the same framework, this paper contributes to the consumer learning literature in two ways.

First, I study consumer learning in a homogeneous goods market, whereas much of the consumer learning literature focuses on learning about individual-specific fit with product attributes in differentiated product markets. An interesting related paper is Spence (2014), which uses semesters enrolled as an experience measure in the college textbook market. With detailed survey data, Spence shows that more experienced students have a lower online search cost after controlling for individual preferences and characteristics. Given that I observe individual consumers' purchase history, my data allows me to thoroughly examine what consumers learn from each purchase.

Second, in addition to a Bayesian learning model, I conduct a model-free test to provide direct evidence that learning occurs in this homogeneous goods market. Given that price is the overwhelmingly most important feature that differentiates sellers in a homogeneous goods market, learning can be directly reflected by the prices that a consumer pays over time and therefore can be tested without any modeling assumptions. Similar in the reduced-form manner, previous studies have shown that experienced consumers have different market behaviors from inexperienced consumers. For example, Jin and Kato (2006) finds that inexperienced consumers are more likely to purchase ungraded baseball cards with high claims on eBay, and Ackerberg (2001) and Blake et al. (2015) show that advertising has significant effects on new and infrequent consumers but minimal effects on experienced consumers.

This paper also draws from the empirical search literature. Due to the lack of consumer

[2]See Ching et al.(2013) for a summary of the empirical works on consumer learning.

data, empirical works before early 2000's study costly search from the perspective of equilibrium prices. Several studies document the widespread price dispersion exists in different markets online and offline (e.g. Smith and Brynjolfsson, 1999; Baye et al., 2004; Eckard, 2004). A number of other authors provide empirical evidence that the observed price difference is consistent with the prediction of a costly consumer search model (e.g. Sorensen, 2000; Brown and Goolsbee, 2002).[3] With the prevalence of Internet shopping, individual-level data on search and purchase behaviors has become available to researchers in the last decade, leading to a large development in the empirical search literature. Several recent papers use detailed consumer search data to study how consumers conduct searches (e.g., De los Santos et al., 2012; Blake et al., 2016; and Bronnenberg et al., 2016). Building upon the canonical theoretical models of consumer search,[4] different structural models have been developed to estimate the search cost that consumers incur in different markets (e.g., Hortaçsu and Syverson, 2004; Hong and Shum, 2006; Kim et al., 2010; Honka, 2014).

While search and learning are two well-developed streams of literature in the study of consumer behaviors, only a few works lie in the intersection of the two themes. Search is not explicitly modeled in the Bayesian learning models, and at the same time learning is not incorporated into search models. Koulayev (2013) and De los Santos et al. (2015) are two exceptions; they postulate that consumers update the information acquired during the search process in a Bayesian fashion. While both of the studies focus on consumer learning within a purchasing decision, this paper investigates how consumers learn from their previous purchase experience. More importantly, I extend their models in two ways: first, in addition to the Bayesian updating, I relax the critical assumption that consumers know the actual price distribution prior to the search; second, the model allows consumers to draw inferences about prices based on the merchant-specific information that they observe before price searches.[5] Both of the methodological differences are shown to

[3]See Baye et al. (2006) for a survey on both early theoretical and empirical works on costly search and price dispersion.

[4]They are the fixed sample search model proposed by Stigler (1961) and the sequential search model developed by McCall (1970) and Mortensen (1970).

[5]That is, consumers may have different price beliefs for different merchants.

be empirically important for the setting.

1.3 Market Background

According to a Washington State wireless consulting company, 70% of Chinese cell phone users are on a prepaid (also called pay-as-you-go) plan, which requires users to have credits in advance of their service use. In the past, to deposit money into an account, the only way was to purchase cell phone refill cards from local stores. In recent years, an online market for cell phone refill cards has emerged, transforming them into virtual products – the face value of a card is automatically transferred to a designated cell phone number within minutes.

In this study, I focus on the online cell phone card market in one of the largest online marketplaces in China, owned by the e-commerce giant, Alibaba.[6] The e-commerce platform was initially created for small businesses to trade at the whole-sale level, but now it also provides business-to-consumer sales services for a variety of consumer products. Because of its nature, sellers on this marketplace are mostly small firms, alleviating the concern for sellers' credibility that arises in many Consumer-to-Consumer online markets.[7] Cell phone refill cards in this market are available in seven different face values, ranging from 1 yuan to 100 yuan, for the big three national cell phone carriers in China.[8] Consumers can purchase multiple refill cards with the same value in one transaction and use them nationwide toward any wireless service.

Purchasing a refill card on this e-commerce website follows steps similar to those in making a purchase on many other online shopping sites. Consumers can search for a specific type of cell phone refill card, defined as a unique combination of carrier and face

[6] Cell phone refill cards are sold on other online e-commerce platforms as well.

[7] Importantly, payment is first kept in a third-party account until the buyer verifies the purchase, which means that sellers have no incentive to default in this marketplace.

[8] They are Yidong (China Mobile), Liantong (China Unicom), and Dianxin (China Telecom).

value, by typing this information into a search box on the top of the home page, as shown in Figure 1-1.

Figure 1-1: Home Page of the Online Marketplace

The website will then guide consumers to the category page, which lists all the sellers who sell this product, their prices, the quantity they have sold, and the number of consumers who bought the product from them in the past. Sellers on this page are initially sorted by a default algorithm, but consumers can easily resort sellers by their price or popularity.[9] Figure 1-2 is a screenshot of the category page returned by searching 50 yuan Yidong (China Mobile) refill cards.

After browsing through the category page, consumers can go to the product page of any seller to learn more information. Figure 1-3 is an example of a product page, which lists the price of the refill card, the city where the seller is located, and the aver-

[9]The platform does not explicitly state the metrics used for the default ranking.

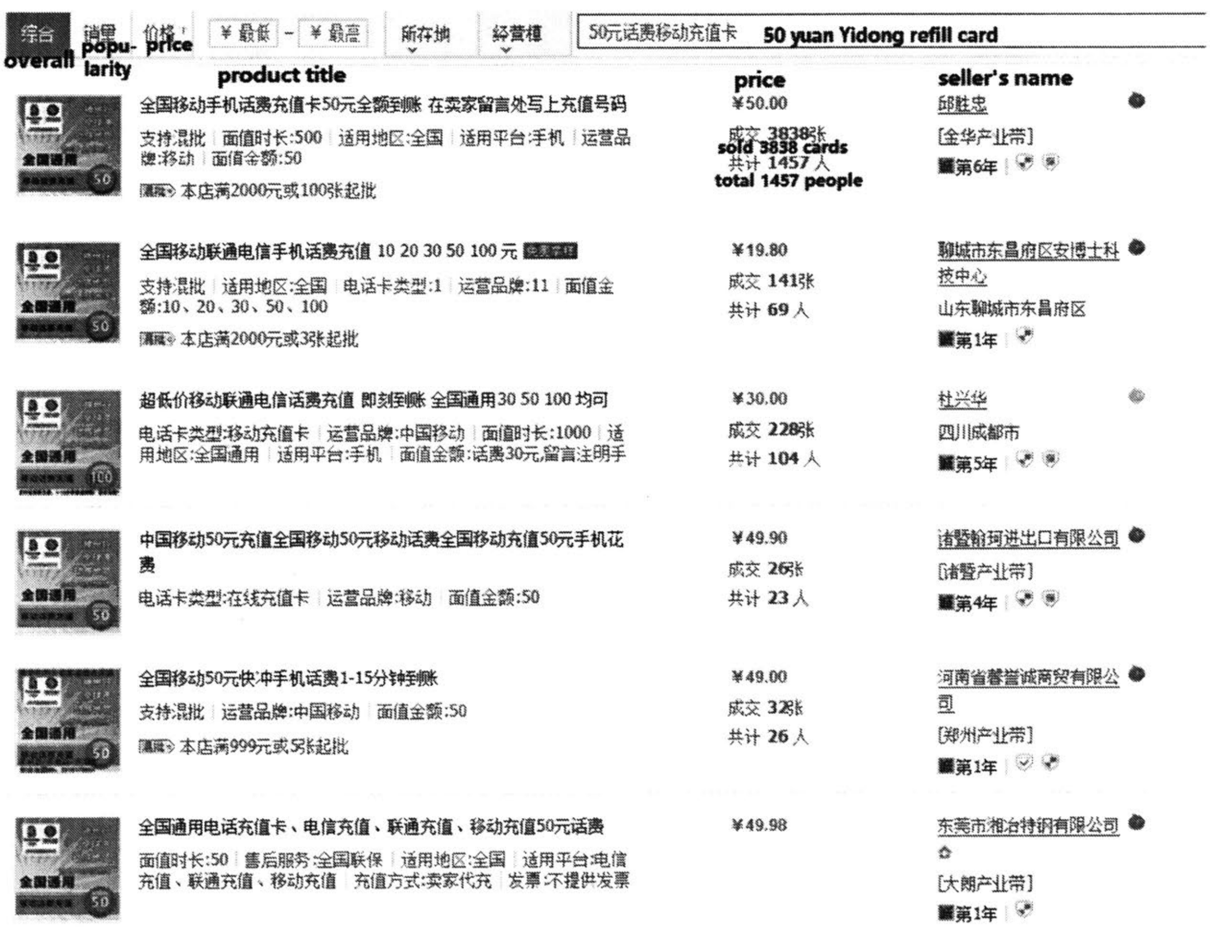

Figure 1-2: Product Category Page of 50 Yuan Yidong Refill Cards

age ratings.[10] During checkout, consumers leave their cell phone number in a box, then a program automatically takes the information and transfers the funds within a few minutes.

Figure 1-3: Product Page of a 100 yuan Refill Card

In comparison to many consumer product markets, the online cell phone refill card market has some unusual market characteristics; but exactly because of its unique aspects, this market is an empirical setting as close to an ideal homogeneous goods market as possible. Specifically, four important market features make it an exceptionally suitable setting to study empirical questions about homogeneous goods markets. First of all, unlike many physical goods markets, the online cell phone refill card market has no quality difference across sellers. Second, the instant online delivery implies that services provided by sellers are essentially the same. Third, unlike large markets with vendors who have different reputations,[11] sellers in this market are small firms with no name recognition and share similarities in a variety of dimensions since they are on the same platform. Lastly, market boundaries are often not clear in many markets in the sense that consumers may have limited information over seller availability and therefore have zero probability of searching some sellers' products.[12] In contrast, once a consumer visits this online marketplace, all

[10]Also on the product page, but not shown in the figure, are individual consumers' ratings on a scale of 1 to 5 and detailed information on every past transaction of this product, including transaction price, number of cards purchased, and time of transaction.

[11]For example, Chevalier and Goolsbee (2003) find that consumers prefer Amazon to Barnes and Noble despite the many similarities of the two sites.

[12]There is a significant body of literature on consideration set due to limited information on product

sellers are potentially in his consideration set.

Since little difference exists in sellers, and the products and services that they provide, price should be the overwhelmingly most important element differentiating sellers. However, the "law of one price" fails again even in this extremely competitive market, in which consumers can easily sort sellers by their price from low to high on the category page. With nearly zero marginal search cost, this finding is surprising. However a closer look at sellers' prices reveals the reason behind this: sellers engage in price obfuscation by making the price on the category page lower than the actual price. What enables them to use obfuscation is a special feature of the marketplace: to facilitate whole-sale level trades, this formerly business-to-business trading platform allows sellers to set several prices for a listing to offer quantity discounts.[13] In the context of the cell phone refill card market, this means that the price of a refill card may vary depending on the number of cards that a consumer purchases in a transaction. If multiple prices are set for a product, only the lowest price is listed on the category page and used for sorting, making the price comparison not as trivial as it seems.

Around three-fifths of the sellers exploit the marketplace's loophole by showing a low price on the search result page and actually charging a higher price. For the sellers who do not engage in price obfuscation, their products are sold at the price advertised on the category page for any quantity. However, consumers are unable to identify which sellers implement obfuscation techniques and have to go to each individual product page to discover the true price that applies to them.

offerings. See Goeree (2008) for an example.

[13]However, whole-sale level trades are not possible for cell phone refill cards because the value purchased in a transaction will be automatically sent to a single cell phone account. Funds, once in an account, cannot be redistributed to other cell phone accounts.

1.4 Data

The dataset consists of product data and transaction data over a six-month period, from September 24, 2013 to March 24, 2014. For each refill card product, the data includes daily prices, the corresponding quantity requirements, product types (i.e. face value and carrier), seller's name, consumer ratings, and the number of refill cards that the seller has previously sold. A total of 357 product-seller pairs were available during the data period with 53 entering and 17 exiting the market. For each of the big three national cell phone carriers, seven different face values were available: 1, 5, 10, 20, 30, 50, and 100 yuan. Table 1.1 summarizes the price and the number of sellers in each market, broken down by value and carrier.

Table 1.1: Summary Statistics of Product Price and Number of Sellers

Face Value	Price				Number of
Carrier	Mean	Stdev	Min	Max	Sellers
1-yuan Card					
Yidong	1.48	0.32	0.98	2	8
Liantong	1.60	0.53	0.98	2	3
Dianxin	1.60	0.53	0.98	2	3
10-yuan Card					
Yidong	10.06	0.17	9.5	10.5	32
Liantong	10.03	0.18	9.5	10.3	16
Dianxin	10.03	0.18	9.5	10.3	16
20-yuan Card					
Yidong	20.07	0.28	19.8	21	35
Liantong	20.01	0.17	19.8	20.5	12
Dianxin	20.02	0.16	19.8	20.5	13
30-yuan Card					
Yidong	29.97	0.13	29.8	30.5	30
Liantong	29.95	0.07	29.8	30	17
Dianxin	29.95	0.07	29.8	30	17
50-yuan Card					
Yidong	49.92	0.40	49.5	52	44
Liantong	49.93	0.53	49.5	52	19
Dianxin	49.93	0.53	49.5	52	19
100-yuan Card					
Yidong	99.71	0.56	98.8	102	31
Liantong	99.64	0.73	98.5	102	19
Dianxin	99.57	0.76	98.5	102	20

Notes: Price here is the market price when one card is purchased in a transaction. Only one seller existed in the 5-yuan market and is therefore not presented.

Like gift cards, cell phone refill cards are typically sold around the face value. Cards with higher face values generally have better value for money. Prices are similar across carriers, but the number of sellers for one carrier is significantly more due to a larger demand.[14] Figure 1-4 displays the price histograms for the most popular carrier, Yidong.[15] With the exception of 1-yuan refill card market, the mode of all the price distributions falls exactly at the face value. Overall, price dispersion is fairly small; price range varies between 0.2 yuan and 3.5 yuan, which is slightly smaller in magnitude than the price dispersion found in the eBay auction market for retail gift cards (Chiou and Pate, 2010).[16] During the data period, some price adjustments took place. Among them, thirty-three were price increases and five were price decreases. The average ratings for sellers are all above 4.5 out of 5 stars, indicating that consumers are satisfied with their purchase experience and more importantly, sellers are minimally differentiated in terms of ratings as well.

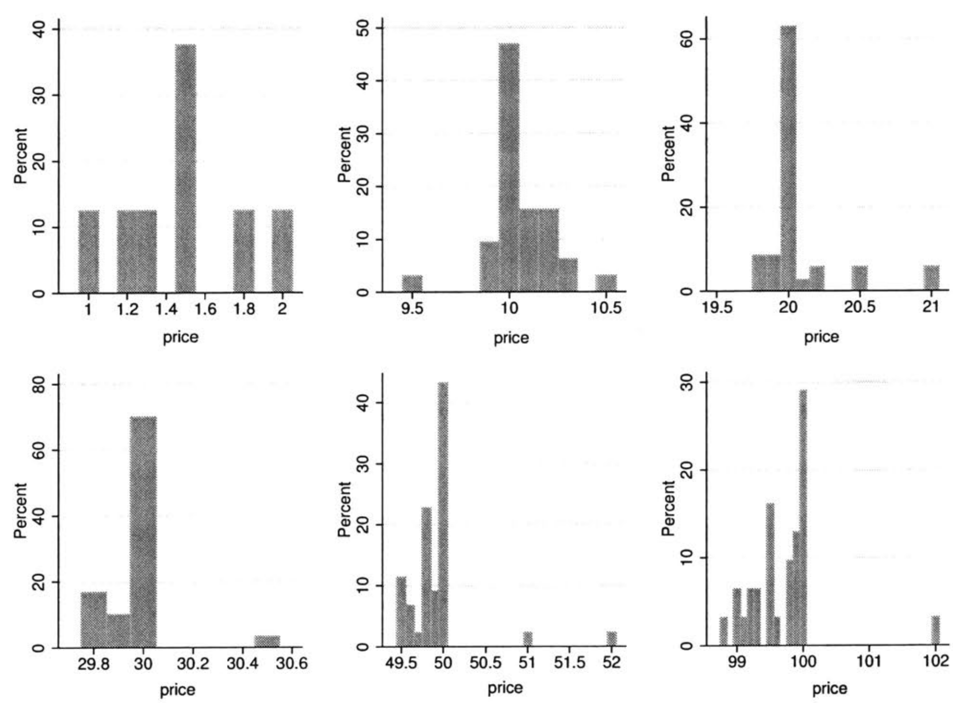

Figure 1-4: Price Histogram of China Mobile Refill Cards

[14]Yidong (China Mobile) accounted for 70% of the market share in China's domestic mobile services in 2010. Liantong (China Unicom) and Dianxin (China Telecom) have 20% and 10% market shares, respectively. See Wikipedia page for china Mobile.

[15]The other two carriers have a similar pattern.

[16]Chiou and Pate (2010) finds that the price range varies between 4% and 8% of the mean price in the auction markets for Best Buy, Home Depot and Wal-Mart gift cards on eBay.

For each transaction, the data consists of consumers' unique user ID, type of card purchased, seller who sold the product, number of phone cards bought in the transaction, price paid, and time of purchase. With unique user ID and transaction time, consumers' purchase history can be traced, providing a precise measure for consumers' level of market experience during the data period.[17]

During the six months, 4,116 transactions were made by 2,794 consumers, totaling more than 200,000 yuan in revenue, which is equivalent to approximately 35,000 dollars. About half of the transactions were made by repeat customers, with someone purchasing as many as 27 times in the half year. Figure 1-5 shows the percent of transactions made by consumers with different numbers of purchases. Table 1.2 summarizes price, quantity, and value purchased in the transaction data.[18] Transaction prices spread across the price range for every face value but have a smaller standard deviation than market prices in Table 1.1. The mean transaction price is smaller than the average market price for all except 10-yuan cards. The average value purchased in one transaction is 57 yuan and the average number of cards bought in a purchase is 1.6 cards.

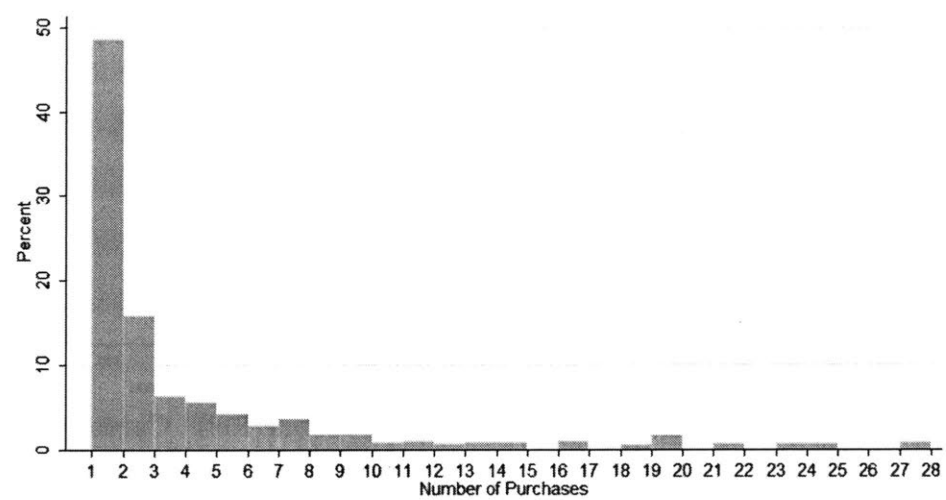

Figure 1-5: Percent of Purchases by Consumers with Different Levels of Experience

[17]See Section 1.5.3 for a discussion of how the data truncation does not affect the empirical analysis.
[18]Because market prices for refill cards with the same face value are very similar across carriers, I group transactions with the same face value together for the summary statistics.

Table 1.2: Summary Statistics of Transaction Data

	Mean	Stdev	Min	Max	Obs.
Transaction Price					
1-yuan card	1.30	0.17	0.98	1.80	412
10-yuan card	10.12	0.09	9.90	10.30	528
20-yuan card	19.91	0.13	19.80	21	196
30-yuan card	29.90	0.11	29.60	30	123
50-yuan card	49.76	0.19	49.00	52	1,696
100-yuan card	99.18	0.45	98.80	100	558
Number of Cards	1.58	1.80	1	50	4,116
Total Face Value	57.07	71.12	1	1,150	4,116

Table 1.3 defines the variables used in the empirical analysis in Section 1.5 and Section 1.6. Several measures for price and consumer experience are exploited to examine the relationship between market experience and prices robustly. A few other control variables are included as well.

Table 1.3: Definition of Variables

Variables	Definition
Price Measure	
Price	Price per card
StandPrice	Standardized price between 0 and 1 defined in Eq. (1)
Cheapest	Dummy equaling 1 if price is the cheapest possible
Experience Measure	
JthPurchase	Dummy equaling 1 if this is the consumer's j^{th} purchase
JTotalExp	Dummy equaling 1 if consumer purchased a total of j times
PastExp	Number of past purchases the consumer had at the time
Others	
Quantity	Number of cards purchased in a transaction
Type	Dummy indicating the type (i.e. value and carrier) purchased
Week	Dummy indicating the week that a transaction takes place
X	Log of the number of cards sold by the seller

1.5 Reduced-form Evidence

In this section, I first provide model-free evidence of consumer learning by examining the prices that consumers paid over time, and then present a reduced-form finding that suggests learning occurs in a passive way in this homogeneous goods market. The section is concluded with a discussion of how data truncation has no impact on the results of the analysis.

1.5.1 Price Improvement

Consumer learning, the process of acquiring information and skills though experience, is often difficult to test because the outcome variables that reflect learning are usually individual-specific and multidimensional. But in a homogeneous goods market, price is the single outcome. Therefore, investigating if consumers learn from their experience is equivalent to examining whether consumers purchase from cheaper sellers.

For a robust analysis, I use three different price measures to better compare prices that consumers paid for different types of refill cards, at different points in time, and for different quantities. First, the most direct measure for price is the transaction price itself. But, since prices are different across markets, change over time, and vary based on the number of cards purchased in one transaction, a more standardized price is also considered. The second price measure, $StandPrice_{it}$, is defined as:

$$StandPrice_{it} = (Price_{it} - LowP_{it})/(HighP_{it} - LowP_{it}), \tag{1.1}$$

where $HighP_{it}$ and $LowP_{it}$ are the highest and lowest market price possible at the time of transaction given the number of cards that consumer i purchased in transaction t, respectively. By construction, $StandPrice_{it}$ is a number between zero and one and measures the relative position of the transaction price in the price range. For the third price measure, I hold consumers to a more stringent standard by examining whether the price that a consumer pays is the lowest price possible at the time of purchase given the number of cards he purchased. This is denoted by the dummy variable, $Cheapest_{it}$.

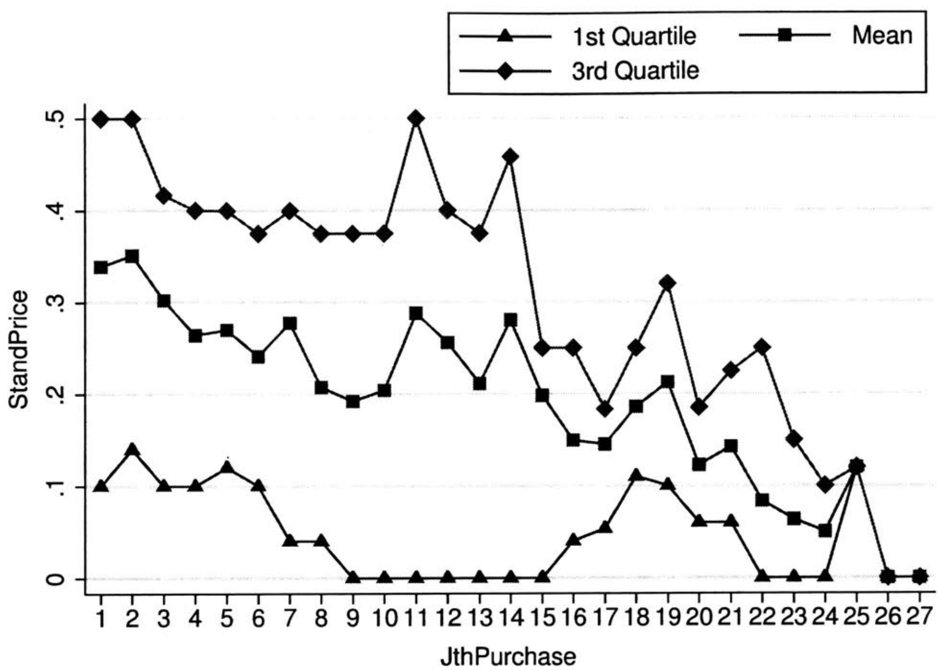

Figure 1-6: Prices Paid throughout Purchase History

Before turning to a formal analysis, I first demonstrate the relationship between price and market experience graphically. Figure 1-6 plots the 25%, mean, and 75% of the standardized prices that consumers paid throughout their purchase history in a chronological order, where the number n in the x-axis represents every consumer's nth purchase observed in the data.[19] A clear decreasing pattern is shown with all three price measures, suggesting a strong negative association between prices and market experience, which could potentially be the evidence of consumer learning.[20]

[19]There is only one consumer who purchased more than 24 times in the data period, and therefore the 25%, mean, and 75% of the standardized prices paid in each of 25^{th}, 26^{th}, and 27^{th} purchases are the same.

[20]This evidence, however, does not necessarily imply that an individual consumer pays a lower price over time because if shoppers are intrinsically different in a way that they always find better deals, then the same pattern depicted in Figure 1-6 can also be observed with the infrequent shoppers not present later in the purchase history. In other words, this can be a result of selection bias.

To control for the intrinsic differences across consumers, I then directly estimate the relationship between price and experience in a regression, allowing the most flexible pattern with the following equation:

$$\begin{aligned} PriceMeasure_{it} &= \beta_0 + \sum_{j=1}^{27} \beta_j JthPurchase_{ijt} + \delta Quantity_{it} \\ &+ \sum_{k=1}^{21} \lambda_k Type_{ikt} + \sum_{l=1}^{26} \zeta_l Week_{ilt} + \alpha_i + \epsilon_{it}, \end{aligned} \tag{1.2}$$

where $PriceMeasure_{it}$ is one of the three price measures described above: $Price_{it}$, $StandPrice_{it}$, and $Cheapest_{it}$. $JthPurchase_{ijt}$ is a dummy variable indicating if transaction t is consumer i's j^{th} purchase in the data, where j goes from 1 to 27. The number of cards that consumer i purchases in transaction t, denoted as $Quantity_{it}$, is also included as an explanatory variable for two reasons. First, the market allows sellers to offer quantity discounts, which give rise to the possibility that the larger the number of refill cards purchased in a transaction, the cheaper the price per card.[21] Second, consumers may have greater incentives to find cheap sellers when they expect to buy several refill cards in one purchase, similar to the intuition in Sorensen (2000).[22] $Type_{ikt}$ and $Week_{ilt}$ are dummies for the type of phone cards purchased in transaction t and the week that transaction t takes place, respectively. Most importantly, I also include consumer fixed effects, α_i, to capture the potential differences between frequent shoppers and infrequent shoppers as a control for the potential selection bias.[23] In the regression, β_1-β_{27} are the parameters of interest because they reveal how the transaction prices change at different points in a consumer's purchase history. It is worth noting that the use of dummy

[21]Note that this explanation is not relevant for second and third price measures because quantity purchased is already factored into the price measures.

[22]Sorensen(2000) finds that relative to one-time prescriptions, the repeatedly-purchased drugs have a significant reduction in price dispersion and profit margins, and he argues that this finding is consistent with the observation that consumers' incentives to price-shop depends on the characteristics of the drug therapy.

[23]It could be the case that consumers who frequently come back to shop in this market are the ones who find low prices earlier in their purchase history or the ones who always enjoy searching for low prices. This scenario will be captured by a smaller α, which means that the low prices found later in the purchase history will not be counted as price improvements.

variable, $JthPurchase_{ijt}$, is the most flexible way to estimate the relationship between price and consumer experience because it does not impose any functional form assumption.

Regression results for $Price_{it}$, $StandPrice_{it}$, and $Cheapest_{it}$ are presented in Columns (1)-(2), Columns (3)-(4), and Columns (5)-(6) of Table 1.4, respectively. In Column (1), the coefficients are all negative and decrease gradually from top to bottom of the column. This shows that consumers purchase from cheaper sellers progressively. Relative to the price paid in a consumer's first purchase, the price improvement is around 2 to 5 cents in his first few purchases and is then doubled and even tripled to nearly 20 cents as the consumer accumulates more market experience. Figure 1-7 depicts a visual representation of the regression results, illustrating the downward price trend in a consumer's purchase history.

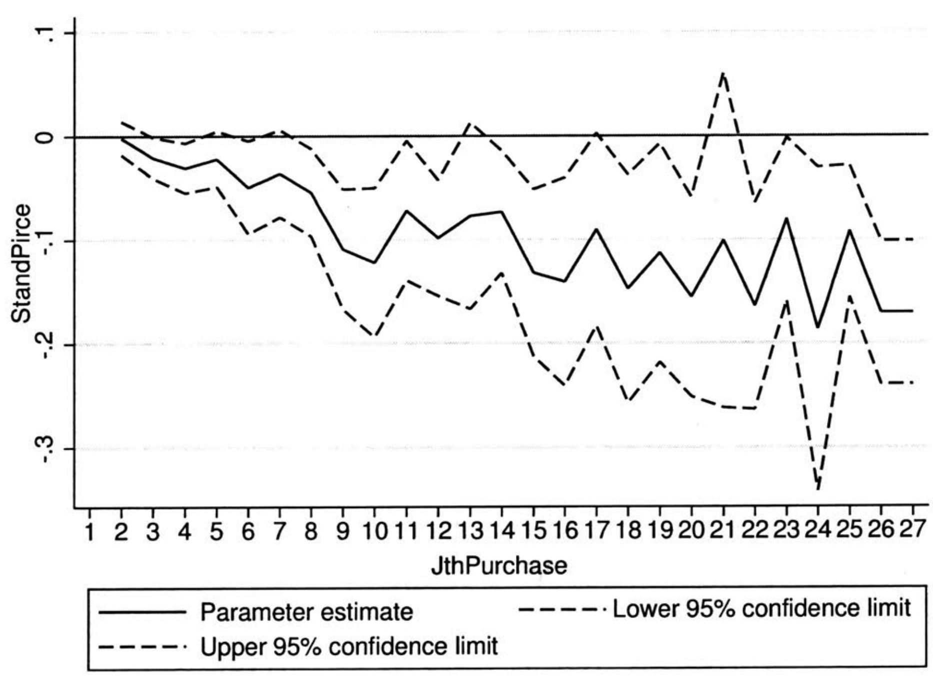

Figure 1-7: Regression of Price on Market Experience

As shown in Column (3) of Table 1.4, the same price pattern is found using the standardized price measure: the prices that a consumer pays move closer toward the lowest

Table 1.4: Impact of Experience on Purchase Prices

	Price		*StandPrice*		*Cheapest*	
JthPurchase	Coeff. (1)	SE (2)	Coeff. (3)	SE (4)	Coeff. (5)	SE (6)
2	-.0021	.0082	-.0072	.0083	.0253*	.0128
3	-.0209*	.0100	-.0336**	.0115	.0578**	.0185
4	-.0307*	.0122	-.0394**	.0137	.0530*	.0223
5	-.0220	.0136	-.0144	.0157	.0120	.0269
6	-.0495*	.0228	-.0330	.0221	.0713*	.0316
7	-.0361†	.0214	-.0248	.0243	.1075**	.0358
8	-.0543*	.0215	-.0510*	.0251	.0956*	.0433
9	-.1091***	.0295	-.0932**	.0270	.1601**	.0488
10	-.1219***	.0367	-.0960**	.0314	.1587**	.0561
11	-.0717*	.0342	-.0332	.0479	.1947**	.0600
12	-.0982***	.0284	-.0663	.0604	.2189**	.0653
13	-.0769†	.0458	-.0825*	.0405	.1725*	.0702
14	-.0731*	.0301	-.0458	.0405	.2247**	.0750
15	-.1319***	.0411	-.1468**	.0510	.2586**	.0821
16	-.1406**	.0512	-.1453**	.0541	.1336	.0821
17	-.0902†	.0473	-.1386***	.0395	.1341	.0917
18	-.1474**	.0562	-.1976***	.0617	.1765†	.0916
19	-.1126*	.0538	-.1150†	.0644	.1962*	.0978
20	-.1552**	.0491	-.1233*	.0578	.2795*	.1288
21	-.1008	.0822	-.0968	.0622	.2835*	.1288
22	-.1640***	.0508	-.1487**	.0530	.5929***	.1486
23	-.0801*	.0399	-.0365	.0885	.1441	.1484
24	-.1863*	.0796	-.1953†	.1039	.2842	.1830
25	-.0919**	.0325	-.0922*	.0449	-.0255	.2560
26	-.1704***	.0352	-.2810***	.0485	.7359**	.2562
27	-.1704***	.0352	-.2810***	.0485	.7359**	.2562
Observations	4116		4116		4116	

† $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Notes: Consumers' first purchases are omitted. The last two rows have the same estimated results because only one consumer purchases more than twenty-four times and the consumer makes the same choice in the 26^{th} and 27^{th} purchase.

possible prices as he gains market experience. Relative to the first transaction price, a consumer's price improvement exceeds 10%, and can be as high as 28%, of the price range after making fourteen or more purchases in this market.

Even with the most stringent price measure, a similar result, presented in Columns (5) and (6), shows that the likelihood of buying from the cheapest seller increases from single digits to double digits as a consumer becomes more experienced in this cell phone refill card market.[24] Together the three regression results indicate that learning does occur in this homogeneous goods market with the strong evidence that a consumer purchases from cheaper sellers over time.

1.5.2 Active Learning vs. Passive Learning

Having documented evidence of consumer learning, I now examine whether consumers learn about market information actively or passively. In other words, do consumers who make multiple purchases in this market have the intention to find cheap sellers early on, expecting that this will not only benefit the current purchase but also future ones? Or does learning happen passively through market interactions? To investigate this, I report estimates of the standardized price in consumer i's first purchase, $StandPrice_i^{first}$, from the following equation:

$$\begin{aligned} StandPrice_i^{first} &= \beta_0 + \sum_{j=1}^{21} \beta_j JTotalExp_{ij} + \delta Quantity_i \\ &+ \sum_{k=1}^{21} \lambda_k Type_{ik} + \sum_{l=1}^{26} \zeta_l Week_{il} + \epsilon_i. \end{aligned} \tag{1.3}$$

$JTotalExp_{ij}$ is a dummy variable indicating whether consumer i purchased j times in the data period. The rest of the explanatory variables are the same as in Equation (1.2).

[24] A logistic regression is performed and yields similar results. For the purpose of a direct interpretation, I report the result for linear regression.

Table 1.5: First vs. Last Purchase Price

	$StandPrice^{first}$		$StandPrice^{last}$	
$JTotalExp$	Coeff. (1)	SE (2)	Coeff. (3)	SE (4)
2	-.0018	.0014	-.0055***	.0008
3	.0216***	.0016	-.0066	.0039
4	.0310***	.0069	-.0121	.0081
5	-.0296***	.0059	-.0122†	.0066
6	.0182**	.0059	-.0305***	.0064
7	-.0467***	.0104	-.0304**	.0104
8	-.0003	.0082	-.0067	.0175
9	.0132	.0142	-.0575***	.0061
10	.1857***	.0197	-.1643***	.0224
11	.0416†	.0214	-.1586***	.0246
12	.4386***	.0191	-.0215	.0190
13	.0760	.0852	-.1747***	.0394
14	.1024**	.0262	-.1309***	.0211
16	-.0243	.0256	-.2535***	.0344
18	-.0533***	.0080	-.3199***	.0413
19	.0035	.0135	.1063***	.0091
21	.1168***	.0211	-.0633***	.0109
23	.1658***	.0313	-.1316**	.0109
24	.2513***	.0242	-.1618***	.0096
27	-.0765†	.0405	-.2391***	.0419
Observations	2794		2794	

† $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Consumers with one purchase are omitted

Columns (1) and (2) of Table 1.5 report the coefficients and standard errors, respectively. If consumers who make repeated purchases in this market are actively learning, a price advantage should be observed in the beginning of their purchase history. However, nearly half of the twenty consumer groups who purchased at least twice in this period paid a significantly higher price in their first purchase, compared to one-time shoppers. They even include consumers who purchased over twenty times. Moreover, changing the specification of market experience to a single variable, which denotes the number of total purchases a consumer has made, gives the result that consumers with more purchases actually paid a higher price in their first transaction with a significance level of .1.[25]

[25]Regression on the prices paid in consumers' second transactions shows no relation with number of

In stark contrast to the first transaction prices, prices paid in the last transactions have a strikingly different relation with consumers' level of total market experience. Keeping the explanatory variables the same, I estimate Equation (1.3) using the standardized price in consumers' last purchases as the dependent variable. Again, one-time consumers are the baseline group. Note that since first purchase is also counted as the last purchase for consumers who purchased only once, the baseline prices are the same in both regressions. Results, reported in Columns (3) and (4) of Table 1.5, show that the coefficients generally become more negative and statistically significant for consumers who made more purchases over the period. The two different patterns together suggest that price advantage stems from market experience in a passive way.

1.5.3 Discussion of Data Truncation

Since the data only spans over half year, this leads to data truncation: consumers' purchases made before the data period are not observed. If some consumers did make purchases earlier, then price improvement found in the data would have happened later in a consumer's purchase history and the impact of experience could potentially be underestimated. However, it's necessary to emphasize that the data truncation does not affect the result that during the data period market experience helps consumers improve on their purchase decision. Moreover, findings in Section 1.5.2 show that consumers who shopped more during the half year did not necessarily pay a lower price in their first transaction relative to one-time shoppers, suggesting that there is no clear evidence that frequent shoppers already had a significant amount of market experience before the data period.

---

total purchases.

1.6 Empirical Model

Results of the reduced-form analysis show that consumers learn from market experience progressively with the evidence that they purchase from cheaper sellers over time. A natural follow-up question is the following: what do consumers learn from their market experience that lead to their improved outcomes? In this section, I incorporate several aspects of learning into a flexible structural search model, which can be roughly divided into two categories. First, I consider whether experienced consumers are better informed about market information (i.e., prices and sellers available in the market). Second, the model accommodates the possibility that consumers may improve their search and purchase behaviors as a result of market experience. I then conclude the section with a discussion of model identification.

1.6.1 *Learning about Prices and Sellers*

Imagine a scenario under which a first-time consumer, who initially believes that refill cards are sold at their face value, finds a better deal. This purchase experience is likely to change his belief about market prices and affect his future purchases in this market. To incorporate the possibility that previous market participation has an impact on consumers' price beliefs, a critical assumption in many empirical search models needs to be relaxed, that is, consumers know exactly the empirical price distribution prior to search and therefore never change their beliefs regardless of the prices they discovered.

Similar to Koulayev (2013) and De los Santos et al. (2015), I assume consumers' price beliefs follow a Dirichlet distribution – the multivariate generalization of beta distribution – because it is the conjugate prior of the categorical distribution. The Dirichlet distribution is parametrized by a vector of positive numbers, $a_1, ..., a_n > 0$, known as the concentration parameters, and defined on a support $x_1, ..., x_n$ with $x_k \in (0, 1)$ and $\sum_1^n x_k = 1$, which is the (n-1)-dimensional simplex. It is essentially a distribution over a discrete distribution. The probability density is given by

$$f(x_1, ..., x_n) = Dir(a_1, ..., a_n) = \frac{\Gamma(\sum_{k=1}^{n} a_k)}{\prod_{k=1}^{n} \Gamma(a_k)} \prod_{k=1}^{n} x_k^{a_k - 1}, \tag{1.4}$$

and the expected value and variance of x_k, $\forall k \in \{1, ..., n\}$, have the following form:

$$\mathbf{E}[x_k] = \frac{a_k}{a_0}, \quad \mathbf{Var}[x_k] = \frac{a_k}{a_0}(1 - \frac{a_k}{a_0})\frac{1}{1 + a_0}, \quad a_0 \equiv \sum_{k=1}^{n} a_k. \tag{1.5}$$

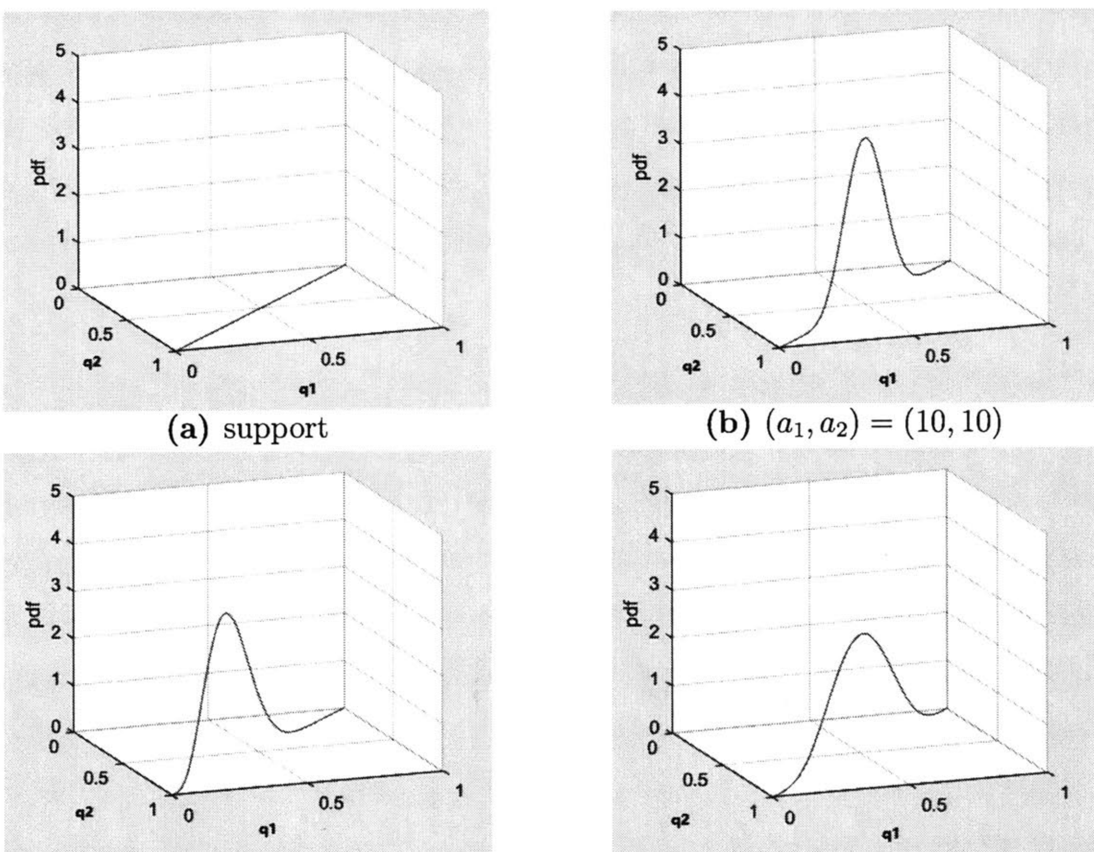

(a) support **(b)** $(a_1, a_2) = (10, 10)$

(c) $(a_1, a_2) = (5, 10)$ **(d)** $(a_1, a_2) = (5, 5)$

Figure 1-8: Examples of Dirichlet Distribution of Order 2

To illustrate how to apply a Dirichlet distribution to consumer's price belief, consider a simple case in which a consumer believes that there are only two possible prices, p_1 and p_2. However, he is not sure about the exact probability that p_1 or p_2 will be realized, and therefore perceives the corresponding likelihoods as random variables q_1 and q_2. By the definition of probability, it is clear that $0 < q_1, q_2 < 1$ and $q_1 + q_2 = 1$. It follows that (q_1, q_2) is any point on the line segment in Figure 1-8(a). Suppose that the price belief, (q_1, q_2), follows a Dirichlet distribution. Figure 1-8(b), 1-8(c), and 1-8(d) illustrate three examples of how the distribution over (q_1, q_2) looks like with different concentration parameters. With $a_1 = a_2$, as in (b) and (d), the distributions are bell-shaped curves with the mean in the middle where $q_1 = q_2 = .5$. Note that the distribution is more dispersed in (d) because a's in (b) are larger than a's in (d). When $a_1 < a_2$, as in Figure 1-8(c), the distribution is shifted toward the points in the support where q_2 is larger than q_1, which demonstrates why a's are given the name – concentration parameters.

The above examples illustrate how price beliefs can be represented by Dirichlet distributions with different concentration parameters in the case of two possible prices. I now turn to the refill card markets which follow the same concept but have more potential market prices. Let S be the set of available market prices for a type of refill card,

$$S = \{P_1, ..., P_n\}, \tag{1.6}$$

where $P_1, ..., P_n$ are sorted in the ascending order.[26] Similar to the example above, the support, $(q_1, ..., q_n)$, is a (n-1)-dimensional simplex. When consumers make refill card purchases, price beliefs, $(q_1, ..., q_n)$, are updated in a Bayesian fashion in the following way: suppose after a number of purchases in the market for a specific type of refill card, a consumer has encountered prices, $P_1, ..., P_n$, $b_1, ..., b_n$ times, respectively. By the property of Dirichlet distribution, the posterior probability density is also a Dirichlet distribution with concentration parameters $a_1 + b_1, ..., a_n + b_n$:

$$f(q_1, ..., q_n | b_1, ..., b_n) = Dir(a_1 + b_1, ..., a_n + b_n). \tag{1.7}$$

[26] I assume that they are all the prices between the minimum and the maximum market prices with a 10-cent gap between the two adjacent prices. For the 50-yuan and 100-yuan refill card markets, the maximum market price is set to be the face value because the general belief is that refill cards with large face value are sold no more than their face value at any place. This results in removing three price outliers, 51, 52 and 102 yuan (see Figure 1-4) and has little effect on the price belief updating process because only one consumer in the data paid 52 yuan to purchase a 50-yuan refill card once and no one else has purchased from sellers who charge these prices. The lowest price for the 10-yuan refill cards is set to be 9.9 yuan instead of the lowest market price, 9.5 yuan, which only lasted for a few days and had no sales.

It follows that the updated expected value of q_k, $\forall k \in \{1, ..., n\}$, can be expressed as:

$$\mathbf{E}[q_k|b_1, ..., b_n] = \frac{a_k + b_k}{\sum_{l=1}^{n} a_l + b_l}. \tag{1.8}$$

Note that holding $\frac{a_k}{a_0}$ constant for all k's, the larger the a_k's, the smaller the variance, and therefore the more weight on the prior belief.[27]

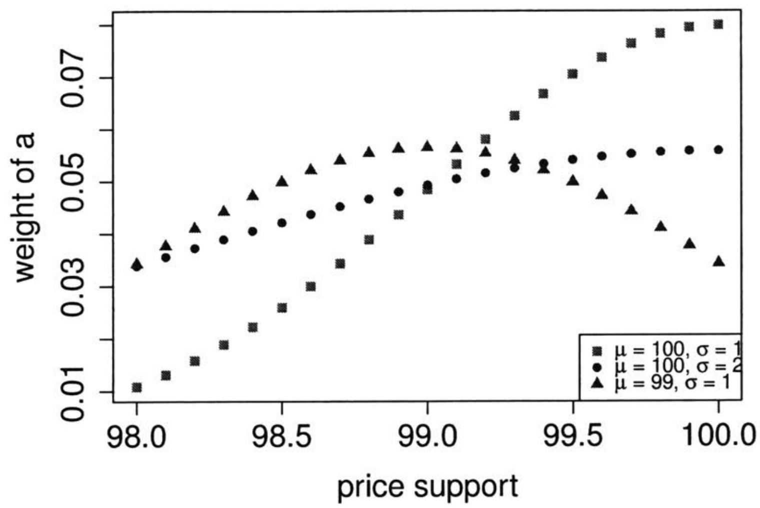

Figure 1-9: Relative Weight of Concentration Parameters

A methodological departure of this paper from the two aforementioned studies is that I estimate concentration parameters, which determine the prior belief, instead of assuming they follow the empirical distribution, because the latter, to some extent, defeats the purpose of allowing flexible beliefs. The advantage of this approach is that consumers' prior belief is estimated in a flexible way, which is more realistic than imposing a distribution assumption, and also provides more insights into how priors may be biased. However, it also comes with the cost of computational burden. So to reduce the number of parameters to be estimated, I assume that $a_1, ..., a_n$, after normalization, have the shape of a truncated normal-distribution with an estimated mean, μ, and a standard deviation,

[27] See Figure 1-8(b) and 1-8(d) for an illustration.

σ. Figure 1-9 plots the relative weight of the concentration parameters, $\frac{a_k}{a_0}$, under three different combinations of μ and σ. Given the expected value equation in Equation 5, the plot also depicts the mean prior belief. The three different examples show that the two parameters can characterize various monotone and uni-modal forms of mean prior beliefs. A larger σ is associated with an expected belief that prices are more likely to have equal or similar probability because a_k's are similar, like examples (b) and (d) in Figure 1-8. A smaller μ is associated with a belief that cheaper prices are more likely to be found. As the weight of prior belief is determined by the magnitude of all a_k's, a multiplier, $\frac{\gamma}{1-\gamma}$, is also estimated and multiplied to every a_k, where $0 < \gamma < 1$ measures the weight on prior belief and $1 - \gamma$ is the weight on prices discovered. It follows that a_k has the following specification:

$$a_k = \frac{\gamma}{1-\gamma} * f(P_k|\mu, \sigma) \quad \forall k \in \{1, ..., n\}, \tag{1.9}$$

where f is the probability density function of normal distribution.

A second methodological difference is that I allow consumers to form different price beliefs based on available information. In the case of the refill card market, consumers may have different price beliefs for different sellers depending on the information that they observe on the category page – how many refill cards a seller has sold before. This relaxes a strong assumption in the existing literature: given merchant-specific information, consumers form the same price belief for all merchants. To allow different priors, parameters govern the concentration parameters are modeled as functions of seller's popularity as follows:

$$\mu_j = \mu(1 + mX_j), \quad \sigma_j = \sigma(1 + sX_j), \quad \gamma_j = \gamma(1 + gX_j), \tag{1.10}$$

where X_j denotes log of the number of cell phone cards sold by seller j and m, s, g are parameters to be estimated.[28] As seen in Figure 1-9, holding everything else constant, a smaller μ and a larger σ represent a belief with higher expected probability on cheaper prices. A larger γ translates into more weight on the prior. While prior price beliefs may

[28] μ is separately estimated for each face value and σ is assumed to be a function of the face value: $\sigma = \sigma_0 + \sigma_1 FaceValue$ following the pattern of standard deviation in Table 1.1.

differ depending on sellers' popularity, the price updating process for all sellers follows equation (8). Going back to the simple example with two possible prices, p_1 and p_2, suppose that the concentration parameters representing the prior beliefs for seller 1 and seller 2 are $(a_{11}, a_{12}) = (5, 5)$ and $(a_{21}, a_{22}) = (10, 10)$, respectively. Then after the consumer encountered price p_1 once, his new updated price beliefs become $(a_{11}, a_{12}) = (6, 5)$ for seller 1 and $(a_{21}, a_{22}) = (11, 10)$ for seller 2.

In addition to a different belief about available prices, consumers with market experience may even remember the sellers and prices in their previous purchases. I assume that consumer i perfectly recalls the seller and price information with probability, p_i^{recall}, in the following form:

$$p_i^{recall} = 1 - exp(-\theta_s PastExp_i), \tag{1.11}$$

where $PastExp_i$ denotes consumer i's number of past purchases at the time of the transaction and $\theta_s > 0$.

1.6.2 Improving Search and Purchase Behaviors

Given the model specifications of consumers learning about market information, I now turn to describe the structural search model. First, consumer i's utility of purchasing from seller j is given as:

$$u_i^j = -\beta_i Price_j Quantity_i + \epsilon_{ij}, \tag{1.12}$$

where $Price_j$ is the price of seller j, and $Quantity_i$ is the number of refill cell phone cards that consumer i purchases in this transaction.[29] The price coefficient β_i is assumed to be a function of consumer i's past purchase experience, $PastExp_i$: $\beta_i = \beta_0(1+\beta_1 PastExp_i)$, to allow for systematic changes in price sensitivity as consumers gain market experience. The idiosyncratic utility draw, ϵ_{ij}, follows a Type 1 Extreme Value distribution and is i.i.d. across consumers, sellers, and transactions.[30] Consumers know $Quantity_i$ and ϵ_{ij}

[29]Transaction index is omitted for clarity.

[30]This term can be interpreted as the errors in individuals' assessments of utility.

before they click into an individual product page, but are uncertain about prices unless they have market experience and can perfectly recall the prices and sellers in their earlier purchases. In a sequential search model, once a price is discovered, consumers make the decision of whether to keep looking for a better deal or terminate the search by purchasing from one of the searched sellers. This decision is determined by the cost and benefit of an additional search: given that the best alternative that consumer i has found so far is utility u^* and he has paid prices, $P_1, ..., P_n$, $b_{i1}, ..., b_{in}$ times in his previous purchases, consumer i continues to search seller j's price if and only if the the net gain from the search, $G_i^j(u^*)$, is greater than zero:

$$\begin{aligned} G_i^j(u^*) &\equiv max(\mathbf{E}[u_i^j], u^*) - u^* - c_i \\ &= \sum_{u_{ik}^j > u^*} (u_{ik}^j - u^*)\mathbf{E}[q_k | b_{i1}, ..., b_{in}] - c_i \\ &= \sum_{u_{ik}^j > u^*} (u_{ik}^j - u^*) \frac{a_k + b_{ik}}{\sum_{l=1}^n a_l + b_{il}} - c_i > 0, \end{aligned} \tag{1.13}$$

where k is the index of the possible utilities (i.e. prices) of seller j and c_i is the cost that consumer i incurs from a price search. To incorporate the possibility that consumers may find it easier to navigate the websites and look for the necessary information after more market experience, the search cost is also modeled as a function of the consumer's level of market experience with the following form:

$$c_i = c_0(1 + c_1 PastExp_i). \tag{1.14}$$

In Equation (1.13), the expected benefit from the search is the expected extra utility of purchasing from seller j given that the current best alternative is u^*. This can be rewritten in a closed-form expression by the posterior mean belief in Equation (1.8). To smooth the search probability, I add an i.i.d. mean-zero stochastic term, η_i, which follows the Type-1 Extreme Value distribution with scale parameter, σ_η, to the search cost, c_i. Given this idiosyncratic error term, the probability that consumer i, whose best alternative is $u^*_{S_i}$ in his search set, S_i, will search seller j is:

$$Prob^{Search}_{ij|u^*_{S_i}} = exp(-exp(\frac{-G^j_i(u^*_{S_i})}{\sigma_\eta})). \tag{1.15}$$

I assume that consumers first search the seller with the largest expected net gain among sellers who have not been searched yet.[31] Therefore the next seller to search can be expressed mathematically as:

$$argmax_j\{G^j_i(u^*_{S_i})|j \in NS_i\}, \tag{1.16}$$

where NS_i denotes the set of sellers who consumer i has not searched yet and $u^*_{S_i}$ is the highest utility consumer i has found. Given that consumer i has searched all the sellers in set S_i, his (conditional) probability of purchasing from seller $j \in S_i$ is given by:

$$Prob_{ij|j \in S_i} = Prob(u^j_i > u^h_i \quad \forall h \neq j \in S_i). \tag{1.17}$$

To obtain the unconditional purchase probability, let $Q^k_i = \{s^{(1)}, ..., s^{(k)}\}$ denote consumer i's optimal search sequence of the first k sellers, where $s^{(h)}$ represents the h^{th} searched seller. Then the purchase probability of the j^{th} seller in the search sequence is:

$$Prob_{ij} = \sum_{j \leq h \leq N} \prod_{l<h} Prob^{Search}_{is^{(l+1)}|u^*_{Q^l_i}} (1 - Prob^{Search}_{is^{(h+1)}|u^*_{Q^h_i}}) Prob_{is^{(j)}|s^{(j)} \in Q^h_i}, \tag{1.18}$$

where the product of search probabilities is the probability that search is terminated at seller $s^{(h)}$, and the last term is the conditional purchase probability of seller $s^{(j)}$ with search set Q^h_i. This is then summed over all possible search sequences that terminate after seller $s^{(j)}$ has been searched. It follows that the sample log-likelihood function is:

$$LL = \sum_t log(p^{recall}_{it} * Prob^{recall}_{ijt} + (1 - p^{recall}_{it}) * Prob^{norecall}_{ijt}), \tag{1.19}$$

[31] Weitzman (1979) introduces the notion of reservation utility which is a fixed reward that makes a consumer indifferent between searching or not searching a seller's price, and he shows that the optimal search sequence in a sequential search model is to search the seller with the largest reservation utility. This optimal search strategy can be easily accommodated into the model, but for computational tractability I assume that consumers search according to the expected net gain.

where $Prob_{ijt}^{recall}$ and $Prob_{ijt}^{norecall}$ are purchase probabilities in Equation (1.18) under the case when consumer i recalls the sellers and prices from his previous purchases and the case when he does not, respectively, and p_{it}^{recall} follows from Equation (1.11). Parameters are then estimated by maximizing the simulated log likelihood with 100 ϵ draws for each transaction[32]. Standard errors are obtained by the estimates of 100 bootstrap samples based on the original data with replacement.

1.6.3 Identification

Price belief parameters are identified from comparing the dynamic purchase patterns of consumers with the same amount of purchase experience but different price history. For example, consider two individuals: consumer A purchased a 50-yuan refill card and consumer B bought a 100-yuan refill card on this platform. Both of them come back to purchase a 100-yuan refill card in their second purchase. Then the potential price difference in their second purchase reflects their different price beliefs for the 100-yuan refill card market, resulting from consumer B's previous 100-yuan refill card purchase experience, for the reason that the price coefficient and search cost are modeled to be the same for consumers with the same level of previous market experience. Therefore, the previous price paid by consumer B, together with their prices paid in the second purchase, help to identify the prior price belief of the 100-yuan refill card market and the weight of prior. Identification of recall probability follows a similar argument: purchase patterns of consumers who paid the same prices but purchased from different sellers help to identify the likelihood that consumers remember their previous purchase information.

Identification of the price coefficient and search cost parameters depends on exogenous market changes, such as price variations across markets and over time and sellers' entry and exit. For instance, when a new seller enters the market, it changes the purchase probability of all existing sellers, and how it affects the probability is governed by the price coefficient and search cost. Moreover, purchases made by consumers who buy the same

[32]I use logit-smoothed AR simulator to smooth the purchase probability in Equation (1.17).

cell phone refill cards but purchase different number of cards help to separately identify price coefficient and search cost because it exogenously changes the benefit of search while keeping search cost fixed.

1.7 Results

This section first presents the parameter estimates for consumer beliefs, price coefficient and search cost and is then followed by a discussion of the implications for consumer search behaviors. I conclude the section with a discussion of the model's fit and a counterfactual analysis.

1.7.1 Beliefs

Consumers' prior price beliefs are parameterized by the mean and standard deviation of the underlying normal distribution. Estimates, presented in Table 1.6, show that a seller's previous sales record, measured by log of the number of cell phone cards sold, plays an important role in forming the price belief. Consumers believe that popularity signals low prices.[33] The parameter estimate for the coefficient of seller popularity in the mean, m, suggests that mean of the underlying normal distribution decreases by 1.5% for one log increase in the amount of a seller's past sales.[34] The positive coefficient of seller popularity in the standard deviation, s, indicates that price beliefs for more popular sellers spread out more. In other words, popular sellers are expected to have extreme prices with higher probability. As an additional search is only beneficial when there is a chance of finding a price lower than the current best price, a seller whose price is believed to

[33]Market price does have an inverse relation with popularity. Regression results suggest that one log increase in the number of refill cards sold in the past is associated with a 2.7-cent decrease in price. The fact that consumers rely on other consumers' earlier choices to make decisions can be viewed as another type of consumer learning – peer learning.

[34]In the case of 100-yuan refill card market, the mean shifts to the left by 2.25 yuan with one log increase in previous sales.

have a more dispersed distribution is more likely to be searched, holding everything else constant. Figure 1-10 shows the empirical price distribution and estimated mean prior price beliefs for refill cards with different values. I illustrate the priors for two different sellers in each market: one who has sold 10 refill cards and the other with a sales record of 500 phone cards.[35] Overall, priors are biased upward, especially for refill cards with large face values or sold by less popular sellers. This is consistent with a survey result that inexperienced consumers' price expectation is higher than actual prices in the online textbook market (Mastumoto and Spence, 2014).

Table 1.6: Price Belief Estimates

Variable	Coeff.	SE
Mean of Normal Distr.		
μ – 1-yuan card	1.60	.15
μ – 10-yuan card	10.82	.10
μ – 20-yuan card	22.46	.48
μ – 30-yuan card	33.30	.91
μ – 50-yuan card	57.41	.13
μ – 100-yuan card	114.86	.60
m – Seller Popularity	-.015	.003
Std. Deviation		
σ_0 – Constant	.014	.035
σ_1 – Face Value	.005	.001
s – Seller Popularity	.094	.028
Weight of Prior		
γ – Constant	1.000	.000
g – Seller Popularity	-.003	.019
Prob. of Seller Memory		
θ_s – Past Exp.	4.53	.021

The constant term in standard deviation is rescaled by multiplying 1000 to the variable.
Standard errors of estimates are computed using 100 non-parametric bootstraps.

[35]At the beginning of the data period, the average number of refill cards sold by a seller is 18.5 with a standard deviation of 72. The most popular seller sold 799 refill cards before the data period.

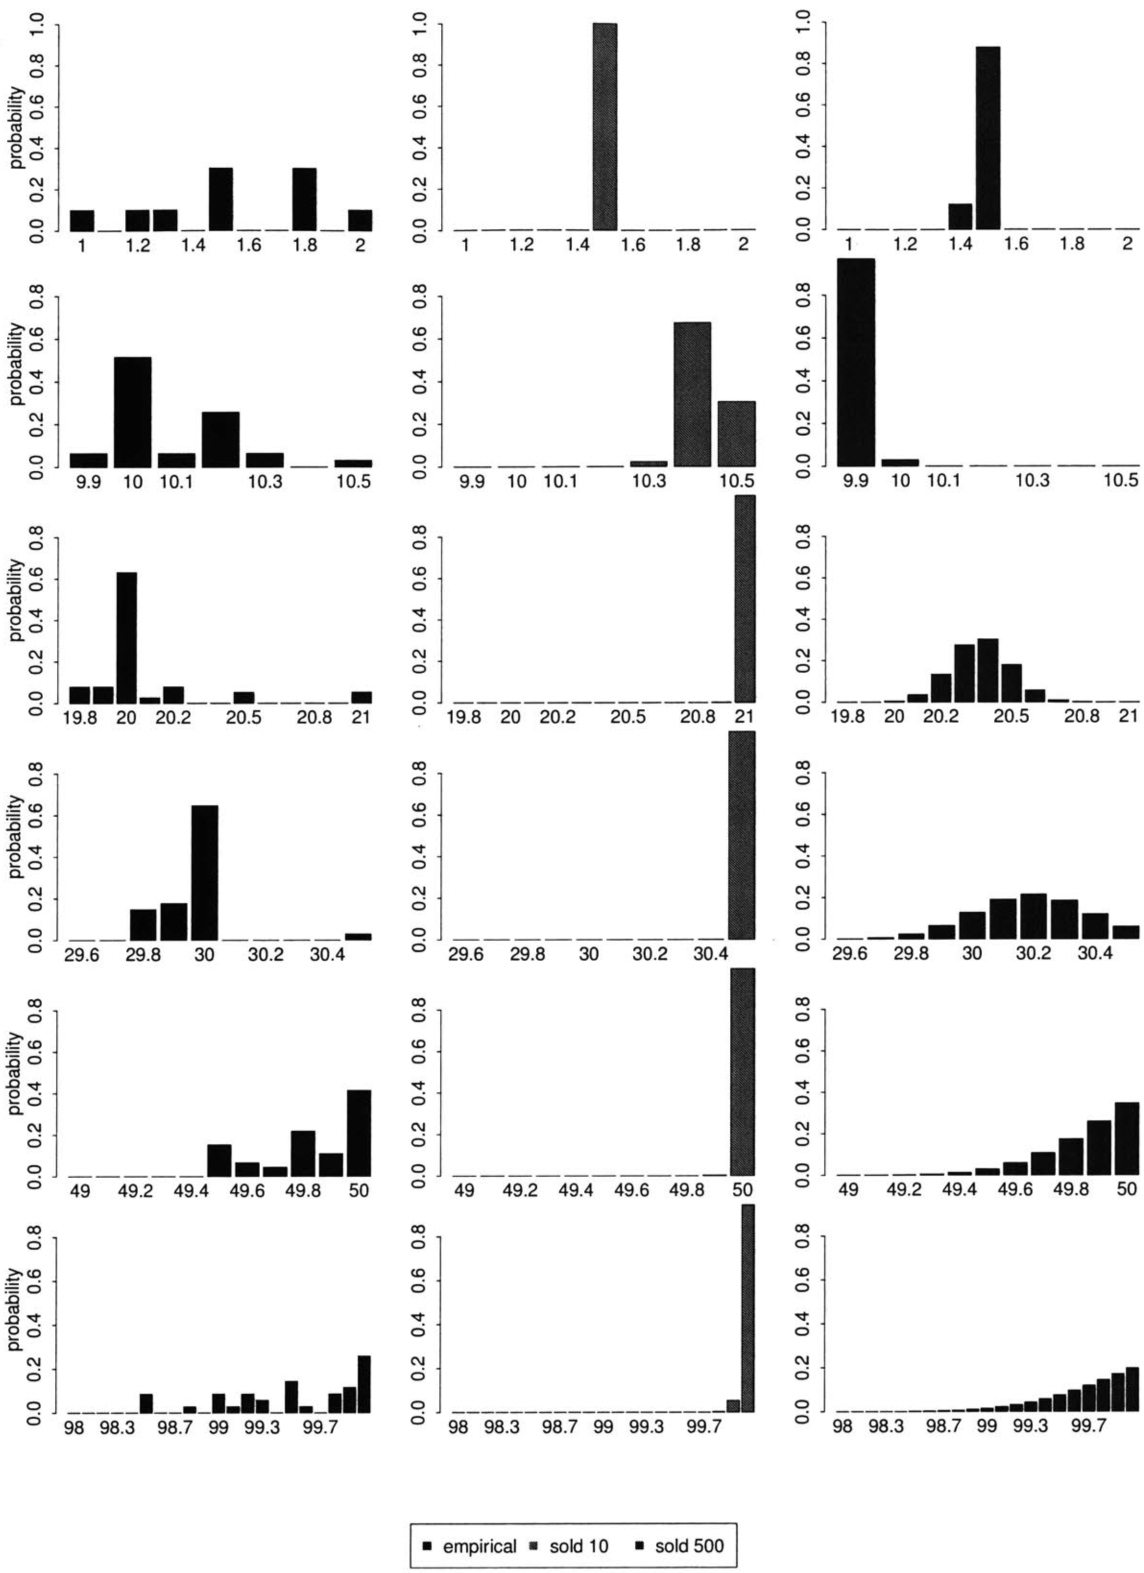

Figure 1-10: Empirical Price Distribution and Prior Price Belief

The parameter estimate for weight on prior, γ, is nearly one, meaning that updating plays very little role in this market. As it turns out, consumers learn about market information in the form of remembering the sellers and prices from previous purchases, as indicated by the parameter estimate for probability of recall, θ_s.[36] The findings imply that even though consumers remember their transaction prices, they only associate the prices with the corresponding sellers with whom they made the purchases and do not change their beliefs for other sellers. One possibility for the the perfect recall is perhaps the easy access to consumers' order history in this online marketplace.

1.7.2 Search Cost and Price Sensitivity

Next we turn to price and search cost parameters, which investigate whether consumers improve their search and purchase decisions. Table 1.7 presents the parameter estimates. The price coefficient results show that consumers become more price sensitive as they gain experience.

Table 1.7: Price Coefficient and Search Cost Estimates

Variable	Coeff.	SE
Price Coeff.		
β_0 – Constant	1.92	.13
β_1 – Past Exp.	.08	.01
Search Cost		
c_0 – Constant	.02	.01
c_1 – Past Exp.	-.0007	.36
σ_η – Variance	1.31	.35

Standard errors of estimates are computed using 100 nonparametric bootstraps.

The estimate for β_1 indicates that one additional past purchase corresponds to an 8% increase in price sensitivity. One explanation for this finding is that after a few purchases consumers may come to the realization that they could save a substantial amount of

[36]With only one previous purchase, the probability that a consumer perfectly recalls seller and price information from the previous purchase is 1-exp(-4.53)=.99. This probability becomes approximately one for consumers with more experience.

money by making some small savings in every purchase. On the other hand, search cost is fairly small and not affected by market experience. The cost associated with making a search is on average one cent.

1.7.3 Implications for Search Behavior

Translating these estimates into predicted search and purchase behaviors, I find that: first, as a result of the upward-biased price belief, consumers on average only search 1.6 sellers with the number of searches positively associated with experience due to an increase in price sensitivity. Second, consumers believe that popular sellers are more likely to offer better deals. However, with the noise in the utility function, their search order does not strictly follow the popularity of sellers. The randomness in the order they search is greater for inexperienced consumers because they are less sensitive of prices and therefore care less about finding the potentially cheapest seller to search. Last, the finding that consumers are able to perfectly recall information from previous transactions implies that experienced consumers can make more effective searches because they have already sampled a subset of sellers and can avoid the expensive sellers with whom they have previously encountered.

1.7.4 Model Fit

Since the goal of the empirical model is to provide explanations for experienced consumers' price advantages, it is important to ensure that the model captures the relationship between price and market experience found in the reduced-form analysis. To examine this, I plug in the parameter estimates to obtain the model predicted price and then estimate Equation (1.2) using the predicted prices as the dependent variable. Columns (3) and (4) in Table 1.8 present the regression results. The coefficients, going down from top to bottom in Column (3), indicate that consumers are predicted to generally pay a statistically significantly lower price over time. Price improvement increases from 2 cents to around 20 cents. The point estimates and standard errors in Columns (1) and (2) in the same table are estimates of Equation (1.2) copied from Columns (1) and (2) in Table 1.4. A comparison of the two regression results suggests that the magnitude of the predicted price improvement at different points in a consumer's purchase history is similar to that of the observed ones, implying that the empirical model does a fairly good job at matching consumer prices and capturing the relationship between prices and consumer experience.

1.7.5 Counterfactual Analysis

In Section 1.7.1 and Section 1.7.2, price sensitivity and information recall are shown to be the two important distinctions between experienced and inexperienced consumers. To examine the role each factor plays throughout the purchase history, I calculate simulated prices for two scenarios: one is that consumers cannot recall the exact information from their previous purchases (i.e., $\theta_s = 0$), and the other case is that consumers not only have no memory about past transactions but also remain as price insensitive as they are in their first purchase (i.e., $\theta_s = 0, \beta_1 = 0$).

I then regress simulated prices on the same exogenous variables in Equation (1.2). Columns (5) and (6) in Table 1.8 present results for the first scenario. Given that consumers do not remember the detailed information on prices and sellers from their previous purchases, coefficients of the experience dummy variables in Column (5) are still mostly negative, but their absolute value and significance level are noticeably smaller. Specifically, the price improvement is less than 5 cents until a consumer's 18th purchase, whereas the predicted price in Column (3) suggests that with only seven purchases, price advantage should be at least 5 cents when consumers can recall market information in their earlier purchase experience. With a considerable amount of market experience – more than fifteen purchases, consumers in general still enjoy a statistically significant price advantage that is about a third to a half of its usual size in Column (3). When consumers not only forget all the market information they acquired but also keep price sensitivity constant over time, the remaining price advantages of the very experienced consumers all

Table 1.8: Observed vs. Predicted Price

	ObservedPrice		*PredictedPrice*		*PredictedPrice* $\theta_s = 0$		*PredictedPrice* $\theta_s = 0, \beta_1 = 0$	
N^{th} Purchase	Coeff. (1)	SE (2)	Coeff. (3)	SE (4)	Coeff. (5)	SE (6)	Coeff. (7)	SE (8)
2	-.0021	.0082	.0018	.0051	.0063*	.0025	.0101***	.0025
3	-.0209*	.0100	-.0172*	.0077	.0035	.0037	.0115**	.0036
4	-.0307*	.0122	-.0206†	.0108	.0044	.0047	.0160**	.0048
5	-.0220	.0136	-.0269*	.0126	-.0119†	.0066	.0035	.0065
6	-.0495*	.0228	-.0514***	.0146	-.0132†	.0069	.0030	.0065
7	-.0361†	.0214	-.0335*	.0165	-.0013	.0098	.0217*	.0089
8	-.0543*	.0215	-.0590***	.0168	-.0194**	.0062	.0059	.0070
9	-.1091***	.0295	-.0964***	.0231	-.0225**	.0085	.0063	.0067
10	-.1219***	.0367	-.0909**	.0323	-.0066	.0119	.0352**	.0114
11	-.0717*	.0342	-.1066**	.0348	-.0297*	.0126	.0162	.0143
12	-.0982***	.0284	-.0738*	.0318	-.0282*	.0114	.0089	.0063
13	-.0769†	.0458	-.1036*	.0405	-.0072	.0130	.0358**	.0113
14	-.0731*	.0301	-.1446***	.0385	-.0144	.0179	.0171	.0173
15	-.1319***	.0411	-.1390***	.0368	-.0448*	.0179	.0227†	.0127
16	-.1406**	.0512	-.0990*	.0493	-.0213	.0311	.0212	.0279
17	-.0902†	.0473	-.0763†	.0401	-.0450**	.0163	-.0018	.0090
18	-.1474**	.0562	-.0871†	.0473	-.0515*	.0211	.0020	.0151
19	-.1126*	.0538	-.0977*	.0415	-.0301†	.0169	.0377*	.0170
20	-.1552**	.0491	-.1378*	.0567	-.0594†	.0310	.0077	.0201
21	-.1008	.0822	-.1200*	.0512	-.0345*	.0172	.0246	.0203
22	-.1640***	.0508	-.1573**	.0497	-.1183***	.0220	.0290	.0223
23	-.0801*	.0399	-.1629***	.0098	-.0816***	.0086	.0058	.0210
24	-.1863*	.0796	.0871	.2621	.0680	.0691	.0716	.0686
25	-.0919**	.0325	-.2223***	.0200	.0130	.0100	.0494***	.0103
26	-.1704***	.0352	-.1809***	.0225	-.0794***	.0117	.0839***	.0114
27	-.1704***	.0352	-.1809***	.0225	-.0532***	.0117	.1470***	.0114
Observations	4116		4116		4116		4116	

† $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

consumers' first purchases are omitted

disappear, as shown by estimation results of the second scenario presented in Columns (7) and (8).

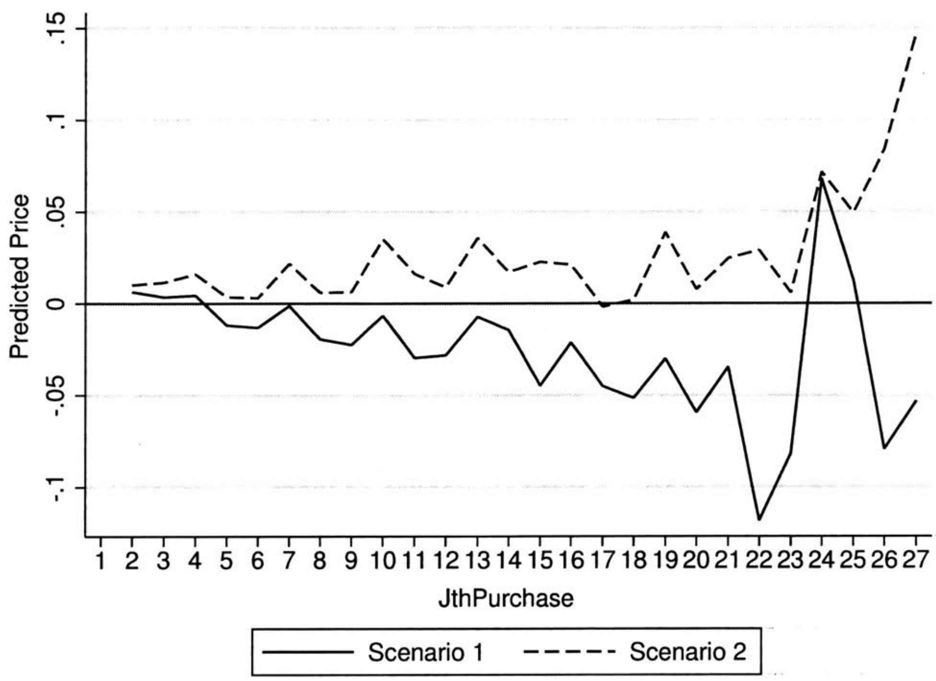

Figure 1-11: Counterfactual Price History

To better compare prices in the two scenarios, Figure 1-11 plots the results in the two scenarios. Together they suggest that consumers' initial price improvements are mostly attributed to the market information acquired in earlier purchases, while price sensitivity plays an increasingly important role, contributing between 30% and 50% to the price advantages in consumers' later purchases. The finding that price sensitivity is the ultimate driver for price improvements is interesting and also intuitive: in a market where sellers are minimally differentiated and sellers' popularity is perceived as a signal of low prices, sellers, once establish a good sales record, have an incentive to increase their prices, in which case consumers' knowledge of their previous sellers and prices may no longer be so useful.

1.8 Conclusion

This paper takes advantage of a very special product market which eliminates potential unobserved seller and product heterogeneities exist in many other markets for seemingly homogeneous goods. Fortunately, I am able to observe every single consumer's purchases made in the market over a period of several months, with many of the them being repeat purchases. With data on the purchase history of individual consumers, I find that market experience significantly changes consumers' search and purchase behaviors, making them savvier shoppers. Reduced-form evidence implies that price improvements stem from market experience and occur in a passive way in this market.

Based on these findings, I incorporate consumer learning into a structural search model and extend recent works by relaxing the distribution assumption imposed on prior belief and allowing consumers to form different prior beliefs for different sellers based on seller-specific information. These methodological differences are shown to be empirically important because prior price beliefs are found to be biased upwards, suggesting that the rational expectation assumption will lead to biased estimates. The estimates of the structural model also indicate that consumers have different prior price beliefs for different sellers: they view sellers' popularity as an important signal of low prices.

My results from the structural estimation shed light on how consumers learn from their experience in homogeneous goods markets, which is a topic that has board implications for firms and policy makers and yet has not been well understood. Specifically, I find that experience helps consumers to become familiar with a group of sellers and their prices, which is an important market advantage because knowing the market information helps consumers to make more effective price searches. Admittedly, the extent to which consumers remember their previous sellers and prices varies across markets: the access to an order history in many online marketplaces facilitate seller and price recall. Nevertheless, the counterfactual analysis of this paper suggests that ultimately what makes consumers' price advantage more sustainable is consumers' increasing price sensitivity.

Chapter 2

Do Online Reviews Affect Quality? Evidence from Yelp.com

2.1 Introduction

The Internet has greatly reduced the cost of communication between people. Complementing the offline word-of-mouth communication, online consumer review platforms, such as Yelp and TripAdvisor, are now an important source of information on product quality to consumers. For instance, Yelp, a crowd-sourced review site for local businesses, had grown into an online community with 168 million unique monthly visitors and 108 million reviews by the end of 2016.[1] This new phenomenon attracts researchers' attention and has led to a significant body of literature studying how consumers react to merchants' online feedback in terms of purchase probability and willingness to pay (e.g. Chevalier and Mayzlin, 2006; Resnick et al., 2006; Jin and Kato, 2006).[2] A majority of these studies finds that sales and prices are positively associated with online reputation. In the case of Yelp.com, Luca (2011) shows that a one-star increase in Yelp rating raises revenue by 5%-9% for independent restaurants.

[1] `https://www.yelp.com/factsheet`

[2] See Tadelis (2016) for a summary of theories behind the reputation mechanisms and empirical works on impact of the feedback systems on online platforms.

Given the substantial evidence that reviews affect demand, this paper investigates whether this new form of quality disclosure – online reviews – also incentivizes the supply side to improve quality in the restaurant industry. My analysis exploits a plausible source of variation in the degree to which restaurants are affected by online reviews: the amount of tourism business in the area where restaurants are located. Since tourists have little local information, we can think of them as customers with high search costs.[3] Reviews available on the Internet can potentially decrease their search costs. Indeed, consumer surveys find that tourists rely more than locals on reviews to inform them about local choices and reputations.[4] Therefore, all else being equal, a restaurant in a dense tourist area is affected by its online reputation more than is a restaurant in a residential area, where customers can also learn about the quality of a restaurant based on their own dining experience, from talking to friends, or through local media.

In an attempt to control for all the other characteristics that may affect restaurant quality, I only include chain-affiliated restaurants in the analysis. My empirical strategy is to first compare the quality of restaurants affiliated with the same chain but located in areas with different degrees of tourism business, in the early days when online reviews did not have much impact. With the baseline quality comparison, I then further implement a difference-in-differences strategy, in which I examine the within-chain differences in quality change of units affected by online reputation to different extents, over the period when online review sites, such as Yelp, soared in popularity. Specifically, the two differences are: the within-chain difference in locations – tourist areas versus residential areas, and the difference over time – a continuous version of the "before and after treatment".

The research design is implemented in the following ways. First, with eleven-year Yelp review data dating back to 2005, when there were only 12,000 reviewers on Yelp, the early ratings in the data are used as the baseline quality measure reflecting restaurant quality

[3]George Stigler and Peter Diamond refer to them as a class of consumers lacking in knowledge about local markets.

[4]Forinstance, seehttp://searchengineland.com/case-study-33-of-local-search-business-is-from-non-locals-and-7-tips-for-capturing-that-traffic-253697

when online reviews had little impact on demand.[5] Second, I use individual ratings in the data to trace out the change in restaurant quality over time. Third, I restrict the sample to restaurants with more than one location in Las Vegas – one of the most popular tourist destinations that has one area concentrated of hotels and casinos called the Strip; I use the distance between a restaurant and the Las Vegas Strip as a measure of its degree of tourism business.

Since businesses would have few reputational incentives when most of their customers are tourists, who lack local information, we would expect that within a restaurant chain, units closer to a tourist attraction would have worse or at most similar quality to units in residential areas. Along these lines, the common trend assumption of my analysis is that restaurants closer to the Strip area exhibit a similar or decreasing pattern in Yelp ratings relative to restaurants that are affiliated with the same chain and are located further away from the Strip, when online reputation is nonexistent or ineffective.

The estimated results suggest that the smallest chains did not initially have a low quality problem at their locations close to the Strip before online reviews had much impact, and consequently, I do not find evidence that online reviews have had any impact on them later on. On the other hand, consistent with prior empirical findings that the degree of repeat customers affects the strength of reputational incentives (Jin and Leslie, 2009), I show that for larger chains, their units closer to the Strip had a lower overall customer satisfaction, when online reviews were not widely used. Note that this finding also supports the aforementioned assumption to some extent. Furthermore, for chains with a moderate size – from a few dozen to a few hundred, their units closer to the Strip improved significantly more in ratings, during the years when Yelp gained popularity. Taken together, the systematic within-chain difference in baseline quality and quality improvement for the mid-sized chains provides evidence that online reputation is an effective

[5] According to Yelp's Wikipedia page, the number of reviewer grew to 100,000 in 2006, and Yelp grew from six million monthly visitors in 2007 to 16.5 million in 2008. This number had increased to 121 million by the end of 2016.

reputation mechanism for chains with an average size. In contrast, no such evidence is found for the large chains that have more than a few hundreds of locations. The results indicate that the large chains actually experienced a significantly steeper decline in Yelp ratings at their units located closer to the Strip.

The rest of the paper is organized as follows. In Section 2.2, I provide a summary of the related literature on firms' responses to reputational incentives and the relation between degree of repeat customers and strength of reputational incentives. In Section 2.3, I present the summary statistics of the Yelp data set and define the variables used in the analysis. In Section 2.4, I describe the research design in detail. Section 2.5 presents the estimated results and a discussion on them. Section 2.6 concludes the paper.

2.2 Literature Review

2.2.1 Supply-side Responses to Reputational Incentives

Since we can think of online ratings as akin to required disclosure of some quality measure, this paper adds to the stream of literature that studies the impact of mandated quality disclosure on firms' behaviors. Several prior studies find a positive impact on quality. Jin and Leslie (2003) provide evidence that mandating restaurants to display hygiene quality grade cards cause restaurant health inspection scores to increase. Along the same lines, Jin and Leslie (2009) show that franchised units of a given chain in general have lower hygiene scores than company-owned units, and interestingly, this difference disappears in the presence of grade cards. They conclude that this is evidence of franchisee free-riding on chain reputation.[6] Similarly, Bennear and Olmstead (2008) examine how Massachusetts drinking water suppliers responded to the Safe Drinking Water Act, which mandated disclosure of contaminant levels. They find that total violations and severe

[6]To some extent, this result also supports the view that online reviews have a weaker effect on large chains since units of a given chain share the same reputation.

health violations decreased by more than 30 percent after larger utilities were required to mail consumer confidence reports directly to customers.

While a number of researchers find that quality disclosure provides firms strong incentives to improve quality in various markets, others also point out that it may encourage gaming activities, especially when the measured quality does not cover all the dimensions of quality. Considerable evidence is found in education and healthcare (e.g. Jacob and Levitt (2003); Lu (2012)).[7] In the same vein, McCarthy and Darden (2017) find that quality disclosure leads firms with good ratings to charge a higher premium in the health insurance market.

2.2.2 Reputational Incentives and Repeat Customers

In addition to changes in reputational incentives due to government intervention, several studies also present empirical evidence of the positive association between the frequency of repeat customers and the strength of reputational incentives. This topic is most studied in the franchise business, in which franchisees and franchisors may not have aligned interests depending on the level of repeat business. As Robert Martin puts it, "in order to preserve the brand-name capital, which is the principal asset in many cases, the franchisor must monitor the franchisee's quality choice. The need to monitor the franchisee's quality choice declines as the proportion of repeat customers increases. The market will monitor the quality choice when transient customers are rare."

Brickley and Dark (1987) provide the descriptive evidence that is consistent with this view. Investigating the effect of the level of repeat customers on the ownership structure of franchise companies, they find that franchising appear to be more frequent in industries with more repeat customers, which include lawn-care companies, automotive services, and pet stores, as opposed to restaurants, hotels and motels, and auto rentals.

[7] See Dranove and Jin (2010) for an excellent summary of this literature.

Similarly, Brickley (1999) finds that area development agreements are significantly more likely to be used by franchisors involved in industries with no repeat customers. With a more direct approach, Jin and Leslie (2009) find that "regional variation in the degree of repeat customers affects the strength of reputational incentives for good hygiene at both chain and nonchain restaurants". In a more controlled setting, List (2006) presents evidence that enhanced consumer learning causes sellers to provide higher quality products in a series of laboratory and field experiments. This paper contributes to the literature by not only empirically showing a positive relation between the level of repeat business and restaurant quality, but also presenting the evidence that the new form of quality disclosure – online reviews – can significantly enhance reputational incentives in areas with few repeat customers.

2.3 Data Summary

The data used in the study is from the publicly available Round Seven Yelp Dataset Challenge, which encourages academics to study Yelp-related projects.[8] It contains 4.1 million Yelp user reviews for around 14,400 businesses in several major cities around the world, as well as detailed information on business attributes and reviewer characteristics. In the study, I focus on all reviews for restaurant chains that have more than one location in the Greater Las Vegas area for reasons explained in Section 2.4.[9] This reduces the sample size to 2,340 restaurants affiliated with 369 chains, and a total of 152,085 reviews.

For each restaurant, I observe the business ID, name of the restaurant, address, price range, type of cuisine, number of reviews, average rating rounded to the nearest half star on a scale of one to five, whether it is still in business, and various restaurant attributes, such as business hours, whether it takes reservation, and noise level. Table 2.1 presents the summary statistics of some key restaurant attributes. As expected, the majority of

[8] `https://www.yelp.com/dataset_challenge`

[9] The restaurant category consists of restaurants, coffee shops, and bars.

restaurants chains are inexpensive; a meal in at least half of the restaurants costs less than $10 per person according to Yelp.[10] A restaurant generally has a few hundred Yelp reviews and an average rating of 3.1 stars. The average number of years that a restaurant has been active on Yelp is 4.4 years, with a few restaurants available throughout the entire eleven-year data period. Nearly 90% of them were still in business by the end of the data period.

Table 2.1: Summary Statistics of Restaurant Attributes and Yelp User Characteristics

	Mean	Stdev	Min	Median	Max
Restaurant Attributes					
Price Range	1.4	0.5	1	1	4
Number of Reviews	381	250	1	330	824
Average Rating	3.1	0.7	1	3	5
Duration (years)	4.4	2.4	0	4.4	10.8
Is Open	0.9	0.3	0	1	1
Distance to Strip (miles)	5.2	3.2	0	5.4	19.5
Size of Chain	676	3,626	2	500	44,690
User Characteristics					
Number of Previous Reviews	133.4	278.1	0	33	8529
Average of Previous Ratings	3.7	0.7	0	3.8	5
Yelping Since	Sep 2011	2.1 yrs	Oct 2004	Sep 2011	Dec 2015

In addition to the Yelp dataset, I obtain the shortest distance from a restaurant to the Strip – the center of Las Vegas tourist activity. Figure A-1 in the appendix is a map of Las Vegas, in which the Strip is represented by the black line segments in the middle. The mean distance to Strip is 5.2 miles. To collect data on the size of a chain, I use each restaurant chain's Wikipedia page and its own webpage for reference.[11] About 26% of the chains are national or multinational chains that have units outside the state of Nevada; the rest are regional chains with a few units in the Greater Las Vegas area. Although there are fewer national chains than local chains, national chains own 1,512 units out of

[10] On Yelp, a dollar sign ($) indicates cost per person less than $10, a two-dollar sign ($$) corresponds to a price range between $11 and $30, a three-dollar sign ($$$) suggests a meal per person is between $31 and $60, and a four-dollar sign indicates it costs at least $61 per person.

[11] The size of a chain is defined as the number of locations worldwide, including both company-owned units and franchised units.

the total 2,340 units in the data. Comparing the average rating of restaurants affiliated with national chains and local Nevada chains, the histogram in Figure 2-1 shows that regional chains' units have slightly better Yelp rating than units affiliated with national chains.

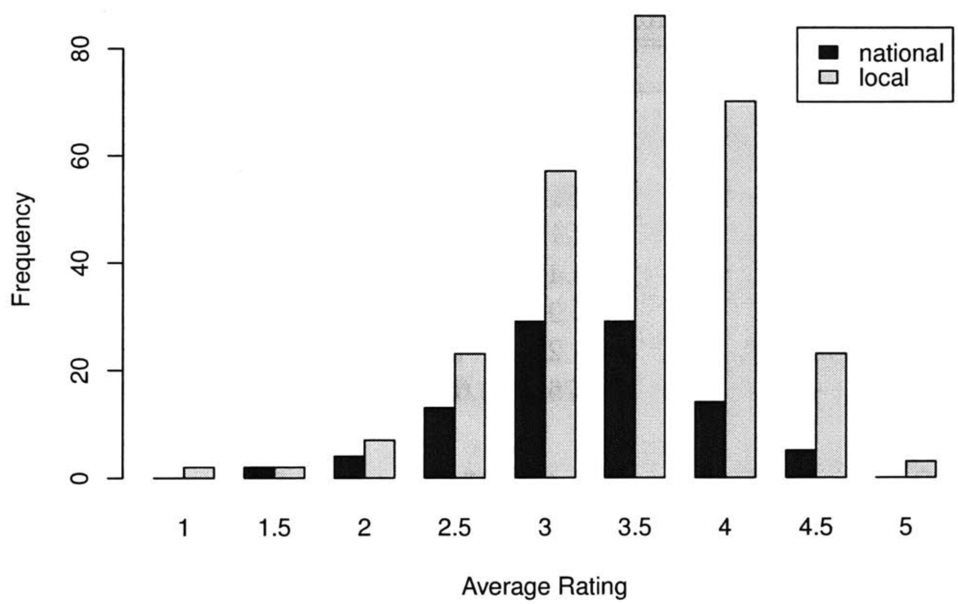

Figure 2-1: Histogram of Average Ratings of Regional Chains and National Chains

The average size of a restaurant chain is 676 units. This large number is driven by a few multi-national chains, like Subway, McDonald's, and Starbucks, that have tens of thousands of locations worldwide. To help visualize where units of small and large chains are located, Figure A-2 plots all the restaurants on a map with different colors indicating the log size of the affiliated chain and a blue diamond-shape symbol denoting the center of Las Vegas Strip. The figure demonstrates that a given neighborhood often has a good mix of restaurants affiliated with chains with different sizes. Furthermore, it turns out that units affiliated with a chain are also scattered in terms of their distance to the Strip. Presented in Figure 2-2 is the histogram of the mean absolute difference in distance to the Strip for any two randomly chosen restaurants within a chain, in which each observation is a restaurant chain.

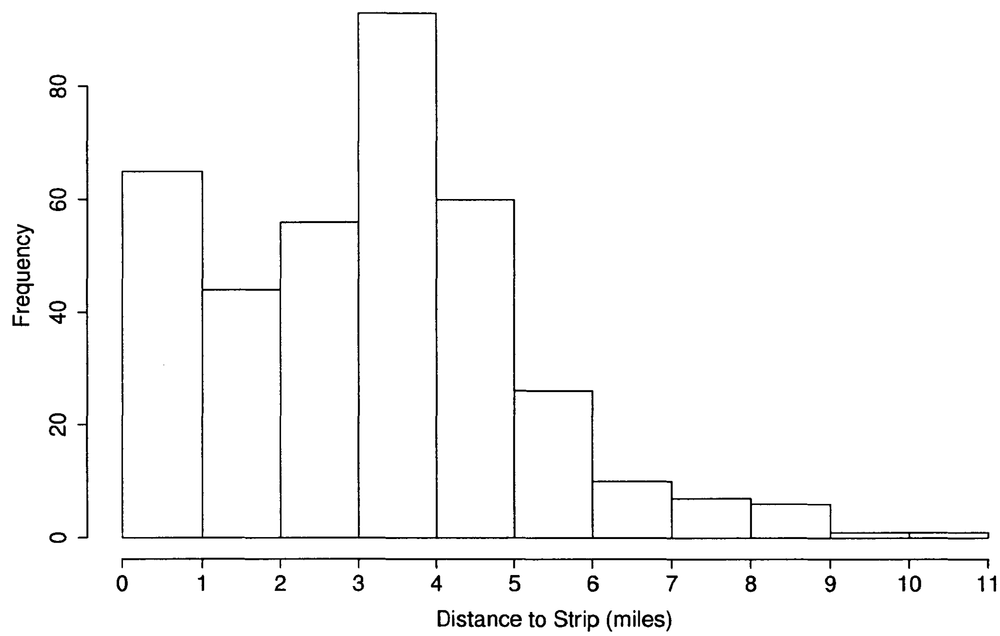

Figure 2-2: Histogram of Mean Absolute Difference of Distance to Strip Within a Chain

Moreover, variation in average rating is also common within a chain. The mean of absolute difference in average rating between two randomly chosen units affiliated with the same chain is 0.48 star on average, compared to a mean absolute difference of 0.83 star between two randomly chosen restaurants. Figure 2-3 presents the histogram of the mean of the absolute difference in rating between two units within a chain.

For each review, I observe the reviewer's user ID, the business ID of the restaurant being reviewed, the rating on a scale of one to five, the date of the review, the average rating of the user's past reviews, the number of reviews that the user wrote in the past, the month and year of the user's first Yelp review, and the number of times this review has been voted funny, useful, or cool by others. Summary statistics of some important user characteristics are presented in the second panel of Table 2.1. A few observations are as follows. First, Yelp reviewers are very willing to offer their own opinions – the

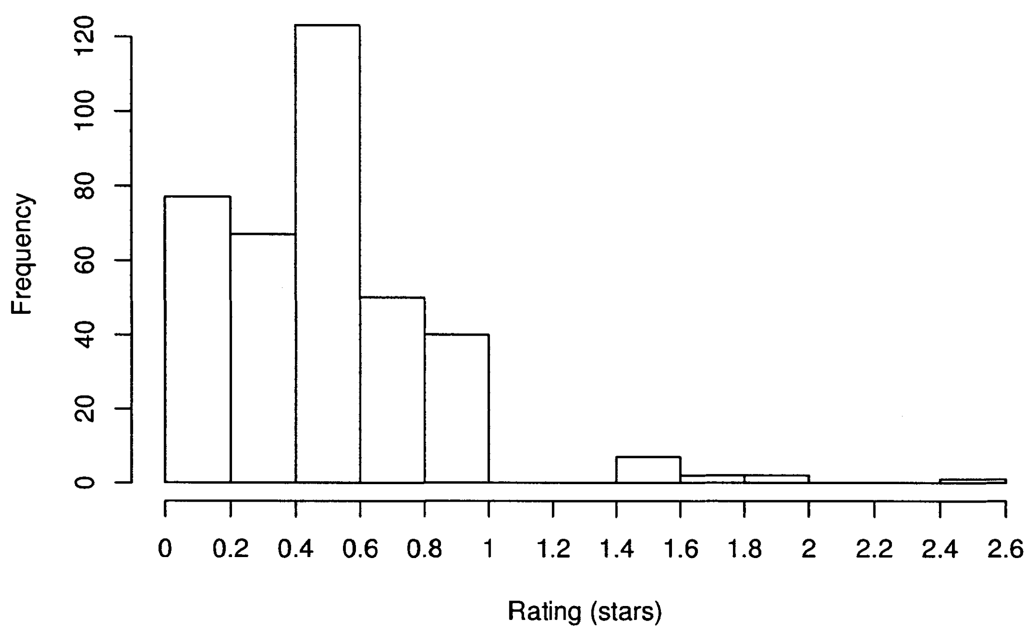

Figure 2-3: Histogram of Mean Absolute Difference of Average Rating Within a Chain

average number of reviews a user has posted is 133.4 reviews with a standard deviation about twice as large as the mean. Second, Yelp users are not too stringent in ratings; the average ratings that Yelp users give have a mean of 3.7 stars with a modest difference across users. Third, the earliest time users started to post Yelp review dates back to October 2004, the inception of Yelp, and about half of the users in the data have written a review before September 2011.

Table 2.2 provides the definition of the variables used in Section 2.4 and Section 2.5. The majority of the variables represent different restaurant attributes. The third to last is the variable for the mean of a user's previous ratings on Yelp. The last two variables are used to account for the time of a review.

Table 2.2: Definition of Variables

Variables	Definition
Duration	Number of years between the restaurant's first and last review
Price	Price range of a restaurant (See Footnote 10 for details)
AverageRating	Average rating of a restaurant
Distance	Distance between a restaurant and the center of the Strip
Size	The number of locations a restaurant chain has worldwide
National	Dummy equaling 1 if the chain has locations outside Nevada
ReviewerAverage	Average ratings of a Yelp reviewer's previous reviews
Date	The date of the review
TimePassed	Number of years between this review and the restaurant's first review

2.4 Empirical Strategy

Tourists are less familiar with local businesses than natives. As a consequence, without any effective form of quality disclosure, it is expected that restaurants in areas with fewer local repeat customers have weaker incentives for high quality. Jin and Leslie (2009) provide evidence that supports this view. Now that online reviews play an important role in consumers' everyday decisions, the businesses that traditionally had little reputational incentives are now affected the most by their online reputation because surveys find that tourists rely more on peer review sites, such as Yelp, for restaurant recommendations. Therefore, location is a plausible source of variation in the degree to which restaurants are affected by online reputation. All else equal, a restaurant in a residential location is affected by its online reputation less than a restaurant in a tourist location.

Despite the potential impact of reputational incentives on quality, from both online reviews and local word-of-mouth communication, restaurant quality is predetermined by many intrinsic restaurant characteristics that are hard to control. To take all the restaurant characteristics that may affect quality into account, my approach is to only focus on chain-affiliated restaurants to utilize chain fixed effect and investigate within-chain difference across units located in areas with different proportions of nonlocal customers.

With the exponential growth of the online review industry in the past decade, a natural approach is a difference-in-differences estimation, in which I compare the quality improvement of restaurants that are affected by online reputation to different extents over the years as online peer review sites gain popularity. Specifically, accounting for common trend of a chain and common trend of a zip-code area, I investigate whether the quality of units in a tourist location improve significantly more than units located in a residential area within the same chain. The analysis hinges on the assumption that without effective online reputational incentives, units close to a tourist attraction should have similar, if not worse, quality trend relative to units affiliated with the same chain but located in areas with more repeat business.

An ideal empirical setting to implement this approach is a city with one single well-defined tourist area so that the level of tourism business radiates from the tourist attraction.[12] Las Vegas is such a city. It has a 4.2-mile stretch, called the Strip, known for its concentration of casinos and resort hotels.[13] Figure A-3 is a map of Las Vegas made by a data artist Eric Fischer who uses geotagging data from Flickr, a popular image hosting website, to map out the areas in which locals and tourists take photos. In the map, photos taken by tourists are shown as red dots, while photos taken by locals are shown as blue dots. The yellow dots are cases where Fischer couldn't determine whether the photographer was a tourist or a local. In the middle of the figure, there is a very distinct red line connected by many dots going through the center of the Strip, denoted by a black diamond-shape symbol. The very concentrated red dots precisely trace out the area of Las Vegas Strip, while blue dots are scattered all over the map. For comparison, Figure A-4 displays a map of where tourists and locals take photos in Boston, a city that attracts many tourists as well. In contrast to the map of Las Vegas, the red dots in the map of Boston are congregated in several regions mixed with blue dots, making it difficult to determine the proportion of tourists in a given area.

[12]Brickley and Dark (1987) propose another example of where to expect fewer repeat customers – restaurants near freeway exits.

[13]`https://en.wikipedia.org/wiki/Las_Vegas_Strip`

Because of the unique geographic pattern in Las Vegas, tourists are mostly concentrated on the Strip and locals are spread out over the Greater Las Vegas area. I use the distance between a restaurant and the Strip as a measure of the degree of tourism business in the area, in which a closer location corresponds to a more touristy area. Figure 2-2 in the Data Summary section demonstrates that there exists a large amount of variation in the degree of tourism business across different units within a chain. That is, some are closer to the popular tourist destination, the Strip; others are further out in the suburb with mostly local customers.

A quality measure is tricky to obtain for two reasons. First, peer review sites may suffer from fraud or biased reviews, while restaurant guides by food critics only cover a small subset of restaurants.[14] Second, restaurant quality has multiple dimensions, which include hygiene, food quality, taste, services, and perhaps even atmosphere, not to mention that these measures are sometimes subjective. Without a better measure, I use Yelp ratings as the restaurant quality measure for the following two reasons. First, Yelp is one of the most influential crowd-sourced review sites that customers turn to when they need restaurant recommendations.[15] Second, despite the allegation of review manipulation on Yelp, Yelp rating is a good indicator of customers' overall satisfaction and reflects restaurant's quality.[16]

To further validate that Yelp rating is a fair representation of restaurant quality, Table 2.3 presents the variance decompositions of restaurants' average rating. The number of active years on Yelp and the price range dummies together only explain 0.8% of the average rating variation, as shown in Columns (1) and (2). Columns (3) and (4) report

[14]See Luca and Zervas (2016) for an interesting investigation of the credibility of Yelp reviews.

[15]In addition to Luca's finding, a 2012 study by Anderson and Magruder show that an increase from 3.5 to 4 stars on Yelp resulted in a 19 percent increase in the chances of the restaurant being booked during peak hours.

[16]Luca (2011) argues that if gaming is present on Yelp.com, then one would expect ratings to be clustered just above the cutoffs so that they could be rounded up to the next half star. He uses McCary density test to rule out this possibility. Moreover, Luca and Zervas (2016) find that although not all reviews are credible and Yelp filters 18% of its reviews, chain restaurants are less likely to commit review fraud.

that including type of cuisine dummies explains nearly an additional 14.5% of the average rating variation and makes the coefficients on price dummies no longer significant. Moreover, results in Columns (5) and (6) suggest that the 5-digit zip code fixed effects explain another 3.3% of the average rating variation. More importantly, the last tow columns of Table 2.4 show that chain fixed effects, together with the number of active years on Yelp, explain 46.5% of the average rating variation. This suggests that Yelp does a fairly good job at capturing the systematic differences across restaurants.

Table 2.3: Variance Decomposition of Average Rating

	Coeff. (1)	SE (2)	Coeff. (3)	SE (4)	Coeff. (5)	SE (6)	Coeff. (7)	SE (8)
Duration	0.018^{**}	0.006	$0.011^{\dagger}$	0.006	0.010	0.006	0.019^{**}	0.006
Price ($$)	0.103^{***}	0.033	-0.010	0.037	0.002	0.037		
Price ($$$)	0.279^{*}	0.140	-0.044	0.145	0.073	0.145		
Price ($$$$)	0.376	0.305	0.101	0.316	0.223	0.312		
Type of cuisine FE			Yes		Yes		Yes	
Zip code FE					Yes		Yes	
Chain FE							Yes	
Adjested R^2	0.008		0.145		0.178		0.465	
Observations	2340		2340		2340		2340	

$\dagger$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Restaurants below $10 are omitted
The price dummies drop out of the last specification due to collinearity with chain fixed effect.

Since Yelp has been growing rapidly in both the number of users and reviews since its inception in October 2004, I use the change in ratings to measure restaurant's quality improvement as a response to the increasing impact of online reputation.[17] In the data, there is a fair amount of changes in restaurants' ratings over time. Simply regressing individual ratings of a restaurant on the time of ratings gives the slope of the rating trend. Figure 2-4 displays the histogram of the average change in ratings per year for restaurants that have been active on Yelp for at least one year and receive more than ten reviews per year. It shows that more restaurants have gone through a decline in ratings than an increase, although the majority of them experience mild changes.

[17] Figure A-5 is the number of cumulative reviews on Yelp since 2005.

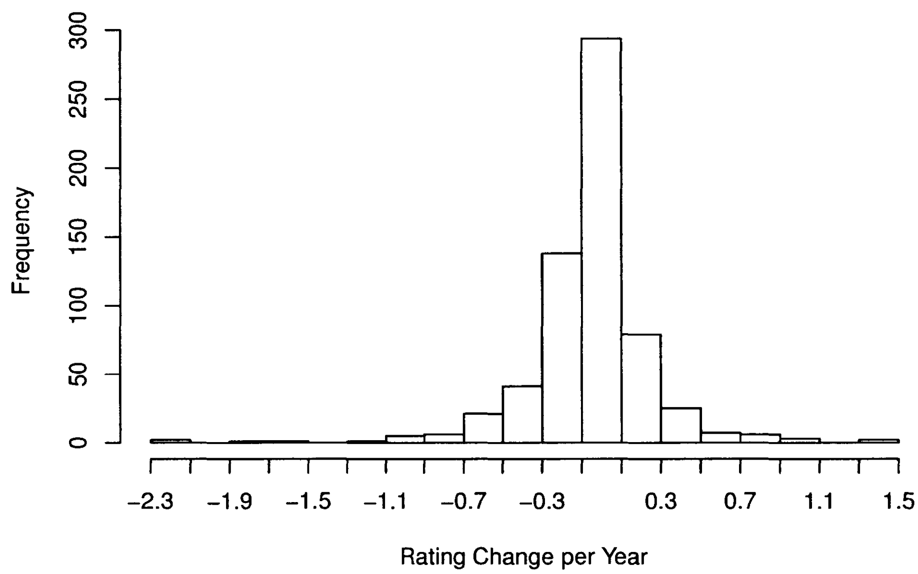

Figure 2-4: Histogram of Change in Average Rating

The measure of quality improvement – slope of ratings – has some desirable features. First, by using individual reviewers' ratings, I can make use of the available data on a Yelp user's average rating of previous reviews to control for the stringency in ratings that might vary significantly across individuals. In a paper that purposes an optimal way to aggregate customer rating, Dai et al. (2012) show that reviewer fixed effects alone account for 23.3% of the total variations in Yelp ratings. Furthermore, the trend of individual ratings over a long period of time is much more likely to be immune to review manipulation.

To formalize the approach, I run a DD regression with the following equation:

$$\begin{aligned} Rating_{ijrt} = \beta_0 &+ \beta_1 TimePassed_{it} + \beta_2 Distance_i TimePassed_{it} \\ &+ \beta_3 Size_j TimePassed_{it} + \beta_4 Distance_i Size_j TimePassed_{it} \\ &+ \beta_5 ReviewerAverage_r + \sum_{z=1}^{63} \gamma_z Date_t + \sum_{j=1}^{369} \eta_j Date_t + \epsilon_{ijrt}, \end{aligned} \tag{2.1}$$

where $Rating_{ijkt}$ is the rating, on a scale of one to five, that reviewer r gives for restaurant i affiliated with chain j at time t. For the explanatory variables, first I control for the common trend of restaurants in the same zip code by including 63 coefficients, γ's, and the common trend of restaurants in the same chain by including 369 coefficients, η's. Second, I account for the restaurant's number of years on Yelp at time t, $TimePassed_{it}$, and I interact this time variable with the distance variable denoting the miles between the restaurant and the Strip, $Distance_i TimePassed_{it}$, with the size of the affiliated chain j, $Size_j TimePassed_{it}$, and with both variables together, $Distance_i Size_j TimePassed_{it}$.[18] Last, I also include reviewer r's average ratings on Yelp, $ReviewerAverage_r$. The parameters of interest are β_2 and β_4, which indicate whether units closer to the Strip have improve their quality more.

2.5 The Effect of Online Reputation

2.5.1 Level of Tourism Business and Early Yelp Ratings

As indicated by the number of cumulative reviews on Yelp since 2005 in Figure A-5, online reviews did not immediately start its exponential growth in the early days. The number of users and reviews in the first few years is dwarfed by the current numbers, perhaps not only because of the huge success of the industry in recent years, but also due to the rapid growth in the use of smartphones. Therefore, I first use early Yelp ratings to study the relationship between restaurant quality and level of tourism business within a chain before online reviews had much impact.

Specifically, I restrict the sample to restaurants that had Yelp reviews before 2008 and

[18] I also control for the size of the affiliated chain because the impact of online reviews may depend on the size of the chain. Luca (2011) finds that Yelp has a significant effect on the demand of independent restaurants but not on that of chain-affiliated restaurants.

calculate their average ratings at the end of 2008.[19] And I report estimates of the average rating of restaurant i affiliated with chain j, from the following equation:

$$\begin{aligned} AverageRating_{ij} = \beta_0 + \beta_1 FirstReview_i + \beta_2 Distance_i + \beta_3 Distance_i Size_j \\ + \alpha_j + \epsilon_{ij}, \end{aligned} \tag{2.2}$$

where I control for the time of restaurant i's first review, $FirstReview_i$, the distance from restaurant i to the Strip, $Distance_i$, the distance interacted with the size of the affiliated chain j, $Distance_i Size_j$, and chain fixed effects, α_j.

The estimates are presented in Columns (1) and (2) of Table 2.4. The coefficient of the variable measuring the distance to the strip is insignificant, while the coefficient of the interaction term is statistically significantly positive. As in Jin and Leslie (2009), we expect quality to be lower closer to the Strip, where there are fewer repeat customers. My results indicate that although this is true for larger chains, the quality of the smallest chain's unit is independent of its distance from the Strip even during the period when online reputation did not have much impact. A plausible reason could be that such chains might have all units owned by the same person and/or internalize how low quality would affect the chain's reputation.

Furthermore, I divide the sample into three subsamples: the smallest one-third of chains, the middle one-third of chains, and the largest one-third of chains. The corresponding chain sizes of the three groups are: less than 5, between 5 and 150, and greater than 150. I then run the same regression as in Equation (2.2), except that I omit the interaction term between distance and size. Columns (3) through (8) report the estimated coefficients. Although none of them is significant perhaps due to the small number of observations, the signs of the coefficient of distance in the three different specifications

[19]In total, there are 301 restaurants in the data that had reviews before 2008. If the cutoff is changed to 2007, the sample is further reduced to only 78 restaurants.

Table 2.4: Average Yelp Ratings Before 2009

	Coeff. (1)	SE (2)	Coeff. (3)	SE (4)	Coeff. (5)	SE (6)	Coeff. (7)	SE (8)
FirstReview	-0.110	0.104	0.355	0.255	-0.236	0.155	-0.103	0.167
Distance	-0.005	0.024	-0.046	0.053	-0.009	0.037	0.048	0.038
Distance * Size	0.065*	0.029						
Full Sample	Yes							
Smallest $\frac{1}{3}$ of Chain			Yes					
Middle $\frac{1}{3}$ of Chain					Yes			
Largest $\frac{1}{3}$ of Chain							Yes	
Adjested R^2	0.24		0.23		0.35		0.19	
Observations	301		120		80		101	

† $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
The variable *Distance* ∗ *Size* is rescaled by dividing 10,000 from the variable

are consistent with the results found using the full sample: for units of large chains, the closer to the Strip, the lower the quality; for small chains, there is no such association. It is worth noting that to some extent, these findings support the assumption that without the impact of online reviews, quality trend of units with high degree of repeat business is similar or stronger relative to units located in areas with more one-time customers.

2.5.2 Level of Tourism Business and Rating Trend

Next I turn to investigate whether ratings have improved the most in areas where online reviews play the biggest role. Specifically, I estimate equation (1) in which each observation is weighted by the inverse of the total reviews that the restaurant receive on Yelp so that every restaurant receives equal weight. The estimates are reported in Columns (1) and (2) of Table 2.5. Coefficient β_2 is insignificant, suggesting that for the smallest chains, the change in ratings is not systematically different for units located within different distances to the Strip. On the other hand, coefficient β_4 is estimated to be 0.0026 with a p-value less than 0.001. This indicates that for units affiliated with a chain that has 10000 locations worldwide, the units that located 10 miles away from the Strip improve at a rate of 0.026 star per year faster than the units on the Strip. In addition, consistent with earlier findings, the mean of a reviewer's earlier ratings is significantly, positively associated with one's current rating.

As shown in Section 2.5.1, the within-chain difference in restaurant quality varies systematically across chains with different sizes even before the widespread use of online reviews. Therefore, I further divide the sample into three subsamples according to the size of a chain, in which the smallest one-third of the chains have no more than 13 units, the middle one-third have between 14 and 2009 units, and the largest one-third of chains have at least 2010 units. Columns (3) through (8) report the estimates of equation (1) under the three subsamples.

The estimated coefficients in Columns (3) and (4) show that for small chains, their units further away from the Strip improve significantly more in ratings over the data period. In contrast, units of mid-sized chains have the opposite relation: the closer to the Strip the faster increase in Yelp ratings, as indicated by the negative and statistically significant coefficient of $Distance * TimePassed$ in Columns (5) and (7), respectively. For example, the coefficient of β_2, -0.0043, reported in Column (5), can be interpreted as that for units affiliated with chains that have 14 locations, one mile closer to the Strip is associated with a 0.0043 star per year increase in ratings. Therefore, over a ten-year period, a unit located ten miles away from the Strip improves almost half star more in customer ratings than a unit that is affiliated with the same chain and located on the Strip. But note that the negative association between distance to the Strip and slope of ratings no longer holds for chains with more than 353 locations, since the coefficient β_4 is estimated to be statistically significantly positive in both cases. That is, for large chains, their Strip units have a faster decline in ratings than off-the-Strip unites. Moreover, the negative coefficient of $Size * TimePassed$ in Columns (5) and (7) suggest that for chains with at least 14 locations, units of the smaller chains in general experience a stronger improvement in Yelp ratings, holding everything else constant.

Table 2.5: Yelp Ratings, Size of Chain and Distance to Strip

	Coeff. (1)	SE (2)	Coeff. (3)	SE (4)	Coeff. (5)	SE (6)	Coeff. (7)	SE (8)
TimePassed	-0.0338**	0.0118	-0.0092	0.0184	0.0073	0.0317	0.0138	0.0505
Distance * TimePassed	0.0011	0.0007	0.0060***	0.0019	-0.0043*	0.0021	-0.0135*	0.0060
Size * TimePassed	0.0037	0.0030	-19.3235	31.4162	-0.4868**	0.1536	-0.0476***	0.0104
Distance * Size * TimePassed	0.0026***	0.0006	-5.0996	4.9380	0.1220***	0.0246	0.0921***	0.0021
ReviewerAverage	0.0093***	0.0000	0.0088***	0.0001	0.0096***	0.0001	0.0102***	0.0002
Full Sample	Yes							
Smallest $\frac{1}{3}$ of Chain			Yes					
Middle $\frac{1}{3}$ of Chain					Yes			
Largest $\frac{1}{3}$ of Chain							Yes	
Adjested R^2	0.39		0.33		0.35		0.46	
Observations	152085		101164		38694		12227	

† $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The variables $Size * TimePassed$, $Distance * Size * TimePassed$ are rescaled by dividing 10,000 from the variables

All specifications include common trends for chains and zip-code areas

The smallest one-third of chains are chains with no more than 13 locations. The middle one-third of chains are chains that have more than 13 locations but fewer than 2010 locations. The largest one-third of chains have at least 2010 locations

2.5.3 *Discussions of Results*

My findings suggest that even before online reviews gain popularity, small chains did not have problems with low quality at their units close to the Strip. Consistent with the result, I also do not find evidence that online reputation has any impact on their units during the periods when online reviews become an important source of information. Since the majority of the sample are restaurants affiliated with small chains, the results using full sample also do not seem to support the hypothesis that online reputation has a significant impact on restaurant quality.

However, for mid-sized chains, I present evidence that online reputation is an effective mechanism that fixes the weak incentive problem in areas with high degree of tourism business. This is implied by the results that (1) their Strip locations have a worse average rating than the off-Strip locations within the same chain before online reviews had much impact, but (2) as online reviews gain popularity, their units closer to the Strip improve significantly more on customer satisfaction.

On the other hand, large chains with more than a few hundreds of units also experienced the weak reputational incentives at their Strip locations in the early days of Yelp, but these units continue to decline even more in ratings than the off-the-Strip units as online reviews play a larger role in consumers' decision. These findings suggest that restaurants affiliated with large chains are not affected by their online reputation. Similarly, Luca (2011) finds that Yelp reviews have a significant effect on the revenue of independent restaurants but not on that of chain-affiliated restaurants, although it's not clear whether small local chains with more than one unit are also considered as restaurant chains by Luca (2011). The intuition is that since restaurants within a chain share similarities in many dimensions, from the standardized menu and price to the look of a place and even employee's uniforms, a dining experience with one restaurant in a chain is likely to give customers a fairly good sense of how they will like other restaurants affiliated to the same chain. As a consequence, many customers, including tourists, are very familiar with large multi-national chains, like McDonald's, such that they no longer rely on these chains' online reviews to make their dining choices.

In summary, for chains with extreme sizes, online reputation has little impact on their quality, possibly because very small chains internalize the reputation externality across all units, and the demand for large chains is independent of their online reviews. In contrast, for chains with a moderate size, I find evidence that online reputation is very effective in terms of encouraging their locations that traditionally have very weak reputational incentives to improve quality.

2.6 Conclusion

Online reviews are now an important source of information that factors into consumers' everyday decisions. With a substantial evidence on its significant impact on the demand side, this paper, to the best of my knowledge, is the first study that attempts to look for evidence that online reviews provide effective incentives for the supply side to improve quality.

In this paper, I develop the argument that the degree of tourism business in the area where a restaurant is located is an useful source of variation to identify the effect of online reviews on restaurant quality. Exploiting the location variation within restaurant chains, I find that for chains with a moderate size, their units closer to a popular tourist area initially had a significantly lower customer satisfaction on Yelp when online reviews did not have much impact, but, strikingly, these units with an initial low quality problem improve significantly more relative to the units in residential areas that have the same chain affiliation, over the years when online reviews play a more important role. These findings support the view that online reviews have a positive impact on restaurant quality for the mid-sized chains. At the same time, I do not find such evidence for either very small chains or large chains, possibly due to differences in incentive mechanisms.

Chapter 3

Obfuscation and Experience: Empirical Evidence from on Online Market

3.1 Introduction

Relative to visiting traditional brick-and-mortar stores, the Internet greatly reduces search cost by making information available to consumers in seconds, not to mention the price-comparison feature that many third-party sites and e-commerce websites now provide. In the early days of Internet commerce, many expected that the low online search cost would intensify price competition, leading to lower prices and smaller or eventually no price dispersion in the onine market. However, researchers find mixed results.[1] While Brynjolfsson and Smith (2000) show that price dispersion in the book retailing industry was smaller online than in the brick-and-mortar stores, Lee and Gosain (2002) find the online and offline price dispersions to be similar in the music CD industry. Moreover, investigating how price dispersions evolve in eight different product categories on the Internet, Pan et al. (2003) conclude that price dispersion is persistent even as the Internet markets mature by showing that price dispersion declined between 2000 and 2001 and increased from 2001 to 2003. Along the same lines, average prices were not found to be lower online (e.g. Tang and Xing (2001); Clemons et al. (2002); Clay et al. (2003)).

[1] Baye et al. (2006) summarize the empirical studies comparing online and offline prices.

Of course, as long as marginal search costs are still present on the Internet, they are a plausible explanation for the observed online price dispersion. A stream of research also considers limited information to be a potential cause of consumers' suboptimal purchase decision (Georee 2008; Eliaz and Spiegler, 2011). But one would predict that the emergence of price comparison features on e-commerce sites or third-party sites would create a scenario in which online prices converge to the "law of one price", since marginal search cost is zero in this environment.[2]

However, a number of prior studies investigating the relationship between market prices and the use of price comparison sites find different results. Brown and Goolsbee (2002) show that as the use of Internet comparison shopping sites spreads, the price of offline life insurance falls by 8-15 percent. Using 4 million daily price observations for over 1000 electronics products on a price comparison site, Baye and Morgan (2004) find that the level of online price dispersion did not decrease as the use of the price comparison site increased, but was negatively correlated with the number of firms listing prices. Tang et al. (2010) show that 1% increase in the use of price comparison sites, which gather multiple vendors' price information, is associated with a $0.41 price decrease in the online book market.

In an attempt to understand why consumers fail to choose the best price in such a frictionless environment, one strand of literature studies firms' pricing strategies that lead to consumers' suboptimal choices. In the theoretical literature, several papers predict that price obfuscation is the optimal strategy because firms are more profitable with more market power; however, it substantially reduces consumer surplus (Carlin 2009; Wilson 2010; Ellison and Wolitzky 2012). Prior empirical studies indeed find evidence that consumers are worse off when the complexity of price structure increases by conducting various experiments. (e.g. Chakravati et al. (2002), Xia and Monroe (2004),

[2]The price comparison feature generally presents a sorted list of prices that different merchants charge for the same product.

Bertini et al. (2009))[3] Ellison and Ellison (2009) is the first paper to directly show how obfuscation mechanism works in an empirical setting and its equilibrium effects. They document obfuscation practices among a group of computer parts retailers "who operate in an environment where a price search engine plays a dominate role". Specifically, when the price comparison site Pricewatch sorted the list only on the basis of merchants' item price, shipping fees were much higher than later when retailers were mandated to report shipping charges and shipping-inclusive prices were used for sorting. Another more subtle but effective obfuscation strategy found on Pricewatch is that some small firms intentionally create a low-quality product and set a very low price to achieve a better rank on the Pricewatch list. Once attracted to the retailer's website, consumers are presented with evidence of how the low-price product is inferior than the other products the firm sells or even offered with several upgrades by default. Furthermore, Ellison and Ellison (2009) find a striking result that offering a low-quality product at a low price improves the sales of the higher quality products of the retailer.

This paper contributes to the empirical price obfuscation literature in two ways. First, it documents the different price obfuscation techniques implemented by different groups of retailers who vary according to experience but are otherwise undifferentiated. Specifically, I show that sellers, who evolve in a bait pricing strategy, actually charge the highest prices and are the most experienced in the market; new sellers tend to engage in mixed bundling practices, in which they combine several similar products into one listing to appear popular. Second, with individual-level consumer purchase data, I follow individual consumers over time and study the effectiveness of different pricing practices on different consumers. My result shows that the more experienced the consumer, the less likely he will be exploited by obfuscation.

In Section 3.2, I describe the online marketplace, the different price obfuscation techniques on the trading platform, and the two potential reasons why consumers make sub-

[3] Ahmetoglu et al. (2014) provide a nice summary of several common pricing practices and the empirical evidence on their effectiveness on consumer behaviors.

optimal choices in this empirical context. Section 3.3 presents the data and summary statistics. Section 3.4 contains an analysis of the price comparison across different obfuscation types. Section 3.5 shows the relation between seller experience and the choice of obfuscation strategy, and the relation between consumer experience and the likelihood of being misled by price obfuscation. Section 3.6 concludes the paper.

3.2 Empirical Setting

3.2.1 The Online Marketplace

The e-commerce site Alibaba.com is owned by the Chinese e-commerce giant Alibaba and named after the parent company. It is a trading platform designed primarily for whole-sale level trades between small businesses and is now the world's largest business-to-business trading platform, according to Wall Street Journal in 2014. In recent years, with the tremendous growth of online shopping in China, the e-commerce website also caters to individual consumers.

In this study, I focus on a particular product in this online marketplace – the cell phone refill cards. They are stored-value cards used for prepaid cell phone accounts, and can be purchased both online and offline at newsstands or convenience stores. The advantage of purchasing refill cards online is that funds are automatically delivered to the recipient's account by software without physical cards changing hands. Commonly used among cell phone users in China, refill cards are available in several different face values with a wide range. On Alibaba.com, consumers can choose from seven different face values: 1, 5, 10, 20, 30, 50, and 100 yuan, for each of the three wireless carriers in China: China Mobile, China Unicom, and China Telecom.

To purchase a refill card on Alibaba, consumers first enter the specific type of cell phone refill card, defined as a unique combination of carrier and value, that they are looking for in the search box on the home page of this e-commerce website, as depicted

Figure 3-1: Main Page of the Online Marketplace

in Figure 3-1. They will then be guided to the category page, which lists all the sellers who sell the product, their prices, past sales information, and sellers' information. Sellers on this page are initially sorted by a default algorithm, but consumers can easily resort sellers by their price or popularity with one click on the icons in the top-left corner of the category page.[4] Figure 3-2 is a screenshot of the category page returned by searching 50 yuan China Mobile refill cards. With a closer look, one might be surprised by the prices that some sellers charge for a 50-yuan refill card because they seem to be too unrealistically low for the face value. I will explain the reasons in Section 3.2.2.

When consumers are interested in learning more about a seller's product, they can click on it and go to the seller's product page. Figure 3-3 is an example of a product page, which lists the price, the quantity requirement, average rating, the product's origin, and available quantity. Also on the product page, but not shown in the figure, are individual consumers' ratings on a scale of 1 to 5 and detailed information on every past transaction of this product, including transaction price, number of cards purchased, and time of transaction. During checkout, there is a box where consumers have to leave the recipient's cell phone number, then a program automatically takes the information and

[4]The platform does not explicitly state the metrics used for the default ranking.

Figure 3-2: Product Category Page of 50 Yuan Yidong Refill Cards

transfers the funds within a few minutes.

Figure 3-3: Product Page of a 100 yuan Refill Card

3.2.2 Evidence of Price Obfuscation

As noted above, some of the prices advertised on the category page appear to be too low. Indeed, after browsing the product page, consumers will notice that prices shown on the category page are lower than the actual prices in many cases. What enables sellers to engage in price obfuscation like this is a special feature of the online marketplace: to facilitate whole-sale level trades, this formerly business-to-business trading platform allows sellers to set multiple prices for a listing to offer quantity discounts. In other words, the price of a refill card of a given value may vary depending on how many such cards consumers purchase. Furthermore, sellers are allowed to combine similar products into one listing on Alibaba, perhaps to facilitate product comparison or offer add-ons. But when a listing has several prices, only the lowest price is advertised on the category page and used for sorting. Therefore even with the price-sorting feature, price comparison is not as trivial as it seems.

Specifically, I observe three distinct types of price obfuscation in this market:

I refer to the first type as bait pricing because sellers who implement this strategy offer several unit prices and require consumers to purchase an unrealistically large number of refill cards in one purchase to qualify for the lowest price. As a result, no one in the market ever meets the requirement.[5] Figure 3-4 is a screenshot of the prices shown on the product page of a seller who implements the bait-pricing strategy. The first row shows the quantity and the second row is the corresponding price per card.

Figure 3-4: Product Page of the Bait-pricing Type

In this example, consumers first see this seller's 100-yuan refill card listed for 98 yuan on the category page; however, they will soon learn on the product page that to qualify for this price they must purchase at least twenty 100-yuan refill cards in one transaction. This would cost a consumer at least 1,980 yuan,[6] an amount that is about twenty times Chinese cell phone users' average monthly spending on cell phone services.[7] As it turns out, no one ever spent this much in a transaction in the data.[8]

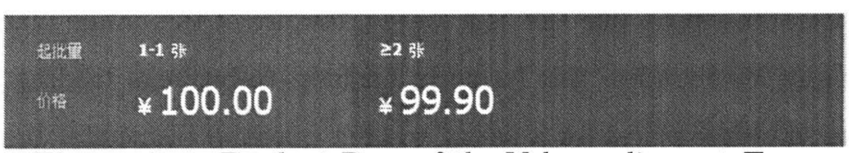

Figure 3-5: Product Page of the Volume-discount Type

I call the second type of strategy volume discount. Sellers who use this strategy also

[5] Whole-sale purchases are not possible in this market because the value purchased is automatically transferred to a cell phone account and is then nontransferable.

[6] This amount is approximately equivalent to $300.

[7] According to a report by Nielsen in 2013, Chinese cell phone users spent an average of 100 yuan per month on cell phone services. See http://www.nielsen.com/content/dam/corporate/uk/en/documents/Mobile-Consumer-Report-2013.pdf.

[8] The total face value bought in 98% of the purchases in this market are less than 100 yuan.

have more than one price for a product, but they do intend to offer discounts when multiple cards are purchased. In other words, they set a reasonably low minimum quantity requirement for the cheapest unit price so that at least one consumer in the data qualifies for the price. Figure 3-5 illustrates an example of this type, in which consumers, who are interested in purchasing only one card, pay exactly the face value, and those who buy more than one card receive a 10-cent discount for each card they purchase. But note that for consumers who only want to purchase one card, the advertised price of the volume-discount type is still higher than the actual price.

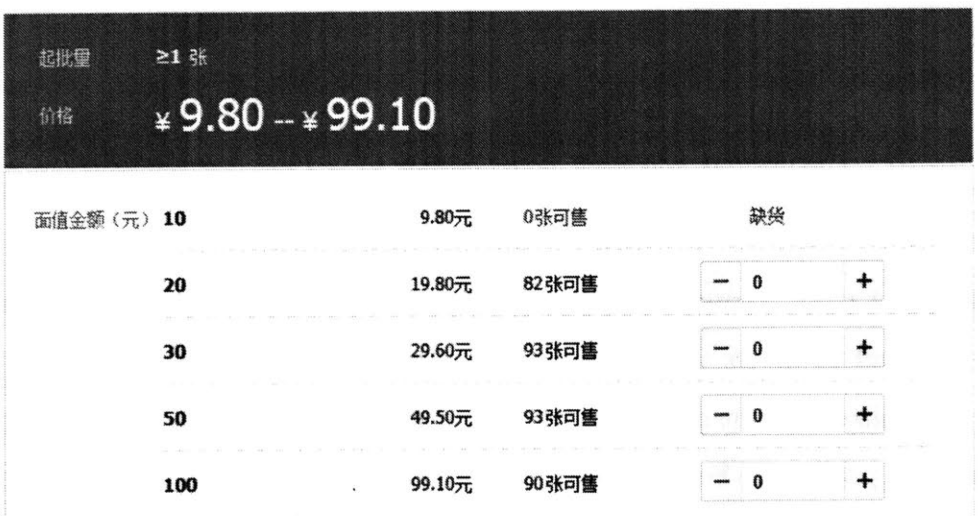

Figure 3-6: Product Page of the Mixed-bundling Type

Sellers of the third type of price obfuscation combine refill cards with different face values into one product page, as shown in Figure 3-6. I refer to this type of obfuscation as the mixed bundling type. In this example, refill cards of five different face values are all in one product page for consumers to choose from, but the 10-yuan cards are in fact listed as out of stock, which is another practice of adding complexity to prices.[9] Consumers who search for any of the five face values will see this seller's price listed as 9.8 yuan on the category page. The advertised price of the mixed-bundling type products may appear too low to be the real price and lead to suspicion, instead of making consumers mistakenly believe that this is the price that they will pay. But one important advantage

[9]Arnold and Saliba (2002) make the point that lack of inventories raises search costs.

of this approach is that it combines the sales of all the products in a listing and makes the products appear to be more popular on the category page than sold separately with each product on its own listing.

While many sellers engage in price obfuscation, two-fifths of the sellers do not use any obfuscation strategies. Their products are sold at the price advertised on the category page for any quantity. However, consumers are unable to identify which sellers implement obfuscation techniques and have to go to each individual product page to discover the true price that applies to them.

3.2.3 Perceived Quality Difference or Price Confusion?

In an excellent survey on the empirical and theoretical studies of why consumers fail to choose the best price, Grubb (2015) summarizes that "there are at least two reasons consumers may fail to do so: First, consumers may be confused about product quality. They may not realize that the goods are homogeneous, and may attribute imaginary quality differences to products. Second, consumers may be confused by complex prices and may be unable to identify the lowest price when comparing two quotes."

There are several scenarios in which consumers may perceive the same product sold by two different sellers to be different. First, physical goods may have subtle differences in their physical condition or other quality-related aspects. Second, what is more prominent is the difference in how consumers perceive different merchants and their relationship with the retailers, such as reputation, perceived service quality, whether they are members of the retailer's loyalty program, and so on. Even for the two well-established online book retailers – Amazon and Barnes & Noble, Chevalier and Goolsbee (2003) find that consumers prefer Amazon to Barnes & Noble, despite the many similarities of the two sites. Last, full information on the available choices is often assumed in many prior empirical studies, but the reality is that consumers are often not aware of the availability of a particular seller due to limited information, in which case the comparison does not even take place.

The online cell phone refill card market on Alibaba.com has some unusual market characteristics; but exactly because of its unique aspects, this market is an empirical setting that eliminates as many potential perceived differences across sellers as possible. Specifically, several important market features alleviate the concern that consumers may view purchasing a cell phone refill card from different sellers differently. First of all, unlike many physical goods markets, the online cell phone refill card market has no quality difference across sellers because the product is essentially a commodity.[10] Second, the instant online delivery implies that services provided by sellers are essentially the same. Third, unlike large markets with vendors who have different reputations, sellers in this market are small firms with no name recognition and share similarities in a variety of dimensions since they are on the same trading platform. Lastly, market boundaries are often not clear in many markets to researchers in the sense that consumers may have limited information on seller availability and therefore have zero probability of searching some sellers' products. In contrast, once a consumer visits this e-commerce website, all sellers are potentially in his consideration set. Because of these unusual market features, I argue that it is highly unlikely that consumers' perception of different sellers and their products and services is the cause of not purchasing from the cheapest seller; what plays a significant role in this homogeneous goods market with a price-sorting feature is rather price obfuscation.

3.3 Data and Descriptive Statistics

The data consists of cell phone refill card transactions on Alibaba.com between September 24, 2013 and March 24, 2014, and listing information during this period. For each

[10]Consumers are well protected from fraud on Alibaba.com since the payment is first kept in an escrow account and will not be transfered to the seller's account until it is confirmed that no dispute occurs after the product delivery.

listing, I observe the unique numeric product ID where a smaller number is associated with an earlier listing, daily price(s), the corresponding quantity requirements, the type(s) of the refill card(s), and other information, including the average rating, the name of merchants, and information on the past sales. A total of 357 listings were available at some point during the data period with 53 entering and 17 exiting the market.

Table 1 presents the summary statistics of listing prices and number of listing that falls into each type of obfuscation, broken down by the face value and carrier of refill cards. The average market prices are around the face value with the unit price decreasing as the face value of a card increases. For different types of refill cards, the price range goes from 0.2 yuan to 3.5 yuan and is typically greater for larger face values. Among refill cards with the same face value, a larger number of sellers is typically associated with a smaller standard deviation in listing prices: China Mobile is the most popular carrier and has the most listings for any given face value, and in four out of the six face values, China Mobile refill cards have the smallest price dispersion measured by standard deviation. This is consistent with the findings in prior empirical studies on the relationship between price dispersion and number of firms (e.g. Borenstein and Rose (1994); Barron et al. (2004)). The number of listings that falls into each type of obfuscation is presented in the last four columns of Table 3.1. More sellers tend to engage in price obfuscation in refill card markets with large face values. The mixed-bundling type is generally the most commonly used obfuscation strategy, and the Volume-discount type is rarely used by sellers.

Table 3.1: Summary Statistics of Product Price and Obfuscation

Face Value	Price				# of Sellers	Number of Listings Using Obfuscation			
Carrier	Mean	Stdev	Min	Max		No obfus	Bait-pricing	Volume-discount	Mixed-bundling
1-yuan Card									
Mobile	1.48	0.32	0.98	2	8	7	1	0	0
Unicom	1.60	0.53	0.98	2	3	2	1	0	0
Telecom	1.60	0.53	0.98	2	3	2	1	0	0
10-yuan Card									
Mobile	10.06	0.17	9.5	10.5	32	24	4	2	2
Unicom	10.03	0.18	9.5	10.3	16	13	1	0	2
Telecom	10.03	0.18	9.5	10.3	16	13	1	0	2
20-yuan Card									
Mobile	20.07	0.28	19.8	21	35	16	4	1	14
Unicom	20.01	0.17	19.8	20.5	12	2	0	0	10
Telecom	20.02	0.16	19.8	20.5	13	2	1	0	10
30-yuan Card									
Mobile	29.97	0.13	29.8	30.5	30	7	3	1	19
Unicom	29.95	0.07	29.8	30	17	4	1	1	11
Telecom	29.95	0.07	29.8	30	17	4	2	0	11
50-yuan Card									
Mobile	49.92	0.40	49.5	52	44	16	7	1	20
Unicom	49.93	0.53	49.5	52	19	5	2	0	12
Telecom	49.93	0.53	49.5	52	19	4	3	0	12
100-yuan Card									
Mobile	99.71	0.56	98.8	102	31	8	6	0	17
Unicom	99.64	0.73	98.5	102	19	6	1	1	11
Telecom	99.57	0.76	98.5	102	20	7	2	0	11

Notes: Price here is the market price when one card is purchased in a transaction.
Only one seller existed in the 5-yuan market and is therefore not presented.

For each transaction, buyer's unique user ID, price paid, product ID, number of cards purchased, time of purchase are observed. The unique user ID allows me to trace the purchase history of each consumer over the data period. In summary, 4,116 cell phone refill card purchases were made by 2,794 buyers. About 82% of the buyers shopped only once and 10% shopped twice in the 6 months. The remaining 8% buyers shopped between 3 and 27 times.

In Table 3.2, which presents the summary statistics of transaction data, the mean transaction prices are around the face values and slightly smaller than the average market prices, except for 10-yuan refill cards.[11] The standard deviation, minimum and maximum of transaction prices show that equilibrium prices spread out over the price range and the majority of consumers fail to buy at the lowest price in this market. The average quantity a consumer purchased in a transaction was 1.6 cards with a standard deviation of 1.8 cards. In about 20% of the transactions, consumers purchased more than one phone card of the same value, and in roughly half of them, a consumer bought at least four cards at once. The average total amount paid in one transaction was 57 yuan with a standard deviation of 71 yuan.

Table 3.2: Summary Statistics of Transaction Data

	Mean	Stdev	Min	Max	Obs.
Transaction Price					
1-yuan card	1.30	0.17	0.98	1.80	412
10-yuan card	10.12	0.09	9.90	10.30	528
20-yuan card	19.91	0.13	19.80	21	196
30-yuan card	29.90	0.11	29.60	30	123
50-yuan card	49.76	0.19	49.00	52	1,696
100-yuan card	99.18	0.45	98.80	100	558
Number of Cards	1.58	1.80	1	50	4,116
Total Face Value	57.07	71.12	1	1,150	4,116

The definition of variables used in the analysis in the subsequent sections is presented in Table 3.3. The variables include several price measures, dummies for obfuscation types,

[11]Because market prices for refill cards with the same face value are very similar across carriers, I group transactions with the same face value together for the summary statistics.

Table 3.3: Definition of Variables

Variables	Definition
Price Measure	
Price	Price per card
AdPrice	Price shown on the category page
Cheapest	Dummy equaling 1 if price is the cheapest possible
PriceDiff	Difference between transaction price and advertised price
Experience Measure	
JTotalExp	Dummy equaling 1 if consumer purchased a total of j times
PastExp	Number of past purchases the consumer had at the time
Obfuscation Type	
NoObfusc	Dummy equaling 1 if seller does not use price obfuscation
VolumeDiscount	Dummy equaling 1 if seller uses volume-discount type obfuscation
BaitPricing	Dummy equaling 1 if seller uses bait-pricing type obfuscation
MixedBundling	Dummy equaling 1 if seller uses mixed-bundling type obfuscation
Other	
Index	Order of listings from the oldest to the newest (1-357)

consumer's experience measures and a seller's experience measure.

3.4 Price and Obfuscation

Since only the lowest price of a listing is advertised on the search result page, it is natural to expect that sellers may have an incentive to take advantage of the loophole by showing a low price on the category page to attract consumers to their product page and charging relatively high prices. The data confirms with the hypothesis: the obfuscation type that advertises the reasonably lowest price on the category page is actually the most expensive among all.[12]

Controlling for product category fixed effects (i.e. a unique combination of carrier and value), I first compare the prices advertised on the category page across all types except for the mixed-bundling type. The baseline group is the bait-pricing type, which

[12]Mixed-bundling type is excluded from the analysis because the price shown on the category page is always much lower than the face value of a refill card and is highly unlikely to convince that this is the price of the product that they intend to buy.

consumers in the data never meet the quantity requirement for their lowest price (i.e. advertised price). The regression estimates reported in Column 1 of Table 3.4 show that the advertised prices of products without price obfuscation are 0.35 yuan higher than the bait-pricing type products. Advertised prices of the volume-discount type also appear to be higher relative to the bait-pricing type, but the estimate is not statistically significant.

Column 2 of Table 3.4 presents the estimation results of the same regression weighted by sales during the data period. The estimated coefficients indicate that among all purchases, the bait-pricing type products have the lowest advertised price, although the price gaps with other types are smaller relative to the unweighted results in Column 1. This finding is intuitive since for sellers who implement the bait-pricing strategy, it is costless for them to advertise a price that they never need to fulfill.

A more striking result is found when comparing all the transaction prices across different obfuscation types. The estimates of regressing actual transaction price on the obfuscation type are presented in column 3 of Table 3.4. Contrary to the advertise price, the bait-pricing type products have the most expensive equilibrium prices, which are on average between 0.1 and 0.25 yuan higher than that of other types. Surprisingly, the mixed-bundling type products turn out to be the cheapest type. Relative to the small price dispersion shown in Table 3.1, these price differences are quite substantial in the market.

For a robustness check, I use a more stringent price measure to compare prices of different obfuscation types within a market; *Cheapest* is a dummy variable indicating whether a product was ever the cheapest among all possible choices in the same category at some point during the data period.[13] Column 1 of Table 3.5 presents the mean and standard deviation of the variable, *Cheapest*, for each obfuscation type.[14] Products of the no-obfuscation type and mixed-bundling type are 30% more likely to be the cheapest

[13]Here I compare the unit price when only one refill card is purchased because in 80% of the transactions consumers only purchased one card.

[14]The volume-discount type products are excluded from the analysis because there were only seven of them and none of them were ever the cheapest choices.

Table 3.4: Price and Obfuscation Type

	(1) AdPrice	(2) AdPrice	(3) Price
NoObfusc	0.353**	0.270***	-0.099***
	(3.08)	(17.43)	(-10.97)
VolumeDiscount	0.322	0.091***	-0.121***
	(1.28)	(5.68)	(-12.93)
MixedBundling	—	—	-0.249***
			(-23.80)
Observations	175	175	4116

t statistics in parentheses

† $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Category fixed effect is included in all regressions

The omitted group is the bait-pricing type products

choices in the market compared with the bait-pricing type products, which are almost never the best choices indicated by a mean close to zero.

Table 3.5: Lowest Price and Obfuscation Type

	(1)	(2)	(3)
	Cheapest		
		difference	
	mean	NoObfusc	BaitPricing
NoObfusc	0.344		
	[0.061]		
BaitPricing	0.048	0.297***	
	[0.048]	(0.078)	
MixedBundling	0.368	-0.024	-0.321***
	[0.079]	(0.100)	(0.093)

Standard deviations are in brackets and standard errors are in parentheses

† $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Results of pairwise t-tests on the mean are reported in Columns 2 and 3 of Table 3.5. Again, I find that the bait-pricing type products rarely have the best prices compared with products of the no-obfuscation type and the mixed-bundling type, consistent with the regression results in Column 3 of Table 3.4. Overall, the findings that products of bait-pricing type have the lowest advertised price, but are able to induce consumers to pay more than other types suggest that price comparison is not trivial in this market even with the presence of a price-sorting feature. Furthermore, the systematic difference between advertised price and actual price of the bait-pricing type products provides evidence that some sellers intentionally engage in price obfuscation to mislead consumers. The nature of this pricing strategy is the same as what is found on Pricewatch by Ellison and Ellison (2009), where small firms advertise a low price in an attempt to induce consumers to their page.

These findings are also consistent with the lab experiment results found by Kalaycı and Potters (2011). In the experiment, they asked the lab subjects who play the role of pricing managers to choose a number between 0 and 18, which is the number of multiplication and addition operations required to calculate the product prices; lab subjects who play the role of consumers have only fifteen seconds to make a purchase decision. Kalaycı and Potters (2011) find that: (1) consumers make more mistakes when the number of operations increases (i.e. level of obfuscation increases); and (2) both offered and transaction prices increase with obfuscation.

3.5 Experience and Obfuscation

After showing that sellers deliberately increase the complexity of price comparison in this cell phone refill card market, I further explore the relation between seller experience and their choice of obfuscation strategy and investigate the effectiveness of obfuscation techniques on consumers with different levels of market experience.

3.5.1 Seller Experience and Obfuscation Strategy

To investigate the relation between sellers' market experience and the type of obfuscation techniques they implement, I regress a dummy indicating whether a product is of a given type of obfuscation strategy on the index of the product, which is an integer between 1 and 357 with a smaller number associated with an earlier listing. The four sets of regression estimates are presented in Table 3.6. Interestingly, the two most distinctive types found in Section 3.4, which are the type with the most misleading prices and the cheapest type – bait-pricing type and mixed-bundling type, respectively, have a statistically significant relationship with the chronological order of the listing – a proxy for the seller's market experience. Specifically, experienced sellers are more likely to use the bait-pricing technique, whereas newer sellers are more inclined to implement the mixed-bundling strategy.

Table 3.6: Seller Experience and Obfuscation Type

	(1) NoObfusc	(2) BaitPricing	(3) VolumeDiscount	(4) MixedBundling
Index	-0.0001	-0.0007***	-0.00004	0.0009***
	(-0.44)	(-4.46)	(-0.70)	(3.89)
Observations			357	

t statistics in parentheses
† $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
category fixed effects are included in all regressions

It makes sense that with more market experience, sellers become more strategic in increasing profits; sellers implementing the bait-pricing strategy provide the best example of this. But why are less experienced sellers more likely to use the mixed-bundling strategy? From my conversations with industry experts, I learned that since product popularity is generally an important aspect that signals high quality or a competitive price on online trading platforms, combining several products into one listing expedites the process of accumulating popularity and is, therefore, more appealing to sellers new to the market. These insights are convincing explanations for why new sellers in the market find the mixed-bundling strategy more optimal, because it is also consistent with the evidence that mixed-bundling type products have the lowest market prices – another strategy to quickly improve the sales record.

3.5.2 Buyer Experience and Price Obfuscation

Next, I turn to the relationship between consumers' experience and the likelihood of being exploited by price obfuscation.

As shown in Section 3.4, products of the bait-pricing type have the highest transaction prices but advertise the lowest prices. Therefore, a naive approach is to examine whether any systematic association exists between consumer experience and the purchase probability of the bait-pricing type products. I run a logistic regression, which is represented by the following equation:

$$\begin{aligned} BaitPricing_{it} &= \beta_0 + \beta PastExp_i + \delta Quantity_i \\ &+ \sum_{k=1}^{21} \lambda_k Type_{ik} + \sum_{l=1}^{26} \zeta_l Week_{il} + \alpha_i + \epsilon_i. \end{aligned} \tag{3.1}$$

Each observation is a transaction and the dependent variable is a dummy indicating whether in this transaction the consumer purchases a bait-pricing type product whose price is not the cheapest at the time of purchase. Control variables include the number of purchases consumer i has made previously in the data period, the number of cards purchased, consumer fixed effects, product type fixed effects, and week fixed effects. β is the parameter of interest.

Regression estimates are presented in Table 3.7. Column 1 reports the estimated coefficient on consumer experience using all transactions in the data. Columns 2 through 5 are the coefficient estimates of the same regression using subsamples of consumers who have made more purchases in the entire data period. Results in all five regressions consistently show that experienced consumers are less likely to buy from a seller who implements the bait-pricing type price obfuscation strategy. The odds of buying from a seller who implements the bait-pricing strategy decrease around 23% to 30% with one additional previous purchase, and this negative relationship becomes more significant when restricting sample to consumers who purchased more in total during the entire data period.

Table 3.7: Purchase Probability of Bait-pricing Type Products and Market Experience

	Purchase a Bait-pricing Type Product				
	(1) all	(2) JTotalExp>2	(3) JTotalExp>3	(4) JTotalExp>4	(5) JTotalExp>5
PastExp	-0.2668*	-0.2716*	-0.3019*	-0.3021*	-0.3466*
	(-2.26)	(-2.03)	(-2.02)	(-2.03)	(-2.15)
Observations	334	280	253	225	195

t statistics in parentheses

† $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Consumer, category, and week fixed effects are included in all regressions

Another way to analyze who are more likely to be exploited by price obfuscation is to compare the price difference between the advertised price of the product a consumer purchases and the actual price that he actually pays for. Excluding all transactions in which consumers paid the lowest price possible for any seller at the time of purchase or bought a mixed-bundling type product, I regress the price difference on the same set of explanatory variables as in equation (3.1).[15] Table 3.8 reports the estimates of the coefficient on consumer experience for different consumer groups as in Table 3.7. The negative signs indicate that the gap between advertised price and transaction price diminishes as consumers gain market experience, and this reduction becomes more significant, both in magnitude and significance level, when restricting samples to consumers with more total

[15] I also regress a dummy variable which indicates whether a consumer purchases from the cheapest seller at the moment of the transaction on the consumer's number of previous purchases, controlling for consumer, category and week fixed effects. The estimated coefficient of past experience is 0.006 with a p-value less than 0.001.

market participation. The two sets of results in Table 3.7 and Table 3.8 together provide convincing evidence that consumers with less market experience are more likely to fall into the trap of price obfuscation, since they tend to purchase the type of products that have the most misleading prices and pay for products with a larger price gap between the advertised price and the transaction price.

Table 3.8: Price Difference and Market Experience

	PriceDiff (Actual Price - Advertised Price)				
	(1) all	(2) JTotalExp>2	(3) JTotalExp>3	(4) JTotalExp>4	(5) JTotalExp>5
PastExp	-0.0011	-0.0015*	-0.0016*	-0.0017*	-0.0021**
	(-1.26)	(-2.42)	(-2.48)	(-2.44)	(-2.82)
Observations	2900	769	613	464	382

t statistics in parentheses

† $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Consumer, category, and week fixed effects are included in all regressions

3.6 Conclusion

The Internet brings a tremendous amount of convenience to consumers, while the Internet itself also creates opportunities for firms to exacerbate consumer confusion. This paper documents an instance of online sellers engaging in obfuscation strategies to make price comparison difficult on an online marketplace, in which sellers can offer several prices in a listing and only the lowest price is shown on the price comparison page.

I find that the choice of strategic obfuscation practices highly depends on how long a seller has been in the market. Experienced sellers prefer the bait-pricing strategy with a low advertised price and a high actual price, and they are most successfully at targeting the least experienced consumers in the market. New sellers often combine several products into one listing and offer low prices to expedite the process of improving their sales records. Since some of the obfuscation found can be eliminated with the effort of the platform, I would be interested to see more future work trying to understand the dilemma of transparency that many e-commerce websites with a price-comparison feature now face between leaving room for sellers to be profitable and offering decent services to consumers.

Appendix A

Figures

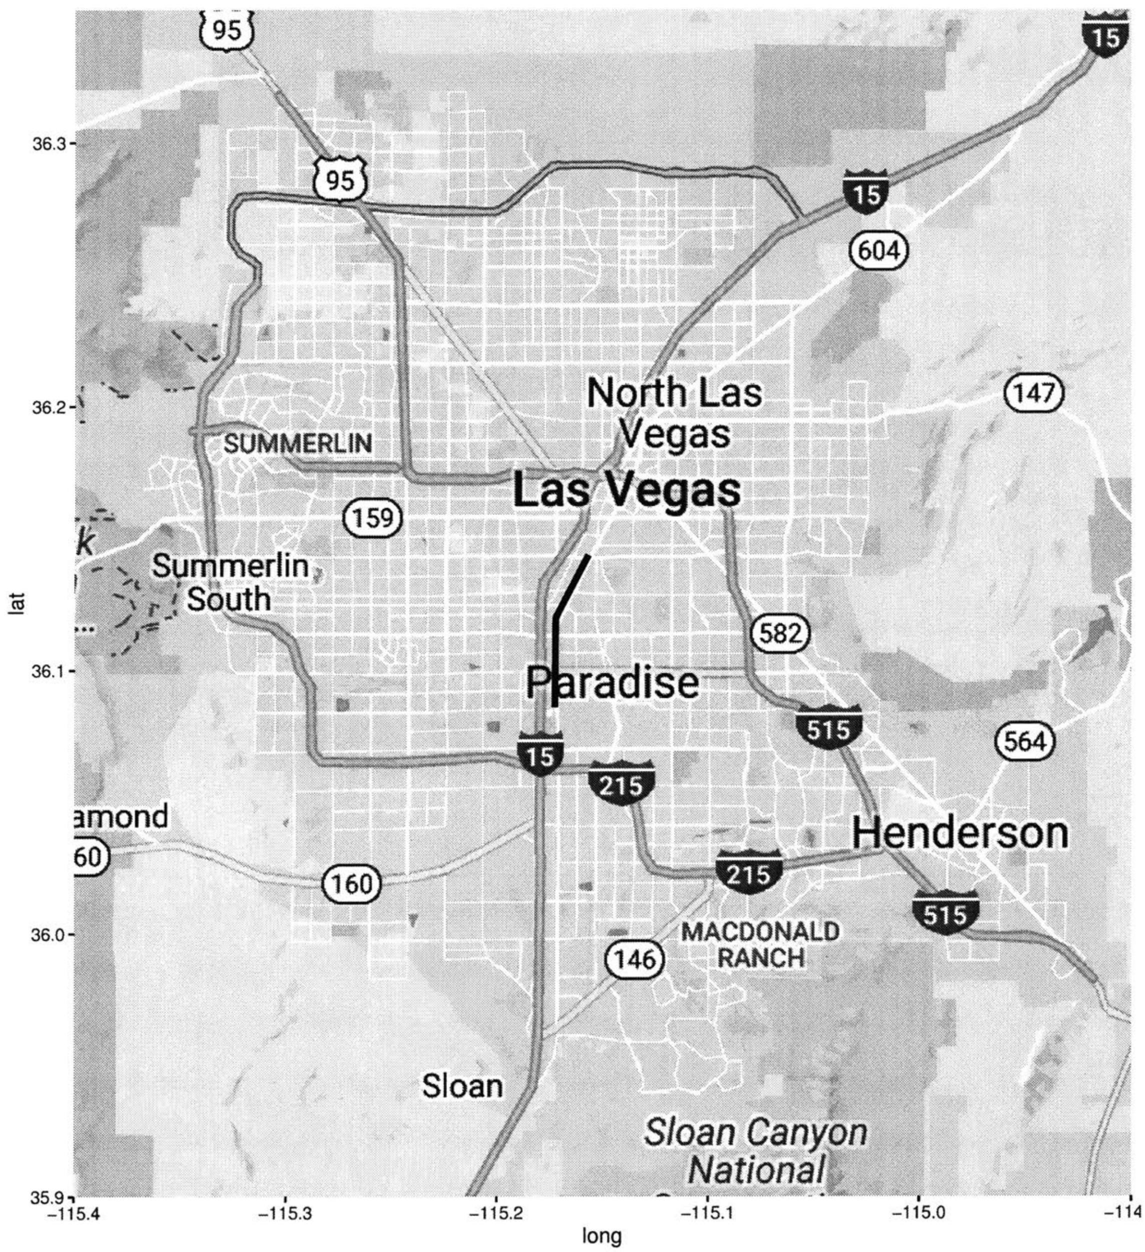

Figure A-1: Las Vegas Strip

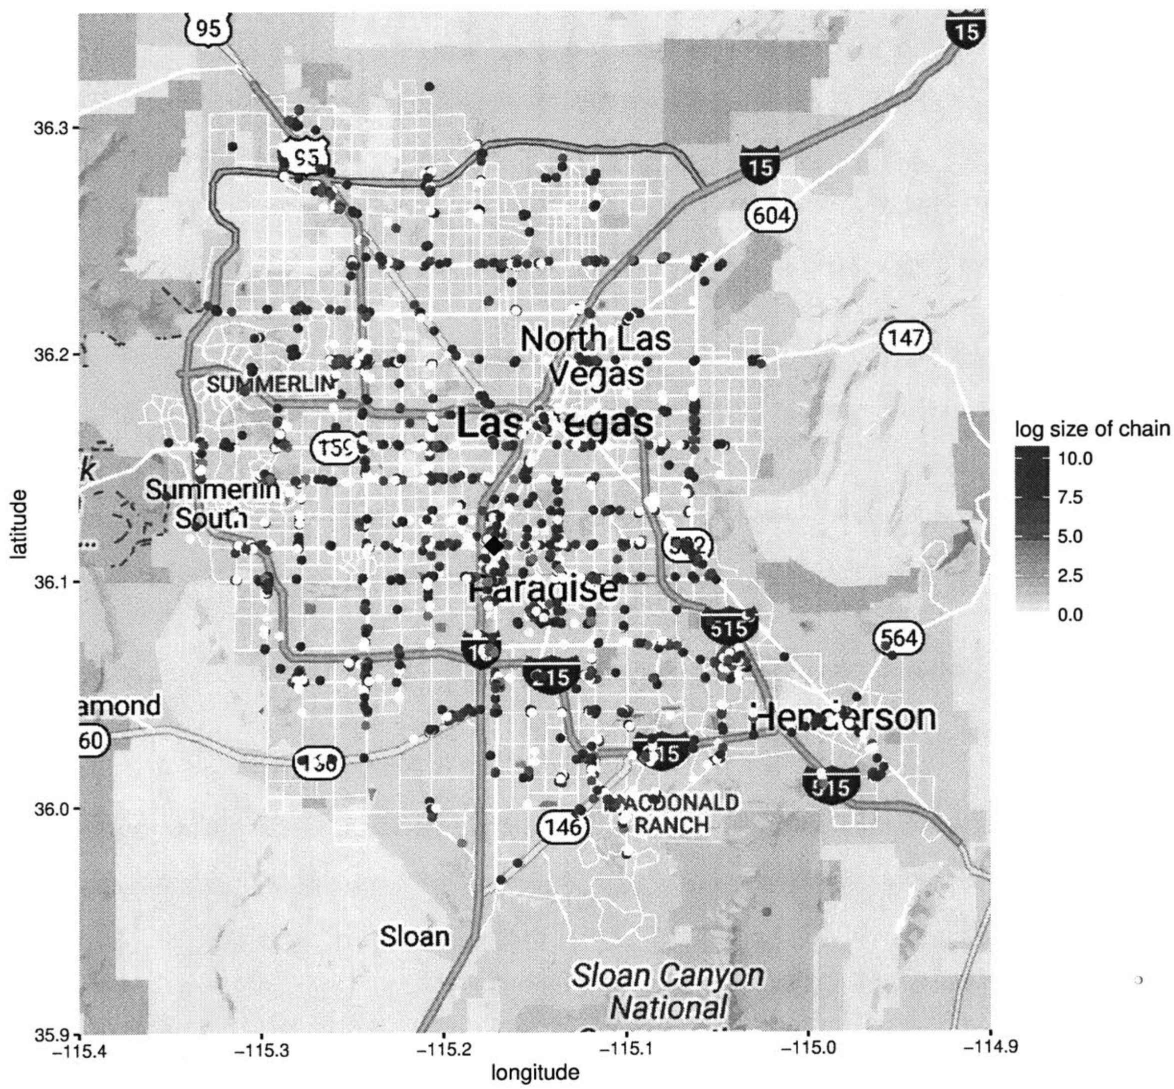

Figure A-2: Map of Chain Restaurants and their Sizes

Figure A-3: Map of Tourists and Locals in Las Vegas (Source: Flickr)

Figure A-4: Map of Tourists and Locals in Boston (Source: Flickr)

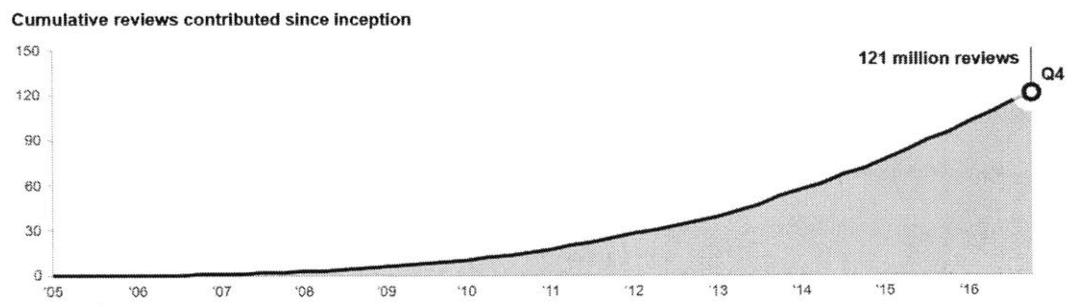

Figure A-5: Cumulative Number of Reviews on Yelp over Time

Bibliography

Ackerberg, D. A., 2001. "Empirically Distinguishing Informative and Prestige Effects of Advertising." *The RAND Journal of Economics*, 32(2):316–333.

Ahmetoglu, G., Furnham, A., and Fagan, P., 2014. "Pricing Practices: A Critical Review of Their Effects on Consumer Perceptions and Behaviour." *Journal of Retailing and Consumer Services*, 21(5):696–707.

Arnold, M. A. and Saliba, C., 2003. ""Price Dispersion in Online Markets: The Case of College Textbooks." *Working Paper.*

Barron, J. M., Taylor, B. A., and Umbeck, J. R., 2004. ""Number of Sellers, Average Prices, and Price Dispersion." *International Journal of Industrial Organization*, 22: 1041–1066.

Baye, M. R. and Morgan, J., 2004. "Price Dispersion in the Small and in the Large : Evidence from an Internet Price Comparison Site." *Journal of Industrial Economics*, 52(4):463–496.

Baye, M. R., Morgan, J., and Scholten, P., 2006. "Information, search, and price dispersion." *Handbook in Economics and Information Systems*, Chapter 6, Hendershott, ed., Elsevier, Amsterdam.

Bennear, L. S. and Olmstead, S., 2008. "The Impacts of the 'Right to Know': Information Disclosure and the Violation of Drinking Water Standards." *Journal of Environmental Economics and Management*, 56(2):117–130.

Bertini, M., Ofek, E., and Ariely, D., 2009. "The Impact of Add-on Features on Consumer Product Evaluations." *Journal of Consumer Research*, 36(1):17–28.

Blake, T., Nosko, C., and Tadelis, S., 2015. "Consumer Heterogeneity and Paid Search Effectiveness: a Large Scale Field Experiment." *Econometrica*, 83(1):155–174.

Blake, T., Nosko, C., and Tadelis, S., 2016. "Returns to Consumer Search: Evidence from eBay." *working paper.*

Borenstein, S. and Rose, N., 1994. "Competition and Price Dispersion in the US Airline Industry." *Journal of Political Economy*, 102:653–683.

Brickley, J. A., 1987. "The Choice of Organizational Form: The Case of Franchising." *Journal of Financial Economics*, 18(2):401–420.

Brickley, J. A., 1999. "Incentive Conflicts and Contractual Restraints: Evidence from Franchising." *The Journal of Law and Economics*, 42(2):745–774.

Brickley, J. A. and Dark, F. H., 1987. "The Choice of Organizational Form: The Case of Franchising." *Journal of Financial Economics*, 18:401–420.

Bronnenberg, B. J., Kim, J. B., and Mela, C. F., 2016. "Zooming In on Choice : How Do Consumers Search for Cameras Online?" *Marketing Science*, 35(3):693–712.

Brown, J. R. and Goolsbee, A., 2002. "Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry." *Journal of Political Economy*, 110(3):481–507.

Brynjolffson, E. and Smith, M., 2005. "Frictionless Commerce? A Comparison of Internet and Conventional Retailers." *Management Science*, 46(4):563–585.

Carlin, B. I., 2009. "Strategic Price Complexity in Retail Financial Markets". *Journal of Financial Economics*, 91(3):278–287.

Chakravarti, D., Krish, R., Paul, P., and Srivastava, J., 2002. "Partitioned Presentation of Multicomponent Bundle Prices: Evaluation, Choice and Underlying Processing Effects." *Journal of Consumer Psychology*, 12(3):215–229.

Che, H., Erdem, T., and Öncü, S., 2015. "Consumer Learning and Evolution of Consumer Brand Preferences." *Quantitative Marketing and Economics*, 3(13):173–202.

Chevalier, J. and Goolsbee, A., 2003. "Measuring Prices and Price Competition Online: Amazon.com and BarnesandNoble.com." *Quantitative Marketing and Economics*, 1(2): 203–222.

Chevalier, J. A. and Mayzlin, D., 2006. "The Effect of Word of Mouth on Sales: Online Book Reviews."

Ching, A. T., Erdem, T., and Keane, M. P., 2013. "Learning Models: An Assessment of Progress, Challenges, and New Developments." *Marketing Science*, 32(6):913–938.

Chiou, L. and Pate, J., 2010. "Internet Auctions and Frictionless Commerce: Evidence from the Retail Gift Card Market." *Review of Industrial Organization*, 36(3):295–304.

Clay, K., Krishnan, R., Wolff, E., and Fernandes, D., 2003. ""Retail Strategies on the Web: Price and Non-Price Competition in the Online Book Industry." *Journal of Industrial Economics*, 50(3):351–367.

Clemons, E., Hann, I. H., and Hitt, L. M., 2002. "Price Dispersion and Differentiation in Online Travel: An Empirical Investigation." *Management Science*, 48(4):534–549.

Crawford, G. S. and Shum, M., 2005. "Uncertainty and Learning in Pharmaceutical Demand." *Econometrica*, 73(4):1137–1173.

Dai, W., Jin, G. Z., Lee, J., and Luca, M., 2012. "Optimal Aggregation of Consumer Ratings : an Application To Yelp.Com." *NBER Working Pape 18567.*

De los Santos, B., Hortaçsu, A., and Wildenbeest, M. R., 2012. "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior." *American Economic Review*, 102(6):2955–2980.

De los Santos, B., Hortaçsu, A., and Wildenbeest, M. R., 2015. "Search with Learning for Differentiated Products: Evidence from E-Commerce." *Journal of Business and Economic Statistics*, forthcoming.

Dranove, D. and Jin, G. Z., 2010. "Quality Disclosure and Certification: Theory and Practice." *Journal of Economic Literature*, 48(4):935–963.

Eckard, W., 2004. "The 'Law of One Price' in 1901." *Economic Inquiry*, 42(1):101–110.

Eliaz, K. and Spiegler, R., 2011. "Consideration Sets and Competitive Marketing." *Review of Economic Studies*, 78(1):235–262.

Ellison, G. and Ellison, S., 2009. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica*, 77(2):427–452.

Ellison, G. and Wolitzky, A., 2012. "A Search Cost Model of Obfuscation." *The RAND Journal of Economics*, 43(3):417–441.

Erdem, T., Keane, M. P., Öncü, S., and Strebel, J., 2005. "Learning About Computers: An Analysis of Information Search and Technology Choice." *Quantitative Marketing and Economics*, 3(2):207–247.

Erdem, T., Keane, M. P., and Sun, B., 2008. "A Dynamic Model of Brand Choice When Price and Advertising Signal Product Quality." *Marketing Science*, 27(6):1111 – 1125.

Goeree, M. S., 2008. "Limited Information and Advertising in the U.S. Personal Computer Industry." *Econometrica*, 76(5):1017–1074.

Grubb, M., 2015. "Failing to Choose the Best Price: Theory, Evidence, and Policy." *Review of Industrial Organization*, 47:303–340.

Hong, H. and Shum, M., 2006. "Using price distributions to estimate search costs." *The RAND Journal of Economics*, 37(2):257–275.

Honka, E., 2014. "Quantifying search and switching costs in the US auto insurance industry." *The RAND Journal of Economics*, 45(4):847–884.

Hortaçsu, A. and Syverson, C., 2004. "Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds." *The Quarterly Journal of Economics*, 119(2):403–456.

Jacob, B. A. and Levitt, S. D., 2003. "Rotten Apples: An Investigation Of The Prevalence And Predictors Of Teacher Cheating." *Journal of Industrial Economics*, 118(3):843–878.

Jin, G. and Kato, A., 2006. "Price, Quality and Reputation: Evidence from an Online Field Experiment." *The RAND Journal of Economics*, 37(4):983–1005.

Jin, G. Z. and Leslie, P., 2003. "the Effect of Information on Product Quality: Evidence From Restaurant Hygiene Grade Cards." *Quarterly Journal of Economics*, 118(2):409–451.

Jin, G. Z. and Leslie, P., 2009. "Reputational Incentives for Restaurant Hygiene." *American Economic Journal: Microeconomics*, 1(1):237–267.

Kalaycı, K. and Portters, J., 2011. "Buyer Confusion and Market Prices." *International Journal of Industrial Organization*, 29(1):14–22.

Kim, J. B., Albuquerque, P., and Bronnenberg, B. J., 2010. "Online Demand Under Limited Consumer Search." *Marketing Science*, 29(6):1001–1023.

Koulayev, S., 2013. "Search with Dirichlet priors: estimation and implications for the price elasticity of demand." *Journal of Business and Economic Statistics*, 31(2):226–239.

Lee, Z. and Gosain, S., 2002. "A Longitudinal Price Comparison for Music CDs in Electronic and Brick-and-mortar Markets: Pricing Strategies in Emergent Electronic Commerce." *Journal of Business Strategies*, 19(1):55–71.

List, J. A., 2006. "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions." *Journal of Political Economy*, 114(1):1–37.

Lu, S. F., 2012. "Multitasking, Information Disclosure and Product Quality: Evidence from Nursing Homes." *Journal of Economics and Management Strategy*, 21(3):673–705.

Luca, M., 2011. "Reviews, Reputation, and Revenue: The Case of Yelp.com." *Harvard Business School Working Paper*, 12-016.

Luca, M. and Zervas, G., 2016. "Fake It Till You Make It: Reputation, Competition, and Yelp Review Frauds." *The Quarterly Journal of Economics*, 62(9).

Matsumoto, B. and Spence, F., 2016. "Price beliefs and experience: Do consumers' beliefs converge to empirical distributions with repeated purchases?" *Journal of Economic Behavior & Organization*, 126(PA):243–254.

McCall, J. J., 1970. "Economics of Information and Job Search." *Quarterly Journal of Economics*, 84(1):113–126.

McCarthy, I. and Darden, M., 2003. "Supply-side Responses to Public Quality Ratings : Evidence from Medicare Advantage." *American Journal of Health Economics*, forthcoming.

Mortensen, D. T., 1970. "Job Search, the Duration of Unemployment, and the Phillips Curve." *American Economic Review*, 60(5):847–862.

Pan, X., Shankar, V., and Ratchford, B. T., 2003. "The Evolution of Price Dispersion in Internet Retail Markets." *Advances in Applied Microeconomics*, 12:85–105.

Resnick, P., Zeckhauser, R., Swanson, J., and Lockwood, K., 2006. "The value of reputation on eBay: A controlled experiment." *Experimental Economics*, 9(2):79–101.

Smith, M. D. and Brynjolfsson, E., 2001. "Understanding Digital Markets: Reviews and Assessments". *Understanding the Digital Economy*, Brynjolfsson and Kahin, eds., MIT Press, Cambridge, MA.

Sorensen, A. T., 2000. "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs." *Journal of Political Economy*, 108(4):833–850.

Spence, F., 2015. "Consumer Experience and the Value of Search in the Online Textbook Market." *working paper.*

Stigler, G. J., 1961. "The Economics of Information." *Journal of Political Economy*, 69 (3):213–225.

Tadelis, S., 2016. "Reputation and Feedback Systems in Online Platform Markets." *Annual Review of Economics*, 8.

Tang, F. F. and Xing, X., 2001. "Will The Growth Of Multi-Channel Retailing Diminish the Pricing Efficiency of the Web?" *Journal of Retailing*, 77:319–333.

Tang, Z., Smith, M. D., and Montgomery, A., 2010. "The Impact of Shopbot Use on Prices and Price Dispersion: Evidence from Online Book Retailing." *International Journal of Industrial Organization*, 28(6):579–590.

Wilson, C. M., 2010. "Ordered Search and Equilibrium Obfuscation." *International Journal of Industrial Organization*, 28(5):496–506.

Xia, L. and Monroe, K. B., 2004. "Price Partitioning on the Internet." *Journal of Interactive Marketing*, 18(4):63–73.