

The Effects of Online Review Platforms on Restaurant Revenue, Survival Rate, Consumer Learning and Welfare

Limin Fang*

January 26, 2019

Abstract

This paper quantifies the effects of online review platforms on restaurants and consumer welfare. Using a novel dataset containing restaurant revenues and information from major online review platforms, I show that online review platforms help consumers learn faster about restaurant quality. The effects on learning show up in restaurant revenues and survival rates. Specifically, doubling consumers' exposure to Yelp, the dominant platform, increases the revenue of a high-quality new independent restaurant by 8-20% and decreases that of a low-quality restaurant by a similar amount. Doubling Yelp exposure also raises the survival rate of a young high-quality independent restaurant by 7-19 basis points and reduces that of a low-quality restaurant by about the same degree. Other platforms have similar effects but in smaller magnitude. In contrast, online review platforms do not significantly affect the revenues or survival rates of chains and old independent restaurants. Building on this evidence, I develop a structural demand model with social learning. Counterfactual analyses indicate that online review platforms improve the welfare for restaurant goers by \$2.5 per person per meal, equivalent to a 12.6% discount on the average meal price. Despite large effects on individual restaurants, online review platforms have little impact on the total industry revenue.

Keywords: consumer learning, online platforms, restaurant quality, dynamic structural models, differentiated products.

*University of British Columbia. Email: limin.fang@sauder.ubc.ca. I would like to thank Victor Aguirregabira, Matthew Osborne, Andrew Ching, Nathan Yang and Ron Borkovsky for great advice, and I am grateful for helpful comments from Paul Ellickson, Keith Head, Ralph Winter, James Brander, Sanghoon Lee, Eduardo Souza-Rodrigues, Avi Goldfarb, Yao Luo, Matthew Mitchell, April Franco, Heski Bar-Isaac, Genet Zinabu, Scott Orr, Demetrios Vakratsas, Laura Lasio and all participants at my seminars at various universities. All errors are mine.

1 Introduction

Over the past two decades, online review platforms, such as Yelp, TripAdvisor, Amazon, and Google, have provided consumers with an enormous amount of information on products. Online reviews have enabled consumers to learn from each other's experiences, have reduced the information asymmetry between consumers and firms, and have helped consumers make better decisions at purchasing products. Given their growing prevalence, it is important to understand the role of these platforms in the economy; in particular, the value of public goods they generate for consumers by producing information. A good understanding of this role will provide the basis for government policies and regulations on online platforms in the digital age¹.

This paper quantifies the effects of online review platforms on firms' revenues, firms' survival rates, consumer learning and welfare. More specifically, I examine the effects of the major online review platforms Yelp, TripAdvisor and Google in the restaurant industry in Texas. I ask two research questions: (1) Do online review platforms help consumers learn faster about restaurant quality? (2) If they do, what are their effects on consumer welfare and industry dynamics? To answer these questions, I collected a novel dataset containing restaurants' revenues, consumers' search intensity for the platforms and the reviews posted on the platforms over the period of 1995-2015. Using this dataset, I test in a reduced-form analysis whether online review platforms speed up consumers' learning process. Online review platforms' effect on learning should show up in restaurant revenues and survival rates. In particular, in areas where review platforms are more widely available, the market share of higher-quality restaurants should rise and that of lower-quality restaurants should fall. These effects should be more pronounced for new independent restaurants, whose quality is more uncertain, but should be more subdued for chain and old established independent restaurants, as their quality is often known to consumers. Furthermore, the changes in market shares should translate into similar patterns in restaurants' survival rates.

I use a triple-difference (3-D) approach that exploits the variation in the timing of penetration and search intensity of the online review platforms across various regions in Texas. I find that doubling consumers' exposure (as measured by search intensity) to Yelp increases the revenue of a high-quality young independent restaurant by 8-20% and decreases that of a low-quality restaurant by about the same amount. Doubling Yelp exposure also raises the survival rate of a high-quality young independent restaurant by 7-19 basis points and reduces that of a low-quality restaurant by a similar level. Other platforms, especially Google, have similar effects but in smaller magnitude. In contrast, online review platforms do not affect the revenues or survival rates of chains and old independent restaurants. Yelp's effect appears to overpower those of other platforms. For example, for the group of restaurants listed on both Yelp and other platforms, Yelp exposure is the only factor that has a significant effect on revenue.

Building on the reduced-form evidence, I develop a novel structural demand model with consumer learning based on the aggregate Bayesian learning model in Ching (2010a) and the constant

¹See European Parliament briefing, *Online Platforms: How to Adapt Regulatory Framework to the Digital Age?* [http://www.europarl.europa.eu/RegData/etudes/BRIE/2017/607323/IPOL_BRI\(2017\)607323_EN.pdf](http://www.europarl.europa.eu/RegData/etudes/BRIE/2017/607323/IPOL_BRI(2017)607323_EN.pdf)

expenditure demand model with differentiated products in Björnerstedt and Verboven (2016). I extend these models by incorporating consumers' exposure to both online review platforms and other information sources into the learning process. Accounting for other information sources is important as it allows for a more accurate assessment of the effect of online reviews. In particular, it ensures that learning still takes place even without online review platforms; learning may be slower, but does not disappear. The complete disappearance of learning is a common assumption in the existing literature, and models based on it are likely to overestimate the effect of online review platforms. This paper shows that without accounting for other information sources, the welfare effect can be overestimated by as much as 30%.

Based on the estimated structural model, I conduct three counterfactual analyses: (1) one where I remove the existence of online review platforms alone; (2) one with no learning, and (3) one with full information. The results from the no-online-review counterfactual show that online reviews improve the welfare of restaurant-goers by \$2.51 per person per meal, equivalent to a 12.6% discount on the average meal price at restaurants. The no-learning counterfactual demonstrates that without other information sources, the welfare would drop by another \$0.76 per restaurant-goer. The full-information counterfactual illustrates that if information is complete, consumer welfare can increase by an additional \$2.16 per person per meal, equivalent to another 11% discount. The difference between the full-information and no-learning counterfactuals reveal the value of information for consumers: a \$5.43 value per person per meal, or 27.3% of an average meal price.

These counterfactual analyses also confirm that online review platforms have opposite effects on the market shares of high- and low-quality young independent restaurants: they increase higher-quality restaurants' shares by about 16% and decrease those of lower-quality restaurants' by 15-21%. Although they affect individual restaurants substantially, the reviews have little impact on the total industry revenue. The substitution of demand occurs almost entirely within the young independent restaurants, with high-quality restaurants drawing demand away from low-quality restaurants.²

1.1 Relevant Literature

Recent research has examined the impact of online reviews through their effects on sales, learning, competition and consumer welfare. For example, Chevalier and Mayzlin (2006) find that higher ratings online lead to an increase in book sales. Zhu and Zhang (2010) investigate how much the effect of online reviews on sales is influenced by product and consumer characteristics, and find that online reviews have a greater impact on the sales of less popular video games and those games whose players are more experienced internet users. Zhao et al. (2013) use a structural learning model to study the effect of online product reviews on consumer purchases of books and find that consumers learn even more from online reviews than from their own experiences of reading similar books in the past. Wu et al. (2015) examine the economic value of online reviews for consumers

²This result is also consistent with an alternative model with costly information about restaurants. Consumers only look up reviews after having decided to go to restaurants. Therefore, reviews affect only substitution within restaurants, but not between restaurants and other eating options.

and restaurants using web browsing data for seven restaurants listed on Dianping.com. Newberry and Zhou (2018) investigate the heterogeneous effect of online reputation for local and national retailers. This paper adds to this literature by showing how consumers' exposure to online reviews can affect both consumer welfare and market dynamics.

This paper is most closely related to Luca (2016).³ Luca employs a regression discontinuity design (RDD) to evaluate the causal effect of Yelp rating changes on restaurant revenues. He finds that an increase in displayed rating by one star raises the revenues of independent restaurants by 5%-9%, but does not affect those of chain restaurants. Luca also discusses whether the way consumers use reviews is consistent with Bayesian learning. This paper differs from Luca's work in a number of aspects: (1) This paper examines not only Yelp but also other online review platforms, including TripAdvisor and Google; (2) This paper focuses on the question of whether online platforms speed up consumer learning by examining the opposite effects of their penetration on high- and low-quality restaurants' revenues, while Luca focuses on the direct effect of online displayed ratings on revenue without emphasizing the channel of the effects; (3) I look at online review platforms' effects not only on revenue, but also on restaurants' survival probabilities; (4) I develop a structural model to quantify the effect of online review platforms on consumer welfare and industry dynamics through counterfactual experiments.

This paper is also closely connected to Lewis and Zervas (2016). Lewis and Zervas (2016) assess the welfare effects of online reviews in the hotel industry, using data on hotel revenues and reviews from a number of crowd-sourcing websites including TripAdvisor and Expedia. Compared to their work, this paper differs in two important dimensions beyond the focus on the restaurant industry: (1) This paper estimates a model with social learning, where the effects of reviews are gradual, whereas their model assumes that the effect is instantaneous.⁴ (2) I incorporate other information sources into consumers' learning process, whereas their study has consumers rely only on online reviews for quality information. In other words, Lewis and Zervas (2016) assume that without online reviews, consumers would never know the true quality of a hotel. That assumption makes the counterfactual in their study equivalent to the no-learning counterfactual in my study. As discussed previously, the welfare difference between the no-learning and the full-information counterfactuals is \$5.43 per person per meal, more than twice the difference between the no-online-review counterfactual and the real world (\$2.51), which is the true effect of online reviews. In this regard, the model in Lewis and Zervas (2016) is likely to overestimate the welfare effect of online reviews.

I find that chain restaurants' market shares are not affected by the increasing popularity of online reviews, in contrast to the results in Luca (2016), Newberry and Zhou (2018) and Hollenbeck (2016), all of which show that online reviews tend to redistribute demand from chains to independent firms. I do not find this result in my study because the average quality of the independent rivals that

³Anderson and Magruder (2012) examine the effect of Yelp ratings on customer flows also using a RDD approach. It is very closely related to Luca (2016), but different in that it does not use revenue data and does not investigate the aspect of consumer learning.

⁴The "instantaneity" assumption makes their model of the real world equivalent to the full-information counterfactual in my study.

chains face in my data is mediocre, and as a result the opposite effects of online reviews on high-quality and low-quality independent restaurants cancel each other out, leaving a net zero effect on the market shares of chain restaurants.

More broadly, this paper contributes to the literature on information, reputation and learning. Examples include the seminar paper Stigler (1961) on the economics of information, Conley and Udry (2010) on social learning, Cai et al. (2009) on observational learning, Acemoglu et al. (2017) on the theory of Bayesian learning from online reviews, Crawford and Shum (2005) on learning and matching, Akerberg (2001 and 2003) on advertising and learning, Jin and Leslie (2003) on how information improves product quality improvement, Berger et al. (2010) on when negative reviews improve sales, Ching, Erdem and Keane (2013 and 2017) on consumer learning models, Anderson (2012) on online reputation and advertising, Newberry (2016) on observational learning in the online music industry, and Ching (2010b) on the effect of learning on dynamic competition in the pharmaceutical industry.

In the context of this broader literature, the contribution of this paper is three-fold: (1) First, the paper is the first that quantifies the welfare value of online review platforms by estimating a demand model with social learning. In particular, it incorporates learning from both online reviews and other information sources, a feature that allows a more accurate account of the effects of online reviews than most models in the existing literature. (2) Second, this paper is the first to document a number of empirical results, including online reviews' effect on survival rates of firms⁵, their opposite effects on high- and low-quality young independent restaurants, zero effect on chain restaurants, and the dominance of Yelp's effect over other platforms'; (3) Third, I develop a novel structural demand model with social learning that utilizes the aggregate revenue data at the product level. The constant expenditure model is particularly convenient for expressing consumers' utilities in terms of revenues instead of quantities. The combination of a social learning model with a constant expenditure model to estimate demand using aggregate revenue data is new.

The paper proceeds as follows. Section 2 describes the data and summarizes the main features of the data. Section 3 discusses the motivating evidence based on the reduced-form analysis. Section 4 introduces the structural demand model with social learning. Section 5 describes the identification and estimation strategy of the structural model, and discusses the findings of the counterfactual analyses. Section 6 concludes by summarizing the key results.

2 Data

My data come from a variety of sources: (1) The first one is the restaurant mixed-beverage revenue data from the Mixed Beverage Tax Information Records held by the Office of the Comptroller of Public Accounts in Texas; (2) The second source includes census data, American Community Survey, visitor and tourism spending data from Texas Economic Development and Tourism, aver-

⁵Anenberg et al. (2018) show using Yelp data that restaurants tend to exit after receiving negative reviews. Different from their paper, I show that consumers' exposure to online reviews helps high-quality restaurants to survive longer and low-quality exit faster.

age daily traffic flow data from the Texas Department of Transportation, consumer expenditure on eating out from the Consumer Expenditure Survey, and wage data from Occupational Employment Statistics; these data inform the market demand and cost characteristics. (3) The third source is the Google Trends data on the search term Yelp and the website TripAdvisor; this data reflects the geographic penetration of Yelp and TripAdvisor over time and are used as the main measure for consumers' exposure to online platforms. My paper focuses on Yelp's effect on restaurants' revenue and survival rates, and uses TripAdvisor as the main control for the penetration of other platforms. (4) The fourth source is detailed restaurant level data collected from the Yelp, TripAdvisor and Google websites, including restaurant overall ratings, pricing, and restaurant cuisine category on each platform. For restaurants listed on Yelp especially, I collected the rating history with time stamp and star rating for each review for each restaurant on Yelp. Below I describe each data source in greater detail.

The first dataset, the restaurant revenue data, is the gross receipts of alcoholic drinks sold at each restaurant in Texas. The volumes of these gross receipts are reported to the State of Texas Comptroller Office on a monthly basis by mixed-beverage permit holders. This dataset started in 1993 and contains information on a restaurant's monthly gross receipts from liquor, wine and beer beverages, restaurant address, owner names and addresses, business location name and address, permit number, permit issue date and out of business date. The dataset covers *all* establishments that hold (or held) a mixed-beverage permit, including short-term events, such as festivals, and non-restaurant entities like convention centers. To obtain a homogeneous set of sample, I select full-service restaurants based on their NAICS code and their names⁶. I exclude primarily night life oriented places such as night clubs or sport bars, and I also remove restaurants at major hotels and airports as they are often monopolies at their locations and are subject to a unique set of regulations. To avoid potential irregularities associated with the earlier effort at collecting this data, I set my sample period to January 1995 - December 2015. This elimination process leaves 15,417 restaurants in my dataset, about 37% of all establishments recorded in the Mixed Beverage Tax Information Records.

In my analysis, I choose the market definition as $\text{zipcode} \times \text{month}$. The geographic area at the zipcode level is appropriate for modeling the competition between full-service restaurants given that most consumers for restaurants are local. To complement the restaurant dataset with market level information, I collected the second dataset related to market demand and costs. They include demographic characteristics, income, tourist spending, traffic volume data, food service industry wage data, and consumer expenditure data. A detailed account of these data sources and how I interpolate some data to each market is shown in Section A of the Online Appendix.

The third dataset, the Google Trends data, tracks the search interest on Google for the term "Yelp" and the TripAdvisor website at the metropolitan area level. The data is monthly, dates from January 2004 to November 2016, and spans across 20 metropolitan areas in Texas⁷. For each metro

⁶Although the mixed-beverage receipt dataset does not include NAICS code, I merged it with another dataset, the Sales Permit dataset, which contains the NAICS code information.

⁷Metropolitan areas are the finest geography for which the data is available. Although these metropolitan areas

area, the data is a time series of a normalized index ranging from 0 to 100. The normalization is carried out by dividing the search volume for the term “Yelp” or the TripAdvisor website for a given metro region at a particular time by the total searches for all terms in that geography at that time. This results in a proportion of the search volume for Yelp or TripAdvisor out of the total search volumes for a given month in a metro area. Among these proportions, the highest one for a metro area is normalized to 100. All the other proportions are scaled accordingly. This normalization is to control for the total population in each metro area, so that the metro area with the highest population does not always rank the first for consumers’ exposure to Yelp or TripAdvisor. For example, Austin ranks the highest by the normalized search interest for Yelp, but it is only the 4th largest metro region by population and the 3rd by the number of restaurants. Through this normalization, the Google Trends search interest can be interpreted as the attention paid to Yelp or TripAdvisor out of a person’s attention on all subjects on Google. This normalized search interest data provides an aggregate measure of consumers’ exposure to Yelp and TripAdvisor and reflects both demand and supply of information on these platforms. It also tracks the geographic penetration of these platforms over time, which provides an important source of variation for identifying the effect of these online platforms on restaurants’ revenues and survival rates.

The fourth data source zooms in at the restaurant level. Of all restaurants included in my dataset, I identified those that are listed on Yelp, TripAdvisor and Google using Yelp API, Google API and automated Bing search on Microsoft Azure. This information includes restaurant overall ratings, pricing and cuisine type on each platform. The data for Yelp and Google were collected in November 2016, and TripAdvisor data were collected in June 2018. The distribution of restaurants across the platforms are illustrated in Figure 1. As shown in the figure, Google covers the highest number of restaurants, and most restaurants that are listed on Yelp and TripAdvisor are also listed on Google. Overall, there are 9,024 restaurants listed on Google, of which 5,307 are also listed on Yelp and 4,726 on TripAdvisor, and 3,591 are listed on all three platforms. TripAdvisor and Yelp alone have about 3,817 overlapping restaurants. In total, 9,900 restaurants have online presence; they account for about 64.2% of all restaurants in the revenue dataset, and 96.5% of all restaurants that were active as of December, 2015, the end of the sample period. Basically by the end of the sample period, almost all restaurants had some type of online presence.⁸

All these platforms provide services to both businesses and consumers. For businesses, they offer business analytics and advertising services, and for consumers, they perform not only as a business directory but also a platform that aggregate consumers’ opinions on a business. For both Yelp and TripAdvisor, anyone can create a profile for a business. Business owners can then choose to claim the business or not.⁹ For Google, since it has the Google Maps GPS system, it covers almost all businesses. This explains why Google has the highest coverage for restaurants in my dataset. All these platforms offer a discrete 1-5 star rating system, and have prices shown in ranges, such as “\$” or “\$\$”. There is also a description of the restaurant cuisine type on each of the platforms.

are delineated by Google, they are very close to the census definition of metropolitan statistical areas.

⁸Note that these platforms retain a record of restaurants that are closed because they draw traffic to their websites.

⁹This is so more for Yelp. TripAdvisor does not offer the “claim” option.

This detailed restaurant-level information allows me to control for restaurant characteristics in my analysis.

For restaurants listed on Yelp, I also collected the rating history for each restaurant, including the time stamp and star rating for each review of each restaurant. I use this data to complement the previous data regarding online review platforms in a number of ways: (1) First, I use the rating history to pinpoint the start date for Yelp’s penetration. As mentioned before, Yelp is the main player of online review platforms in the restaurant industry, and the focus of the paper. Having a precise information regarding its penetration is of particular interest. According to the rating history data, the very first review on Yelp was written on March 29th, 2005 for an Italian restaurant “Carrabba’s Italian Grill” in Plano, Texas. I use this date as the start of online review platforms’ penetration in my analysis since Yelp was the pioneer platform for restaurant reviews. Google and TripAdvisor’s restaurant review services lagged that of Yelp. (2) Second, I deploy this rating history information to construct an additional measure of Yelp’s penetration over time and geographic regions and compare that to the measure from the Google Trends data. As will be discussed later, this alternative measure gives a more concrete interpretation of the normalized index provided by Google Trends, and supports the validity of Google Trends’ normalized index as a measure for platforms’ geographic penetration. (3) Third, I employ this data to construct various measures of restaurant quality. Knowing each review’s date and rating for each restaurant, I can calculate a restaurant’s average rating at a particular time. This average rating can predate the revenue data. For example, I can use restaurants’ first 10 reviews to construct a predetermined quality measure for each restaurant, and then examine how exposure to Yelp affects their revenues after the first 10 reviews date. This measure is more exogenous compared to the overall ratings displayed on each platform at the date of data collection (i.e. November 2016 or June 2018), which postdates the sample period. Though having the advantage of being a more exogenous measure, this “predetermined” measure limits the number of observations on restaurant revenues. It is used alongside the overall ratings collected in November 2016 and June 2018 as robustness checks. (4) Last, the rating history data also informs how restaurants’ quality might have evolved over time. It is used in the examination of other major channels through which these platforms may have an effect on restaurant revenues.

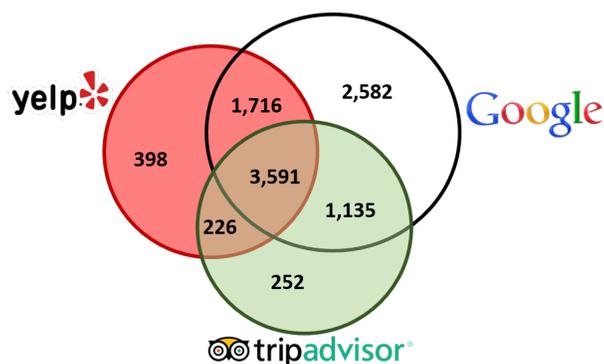


Figure 1: Distribution of Restaurants Listed on Different Platforms

Combining all these sources, I obtain a panel that covers monthly mixed beverage sales for 15,417 restaurants during the period of January 1995 to December 2015, demographic and income information for each zipcode tabulation areas, traffic information based on each outlet’s location, the monthly Yelp popularity measure for each metro area in Texas, the number of monthly reviews and average rating for each Yelp listed restaurant, and the overall rating, price range and cuisine types for all restaurants that are listed on Yelp, Google and TripAdvisor. In total, the dataset includes 1.14 millions of observations in 846 geographic markets, which span over 482 cities and places, 113 counties and 20 metropolitan areas. The key features of this data and some stylized facts exhibited by this data are summarized in the next section.

2.1 Summary Statistics and Measures

Table 1 summarizes key features of the data. All monetary values in this dataset have been adjusted for inflation to December 2000 values. On average restaurants make \$23,421 per month from mixed beverages. The highest amount of revenue is over \$1.67 million. This indicates that there is a large variation in restaurant revenues. Some restaurants have survived a long time, as old as 22 years (or 265 months). There are also many new restaurants that just opened. The average age of a restaurant in my dataset is about 6 years (or 74 months).

These restaurants received a wide range of ratings on Yelp, Google and TripAdvisor. In my study, I differentiate between independent and chain restaurants since consumer learning is likely different for chain restaurants compared to independent restaurants. I define chain restaurants as those restaurants that share the same name and whose total number of outlets at a given month reached 10 or greater in Texas during the sample period. By this definition, there are 3,328 unique chain restaurants in my dataset. As shown in the table, on average, independent restaurants received higher ratings than chain restaurants. The overall average rating for independent restaurants is 3.52 on Yelp (as of November 2016), 3.93 on Google (as of November 2016), and 4.0 on TripAdvisor (as of June 2018), whereas that for a chain restaurant is 2.96 on Yelp, 3.63 on Google and 3.75 on TripAdvisor. Across online review platforms, restaurants’ ratings are also systematically higher on Google and TripAdvisor than on Yelp. Figure 3 in Appendix A illustrates the overall average rating distributions for all three platforms. As shown, rating distributions for all three platforms are skewed to the left, and Google’s and TripAdvisor’s ratings are higher than Yelp’s; their mode ratings are 4.0 compared to 3.5 on Yelp.

Although restaurants are rated differently across different review platforms, these overall ratings are correlated across platforms. Their correlations are shown in Table 2. As can be seen, the correlation between Yelp and Google ratings is higher than that between Yelp and TripAdvisor or between Google and TripAdvisor. The correlation between the first two is 0.72, while those between either of them and TripAdvisor are only 0.58 to 0.59. These pairwise correlations indicate that TripAdvisor is different from the other two platforms. TripAdvisor targets primarily travellers or tourists, whose taste and rating standards for restaurants may be systematically different from those for Yelp and Google, which are used widely by both the local and tourist population. Given these differences in targeted users, I use Yelp and Google ratings as the primary measures for

restaurant quality. I make this decision because the paper’s focus is on consumers’ continuous learning of quality, and tourists, who are one-time visitors, rarely learn accumulatively.

Table 1: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	Number of Obs.
Monthly revenue (\$)	23,421	32,567	0	1,668,320	1,135,220
Number of months a restaurant survived to	74.6	71.3	1	265	15,417
Yelp Average Rating for independent (November 2016)	3.52	0.59	1	5	3,995
Yelp Average Rating for chain (November 2016)	2.96	0.55	1	5	1,935
Google Average Rating for independent (November 2016)	3.93	0.43	1.5	5	6,356
Google Average Rating for chain as of (November 2016)	3.63	0.46	2	5	2,668
TripAdvisor Average Rating for independent (June 2018)	4.0	0.47	1	5	3,617
TripAdvisor Average Rating for chain (June 2018)	3.75	0.51	1	5	1,587
Yelp search interest (Google Trends)	4.5	9.2	0	100	5,040
TripAdvisor search interest (Google Trends)	10.8	15.3	0	100	5,040

Note: All monetary amounts are in December 2000 dollars.

Table 2: Correlation Between Ratings Across Platforms

	Yelp	Google
Google	0.72	
TripAdvisor	0.59	0.58

Table 1 also highlights some key characteristics of the Google Trends data on consumers’ search interest for Yelp and TripAdvisor. During January 1995 to December 2015, the average Yelp search interest for a metro area at a given month is 4.5, and the standard deviation is 9.2, indicating a large variation in Yelp’s popularity across regions and time¹⁰. For TripAdvisor, the search interest is higher than Yelp, with an average of 10.8 and a standard deviation of 15.3. This larger search interest can be explained by the fact that TripAdvisor was founded in February 2000, 4 years earlier than Yelp. Consumers’ exposure to a more established platform, TripAdvisor, is therefore higher than Yelp on average.

The overall trends in consumers’ search interest for both websites for the entire state of Texas are illustrated in Figure 4 in Appendix A. As can be seen, TripAdvisor was active long before Yelp became popular in Texas. We can also see that although Yelp started in Texas in March 2005, its popularity did not pick up until August 2007, and at around 2010, its popularity surpassed that of TripAdvisor¹¹. It should be noted that despite its popularity, TripAdvisor is a platform for both hotels and restaurants, while Yelp concentrates mostly on restaurants. Nonetheless, the penetration of these two platforms over time is highly correlated, especially for years after 2010. This may be due to the increasing trend in the usage of the Internet and the Google search engine.

As mentioned before, besides the Google Trends data, I also use the rating history on Yelp to construct an alternative set of measure for Yelp’s penetration. The rating history data gives the

¹⁰ For periods before the start of Yelp’s penetration, March 2005, the search interest for Yelp is set to 0

¹¹I have also collected the Google Trends data on smaller platforms such as urbanspoon and zomato. However, their search interest is very small in comparison to those for Yelp and TripAdvisor. Therefore, I have not included them in the graph or in the empirical analysis of the paper.

number of reviews written about a restaurant in each month. At the restaurant level, this data is highly endogenous with respect to the restaurant revenue because the higher the number of visitors to a restaurant, a higher the restaurant’s revenue, and the greater the number of reviews written about this restaurant. Given these considerations, I aggregate the number of reviews for each month at the metro region level and then divide them by the number of restaurants in that region to obtain an average number of reviews per month per restaurant. This measure is likely to even out the idiosyncratic shocks at the restaurant or market level, and can be used as an alternative representation of Yelp’s penetration. I plot this measure against time and compare the graph to that of the Google Trends data for major metro regions in Texas (including Austin, Houston, Dallas and Fort-Worth, and San Antonio). I display these two graphs in Figure 2. The left panel in the figure is the Google Trends search interest data, and the right the average number of reviews per months per restaurant in a metro region. As can be seen, these graphs have almost identical shapes and trends, and both show that Yelp is the most popular in Austin than the other major regions. The correlation between these two measures is 0.85. Given this high correlation, I use only one measure in my empirical analysis to show the effect of Yelp on consumer learning and restaurant revenues, and that measure is the Google Trends data. Similarly, to capture the effect of other platforms, I use Google Trends search interest data for other online review platform.

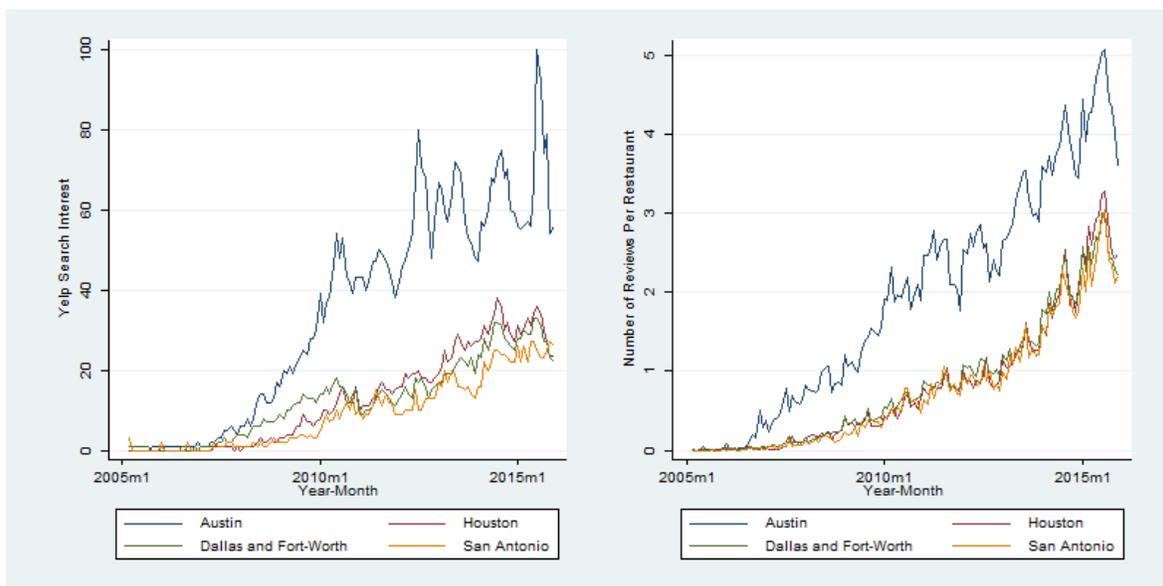


Figure 2: Yelp Search Interest and Number of Reviews per Restaurant for Selected Metro Areas, 2005-2015

3 Motivating Evidence

In this section, I provide motivating evidence through a reduced-form analysis that online review platforms help consumers learn about restaurant quality. The analysis is based on the following idea: When a high-quality independent restaurant first opens, consumers would not know its true

quality, and therefore people would not go there very often. However, as consumers learn about the restaurant’s quality from both their own experiences and others’ experiences through word-of-mouth communication (including reading reviews from online review platforms), they would update their beliefs of the restaurant’s quality and visit there more often. Due to learning, over time, this high-quality independent restaurant’s revenue would go up, and the increase in revenue should be faster in areas where consumers learn faster. Given that online review platforms provide information on restaurant quality and consumers learn from these reviews, high-quality independent restaurants’ revenues should grow faster in areas where online review platforms are more popular. The reverse should be true for lower-quality independent restaurants.

As independent restaurants pass a certain age, however, their quality would have been fully revealed to consumers. Online review platforms’ popularity would no longer have an effect on their revenues. Similarly, for chain restaurants, as they benefit from their headquarters’ centralized advertising and from the chain’s reputation built through multiple locations with uniform service, it is likely that their quality is known to consumers already when they first open. As a result, online review platforms should not affect chain restaurants’ revenues substantially.

Nonetheless, chain restaurants and old established independent restaurants’ revenues could change through the effect of competition induced by online review platforms. For example, if young independent restaurants in the market are mostly of high quality, as online review platforms help consumers learn faster about their qualities, demand would be directed from chain or old established independent restaurants to young independent restaurants. In that case, chain and old established independent restaurants’ revenues could go down with the penetration of online review platforms.

How chain and old established independent restaurants’ revenues would change as a result of online review platforms’ penetration would depend crucially on the average quality of their young independent rivals in the market. If the overall average quality is mediocre, then it is likely that the increase in high-quality young independent restaurants’ revenues would be offset by the decrease in low-quality restaurants’ revenues, leaving a net zero effect on other restaurants in the market. The sign of the overall effect on the revenues of chain and old established independent restaurants is an empirical question.

Based on the above analysis, I test two hypotheses: Hypothesis (1) online review platforms have the opposite effects on high- and low-quality young independent restaurants; Hypothesis (2) online review platforms have an effect on the revenues of chain and old established independent restaurants, but with the same sign for all quality levels. The empirical strategy are described in Subsection 3.1, and results shown in Subsection 3.2. At the end of this section, I also discuss other potential mechanisms through which online review platforms may affect restaurants’ revenues and rule them out using a number of tests.

3.1 Empirical Strategy

To test these effects of online review platforms, I use a triple difference-in-difference (3-DID) approach. I exploit the variation in the penetration of online review platforms across region and time

to tease out their effects on restaurants' revenues. As shown previously in Figure 2, the penetration of the dominant platform Yelp varied across metropolitan regions in Texas. Not only the intensity of the penetration was different, but also the timing. Although the graphs show the penetration of Yelp for only four metropolitan regions, a similar type of variation in time and intensity can be found for other metro regions and for other online review platforms as well. The main idea behind the 3-DID approach is that for a given quality, the change (increase or decrease) in young independent restaurants' revenues should be larger in regions with higher exposure to online review platforms than those with lower exposure. By controlling for restaurant fixed effects and calendar time fixed effects, I compute the change in restaurants' revenues for a given quality level for each region and relate that to each region's exposure to online reviews. This relationship identifies the effect of online review platforms. In addition, I also compare the effects across quality groups, age groups, and chain affiliation.

The econometric specification is as follows:

$$\begin{aligned} \log(Rev.al_{jt}) = & [\theta_y + \theta_{yr}R_j] \times \log(Y_{mt}) \times (1 - D_j^C) + [\theta_{yc} + \theta_{ycr} \times R_j] \times \log(Y_{mt}) \times D_j^C \\ & + \mathbf{X}_{jt}\boldsymbol{\theta}_x + \theta_{nI}n_{jt}^I + \theta_{nc}n_{jt}^C + \theta_t + \theta_{tc} \times D_i^C + \theta_j + \xi_{jt} \end{aligned} \quad (1)$$

where $\log(Rev.al_{jt})$ is the logarithm of restaurants' revenues from alcoholic drinks. I assume that restaurants' alcohol sales account for a fixed proportion of restaurant total revenues. This fixed proportion is restaurant specific and can be accounted for by restaurant fixed effect θ_j . Under this assumption, the changes in restaurants' alcoholic drink sales reflect the fluctuations in restaurants' total revenues. This assumption is in line with the restaurant industry standard that alcohol sales represent roughly a fixed proportion of the total revenue because alcohol drinks and food are often consumed together at full-service restaurants. Empirical support for this assumption can be found in Section B of the Online Appendix. $\log(Y_{mt})$ is the logarithm of the Google Trends measure of exposure to online review platforms in market m at time t . R_j is the quality of restaurant j . This quality is assumed to be time invariant. This assumption will be discussed in greater detail in Section 3.3 and the structural model Section 4. $D_j^C = \mathbb{1}\{\text{restaurant } j \text{ belongs to a chain}\}$ is a chain dummy. \mathbf{X}_{jt} include market demand shifters and restaurant characteristics. n_{jt}^I and n_{jt}^C are the numbers of independent and chain rivals in a market that restaurant j faces at time t respectively. $\boldsymbol{\theta}$ are the coefficients associated with these variables. θ_t and θ_{tc} are calendar fixed effects for independent and chain restaurants respectively. They represent aggregate trends in restaurant revenues due to changes in consumer tastes; for example, over time, consumers may prefer independent restaurants more than chain restaurants or would like to eat in more often instead of dining at restaurants. ξ_{jt} is an i.i.d demand shock for restaurant j at time t and has mean 0.

In this regression, the composite coefficient $\theta_y + \theta_{yr}R_j$ (or $\theta_{yc} + \theta_{ycr}R_j$ for chains) is the percentage change in restaurant j 's revenue when consumers' exposure to online review platforms increases by 100%. This can be seen as the average treatment effect on all restaurants of the quality class R_j .

To differentiate the heterogeneous treatment effects of online review platforms on young and

old restaurants, I divide restaurants into young and old groups based on their entry dates. For restaurants that entered long before online review platforms for restaurants became popular, I labeled them as old established restaurants. More specifically, I use two years before the first Yelp restaurant review was written in Texas (in March 2005) as the cutoff entry date. This date is March, 2003. All restaurants that entered before this month would have been older and relatively established restaurants in the era of online review platforms. Restaurants that entered after the first Yelp review date are defined as young. Note that the choice of the exact cutoff date within the two year window of the first Yelp review does not affect much the results.

As in all DID methods, the key identifying assumption of this analysis is that the trends in restaurant revenues are the same across regions in absence of online review platforms; that is, without online review platforms, regions with different degrees of exposure to online review platforms would show a similar trend in the change in revenues of their restaurants. Factors such as consumers' inherent willingness to learn about restaurant quality in various regions could create variations in restaurant revenue trends that are similar to the effect of online review platform, and if the exposure to online review platforms is correlated with this type of "preexisting" variation, then the estimates of the effects of online review platforms would be biased. To test this kind of endogeneity, I conduct a "pre-trend" analysis or a placebo test, where I move online review platforms' exposure data to earlier dates by 10 years; that is, instead of starting in March 2005, I move online review platforms' penetration to March 1995. Then I ran regression 1 for the sample of restaurant revenues before March 2005. If the measure of exposure to online review platforms picks up other factors that would have the same type of effect on restaurants in absence of online reviews, then we should see the effect of online review platforms being estimated as significant in this "placebo" test. I include the results from this placebo test and other robustness checks in Section B.1 of the Appendix. This test result demonstrates that this type of endogeneity is not of serious concern.

3.2 Online Review Platforms' Effects

This section presents the results for regression 1. I first focus on presenting the effect of Yelp on restaurants listed on Yelp since it is the dominant platform, and then I extend the analysis to include restaurants listed on TripAdvisor and Google.

3.2.1 Yelp's Effects on Revenues

Using the sample of restaurants listed on Yelp, I ran regression 1 separately for old and young restaurants. To control for quality, I use a variety of measures: overall average Google ratings and overall average Yelp ratings as of November 2016, and the "pre-determined" ratings based on the average of restaurants' first 10 reviews.¹² The results are shown in Table 3.

The first two columns represent results for old and young restaurants by using the Google rating as quality measure. The middle two columns are results from using the Yelp ratings, and the last

¹²Note that using the average of the first 10 reviews is without the loss of generality. The results for using the average of first 15 or 20 reviews are very similar.

Table 3: Effects of Yelp Exposure on Revenues

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Revenue	Log Revenue	Log Revenue	Log Revenue	Log Revenue	Log Revenue
Log Yelp (θ_y)	-0.272 (0.218)	-0.575*** (0.143)	-0.0246 (0.137)	-0.357*** (0.0999)	-0.118 (0.115)	-0.474*** (0.156)
Log Yelp×Rating (θ_{yr})	0.0523 (0.0555)	0.155*** (0.0363)	-0.00461 (0.0444)	0.108*** (0.0281)	0.0467 (0.0320)	0.139*** (0.0417)
Log Yelp×Chain Dummy (θ_{yc})	-0.251 (0.170)	-0.0547 (0.0641)	-0.0143 (0.0919)	0.00516 (0.0418)	0.257 (0.286)	-0.151 (0.101)
Log Yelp×Chain Dummy×Rating (θ_{ycr})	0.0625 (0.0413)	0.0225 (0.0174)	-0.00414 (0.0294)	0.00714 (0.0135)	-0.0504 (0.0796)	0.0467 (0.0297)
Controls	X	X	X	X	X	X
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE×Chain Dummy	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant FE	Yes	Yes	Yes	Yes	Yes	Yes
Age Group	Old	Young	Old	Young	Old	Young
Rating	Google Nov. 2016	Google Nov. 2016	Yelp Nov. 2016	Yelp Nov. 2016	Yelp First 10 Reviews	Yelp First 10 Reviews
<i>N</i>	121,083	165,221	145,694	180,278	59,220	65,219

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Estimates for Controls are shown in Section F of the Online Appendix

two are those from using the first 10 ratings. For the first two quality measures, the regressions are run on the entire sample of Yelp listed restaurants, whereas for the “pre-determined” quality measure, the sample is restaurant revenues *after* restaurants received their first 10 reviews on Yelp. As mentioned before, the “pre-determined” ratings are more exogenous compared to the other two quality measures, given that this quality measure pred-dates the revenue sample.

In this table, the first 4 rows of coefficients are the key parameters of interest. They represent the effect of Yelp exposure on restaurant revenues by quality class. The first two rows (θ_y and θ_{yr}) illustrate Yelp’s effects on independent restaurants, and the 3rd and 4th rows show effects on chain restaurants. There are 3 notable features in these regression results: first, it is uniformly true across all regression results that the coefficients associated with chain restaurants are insignificant, implying that Yelp exposure has very little effect on chain restaurants’ revenues. Second, all coefficients associated with Yelp exposure in Columns 1, 3 and 5 are also insignificant. These results are for the sample of old restaurants that entered long before Yelp’s penetration. They suggest that Yelp exposure does not have an effect on old established restaurants. These two features of the results contradict with Hypothesis (2) that online review platforms affect the revenues of chain and old established independent restaurants. However, they are consistent with the analysis that online review platforms influence chain and old established restaurants only through competition effects not learning.

The third feature of the results is that the Yelp coefficients associated with young independent restaurants in the 2nd, 4th and 6th columns are all significant at the 99% confidence level. These significant coefficients suggest that Yelp exposure has a non-trivial effect on young independent restaurants’ revenues. In particular, the intercept parameter θ_y is negative in all 3 columns, indicating that when a restaurant’s quality is low, Yelp exposure has a negative effect on the revenue. The slope parameter θ_{yr} is positive in all 3 regressions, suggesting that the effect of Yelp exposure on revenue increases with the quality level. A negative intercept and positive slope together implies that Yelp has the opposite effect on the revenues of high- and low-quality restaurants. This

“opposite” effect is best illustrated by the composite coefficient $\theta_y + \theta_{yr}R_j$, which represents the effect of Yelp for a given quality level, and they are summarized in Table 4¹³.

Table 4: Effects of Yelp Exposure on Revenue by Quality Level

Star Rating	Google	Yelp	Pre-determined Rating
2	-0.266*** (0.0715)	-0.141*** (0.0468)	-0.197** (0.0784)
3	-0.111*** (0.0380)	-.0325 (0.0259)	-0.0581 (0.0473)
4	0.0433** (0.0202)	0.0757*** (0.0270)	0.0806** (0.0425)
5	0.198*** (0.0447)	0.184*** (0.0487)	0.219*** (0.0697)

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

As can be seen, when Yelp exposure increases by 100%, the revenue of a 2 star restaurant based on Google ratings would decline by 26.6%, while that of a 5 star restaurant would increase by 19.8%. For restaurants in the middle tier quality classes with 3 or 4 star ratings, the effects are more subdued. This pattern persists when we use either Yelp or “pre-determined” ratings. In particular, when we use Yelp ratings, the effect of 100% increase in Yelp exposure on a 2 star rated restaurant is much smaller than that estimated by using Google rating: 14.1% v.s. 26.6%. This stark difference can be explained by the fact that Google ratings are systematically higher than Yelp ratings. Recall from the summary statistics of the data in Table 1, the average of overall ratings on Google is rated 4 star while being 3.5 on Yelp. A 2 star restaurant on Yelp is roughly equivalent to a 2.5 to 3 star restaurant on Google. Therefore, the decline in the revenue of a 2 star restaurant as rated by Yelp is much less dramatic as that of a 2 star restaurant as rated by Google. At the higher end, however, the effects of Yelp exposure on top rated 5 star restaurants are very similar (19.8% v.s. 18.4%) across the two rating systems. This is not surprising since 5 stars is the top tier. The results from using the pre-determined ratings are also very similar to those based on Yelp and Google ratings. These results are consistent with Hypothesis (1).

As robustness checks of these results, I also use Heckman’s correction to control for endogenous exit, time×metro region fixed effects to account for region and time specific factors that may be correlated with the penetration of Yelp, and a placebo test to see if the “pre-trend” of restaurant revenues in various regions were indeed uncorrelated with the cross-section variations in the Yelp search interest data. All of these robustness checks confirm the validity of the results shown in Table 4. The details of these robustness analyses are shown in Appendix B.2 .

To explore further whether consumers take ratings on Yelp as quality signals and react to them, I also conduct a regression discontinuity (RDD) analysis using the Yelp review history data to evaluate the effect of changes in Yelp displayed ratings on restaurant revenues. The RDD takes advantage of the rounding in displayed ratings on Yelp, which rounds the raw average ratings to

¹³Since the sample of new restaurants do not include restaurants with a 1-star rating, I do not report the Yelp effect for that quality class.

the nearest 0.5. The details are presented in Section C of the Online Appendix. The results of the RDD analysis are consistent with the results in this section: changes in displayed ratings on Yelp affect only young independent restaurants’ revenues, and have very little effect on either chain or old established independent restaurants.

In summary, these results are consistent with what we would expect to see in terms of how the revenues of restaurants would change if Yelp helps consumers learn about restaurant quality.

3.2.2 Yelp’s Effect on Survival Rates

Given that exposure to Yelp affects restaurants’ revenues, it should influence restaurants’ survival rates as well. In particular, it should help the market weed out the lower quality restaurants by pushing them to exit faster, and assist the higher-quality restaurants in staying in market longer. To examine how Yelp affects restaurants’ survival rates, I run a linear probability model using almost exactly the same econometric specification as regression 1, only that I replace the dependent variable ($\log(Rev_{aljt})$) with $action_{jt} = \mathbb{1}\{\text{restaurant } j \text{ is active at time } t\}$, and I add in w_{mt} , wage for food service workers, as an additional explanatory variable. In this regression, I use both Google and Yelp ratings as quality measures because the sample associated with predetermined ratings contain a much smaller number of restaurant exits. The regression results are shown in Table 21 of Appendix G. Overall the results are consistent with Yelp’s effect on revenue; that is, Yelp exposure again affects only the survival rates of young independent restaurants but not those of chain or old established restaurants. Here I present in Table 5 the composite parameters that demonstrate Yelp’s effect on the survival rates of young independent restaurants.

Table 5: Effects of Yelp Exposure on Survival Rate

Star Rating	Google	Yelp
2	-0.00293** (0.00144)	-0.00179* (0.00102)
3	-0.00139 (0.000848)	-0.000556 (0.000631)
4	0.000161 (0.000561)	0.000678 (0.000618)
5	0.00171* (0.000930)	0.00191* (0.000992)

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

As can be seen, when Yelp exposure doubles, the survival rate of a 5-star new independent restaurant would increase by 17.1–19.1 basis points, and that of a 2-star restaurant would decrease by 17.9–29.3 basis points. For the two other quality classes, the effects are insignificant. This result confirms that Yelp helps the market weed out bad quality restaurants by pushing them to exit faster.¹⁴

¹⁴As a result of this function of Yelp, the average quality of the stock of restaurants in Texas has increased overtime, and the variance of these restaurants’ qualities has decreased overtime. The average quality of the stock of restaurants as measured by Yelp ratings increased by 0.16% per year or 3.4% over the entire sample period, and the variance

3.2.3 Yelp and Other Platforms

Having examined the effect of Yelp on restaurants listed on Yelp, I extend the analysis to other online platforms and their effects on restaurant revenues. This analysis serves two purposes: (1) one is to examine if the Yelp exposure measure has picked up the penetration of other platforms; by controlling for the penetration of other online review platforms and by comparing the effects of platforms' penetration on restaurants listed on v.s. off Yelp, we can test if Yelp's effects are restricted only to restaurants listed on Yelp. If that is indeed the case, we can conclude that the Yelp exposure measure used in the previous analysis is exogenous. (2) The other purpose of this analysis is to capture the effects of other online review platforms and see if they also have the opposite effects on high- and low-quality restaurants.

For this analysis, I use the sample of all restaurants that are listed on Google. As mentioned previously, Google covers over 9,000 restaurants and has a rating for every restaurant on it. This rating provides a uniform quality measure for restaurants listed across platforms. Recall in Figure 1, there are four types of restaurants according to which platform they are listed on: Type (1), those listed only on Yelp; Type (2), those listed on both Yelp and TripAdvisor; Type (3), those listed only on TripAdvisor; and Type (4), those listed neither on Yelp nor on TripAdvisor. There are at least 1,100 restaurants in each category, providing power for identifying the effects of various online platforms.

As mentioned earlier, to control for the penetration of other platforms, I use the Google Trends data on the search interest for the TripAdvisor website. This measure likely captures not only the penetration of TripAdvisor but also partly the penetration of Google's restaurant review service. Since Yelp was the leader that popularized the usage of online restaurant reviews and the other platforms imitated it, the penetrations of all online review platforms are likely correlated with each other in terms of both timing and geographic locations. In particular, in areas where Yelp is popular, consumers have been exposed to online reviews and may be more prone to using other online review platforms as well. Ideally, we want a direct measure of the penetration of Google's restaurant review service. However, this data is very difficult to obtain.¹⁵ Due to this difficulty, I use TripAdvisor exposure measures to capture partly Google's effects on restaurants.

The econometric specification for this analysis is as follows:

$$\begin{aligned} \log(Rev_{aljt}) = & \sum_{p=1}^4 D_j^p \left([\theta_y^p + \theta_{yr}^p R_j] \times \log(Y_{mt}) \times (1 - D_j^C) + [\theta_{yc}^p + \theta_{ycr}^p \times R_j] \times \log(Y_{mt}) \times D_j^C \right) \\ & \sum_{p=1}^4 D_j^p \left([\theta_g^p + \theta_{gr}^p R_j] \times \log(T_{mt}) \times (1 - D_j^C) + [\theta_{gc}^p + \theta_{gcr}^p \times R_j] \times \log(T_{mt}) \times D_j^C \right) \\ & + \mathbf{X}_{jt} \boldsymbol{\theta}_x + \theta_{nl} n_{mt}^I + \theta_{nc} n_{mt}^C + \theta_t + \theta_{tc} D_i^C + \theta_j + \xi_{jt} \end{aligned} \quad (2)$$

where D_j^p is a dummy for what type restaurant j is in terms of listing platforms, with $p = 1$ denoting Type 1, 2 Type 2, etc. $\log(T_{mt})$ is the logarithm of the Google Trends measure of search interest

decreased by 0.44% per year or 9.1% during the sample period.

¹⁵This is because Google Trends data do not track that particular segment of Google's service, and Google provides very limited amount of information through its API on the number of reviews that restaurants receive at a given time.

for the TripAdvisor website. Parameters θ_g^p , θ_{gr}^p , θ_{gc}^p and θ_{gcr}^p capture the effect of TripAdvisor on the revenues of Type p restaurants. Similarly, the parameters associated with the Yelp effect are also indexed with p . Note that this specification is very similar to regression 1, only that I allow Yelp and TripAdvisor to have different effects on each type of restaurant. I run the regression for only young restaurants that entered after Yelp’s penetration. The results are shown in Table 6.

Table 6: Effects of Exposure to Online Review Platforms on Restaurant Revenue

	Type (1) Log Revenue	Type (2) Log Revenue	Type (3) Log Revenue	Type (4) Log Revenue
Log Yelp (θ_y)	-0.912** (0.3850)	-0.425*** (0.1500)	-0.231 (0.2030)	-0.183 (0.1700)
Log Yelp×Rating (θ_{yr})	0.231** (0.0988)	0.124*** (0.0371)	0.0681 (0.0524)	0.0412 (0.0449)
Log Yelp×Chain Dummy (θ_{yc})	-0.0397 (0.0963)	-0.0655 (0.0876)	0.172 (0.1440)	0.335 (0.3580)
Log Yelp×Chain Dummy×Rating (θ_{ycr})	0.0127 (0.0264)	0.0284 (0.0243)	-0.0466 (0.0393)	-0.0986 (0.1020)
Log TripAd (θ_g)	0.421 (0.4970)	0.101 (0.1680)	0.0761 (0.1840)	-0.318** (0.1270)
Log TripAd×Rating (θ_{gr})	-0.107 (0.1280)	-0.0255 (0.0413)	-0.0073 (0.0449)	0.0825** (0.0322)
Log TripAd×Chain Dummy (θ_{gc})	-0.0293 (0.1120)	0.00693 (0.0576)	-0.387 (0.2770)	-0.282 (0.2680)
Log TripAd×Chain Dummy×Rating (θ_{gcr})	-0.00045 (0.0290)	-0.000861 (0.0156)	0.143* (0.0859)	0.077 (0.0715)
Controls	X			
Time FE	Yes			
Time FE×Chain Dummy	Yes			
Restaurant FE	Yes			
Sample Period	After Yelp’s penetration			
Sample Platform Presence	Yelp not TripAd	Yelp & TripAd	not Yelp but TripAd	neither Yelp nor TripAd
Rating	Google Rating			
N	239,104			

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Effect of Google

	2 Star	3 Star	4 Star	5 Star
Effect on Revenue	-0.153** (0.0754)	-0.0705 (0.0589)	0.012 (0.0226)	0.0945* (0.0556)

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The four columns of the table show the parameters associated with the effects of Yelp and TripAdvisor respectively for each type of restaurant. As can be seen from columns (1) and (2), even after controlling for the penetration of TripAdvisor, the parameters associated with Yelp’s effect is still significant for young independent restaurants, insignificant for chain restaurants. In particular, the estimates of Yelp’s effect on young independent restaurants that are listed only on Yelp (column 1) are much higher in magnitude compared to those that are listed both on

TripAdvisor and on Yelp (column 2), with the intercepts at -0.912 vs. -0.425 and slopes at 0.231 vs. 0.124. These parameters imply that Yelp’s effects are much stronger for restaurants that are listed only on Yelp.

Furthermore, Yelp’s effects are insignificant for restaurants not listed on Yelp, as shown in the first four rows of columns (3) and (4). These results imply that Yelp exposure affects only restaurants listed on Yelp, and it does not pick up the effects of other online review platforms once the exposure to TripAdvisor is controlled for.

For restaurants listed on Yelp, as illustrated in columns (1) and (2), the parameters associated with the effect of TripAdvisor are insignificant. This result suggests that the effect of Yelp overpowers the effects of TripAdvisor’s and other online review platforms’ on restaurants listed on Yelp. Given that Yelp is the dominant platform in the restaurant review service industry, it is very reasonable that there is very limited additional benefit to being listed on other platforms once a restaurant is already listed on Yelp.

For restaurants not listed on Yelp but on TripAdvisor, as shown in column (3), the parameters associated with TripAdvisor’s penetration are mostly insignificant, except for the interaction term between rating and chain, which suggests that the effect of TripAdvisor on chain restaurants’ revenue increases with quality. However, this parameter has a relatively low confidence level, only 90%, resulting in that the composite parameters that summarize TripAdvisor’s effect on chain restaurants’ revenues for a given quality level are insignificant for all quality levels. This weak effect of TripAdvisor on restaurants listed only on TripAdvisor is somewhat puzzling. It may be explained by the quality measure used for this analysis. Recall from Section 2.1 that TripAdvisor’s ratings are correlated with those from Yelp and Google only at about 58%, while the correlation between the other two rating schemes are around 72%. The quality measure used in this regression is the Google rating, which might have lumped together both high- and low-quality restaurants as rated by TripAdvisor into one quality class due to the low rating correlation. As a result, TripAdvisor’s penetration shows very little effect.

For restaurants listed neither on Yelp nor on TripAdvisor, but only on Google, as shown in column (4), the parameters associated with the effect of TripAdvisor are significant at the 95% confidence level for young independent restaurants, and insignificant for chain restaurants. The negative intercept parameter θ_g (-0.318) and positive slope parameter θ_{gr} (0.0825) suggest that TripAdvisor’s penetration has the opposite effect on high- and low-quality restaurants that are listed only on Google. TripAdvisor’s penetration having an effect on restaurants that are *not* listed on TripAdvisor but only on Google confirms that its penetration is correlated with Google’s. In this regard, TripAdvisor’s effects should be interpreted as Google’s effects.

I summarize Google’s effects by restaurant quality levels, and they are shown in Table 7. As can be seen, for a 5 star restaurant, doubling the exposure to Google restaurant review service increases its revenue by 9.5%, whereas for a 2 star restaurant, doubling the exposure would decrease its revenue by 15.3%. These contrasting effects on high- and low-quality restaurants indicate that Google also helps consumers learn about restaurant quality.

Although other online platforms affect restaurants’ revenues through learning, I cannot detect

that they have significant effects on the survival rates of restaurants not listed on Yelp. The linear regression on restaurants' exit in the same fashion as that done in Section 3.2.2 does not detect any significant effect of Yelp and TripAdvisor's penetration on the survival rates of restaurants not listed on Yelp. I thereby do not present those results on survival here.

To summarize, the empirical analysis in this section provides convincing evidence that online review platforms speed up consumers' learning about restaurant qualities. In particular, through learning, Yelp affects both restaurants' revenues and survival probabilities. Doubling Yelp exposure increases the revenue of a high-quality new independent restaurant by 18-27% and decreases that of a low-quality restaurant by about the same amount. Doubling Yelp exposure also raises the survival rate of a new high-quality independent restaurant by 18-29 basis points and reduces that of a low-quality restaurant by a similar level. In contrast, Yelp does not affect chain or old independent restaurants. Other platforms have a similar effect. In particular, if restaurants are listed on both Yelp and other platforms, Yelp's effects dominate. This confirms that Yelp is the most dominate online review platform in the restaurant industry.

3.3 Other Channels of Online Review Platforms' Effects Besides Learning

Although this paper focuses on the effect of online review platforms on restaurants through the channel of consumer learning, there are other mechanisms through which online review platforms can have an impact. Here I briefly discuss other potential channels and argue that they are not very important in the context of my analysis. Specifically, I discuss four main channels: managerial learning, consumer sorting, herding, and online reputation.

For managerial learning, I focus on the mechanism shown by Gin and Leslie (2003) that greater information exposure incentivizes restaurants to improve their quality. In the context of online review platforms, managers might be able to receive more feedback about restaurant quality by reading online reviews and thereby improve their services and food quality over time. If this is true, we should see new ratings that a restaurant receives per month would go up with the restaurant's age. I test this hypothesis by restaurants' rating history on Yelp. I regress the average of new ratings a restaurant receives per month over its age, and find that new ratings do not show increasing or decreasing trends. The details of this analysis are shown in Appendix B.3. Given this result, managerial learning is not a very important channel in the context of my study.

Consumer sorting is another channel through which online review platforms can affect revenues. Due to taste differentials among consumers, it is possible that consumers rank the same set of restaurants differently. Information displayed on online review platforms could help consumers with horizontal matching. Consumers could sort themselves into their own favourite restaurants over time, resulting in that only customers who love a restaurant would visit that restaurant. This type of sorting behavior should also cause a restaurant's rating to increase over time. However, as has been demonstrated previously, there is no increasing or decreasing trend in a restaurant's ratings. This evidence again rules out the channel of consumer sorting.

To investigate the channels of online reputation and herding, I argue that if online reputation and herding are important, then changes in online ratings should affect both young and old independent

restaurants. In theory, both online reputation and herding imply that rating changes on online review platforms would directly affect consumers’ preferences; that is, if a restaurant’s overall rating online becomes higher, more consumers would visit that restaurant regardless whether it is a young or old established restaurant. As shown in the reduced-form analysis, online review platforms do not affect old established independent restaurants. That evidence rules out the channels of reputation and herding.

In short, other channels of managerial learning, consumer sorting, reputation and herding do not seem to play an important role in the effects of online review platforms on restaurant revenues. The empirical evidence in this section points to consumer learning as the main mechanism.

4 Structural Demand Model With Consumer Learning

In this section, I develop a structural demand model with social learning. This structural model builds on the findings in the reduced-form section. The structural model incorporates (1) a Bayesian social learning model based on Ching (2010b) and (2) a constant expenditure model based on Björnerstedt and Verboven (2016). Both components lend themselves very well to modelling consumer choices over differentiated products by using aggregate product level revenue data.

4.1 Model

I assume that consumers’ decisions include two steps: first, consumers choose a restaurant, and then they decide how much to eat at a restaurant. This assumption of deciding the quantity to consume after having chosen a restaurant is the key difference between the constant expenditure demand model and the standard logit demand model, which assumes unit demand. As will be shown later, the constant expenditure model allows me to express consumers’ indirect utility in the form of revenues instead of quantities as are required by the logit demand model. This feature is convenient for using revenue data to model demand.¹⁶ Consumers choose which restaurant to go to in the first step based on the expected indirect utility they would receive at a restaurant. The expected indirect utility can be obtained by taking expectations over consumers’ *ex-post* indirect utility after having dined at a restaurant.

Consumers’ *ex-post* indirect utility is expressed in terms of income, demographic characteristics, and market and restaurant characteristics:

$$U_{ijt} = \mathbf{X}_{jt}\boldsymbol{\theta}_x + \alpha\gamma^{-1}\log(y_{it}) - \alpha\log(p_{jt}) + \tilde{A}_{ijt} + \xi_{jt} + e_{ijt} \quad (3)$$

where \mathbf{X}_{jt} includes restaurant characteristics that are observable to consumers initially and do not require learning as well as market characteristics that capture demand factors such as demographic attributes. $\log(y_{it})$ is the natural log of consumer i ’s income; γ is the budget share of consumer i ’s income spent on food.¹⁷ $\log(p_{jt})$ is the natural log of the average price of a meal at restaurant

¹⁶Although the price data are available for restaurants listed online, the prices are in broad ranges. Dividing revenues by these gross price measures does not give a very good approximation for quantity.

¹⁷The term $\alpha\gamma^{-1}\log(y_{it})$ will cancel out in consumers’ choice decisions.

j , and α is the price coefficient. \tilde{A}_{ijt} is consumer i 's experience from dining at restaurant j ; ξ_{jt} is an aggregate demand shock for restaurant j at time t , observed by consumers, but not by the econometrician; e_{ijt} is an idiosyncratic taste shock that is i.i.d. across consumers, restaurants and time; it follows the extreme value type I distribution. θ_x are coefficients associated with \mathbf{X}_{jt} .

There are two features worth noting in equation 3: (1) the constant expenditure model implies that consumers spend a constant proportion (γ) of their income (y_{it}) on food.¹⁸ (2) prices come into consumers' indirect utility function in the log form instead of the linear form as is in the standard logit demand function.

In equation 3, \tilde{A}_{ijt} represents the component of a restaurant's quality that is observable to consumers only after having eaten at a restaurant. \tilde{A}_{ijt} varies across consumers and time. It is a random draw from restaurant j 's quality distribution:

$$\tilde{A}_{ijt} = A_j + \delta_{ijt} \quad (4)$$

where A_j is the true mean quality of the restaurant j . δ_{ijt} is a random deviation from the mean; it is i.i.d. across consumers, restaurants and time and follows a normal distribution $\delta_{ijt} \sim N(0, \sigma_\delta^2)$.

I assume a varying quality instead of a constant quality for restaurants because (1) there is inherent variability in restaurant meals; even the same dish can taste different from time to time; (2) consumers' dining experience of the same restaurant can also vary across occasions; (3) a constant mean quality is supported by the reduced-form evidence that online reviews do not affect the revenues of old established independent restaurants; if restaurants' mean qualities change over time, then online ratings, which reflect restaurants' underlying quality, should affect the revenue of old established restaurants. This is found not to be the case in reduced-form analysis.

Here I also assume that the mean quality A_j is vertically differentiated and the same for all consumers. This assumption does not preclude horizontal differentiation between restaurants. Rather I assume that restaurants' characteristics that are horizontally differentiated are observable to consumers initially and do not require learning. The most typical characteristics of this type would be cuisine types; some people may prefer certain cuisine types that others dislike. These types of characteristics are included in \mathbf{X}_{jt} .¹⁹ Conditional on \mathbf{X}_{jt} , I assume that restaurants have a vertical differentiation in term of quality (A_j), which requires learning.

Consumers know the value of σ_δ^2 , but do not know the mean A_j , and therefore, they use experience signals A_{ijt} to learn about A_j . The experience signals can come from their own visits and others' experiences shared through online platforms or face-to-face communication. The learning process follows a Bayesian updating process: consumers have an initial prior of A_j , which follows a normal distribution with mean A_j^p and variance σ_A^2 . Consumers incorporate new experience signals to update their beliefs of A_j .

The belief updating process follows Bayes' rule. In particular, let $\eta_{jt+1} \equiv E[A_j|I(t+1)]$ denote the mean of the posterior distribution of A_j given consumers' information set at $t+1$, $I(t+1)$, and

¹⁸This implication is shown in the detailed derivation of the indirect utility function shown in Section D of the Online Appendix.

¹⁹In \mathbf{X}_{jt} , I include interactions terms between demographic characteristics and cuisine types to capture the fact that different demographic groups may prefer different cuisines.

σ_{jt+1}^2 denote the variance of the posterior distribution, Bayes' rule (DeGroot, 1970) gives

$$\eta_{jt+1} = (1 - \beta_{jt})\eta_{jt} + \beta_{jt}\bar{A}_{jt} \quad (5)$$

where η_{jt} is the prior mean and β_{jt} is

$$\beta_{jt} = \frac{\frac{1}{\sigma_{\delta}^2}n_{jt}}{\frac{1}{\sigma_{jt}^2} + \frac{1}{\sigma_{\delta}^2}n_{jt}} \quad (6)$$

in which n_{jt} is the number of experience signals that are released to consumers through their own trials or social media at time t ²⁰, and \bar{A}_{jt} is the sample mean of these experience signals. \bar{A}_{jt} follows a conditional normal distribution: $\bar{A}_{jt}|n_{jt} \sim N(A_j, \frac{\sigma_{\delta}^2}{n_{jt}})$. The precision (or the inverse of the variance) of the posterior is

$$\frac{1}{\sigma_{jt+1}^2} = \frac{1}{\sigma_{jt}^2} + \frac{n_{jt}}{\sigma_{\delta}^2} \quad (7)$$

where σ_{jt}^2 is the prior variance.²¹

As can be seen from equations 5 to 7, the posterior distribution of beliefs depends on \bar{A}_{jt} and n_{jt} . Although they are observed by consumers during the learning process, they are unobservable to me as the researcher. This unobservability necessitates more modeling assumptions: the number of signals n_{jt} is modeled to be proportional to consumers' exposure to online review platforms in a market and the number of consumers who dined at the restaurant in period t :

$$n_{jt} = (\kappa + \lambda_y \log(Y_{mt}) + \lambda_g \log(T_{mt})) \frac{Rev_{jt}}{\gamma \bar{y}_{mt}} \quad (8)$$

where $\frac{Rev_{jt}}{\gamma \bar{y}_{mt}}$ represents the number of customers that visited a restaurant in period t in the constant expenditure model. This representation is derived in equation 33 in Appendix C. Here Rev_{jt} is the revenue of restaurant j at time t , and \bar{y}_{mt} the average income of consumers in market m at time t . $\log(Y_{mt})$ and $\log(T_{mt})$ represent how much consumers are exposed to online review platforms in market m at time t . κ captures the portion of information that consumers receive from other sources such as their own trials, interactions with friends or other media platforms. λ_y and λ_g are the coefficients for the exposure to Yelp and TripAdvisor. Together $\lambda_y \log(Y_{mt}) + \lambda_g \log(T_{mt})$ represent the proportion of signals that come from online review platforms. Both κ and λ 's are assumed to be positive.

The belief updating process is intuitive. The posterior mean η_{jt+1} is simply a weighted average of the prior mean η_{jt} and the average of the new signals. This weight is β_{jt} , which depends on the precision of the prior ($\frac{1}{\sigma_{jt}^2}$) and that of new signals. If at time t , the prior η_{jt} is already precise,

²⁰Ideally, we want to separate consumers' own visits from the experience signals they learn from word-of-mouth. However, this is not possible given the limitation of the aggregate data at the product level.

²¹Note that in equations 5 to 7, I omit the notation i for two reasons: (1) A_j is vertically differentiated and therefore the same to everyone; (2) all consumers in a market are assumed to be exposed to the same set of experience signals. Given this assumption, the society as a whole updates their prior belief together. In this aspect, the learning process can be seen as the diffusion of knowledge on restaurants throughout the society.

then $\frac{1}{\sigma_{jt}^2}$ would be large, and the updating process would put more weight on the prior and less on the new signals \bar{A}_{jt} . If, however, the prior is not precise and $\frac{1}{\sigma_{jt}^2}$ is small, then the updating process would place a greater weight on \bar{A}_{jt} and less on the prior.

Typically as t gets large, the prior η_{jt} will become very close to the true mean A_j and very precise because it has incorporated all the experience signals that have been released to consumers up to period $t - 1$ ($\sum_{l=0}^{t-1} n_{jl}$). At this point, additional new signals would contribute little to consumers' belief update, and η_{jt+1} stays very close to η_{jt} and the true mean A_j . This feature of the model generates the result we see in the reduced-form analysis that online reviews do not affect the learning of old established restaurants.

The belief updating process also generates the phenomenon from the reduced-form analysis that online reviews have the opposite effects on higher- and lower-quality restaurants. Suppose that the initial prior mean A_j^p is very similar across restaurants, then over time, the posterior mean η_{jt+1} of higher-quality restaurants converges to a high number, and that of lower-quality restaurants converges to a low number. This pattern of convergence will show up in restaurants' revenues as an increasing or decreasing trend. The more online reviews there are, the faster the rate of convergence, and the steeper of the shape of the trend, and hence, the opposite effects of online reviews on higher- and lower-quality restaurants' revenues.

Consistent with the reduced-form evidence that online review platforms affect only independent restaurants and not chain restaurants, I assume that consumers learn only about independent restaurants' quality not chains'.²² Therefore, $\eta_{jt} = A_j, \forall t$ for chain restaurants. With this assumption, consumer i 's expected indirect utilities for independent and chain restaurants can be expressed in equations 9 and 10 respectively:

$$E[U_{ijt}|I(t)] = \mathbf{X}_{jt}\boldsymbol{\theta}_x + \alpha\gamma^{-1}\log(y_{it}) - \alpha\log(p_{jt}) + \eta_{jt} + \xi_{jt} + e_{ijt}, \forall j \in J_I \quad (9)$$

$$U_{ijt} = \mathbf{X}_{jt}\boldsymbol{\theta}_x + \alpha\gamma^{-1}\log(y_{it}) - \alpha\log(p_{jt}) + A_j + \xi_{jt} + e_{ijt}, \forall j \in J_C \quad (10)$$

where $I(t)$ denotes the information set at time t , and J_I denotes the set of independent restaurants in a market²³. J_C is the set of chain restaurants in a market. These equations are the result of taking expectation of the indirect utility in equation 3 over \tilde{A}_{ijt} conditional on the information set $I(t)$ at time t .

In additional to chain and independent restaurants, consumers also have an outside option in

²²I test the assumption of no-learning for chain restaurants by examining whether chain restaurants' revenues evolve systematically with age. If consumers learning about quality is important for chain restaurants, their revenues should increase or decrease with age after controlling for competition and demand factors in the market. I construct non-parametrically chain restaurants' revenue-age profile and find that it is very flat, a shape that is consistent with the assumption that most chain restaurants' quality is already known to consumers when they first open. Details are shown in Section G of the Online Appendix.

²³Here I assume that consumers are risk neutral with respect to the restaurant's underlying quality A_{ijt} . I could build in risk aversion by adding a quadratic term A_{ijt}^2 into the indirect utility function in equation 3, and both η_{jt}^2 and σ_{jt}^2 will show up in the expected indirect utility function 9. Including η_{jt}^2 in the indirect utility function implies a non-linear Gaussian structure of the estimating equation and that the likelihood function would no longer have a closed-form solution. The estimation method in that case would be simulated maximum likelihood with sequential filtering that numerically approximates the likelihood function. That method is very computationally burdensome. For this reason, I abstract away from risk aversion in the current estimation.

their choice set, which includes eating at fast food restaurants or at home. The indirect utility for the outside option is normalized to

$$U_{i0t} = \alpha\gamma^{-1}\log(y_{it}) - \theta_t + e_{i0t} \quad (11)$$

where θ_t represents changes in consumer taste towards eating out at full-service restaurants relative to the outside option over time. Note that the term $\alpha\gamma^{-1}\log(y_{it})$ shows up in the indirect utility in all restaurants in the choice set and therefore will cancel out in consumers' choice probability function.

Given the utilities associated with each option in their choice set, consumers choose the option with the highest expected utility.²⁴ A consumer's probability of choosing restaurant j is

$$s_{jt} = \frac{\exp(\Delta_{ijt} - \Delta_{i0t})}{1 + \sum_l \exp(\Delta_{ilt} - \Delta_{i0t})} \quad (12)$$

where Δ_{ijt} denote the relative mean utility from dining at a restaurant j at time t compared to the outside option, with $\Delta_{ijt} \equiv E[U_{ijt}|I(t)] - e_{ijt}$ for independent restaurants, $\Delta_{ijt} \equiv U_{ijt} - e_{ijt}$ for chain restaurants, and $\Delta_{i0t} \equiv U_{i0t} - e_{i0t}$. Here I omit the subscript i in s_{jt} because $\gamma^{-1}\log(y_{it})$ cancels out in this expression.

It is easy to show that under the constant expenditure model, the ratio of consumers' probabilities of choosing one restaurant over another is the ratio of the restaurants' revenues. This derivation is shown in Appendix C. In particular, we can take the log ratio of a restaurant's revenue over the revenue of the outside option to extract a restaurant's mean utility relative to that of the outside option:

$$\log\left(\frac{Rev_{jt}}{Rev_{0t}}\right) = \log\left(\frac{s_{jt}}{s_{0t}}\right) = \Delta_{jt} - \Delta_{0t} \quad (13)$$

If restaurant j is an independent restaurant, then

$$z_{jt} = \mathbf{X}_{jt}\boldsymbol{\theta}_x - \alpha\log(p_{jt}) + \eta_{jt} + \theta_t + \xi_{jt} \quad (14)$$

where $z_{jt} = \log\left(\frac{Rev_{jt}}{Rev_{0t}}\right)$. Equation 14 relates the revenue of independent restaurant j to consumers' mean utility relative to the outside option. It allows me to estimate the structural parameters by using the revenue data.²⁵

Given this main setup of the model, I make other arrangements to exploit the features of the data:

²⁴I assume that consumers are myopic in their decision making; that is, consumers do not experiment by actively trying out new restaurants today so that they can learn about restaurants' quality and make better decisions tomorrow. The myopic assumption is reasonable in the context of social learning: the gain from experimenting with new restaurants should be small if consumers can learn a lot from others' experiences. Also this assumption is consistent with the findings by Erdem and Keane (1996), who compare models with myopic and forward-looking behaviors in a experienced-good category and find that there is little difference between the two models in terms of the structural parameter estimates and model predictions.

²⁵This relationship is analogous to the log quantity ratio reflecting the mean utility of a product in the standard logit model.

1. First and foremost, I forecast restaurants' total revenues from alcohol drink sales. Given that my revenue data are partial revenues from only alcoholic drink sales (denoted by Rev_{aljt}), I must make some assumptions to approximate the total revenues using the partial revenues. I carry out this approximation through the following steps:
 - (a) I assume that the alcohol drink sales are proportional to a restaurant's total revenue with the relationship $\log(Rev_{jt}) = \log(g_j) + \log(Rev_{aljt})$, where g_j is the ratio of total revenue at a restaurant to its alcohol sales, and it is restaurant specific.²⁶
 - (b) Without observing the actual total revenues at each restaurant, I use the number of Yelp reviews that a restaurant receives each month as a proxy for the number of consumers who patronized a restaurant in that month. As mentioned previously,²⁷ the number of consumers who visited a restaurant is proportional to the restaurant's revenues subject to the average income in a market. This relationship allows me to connect the number of Yelp reviews to restaurant total revenues. Then by connecting the number of Yelp reviews to restaurants' alcohol sales, I estimate g_j through restaurants' fixed effects.
 - (c) Since not all restaurants are listed on Yelp, I use restaurants' characteristics, especially their cuisine types, to predict g_j for all other restaurants. The econometric details regarding the estimation of g_j for Yelp listed restaurants and the imputation of g_j for other restaurants are shown in Section B of the online Appendix.
 - (d) Overall, the estimated g_j shows that the average percentage of alcohol sales out of the total revenue differ across cuisine types.²⁸ For breakfast restaurants and cafés, the average percentage of alcohol sales is about 8%; for Asian restaurants, 6%; for bars, 38%; for Tex-Mex, 19%; for European and Mediterranean cuisines, 22%. These percentages are consistent with industry standards for each type of restaurants.²⁹
 - (e) Once I have the estimated and imputed g_j in hand, I calculate restaurants' total revenues based on their alcohol sales and g_j .
2. Second, with the approximated total revenues in hand, I construct the revenues of the outside option (Rev_{0t}) by subtracting all restaurants' revenues in a market from the total market size. As shown in Appendix C, the market size is the total amount of income spent on food per month for all consumers in a market. I compute this number by multiplying the sum of consumers' total income in a market by the fraction of their income spent on food. The fraction is provided by the Consumer Expenditure Survey held by the U.S. Bureau of Labor Statistics.³⁰

²⁶ As mentioned previously, empirical support for this assumption can be found in Section B of the Online Appendix.

²⁷ Also shown in Appendix C.

²⁸ Using restaurant cuisine types described on Yelp and TripAdvisor and judging by restaurant names, I classify all restaurants in my dataset into five main cuisine categories: (1) breakfast and café, (2) Asian, (3) bars, (4) Tex-Mex, and (5) European. A detailed description of these cuisine categories is shown in Table 19 in Appendix G.

²⁹ Note that I excluded bars that focus mainly on night life because they serve different clientele compared to full-service restaurants.

³⁰ Forecasting the total revenue from alcohol drinks based on a constant ratio g_j does not affect the estimation of

3. Third, I estimate the demand equations only for young independent restaurants that have online presence during the sample periods. The demand equation is shown in equation 14. I make this decision because learning for chain restaurants and old established restaurants is not affected by online review platforms, and their mean utilities can always be expressed in terms of revenue ratios. These mean utilities remain unchanged in the counterfactual exercise where the online review platforms are removed. For young independent restaurants that did not have online presence, I treat them the same way. Reducing the estimating sample to only those restaurants that are affected by online review platforms cuts down the estimation burden and does not influence the counterfactual analysis. The sample for this estimation includes 3,995 restaurants.
4. Fourth, I model the initial prior mean of a restaurant’s quality as a linear function of the quality measure R_j : $A_j^p = \theta_r^p R_j$. For the structural analysis, I use Google ratings as the quality measure as it covers the largest number of restaurants. For restaurants that are not listed on Google but on Yelp or on TripAdvisor (including 463 new independent restaurants), I predict their Google ratings based on the correlation between the ratings on these three platforms. Modeling the initial prior mean this way also controls for unobserved time-invariant characteristics of a restaurant that are correlated with the Google rating but do not require learning.
5. Fifth, instead of estimating the true mean quality for each individual restaurant, A_j , I model A_j also as a linear function of the quality measure: $A_j = \theta_0 + \theta_r R_j$. This formulation keeps the structural model parsimonious. Note that the linear function of A_j includes an intercept, but that of A_j^p does not. This is because only one intercept can be identified in addition to the restaurant cuisine and group fixed effects, which are included in \mathbf{X}_{jt} .
6. Sixth, the aggregate demand shock ξ_{jt} is assumed to be serially correlated with an autocorrelation ρ : $\xi_{jt} = \rho \xi_{jt-1} + \varepsilon_{jt}, \varepsilon_{jt} \sim i.i.d.N(0, \sigma_\varepsilon^2)$ This assumption is to capture any unobserved demand shocks that may be persistent. Furthermore, since both consumers’ expectations of quality η_{jt} and ξ_{jt} are unobserved, and η_{jt} is serially dependent, without controlling for serial correlation in ξ_{jt} , η_{jt} may pick up the additional serial correlation in ξ_{jt} and thereby overestimate the importance of learning.
7. Seventh, I use a fifth-order polynomial of the calendar time to approximate the time trend θ_t instead of using a time dummy for each month. This functional form also keeps the structural model parsimonious.
8. Last, restaurants listed on various platforms are given their own learning parameters; that is, I expand the κ and λ parameters to accommodate the different learning speed for the four types of restaurants listed on various platforms: (1) those listed only on Yelp; (2) those listed on both Yelp and TripAdvisor; (3) those listed only on TripAdvisor; (4) those listed on

the main structural parameters, especially the learning parameters, because the identification of these parameters come mostly from the change in restaurants’ revenues with respect to age. The ratio g_j is equivalent to a restaurant fixed effect in the log revenue equation. However, g_j is important for the counterfactual exercise, where we need the estimates of restaurants’ mean utilities to simulate consumers’ choices in the counterfactual world.

neither Yelp nor TripAdvisor, but on Google.

With these modeling arrangements, the estimating equation can be written as

$$\begin{aligned}
z_{jt} &= \mathbf{X}_{jt}\boldsymbol{\theta}_x - \alpha \log(p_{jt}) + \eta_{jt} + \sum_{k=1}^5 \theta_k \tau^k + \xi_{jt} \\
\xi_{jt} &= \rho \xi_{jt-1} + \varepsilon_{jt}, \varepsilon_{jt} \sim i.i.d. N(0, \sigma_\varepsilon^2)
\end{aligned} \tag{15}$$

where τ is the calendar time in month; θ_k is simply the coefficient associated with the k-th order polynomial of τ , and

$$\eta_{jt} = (1 - \beta_{jt-1})\eta_{jt-1} + \beta_{jt-1}A_j + \nu_{jt} \tag{16}$$

$$\eta_{j0} = A_j^p = \theta_r^p R_j \tag{17}$$

$$A_j = \theta_0 + \theta_r R_j \tag{18}$$

$$\nu_{jt} = \beta_{jt-1}(\bar{A}_{jt-1} - A_j) \tag{19}$$

$$\beta_{jt-1} = \frac{\frac{1}{\sigma_\delta^2} n_{jt-1}}{\frac{1}{\sigma_{jt-1}^2} + \frac{1}{\sigma_\delta^2} n_{jt-1}} \tag{20}$$

$$n_{jt-1} = \left(\sum_{a=1}^4 D_j^a (\kappa^a + \lambda_y^a \log(Y_{mt-1}) + \lambda_g^a \log(T_{mt-1})) \right) \frac{Rev_{jt-1}}{\gamma \bar{y}_{mt-1}} \tag{21}$$

$$\nu_{jt} \sim N\left(0, \frac{\beta_{jt-1}^2 \sigma_\delta^2}{n_{jt-1}}\right) \tag{22}$$

where a denotes the platform type of restaurant j as described in point 6 above; D_j^a is a dummy which equals 1 if restaurant j belongs to type a , 0 otherwise.

5 Structural Estimation and Results

This section discusses the identification and structural estimation method. The identification of most structural parameters in equation 15 is straightforward except for the learning parameters. One important issue is how we differentiate η_{jt} , learning, from ξ_{jt} since both are serially correlated. The key source of identification between the two elements comes from the fact that the serial correlation in η_{jt} depends on age, while that in ξ_{jt} does not. In particular, once a restaurant's quality has been learned, η_{jt} stays constant, whereas ξ_{jt} still follows an AR(1) process with an autocorrelation ρ .

Most of the learning parameters are easily shown to be identifiable, except for σ_δ^2 , which cannot be separately identified from the κ and λ parameters, and is normalized to 1. A detailed discussion regarding the identification of each learning parameter is provided in Appendix D. In essence, the initial prior A_j^p is identified by a restaurant's initial revenue when it just opens; the true mean quality A_j is identified by the restaurant's later revenue when the revenue stabilizes.³¹ How fast

³¹These revenues are residual revenues after all other demand factors and competitive effects in the market have been accounted for.

restaurants' revenues converge to their stable level identifies the κ and λ parameters as well as the initial prior variance σ_A^2 .

To estimate equation 15, I use a quasi-maximum likelihood approach. The key estimation issue is that η_{jt} is unobservable to the econometrician, and needs to be integrated out. Because of the Gaussian structure in both η_{jt} and ξ_{jt} , the integrated likelihood function has a closed-form solution. I derive its functional form in Appendix E. To the best of my knowledge, this paper is the first that accounts for the serial correlation in ξ_{jt} in social learning models.

As in all demand models, an important endogeneity concern is the restaurant price. To correct for endogeneity, I use the standard BLP instruments as IVs for prices (Berry et al. (1995)). Specifically, the set of IVs includes the average ratings of the independent and chain competitors and the total number of competitors in a market. The assumption associated with these IVs is that the entry of the competing restaurants is exogenous in the quality dimension. That is, at the time of entry, competing restaurants choose a quality level that is unrelated to η_{jt} and the demand shock ξ_{jt} of restaurant j . Since entry is sticky, once entered, competing restaurants' mean qualities are fixed, and thereby uncorrelated with future demand shocks for restaurant j , but they affect the prices that restaurant j can charge. Under this assumption, these instruments are valid. I carry out this IV approach in two steps. In the first step, I regress restaurant prices on these instruments and obtain a fitted value. The estimating equation in the first step is shown below:

$$\log(p_{jt}) = \theta_{pI}AR_{jt}^I + \theta_{pC}AR_{jt}^{ch} + \theta_r n_{rjt} + \theta_m + \tilde{\epsilon}_{jt} \quad (23)$$

where AR_{jt}^I and AR_{jt}^{ch} are the average ratings of restaurant j 's independent and chain rivals respectively at time t ; n_{rjt} is the total number of competitors, and θ_m is a zipcode fixed effect to capture persistent market characteristics that affect restaurant prices.

In the second step, I plug the fitted log price into equation 15 and conduct the estimation using the quasi-maximum likelihood approach. To adjust the standard errors from the second step to incorporate the errors from the first step estimation, I implement the method introduced in Murphy and Topel (1985)³².

5.1 Structural Estimates

The estimates of the key structural parameters are shown in Table 8. These estimates include the learning parameters and the log price coefficient; the estimates of control variables can be found in Section F of the Online Appendix. As shown in the table, the log price coefficient α is 6.74 and significant. It can be used to construct the price elasticity of demand, which is $1 + \alpha(1 - s_{jt})$ in the constant expenditure model. Since s_{jt} is very small for most restaurants (about 95% of restaurants in the sample have a revenue share of less than 0.071, and the mean of all revenue shares is 0.019), the price elasticity of demand for a restaurant can be approximated by $(1 + \alpha)$, which equals 7.74. This elasticity implies that if a restaurant's price declines by 10%, its demand will increase by 77.4%. This large elasticity of demand is consistent with the common observation that the restaurant industry is highly competitive.

Table 8: Structural Estimates: κ and λ 's

α	6.74*** (0.563)			
σ_A^2	0.399*** (0.0425)			
θ_r^p	0.0828*** (0.00204)			
θ_0	-2.86*** (0.101)			
θ_r	0.817*** (0.0254)			
σ_ε^2	0.638*** (0.00223)			
ρ	0.572*** (0.00230)			
	On Yelp not on TripAd	On Yelp and TripAd	Not on Yelp but on TripAd	Not on Yelp nor on TripAd
κ	2.54e-8 (5.19e-6)	1.74e-5* (9.40e-6)	1.53e-5* (8.42e-6)	2.16e-9 (3.98e-7)
λ_y	1.8e-2*** (6.43e-3)	1.07e-5** (5.53e-6)	2.68e-9 (8.02e-8)	1.9e-5 (1.49e-5)
λ_g	4.95e-5 (5.70e-5)	2.56e-9 (1.74e-7)	6.84e-7 (4.97e-6)	2.75e-4* (1.05e-4)
N	179,834			

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Table 9: Distribution of Information Source

Type	Yelp	Google/TripAdvisor	Other Info Source	No. of Restaurants
(1) On Yelp not on TripAdvisor	99.7%	0.30%	0.0%	987
(2) On Yelp and TripAdvisor	63.0%	0.02%	37.0%	1,555
(3) Not on Yelp but on TripAdvisor	0.0%	12.1%	87.9%	569
(4) Not on Yelp nor on TripAdvisor	5.4%	94.6%	0.0%	884

The bottom four rows of the table display the estimates of the key learning parameters κ and λ s for each group. As can be seen, the κ coefficient is significant only for the second and third groups, including all restaurants listed on TripAdvisor. The Yelp exposure coefficient λ_y is significant only for the first and second groups, including all restaurants listed on Yelp. The TripAdvisor exposure coefficient λ_g is significant only for the fourth group, restaurants listed on Google but not on the other two platforms. The estimates of λ_y and λ_g are consistent with the results shown in Table 6 from the reduced-form analysis: Yelp affects only restaurants listed on Yelp, and the TripAdvisor exposure measure picks up the effect of Google.

The estimates of κ and λ 's can tell us what information source is the most important in the consumer learning process. The ratio $\frac{\lambda_y \log(Y_{mt})}{(\kappa + \lambda_y \log(Y_{mt}) + \lambda_g \log(T_{mt}))}$ can be interpreted as what percentage of information that consumers use to learn about restaurants comes from Yelp. Similarly, $\frac{\lambda_y \log(T_{mt})}{(\kappa + \lambda_y \log(Y_{mt}) + \lambda_g \log(T_{mt}))}$ and $\frac{\kappa}{(\kappa + \lambda_y \log(Y_{mt}) + \lambda_g \log(T_{mt}))}$ capture how much information comes from other platforms and from other information sources respectively. Plugging in the average measures for Yelp and TripAdvisor exposure, we can calculate these percentages for each group. Table 9 presents the result.

As can be seen, for Type (1) restaurants, almost all information for learning comes from Yelp. For Type (2), about 63% of the information comes from Yelp and 37% from other information sources. For Type (3), about 88% of the information comes from other information sources; online review platforms account for only about 12%. For Type (4), almost all information comes from Google, about 95%; Yelp seems to account for some information as well but at a very small level.

Table 9, the table for key structural estimates, also illustrates the coefficients related to consumers' prior beliefs and restaurants' true qualities, including the slope of the initial prior with respect to ratings (θ_r^p) and the intercept and slope of the true mean quality (θ_0 and θ_r respectively). These coefficients implies how much learning changes consumers' expectation of a restaurant's quality. The slope of the initial prior θ_r^p is 0.0828, positive and significant, but small. It indicates that consumers' initial prior belief of quality is almost the same for all restaurants. The slope of the true mean quality θ_r is 0.817, positive, significant and large. It suggests that restaurants' true mean quality increases with Google ratings. The intercept θ_0 is -2.86 , also significant and large, and has a negative sign. The negative intercept indicates that when a restaurant's rating is very low, its true mean quality is much lower than the initial prior.

The most informative form of these parameters is the difference between the true mean quality and the prior, $\theta_0 + (\theta_r - \theta_r^p) * R_j$. It tells us how much consumers have updated their beliefs during the learning process, and whether the posterior true mean quality is higher or lower than the prior for each quality class. These differences are shown in Table 10. As can be seen, for 2 and 3 star restaurants, consumers lower their beliefs during the learning process, and for 4 and 5 star restaurants, consumers raise their expectations. In particular, the changes in consumers' beliefs are large for restaurants that are of either very low or very high quality (2 or 5 stars), whereas those changes are moderate for restaurants of medium quality (3-4 stars). The gap between the initial prior and the true mean quality is asymmetric across quality ratings, with 2 star restaurants

³²Specifically, page 374, equation 15.

experiencing the highest absolute change in consumers’ beliefs, a 1.389 drop, while 5 star restaurants encountering a more moderate increase of 0.812. This asymmetric pattern can be explained by that the average of Google ratings is about 3.93 (as shown in Table 1); therefore, the quality gap between a 2 star restaurant and an average restaurant is much bigger than that between a 5 star restaurant and an average restaurant (1.93 vs. 1.07 stars).

Table 10: Gap Between True Mean Quality and Prior

	2 Star	3 Star	4 Star	5 Star
$\theta_0 + (\theta_r - \theta_r^p) * R_j$	-1.389	-0.656	0.078	0.812

The gap between the true mean quality and initial priors accounts for a large proportion of a restaurant’s true mean utility. As shown in Figure 5 in Appendix A, a histogram of restaurants’ true mean utility, almost all restaurants lie between -10 and 0. The average of the true mean utility is about -4.3. This distribution implies that a drop of 1 in consumers’ expectation of quality during learning can lead to a 23% change in an average restaurant’s true mean utility, and thereby lead to a substantial change in the restaurant’s revenues. As will be demonstrated in the counterfactual analysis, restaurants do experience a large change in their revenue shares during the learning process.

5.2 Counterfactual Analysis

Having estimated the structural parameters, I conduct counterfactual analyses to evaluate the effect of online review platforms on consumer welfare, restaurants’ revenue shares, and the total revenue of the industry. I evaluate three counterfactuals: (1) a “no-online-review” counterfactual, where I remove the existence of online review platforms by setting all the λ parameters to 0, and consumers rely only on other information sources to learn about restaurant quality. (2) a “no-learning” counterfactual, where I shut down learning completely, and consumers make choices only based on their prior beliefs. (3) a “full-information” counterfactual, where I assume that learning is immediate, and consumers know the true mean qualities of restaurants as soon as they open. In all three counterfactual exercises, I simulate consumer choices in the counterfactual worlds and compare them to those in the real world. The resulting change in consumer welfare, restaurant revenue shares and the total industry revenue gives me the effect of online review platforms (in the no-online-review case) and the general effect of information (in the no-learning and full-information case). In particular, comparing the results from the “no-online-review” counterfactual and those from the “no-learning” one gives us the importance of other information sources.

In all of the counterfactual analyses, I hold the supply side constant: I assume that restaurants’ prices do not change, and their entry and exit patterns stay the same with or without information. The assumption of constant prices is innocuous. As mentioned before, the price elasticity of demand is big in this industry, and restaurants do not have a very strong pricing power. Furthermore, it is easy to show that restaurants’ optimal prices are set at $p_{jt} = c_{jt}(1 + \frac{1}{\alpha(1-s_{jt})})$, where c_{jt} is the marginal cost at restaurant j at time t . Since the revenue share s_{jt} is very small for most

restaurants and α is large, the change in optimal prices from the real world to the counterfactual worlds is negligible.

The assumption of exogenous entry and exit is somewhat restrictive. As shown in the reduced-form analysis, online review platforms affect restaurants' survival rates and help the market weed out bad quality restaurants. Without accounting for endogenous entry and exit of restaurants, the estimates of consumer welfare and industry revenues are likely underestimated. This is because without online review platforms, bad quality restaurants would survive longer, and consumers are more likely to go to the bad quality restaurants. To account for endogenous entry and exit accurately, we need to combine the learning process with a model of dynamic entry and exit, which is a relatively large undertaking and is currently being explored in a separate paper. In this paper, I treat entry and exit as exogenous, and we can see the welfare effect as a lower bound.

Compensating Variation: To translate the welfare effect of online review platforms into dollar values, I use the concept of compensating variation. In particular, I use the metric of compensating variation per restaurant-goer. This metric represents how much money should be given to restaurant goers such that their utilities in the counterfactual world would be brought up to the same level as those in the real world. The derivation and rationale behind this metric is shown below.

In the real world, consumer welfare can be written as consumers' expected utility from all choices of food:

$$\begin{aligned} CS_{it} &= \frac{\gamma}{\alpha} \left[\log \left(\exp(\Delta_{0mt}) + \sum_{j \in J_{mt}} \exp(\Delta_{jt}) \right) + \sum_{j \in J_{mt}} s_{jt}(\Delta_{jt}^T - \Delta_{jt}) \right] + C \\ &= \frac{\gamma}{\alpha} \left[\log \left(1 + \sum_{j \in J_{mt}} \exp(\Delta_{jt} - \Delta_{0mt}) \right) + \sum_{j \in J_{mt}} s_{jt}(\Delta_{jt}^T - \Delta_{jt}) \right] + C + \log(y_{it}) - \frac{\gamma}{\alpha} \theta_t \quad (24) \end{aligned}$$

where Δ_{jt}^T is the actual mean utility that consumers receive when visiting restaurant j at time t , which is equivalent to replacing η_{jt} with the true mean utility A_j ; J_{mt} denotes the set of restaurants in market m at time t ³³; C is a constant.

This formula is derived based on the concept of inclusive values, which is a measure of consumer welfare (Anderson et al. (1992), Small and Rosen (1981)). The welfare number here is normalized by $\frac{\gamma}{\alpha}$, such that the unit of the welfare is in terms of \log income. Note that in demand models with complete information, the consumer welfare formula do not include the $\frac{\gamma}{\alpha} \sum_{j \in J_{mt}} s_{jt}(\Delta_{jt}^T - \Delta_{jt})$ component. This component is added in the welfare formula in the model with consumer learning because when consumers choose a restaurant to go to, the choices are based on their expectations of a restaurant's quality η_{jt} , not the actual experience they will receive, which is A_j . In this sense,

³³ J_{mt} includes all full-service restaurants in my data, not just the young independent restaurants that were listed online. Other restaurants, such as chains, old established independent restaurants, and young independent restaurants not listed online, are also included in this calculation. For these restaurants, their true mean utility Δ_{jt}^T is assumed to be the same as Δ_{jt} . This assumption is innocuous for both chain and old established restaurants because their qualities were already known to consumers during the age of online review platforms. For young independent restaurants not listed online, this assumption is more restrictive. However, given that very little is known about their characteristics and that they do not account for a large proportion of restaurants in each market, this assumption is mostly harmless.

consumers' choices in a model with learning are not optimal. As consumers learn about restaurants' quality, their expectations get closer to the true mean quality of a restaurant, and their choices get closer to being optimal. That is how greater availability of information improves consumer welfare in the context of consumer learning. The derivation of this formula can be found in Section E of the Online Appendix.

Similarly in the counterfactual world, the welfare formula is

$$CS_{it}^c = \frac{\gamma}{\alpha} \left[\log \left(1 + \sum_{j \in J_{mt}} \exp(\Delta_{jt}^c - \Delta_{0mt}) \right) + \sum_{j \in J_{mt}} s_{jt}^c (\Delta_{jt}^T - \Delta_{jt}^c) \right] + C + \log(y_{it}) - \frac{\gamma}{\alpha} \theta_t \quad (25)$$

where the superscript c denotes "counterfactual."

To bring consumers' welfare in the counterfactual world to be the same as the one in the real world, we can increase consumers' income y_{it} in the counterfactual world by a multiplier ψ_{it} . That is

$$CS_{it} = CS_{it}^c - \log(y_{it}) + \log((1 + \psi)y_{it}) = CS_{it}^c + \log(1 + \psi_{it}) \quad (26)$$

ψ_{it} is in general very small in this type of model because it is applied to consumers' entire income, not just budget on food. For a small ψ_{it} , we can write

$$\begin{aligned} \psi_{it} &\approx CS_{it} - CS_{it}^c = \frac{B_{mt}\gamma}{\alpha}, \text{ where} \\ B_{mt} &= \log(Rev_{0mt}^c) - \log(Rev_{0mt}) + \sum_{j \in J_{mt}} s_{jt} (\Delta_{jt}^T - \Delta_{jt}) - \sum_{j \in J_{mt}} s_{jt}^c (\Delta_{jt}^T - \Delta_{jt}^c) \end{aligned} \quad (27)$$

The above equation 27 shows that ψ_{it} is market specific and could be replaced with ψ_{mt} . To translate ψ_{mt} into a monetary amount, I multiply it by the total income of consumers in a market and then divide it by the number of consumers who go to those restaurants listed online:

$$CV_{mt} = \frac{\psi_{mt} Y_{mt}}{S_{mt} M_{mt}} \quad (28)$$

where Y_{mt} is the total income of all consumers in a market. M_{mt} is the total population in the market, $S_{mt} = \sum_{j \in JO_{mt}} s_{jt}$ is the total proportion of consumers who patronize the set of young independent restaurants listed online in that market. JO_{mt} denotes the set of young independent restaurants listed online.

The metric CV_{mt} has the interpretation that the welfare gain is equivalent to giving each consumer CV_{mt} dollars every time he or she eats at a young independent restaurant listed online. In particular, in equation 28, the numerator $\psi_{mt} Y_{mt}$ represents how much money should be given to the entire population in a market in order to bring their utility in the counterfactual world to the same level as the one in the real world. Dividing this by the number of population will give us a measure of the welfare benefit per capita. Here I choose the subset of consumers who actually eat at those young independent restaurants listed online instead of the entire population in the market (e.g. M_{mt}). This is because the market size definition in my study is relatively big – all food choices in a person's monthly consumption. While this market size definition has the benefit

of being more stable compared to a definition of all eating-out options,³⁴ it also makes the revenue shares of full-service restaurants relatively small. In particular, the aggregate revenue shares of all young independent restaurants listed online, the group that are generating the welfare gain, are even smaller (with a mean of 0.063).³⁵ Given that at a given time, only a very small portion of the consumers actually eat at these young independent restaurants listed online, the welfare gain generated by these restaurants should be divided by the set of consumers who go to these restaurants. The resulting metric then captures the welfare gain per restaurant-goer.

In addition to this metric, I also calculated the total welfare gain for the entire state of Texas over the sample period using the formula $\sum_{m,t} \psi_{mt} Y_{mt}$, and dividing this total by all restaurant goers in each market over time ($\sum_{m,t} S_{mt} M_{mt}$) gives the overall average of welfare gain per restaurant-goer. The technical details regarding how to obtain these welfare numbers are shown in Appendix F. The subsections below present the results.

5.2.1 Welfare Effect of Online Review Platforms

The results for the welfare analysis are shown in Table 11 below. For the no-online-review counterfactual, the estimated compensating variation CV_{mt} for each market ranges from \$0 to about \$303, averaging at \$5.31. The total welfare gain for Texas amounts to \$110.4 million. The overall state average per restaurant-goer is about \$2.5. Given that the average meal price in Texas is about \$20, this welfare gain is equivalent to a 12.6% discount on each meal.

For the no-learning case, CV_{mt} for each market has the same min and max as in the no-online-review scenario; however, the mean is higher by about 44 cents or 8%. The total welfare for the state is also larger by about \$31.8 million, and the corresponding overall average welfare per restaurant-goer is bigger by 76 cents or 30%. These differences in the welfare effect between the no-online-review and no-learning scenarios illustrate that without accounting for other information sources, consumers' welfare gain can be overestimated by as high as 30%. For the full-information counterfactual, the welfare gain per restaurant-goer for each market ranges from \$0.014 to \$240.4, with an average of \$6.76. The total welfare for the state is about \$94.6 million, and the overall average per restaurant goer is about \$2.16, representing an 11% discount on an average restaurant meal. These results imply that as we move from the real world to a world with full information, there is still a substantial amount of additional welfare to be attained.

5.2.2 Effect on Revenue Shares and Total Industry Revenue

Having examined the effect of online review platforms (or information in general) on consumer welfare, I explore the effects on restaurants' revenue shares and industry revenues in this section.

³⁴The market size definition of all eating-out options, including full-service restaurants, restaurants that do not serve alcohol and fast food places, is likely endogenous. The sheer availability of many restaurants in a market can increase consumers' budget on eating out, and thereby increasing the market size.

³⁵These aggregate revenue shares were particularly small at the beginning of online review platforms' penetration. As online reviews grew more popular, the aggregate shares grew bigger. Nonetheless, on average, the shares of young independent restaurants listed online were small.

Table 11: Welfare Analysis Results

	No-Online-Review	No-Learning	Full-Information
Min CV_{mt} (\$)	0.00	0.00	0.014
Mean CV_{mt} (\$)	5.31	5.75	6.76
Max CV_{mt} (\$)	302.7	302.7	240.4
Total Welfare for Texas (\$ million)	110.4	142.2	94.6
Overall Average Welfare per Restaurant-Goer per Meal (\$)	2.51	3.27	2.16
As % of Average Restaurant Meal Price	12.6%	16.5%	10.9%

All monetary values are in October 2018 dollars.

Table 12 summarizes the average change in restaurants' revenue shares by quality level for each counterfactual, and Table 13 illustrates the changes in total industry revenues.

As shown in Table 12, the patterns of average changes in revenue shares by quality for both the no-online-review and no-learning counterfactuals are very similar to those shown in Table 4 in the reduced-form analysis. For the no-online-review counterfactual, the average change in revenue share for lower quality restaurants is negative, with a decline of 20.7% for 2 star restaurants and 15.4% for 3 star restaurants. The effect on higher quality restaurants is the opposite: although the change for 4 star restaurants is very small, the change for 5 star restaurants is a substantial gain of 15.7%. In the case of no learning, the effects for each quality level is more exaggerated compared to those in the no-online-review case. This is because the no-learning counterfactual captures the effect of other information sources. Nonetheless, the difference between the average changes in revenue shares between the no-online-review counterfactual and the no-learning one is not very big. This small difference implies that online review platforms are the main information source for restaurants.

For the full information counterfactual, the analysis shows that as we move from the real world to a world with full information on quality, lower quality restaurants would suffer an additional 26% to 51% drop in revenue shares, and higher quality restaurants would gain as much as another 43%. These large changes in revenue shares imply that restaurants' revenue shares are very sensitive to information and that there is still a lot more effect of information that could be exploited.

Table 12: Mean Percentage Change in Revenue Share of a Restaurant

Quality	No-Online-Review	No-Learning	Full-Information
2 stars	-20.7%	-21.9%	-51.4%
3 stars	-15.4%	-16.4%	-25.6%
4 stars	0.22%	2.31%	4.94%
5 stars	15.7%	17.1%	43.4%

While the analysis on revenue shares shows that the effect of information has a substantial impact on individual restaurants, in aggregate, information does not seem to affect much the restaurant industry as whole. This result can be gleaned from Table 13, which summarizes the absolute and percentage changes in total restaurant revenues for three groups of restaurants, including (1) all the young independent restaurants listed online; (2) all other full-service restaurants, such as chain and old established independent restaurants; (3) the sum of the first two groups, i.e.

all full-service restaurants as a whole. Breaking down the restaurants into these groups helps us understand better the substitution patterns of demand between the young independent restaurants listed online and those restaurants that are not.

There are two main features of Table 13 worth noting. The first one is that in terms of percentage changes, the effect of information has very little effect on any group of restaurants in almost all counterfactuals. These percentage changes are mostly less than 1%, except for the 3.6% increase for the young independent restaurant group listed online in the full-information case. The reason for such a small change in the industry total revenue may be attributed to that there are both high- and low-quality restaurants in the group of young independent restaurants listed online. While information benefits the high-quality restaurants, it harms the low-quality ones, with the net effect on the entire group and the industry being close to 0. Nonetheless, the 3.6% increase in the total revenues for young independent restaurants in the full-information case is non-negligible. This 3.6% or \$1.28 billion increase suggests that more information (compared to what we have in the real world) will benefit higher-quality restaurants more than lower-quality restaurants, leading to a net increase in revenue for the young independent restaurant group as a whole.

The second notable feature of Table 13 is that the change in total revenue of young independent restaurants listed online gets mostly passed through to the entire industry of full-service restaurants. For example, for the no-online-review counterfactual, almost 70% of the change in group revenue (\$-227.6 million) is passed through the entire industry (\$-158.0 million). This high pass-through rate implies that the main substitution pattern occurs mostly within the young independent restaurants listed online, not between this group and other full-service restaurants. For the no-online-review counterfactual especially, the percentage change in the revenues of other full-service restaurants is only 0.097%, very close to 0. This negligible change explains why the revenues of chain and old established independent restaurants are not affected by online review platforms, one of the main findings in the reduced-form analysis.

Table 13: Change in Total Industry Revenue

	No-Online-Review	No-Learning	Full-Information
Young Independent Restaurants Listed Online (\$ million)	-227.6	164.1	1,284.6
% Change	-0.63%	0.46%	3.56%
Other Full-Service Restaurants (\$ million)	71.7	-87.3	-416.7
% Change	0.097%	-0.12%	-0.56%
All Full-Service Restaurants (\$ million)	-158.0	76.8	867.9
% Change	-0.14%	0.070%	0.79%

All monetary values are in October 2018 dollars.

In summary, the counterfactual analyses show that online review platforms improve the welfare of restaurant goers by \$2.5 per person per meal, which is equivalent to a 12.6% discount on an average restaurant meal. Online review platforms have a very large impact on individual restaurants' revenue shares, but very little effect on the industry's total revenues. This is largely because information provided by online review platforms transfers demand from lower-quality restaurants to higher-quality restaurants, leaving a net zero effect on the demand for other dining choices, such as eating at fast-food restaurants or eating at home. The counterfactual analysis based on

the no-learning scenario reveals that without accounting for other information sources, the welfare effect of online review platforms can be overestimated by as high as 30%. The analysis based on the full-information scenario shows that if we make information on quality instant and complete, consumers' welfare can increase by an additional \$2.16 per person per meal or a 11% discount on an average restaurant meal. The difference between the "full-information" and "no-learning" counterfactuals uncovers the value of information on welfare: a \$5.43 value per person per meal, or 27.3% of an average meal price.

6 Conclusion

This paper quantifies the effects of online review platforms on restaurant revenues, survival rates, consumer learning and welfare. More specifically, I examine the effects of major online review platforms, including Yelp, TripAdvisor and Google, in the restaurant industry in Texas. Using a novel dataset containing the universe of full-service restaurants' monthly revenues in Texas, search interest on Yelp and TripAdvisor, and online review information from Yelp, TripAdvisor and Google for the period of 1995 to 2015, I first test in a reduced-form analysis if online review platforms speed up consumers' learning process. I find that online review platforms have the opposite effects on high- and low-quality restaurants' revenues and survival probabilities. Specifically, doubling consumers' exposure to Yelp, the dominant platform, increases the revenue of a high-quality new independent restaurant by 8-20% and decreases that of a low-quality restaurant by about the same amount. Doubling Yelp exposure also raises the survival rate of a new high-quality independent restaurant by 7-19 basis points and reduces that of a low-quality restaurant by a similar level. Other platforms, especially Google, have similar effects but with smaller magnitude. In contrast, online platforms do not affect the revenues or survival rates of chains and old independent restaurants. These results are consistent with online review platforms' helping consumers learn faster about restaurant quality.

Building on this reduced-form evidence, I develop a novel structural demand model with social learning to quantify the effect on consumer welfare, restaurant revenue shares and total industry revenues. The counterfactual analysis, where I remove the existence of online review platforms, shows that the welfare effect of online review platforms is equivalent to giving restaurant goers \$2.5 per month per person, which is 12.6% of an average meal price at restaurants. The counterfactual analysis also shows although online review platforms affect individual restaurants' revenues substantially, they influence the total industry revenue very little. The substitution pattern occurs almost entirely between young high- and low-quality independent restaurants, not between young independent restaurants and other restaurants or other dining options.

References

- [1] Acemoglu, D., Makhdoumi, A., Malekian, A. and Ozdaglar, A., 2017. Fast and slow learning from reviews. Working Paper No. 24046. National Bureau of Economic Research.

- [2] Akerberg, D.A., 2001. Empirically distinguishing informative and prestige effects of advertising. *RAND Journal of Economics*, pp.316-333.
- [3] Akerberg, D.A., 2003. Advertising, learning, and consumer choice in experience good markets: an empirical examination. *International Economic Review*, 44(3), pp.1007-1040.
- [4] Anderson, S.P., 2012. Advertising on the Internet. *The Oxford handbook of the digital economy*, pp.355-396.
- [5] Anderson, M. and Magruder, J., 2012. Learning from the crowd: Regression discontinuity estimates of the effects of an online review database. *The Economic Journal*, 122(563), pp.957-989.
- [6] Anenberg, E., Kuang, C., Kung, E. 2018. Social Learning and Local Consumption Amenities: Evidence from Yelp. Working paper.
- [7] Berger, J., Sorensen, A.T. and Rasmussen, S.J., 2010. Positive effects of negative publicity: When negative reviews increase sales. *Marketing Science*, 29(5), pp.815-827.
- [8] Björnerstedt, J. and Verboven, F., 2016. Does merger simulation work? evidence from the swedish analgesics market. *American Economic Journal: Applied Economics*, 8(3), pp.125-164.
- [9] Cai, H., Chen, Y. and Fang, H., 2009. Observational learning: Evidence from a randomized natural field experiment. *The American Economic Review*, 99(3), pp.864-882.
- [10] Chevalier, J.A. and Mayzlin, D., 2006. The effect of word of mouth on sales: Online book reviews. *Journal of marketing research*, 43(3), pp.345-354.
- [11] Ching, A.T., 2010a. Consumer learning and heterogeneity: Dynamics of demand for prescription drugs after patent expiration. *International Journal of Industrial Organization*, 28(6), pp.619-638.
- [12] Ching, A.T., 2010b. A dynamic oligopoly structural model for the prescription drug market after patent expiration. *International Economic Review*, 51(4), pp.1175-1207.
- [13] Ching, A.T., Erdem, T. and Keane, M.P., 2013. Learning models: An assessment of progress, challenges, and new developments. *Marketing Science*, 32(6), pp.913-938.
- [14] Ching, A.T., Erdem, T. and Keane, M.P., 2017. Empirical Models of Learning Dynamics: A Survey of Recent Developments. In *Handbook of Marketing Decision Models* (pp. 223-257). Springer, Cham.
- [15] Conley, T.G. and Udry, C.R., 2010. Learning about a new technology: Pineapple in Ghana. *The American Economic Review*, 100(1), pp.35-69.

- [16] Crawford, G.S. and Shum, M., 2005. Uncertainty and learning in pharmaceutical demand. *Econometrica*, 73(4), pp.1137-1173.
- [17] DeGroot, M.H., 2005. *Optimal statistical decisions* (Vol. 82). John Wiley & Sons.
- [18] Dempster, A.P., Laird, N.M. and Rubin, D.B., 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical society. Series B (methodological)*, pp.1-38.
- [19] Dubé, J.P., 2004. Multiple discreteness and product differentiation: Demand for carbonated soft drinks. *Marketing Science*, 23(1), pp.66-81
- [20] Hall B, Jaffe A, Trajtenberg M. 2002. The NBER patent citations data file: lessons, insights, and methodological tools. In *Patents, Citations and Innovations: A Window on the Knowledge Economy*, Jaffe A, Trajtenberg M (eds). MIT Press: Cambridge, MA; 403–460.
- [21] Hendel, I., 1999. Estimating multiple-discrete choice models: An application to computerization returns. *The Review of Economic Studies*, 66(2), pp.423-446.
- [22] Hollenbeck, B., 2016. Online Reputation Mechanisms and the Decreasing Value of Brands. Unpublished manuscript.
- [23] Jin, G.Z. and Leslie, P., 2003. The effect of information on product quality: Evidence from restaurant hygiene grade cards. *The Quarterly Journal of Economics*, 118(2), pp.409-451.
- [24] Jovanovic, B., 1982. Selection and the Evolution of Industry. *Econometrica*, pp.649-670.
- [25] Kalman, R.E., 1960. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1), pp.35-45.
- [26] Lewis, G. and Zervas, G., 2016. The Welfare Impact of Consumer Reviews: A Case Study of the Hotel Industry. Unpublished manuscript.
- [27] Luca, M., 2016. Reviews, reputation, and revenue: The case of Yelp. com. Harvard Business School Working Paper 12-016.
- [28] Luca, M. and Zervas, G., 2016. Fake it till you make it: Reputation, competition, and Yelp review fraud. *Management Science*, 62(12), pp.3412-3427.
- [29] Lucas, R. Jr. 1976. Econometric Policy Evaluation: A Critique, in: K. Brunner and A. Meltzer (eds.), *The Phillips Curve and Labor Markets*, Carnegie-Rochester Conference Series on Public Policy, Volume 1, pp. 19-46.
- [30] Mayzlin, D., Dover, Y. and Chevalier, J., 2014. Promotional reviews: An empirical investigation of online review manipulation. *The American Economic Review*, 104(8), pp.2421-2455.

- [31] Mehta, A., Rysman, M. and Simcoe, T., 2010. Identifying the age profile of patent citations: New estimates of knowledge diffusion. *Journal of Applied Econometrics*, 25(7), pp.1179-1204.
- [32] Stigler, G.J., 1961. The economics of information. *Journal of Political Economy*, 69(3), pp.213-225.
- [33] Murphy, K.M. and Topel, R.H., 1985. Estimation and inference in two-step econometric models. *Journal of Business & Economic Statistics*, 20(1), pp.88-97.
- [34] Newberry, P.W., 2016. An empirical study of observational learning. *The RAND Journal of Economics*, 47(2), pp.394-432.
- [35] Newberry, P.W. and Zhou, X. 2018. The heterogenous effects of online reputation for local and national retailers. Working Paper, Penn State University.
- [36] Suits, D.B., Mason, A. and Chan, L., 1978. Spline functions fitted by standard regression methods. *The Review of Economics and Statistics*, pp.132-139
- [37] Zhao, Y., Yang, S., Narayan, V. and Zhao, Y., 2013. Modeling consumer learning from online product reviews. *Marketing Science*, 32(1), pp.153-169.
- [38] Zhu, F. and Zhang, X., 2010. Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics. *Journal of marketing*, 74(2), pp.133-148.
- [39] Wooldridge, J.M., 1995. Selection corrections for panel data models under conditional mean independence assumptions. *Journal of econometrics*, 68(1), pp.115-132.
- [40] Wu, C., Che, H., Chan, T.Y. and Lu, X., 2015. The economic value of online reviews. *Marketing Science*, 34(5), pp.739-754.

Appendix

A Figures

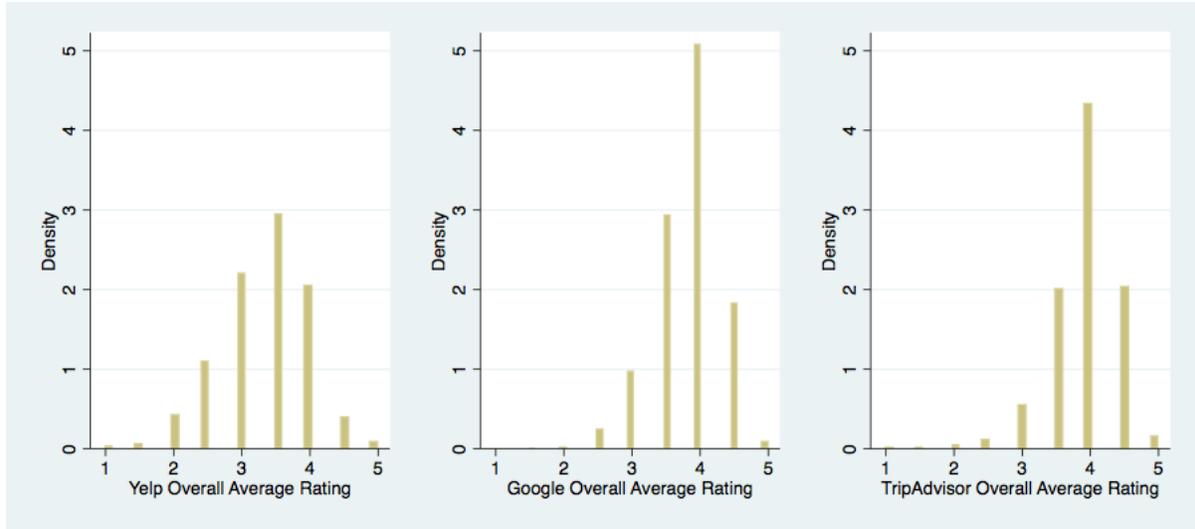


Figure 3: Distribution of Restaurants over Average Ratings

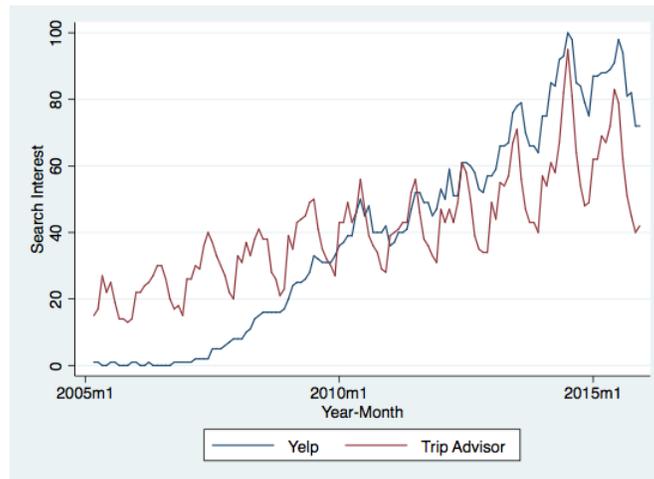


Figure 4: Yelp and TripAdvisor Search Interest Comparison, Texas

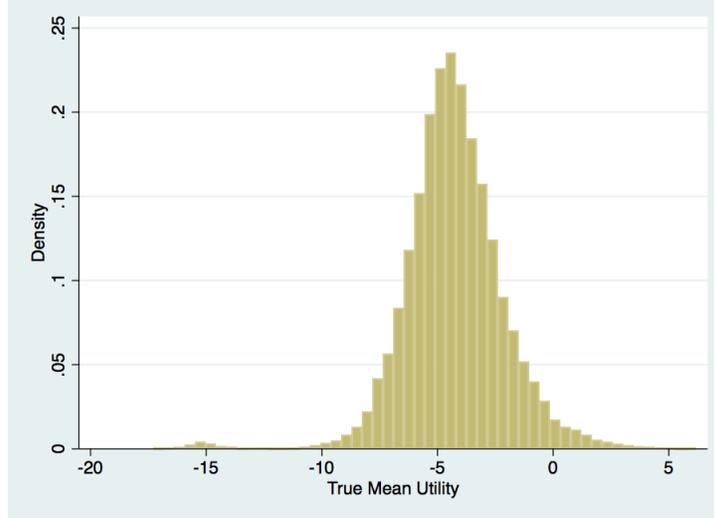


Figure 5: Distribution of Restaurants’ True Mean Utility

B Placebo Test, Heckman’s Correction and Other Robustness Checks

B.1 Placebo Test

As mentioned in Section 3.1, I conduct a placebo test to see if there is variation in the “pre-trend” of restaurant revenues that may be correlated with the exposure online review platforms. To do so, I date Yelp exposure data to 10 years earlier to a start date of March 1995, and then run regression 1 using the sample of restaurant revenues for the period between March 1995 to March 2005. I again use both Google ratings and Yelp ratings as quality measures. The pre-determined ratings cannot be used here because the sample associated with it does not exist before March 2005. The regression results are shown in Table 14. As can be seen, none of the coefficients associated with the Yelp exposure is significant. These results indicate that this source of endogeneity for the exposure measure is not of great concern.

B.2 Heckman’s Correction

It is important to note that the results shown previously do not correct for selection due to endogenous exit. To deal with selection, I again use Heckman’s correction. To control for restaurant fixed effects in the first-stage probit regression, I use the method introduced in Wooldridge (1995), which includes the whole history of control variables as regressors³⁶. The specification for the first-stage probit regression is very similar to equation 1, only that the dependent variable is $action_{jt}$ instead of $\log(Rev_{jt})$ and that I add wage w_{mt} and the history of other independent variables as additional controls. Given the limited number of exits observed for Yelp listed restaurants (about 660 out of 5930 restaurants), I run regression using only Google and Yelp November 2016 ratings as the quality measures as the sample associated with “predetermined” ratings does not have many exits. In addition, as a robustness check, I include $region \times time$ FEs to capture any type of time varying factors that are specific for each region. The raw regression results for using Google and Yelp ratings are shown in Tables 16 and 17 respectively. The main results do not change substantially from the ones that do not control for selection. Below I present the translated results for the effects of Yelp on revenues for given rating classes in Table 15. As can be seen in the table, once we control for selection and $time \times metro$ FEs, the results do not change much, only that once we add $time \times metro$ FEs, the significance

³⁶The Wooldridge method works well only if the number of observations is much larger than the number of time periods in the panel data; otherwise, the number of independent variables could exceed or become very close to the number of observations, rendering the probit regression invalid. In my application, some restaurants have a very long time horizon. To solve the problem, I cut my sample into many 12-month periods and include the history of controls during the 12 months.

Table 14: Placebo Test Results

	(1)	(2)
	Log Revenue	Log Revenue
Log Yelp (θ_y)	0.344 (0.227)	0.113 (0.175)
Log Yelp×Rating (θ_{yr})	-0.0747 (0.0563)	-0.0254 (0.0483)
Log Yelp×Chain Dummy (θ_{yc})	-0.0385 (0.0977)	0.0707 (0.0592)
Log Yelp×Chain Dummy×Rating (θ_{ycr})	0.0224 (0.0274)	-0.0103 (0.0175)
Controls	X	X
Time FE	Yes	Yes
Time FE×Chain Dummy	Yes	Yes
Restaurant FE	Yes	Yes
Sample Period	March 1995-March 2005	March 1995-March 2005
Rating	Google Rating	Yelp Nov 2016 Rating
N	91,963	56,147

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Estimates for Controls are shown in Section F of the Online Appendix.

levels of the estimated effects reduce. Nonetheless, the magnitude of the effects and the significance levels for very high and very low quality restaurants stay almost the same.

B.3 No Evidence of Quality Improvement

I test the hypothesis that restaurants' quality increases with age by examining the average of new ratings restaurants receive each month. The econometric specification is as follows:

$$\begin{aligned} \log(ANR_{jt}) = & \theta_a \log(a_{jt}) + \theta_{ac} \log(a_{jt}) \times D_j^C + \theta_y \log(Yelp_{mt}) + \theta_{yc} \log(Yelp_{mt}) \times D_j^C \\ & + \mathbf{X}_{mt} \boldsymbol{\theta}_x + \theta_{nI} n_{mt}^I + \theta_{nc} n_{mt}^C + \theta_t + \theta_{tc} D_j^C + \theta_j + \xi_{jt} \end{aligned} \quad (29)$$

where ANR_{jt} is the average of new ratings that restaurant j receives in month t . The rest of the parameters and variables are defined in the same way as those in equation 1.

In this regression, θ_a and θ_{ac} are the parameters of interest. Them being positive and significant implies that there is quality improvement and thereby managerial learning. Here I include both calendar time fixed effects and their interactions with the chain dummy to control for aggregate shifts in consumer tastes that affect how Yelp reviewers rate restaurants. Furthermore, I also include restaurant fixed effects to account for restaurant-specific time-invariant characteristics that make them survive longer or exit faster. Failure to include restaurant fixed-effects can create bias in the estimation of the age coefficients θ_a and θ_{ac} since the identification of these coefficients for older ages comes from restaurants who have survived longer and tend to have higher qualities or ratings. The regression results are shown in Table 18. Columns 1 to 3 of the table are regression results without including market characteristics. Column 1 are the results for the entire sample including all restaurants listed on Yelp. Columns 2 and 3 are the results for the sample before and after Yelp penetration respectively. Columns 4 and 5 are regression results by including all market controls variables, and they are for the full sample and the sample with only new restaurants respectively. As can be seen, none of the coefficients associated with age in these results are significant. This implies that there is no significant improvement in restaurants' services or qualities as a result of managerial learning.

Table 15: Effects of Yelp Exposure on Revenue by Quality Level With Heckman Correction

Star Rating	Google	Google	Yelp	Yelp
2	-.260*** (.0705)	-.233** (.103)	-.155*** (.0468)	-.174*** (.0688)
3	-.109*** (.0377)	-.0665 (.0831)	-.0436** (.0251)	-.0502 (.0548)
4	.0413** (.0203)	.100 (.0753)	.0682** (.0273)	.0731 (.0542)
5	.184*** (.0490)	.267*** (.0827)	.180*** (.0504)	.196*** (.0672)
Time FE	Yes	Yes	Yes	Yes
Time FE×Chain Dummy	Yes	Yes	Yes	Yes
Time ×Metro FE	No	Yes	No	Yes
Restaurant FE	Yes	Yes	Yes	Yes
Sample	Entry after Yelp penetration			
<i>N</i>	131,138	131,024	158,817	158,724

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Table 16: Effects of Yelp Exposure on Revenue With Heckman Correction (Google Ratings)

	(1)	(2)	(3)
	Log Revenue	Log Revenue	Log Revenue
Log Yelp (θ_y)	-0.166 (0.102)	-0.560*** (0.140)	-0.567*** (0.158)
Log Yelp×Rating (θ_{yr})	0.0487* (0.0259)	0.150*** (0.0357)	0.167*** (0.0346)
Log Yelp×Chain Dummy (θ_{yc})	-0.0807 (0.0542)	-0.0319 (0.0655)	NA NA
Log Yelp×Chain Dummy×Rating (θ_{ycr})	0.0202 (0.0149)	0.0195 (0.0174)	0.0303 (0.0185)
Controls	X	X	X
Inverse Mills Ratio	-0.177 (0.110)	-0.0891 (0.115)	-0.0793 (0.118)
Time FE	Yes	Yes	Yes
Time FE×Chain Dummy	Yes	Yes	Yes
Time ×Metro FE	No	No	Yes
Restaurant FE	Yes	Yes	Yes
Age Group	Old	Young	Young
<i>N</i>	116,831	131,138	131,024

Standard errors in parentheses are clustered by restaurant.

* p<0.10, ** p<0.05, *** p<0.01

Estimates for Controls are shown in Section F of the Online Appendix.

Table 17: Effects of Yelp Exposure on Revenue With Heckman Correction (Yelp Ratings)

	(1)	(2)	(3)
	Log Revenue	Log Revenue	Log Revenue
Log Yelp (θ_y)	-0.0376 (0.138)	-0.379*** (0.102)	-0.420*** (0.115)
Log Yelp×Rating (θ_{yr})	-0.000344 (0.0446)	0.112*** (0.0290)	0.123*** (0.0287)
Log Yelp×Chain Dummy (θ_{yc})	-0.00818 (0.0940)	0.00535 (0.0428)	NA NA
Log Yelp×Chain Dummy×Rating (θ_{ycr})	-0.00558 (0.0300)	0.00779 (0.0139)	0.00838 (0.0146)
Controls	X	X	X
Inverse Mills Ratio	-0.674 (0.553)	-0.558 (0.484)	-0.548 (0.493)
Time FE	Yes	Yes	Yes
Time FE×Chain Dummy	Yes	Yes	Yes
Time ×Metro FE	No	No	Yes
Restaurant FE	Yes	Yes	Yes
Sample	Old	Young	Young
<i>N</i>	139,311	158,817	158,724

Standard errors in parentheses are clustered by restaurant.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Estimates for Controls are shown in Section F of the Online Appendix.

Table 18: Relationship Between Average New Rating Per Month and Age

	(1)	(2)	(3)	(4)	(5)
	Log Average New Rating Per Month				
Log Age	-0.00291 (0.00320)	-0.0375 (0.0401)	-0.00493 (0.00370)	-0.00287 (0.00321)	-0.00516 (0.00368)
Log Age×Chain Dummy	-0.00595 (0.00673)	-0.0308 (0.0769)	-0.00423 (0.00820)	-0.00536 (0.00669)	-0.00346 (0.00815)
Log Yelp				-0.0121 (0.00751)	-0.00970 (0.0102)
Log Yelp×Chain Dummy				-0.000707 (0.0136)	-0.0227 (0.0207)
Controls				X	X
Time FE	Yes	Yes	Yes	Yes	Yes
Time FE×Chain Dummy	Yes	Yes	Yes	Yes	Yes
Restaurant FE	Yes	Yes	Yes	Yes	Yes
Age Group	All	Old	Young	All	Young
<i>N</i>	144,659	59,899	84,754	144,659	84,754

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Estimates for Controls are shown in Section F of the Online Appendix.

C Derivation of Revenue Shares and Number of Customers

As mentioned in Section 4, s_{jt} is the probability of a consumer choosing restaurant j at time t . Note that s_{jt} is different from the market share in the standard discrete choice models. It represents how many consumers will choose restaurant j out of one unit mass of consumers, but does not represent the quantity sold at restaurant j as in the standard discrete choice demand model. To obtain quantity, one can use Roy's identity to derive a consumer's demand for restaurant j , $d_j(y_{it})$, conditional on her having chosen j :

$$\begin{aligned} d_j(y_{it}) &= -(\partial E[U_{ijt}|I(t)]/\partial p_{jt})/(\partial E[U_{ijt}|I(t)]/\partial y_{it}) \\ &= \gamma \frac{y_{it}}{p_{jt}} \end{aligned} \quad (30)$$

where y_{it} is consumer i 's income at time t . By integrating over the income distribution in a market and multiplying the average individual demand by the market size, we have that the aggregate demand q_{jt} is

$$\begin{aligned} q_{jt} &= \int s_{jt} d_j(y) dF_y(y) M_{mt} \\ &= s_{jt} \int d_j(y) dF_y(y) M_{mt} \\ &= s_{jt} \gamma Y_{mt} / p_{jt} \end{aligned} \quad (31)$$

where M_{mt} is the size of the population in market m at time t ; $F_y(y)$ is the income distribution of the population in the market; Y_{mt} is the total income of consumers in market m at time t , and γ , as mentioned before, is the budget share of consumers' income spent on eating out. This equation implies that s_{jt} can be written as a function of the revenue:

$$\begin{aligned} s_{jt} &= \frac{p_{jt} q_{jt}}{\gamma Y_{mt}} \\ &= \frac{Rev_{jt}}{\gamma Y_{mt}} \end{aligned} \quad (32)$$

where Rev_{jt} is the revenues of independent restaurant j at time t .

In equation 32, γY_{mt} can be seen as the market size, given that

$$1 = \sum_{j \in J_{mt}} s_{jt} = \frac{\sum_{j \in J_{mt}} Rev_{jt}}{\gamma Y_{mt}}$$

where J_{mt} is the set of all restaurants and the outside option in a market. Therefore, s_{jt} does not only represent the probability of choosing restaurant j but also restaurant j ' share of revenues in a market.

For the number of customers who visited restaurant j at time t , denoted by NC_{jt} we can write

$$NC_{jt} = s_{jt} M_{mt} = \frac{Rev_{jt}}{\gamma Y_{mt}} M_{mt} = \frac{Rev_{jt}}{\gamma \frac{Y_{mt}}{M_{mt}}} = \frac{Rev_{jt}}{\gamma \bar{y}_{mt}} \quad (33)$$

D Identification of the Structural Model

For simplicity, here I use only one restaurant platform type to discuss identification issues of the learning parameters. As shown in Section 4, the estimating equation of the structural model is

$$\begin{aligned} z_{jt} &= \mathbf{X}_{jt} \boldsymbol{\theta}_x - \alpha \log(p_{jt}) + \eta_{jt} + \sum_{k=1}^5 \theta_k \tau^k + \xi_{jt} \\ \xi_{jt} &= \rho \xi_{jt-1} + \varepsilon_{jt}, \varepsilon_{jt} \sim i.i.d. N(0, \sigma_\varepsilon^2) \end{aligned} \quad (34)$$

where

$$\eta_{jt} = (1 - \beta_{jt-1})\eta_{jt-1} + \beta_{jt-1}A_j + \nu_{jt} \quad (35)$$

$$A_j = \theta_0 + \theta_r R_j \quad (36)$$

$$\nu_{jt} = \beta_{jt-1}(\bar{A}_{jt-1} - A_j) \quad (37)$$

$$\beta_{jt-1} = \frac{\frac{1}{\sigma_\delta^2} n_{jt-1}}{\frac{1}{\sigma_{jt-1}^2} + \frac{1}{\sigma_\delta^2} n_{jt-1}} \quad (38)$$

$$n_{jt-1} = (\kappa + \lambda_y \log(Y_{mt-1}) + \lambda_g \log(T_{mt-1})) \frac{Rev_{jt-1}}{\gamma \bar{y}_{mt-1}} \quad (39)$$

$$\nu_{jt} \sim N\left(0, \frac{\beta_{jt-1}^2 \sigma_\delta^2}{n_{jt-1}}\right) \quad (40)$$

where ν_{jt} is an error term with mean 0 and a time-varying variance.

Here I discuss in greater detail about the identification of the learning parameters. First, as shown in equations 38 to 40, σ_δ cannot be separately identified from κ and λ 's. Therefore, it is normalized to 1. Second, the identification of other learning parameters, the initial prior mean A_j^p , the initial prior variance σ_A^2 , κ , λ_y , and λ_g can be demonstrated through the following analysis. η_{jt} can be written as follows

$$\eta_{jt} = \frac{A_j^p + \sigma_A^2 N_{jt-1} A_j + \sigma_A^2 (\sum_{l=1}^{N_{jt-1}} \delta_l)}{1 + \sigma_A^2 N_{jt-1}} \quad (41)$$

where $N_{jt-1} \equiv \sum_{l=1}^{t-1} n_{jl}$, i.e. the total number of experience signals received by consumers; δ_l is the random component of the experience signals. It is i.i.d. and follows a standard normal distribution since σ_δ has been normalized to 1.

Then η_{jt} can be decomposed into three terms:

$$\eta_{jt} = \frac{A_j^p}{1 + \sigma_A^2 N_{jt-1}} + \frac{\sigma_A^2 N_{jt-1} A_j}{1 + \sigma_A^2 N_{jt-1}} + \frac{\sigma_A^2}{1 + \sigma_A^2 N_{jt-1}} \left(\sum_{l=1}^{N_{jt-1}} \delta_l \right) \quad (42)$$

It is easier to think about the identification of the parameters for a given rating class first. Suppose we take all restaurants with a rating of 3 in our estimation, then A_j^p and A_j are the prior mean and true mean quality of these restaurants. We can focus on how we can identify A_j^p and A_j and other learning parameters for this rating class. The first and second terms of equation 42 is a weighted average of the initial prior mean A_j^p and the true mean quality A_j . Together they determine the evolution of the mean revenues with respect to age t after all other demand variables have been controlled for. In particular, at age 0, the revenues start with A_j^p , and when a restaurant reaches a certain age and N_{t-1} is large enough, the revenues will converge to A_j . Therefore, restaurants' revenues at age $t = 0$ can identify A_j^p , and the revenues at a much later age identify A_j . How fast the mean revenues converge from A_j^p to A_j identifies $\sigma_A^2 N_{t-1}$. It is important to know if σ_A^2 can be separately identified from N_{t-1} . If so, the variations of $\log(Y_{mt-1})$ and $\log(T_{mt-1})$ in N_{t-1} will identify both κ and λ 's.

The variance of the third term in equation 42 identifies σ_A^2 . To see this, we can write the variance as

$$Var = \frac{\sigma_A^4}{(1 + \sigma_A^2 N_{t-1})^2} N_{t-1} = \frac{\sigma_A^2}{(1 + \sigma_A^2 N_{t-1})^2} \sigma_A^2 N_{t-1} \quad (43)$$

Since $\sigma_A^2 N_{t-1}$ can be identified from the speed of convergence of the revenues, the variance of the third term therefore identifies σ_A^2 .

Once A_j^p and A_j have been identified for each rating class, we can collect these A_j^p and A_j for all rating classes from 1 star to 5 stars and fit them as a linear function of these star ratings. Fitting either A_j^p or A_j as a linear function of rating R_j will result in an intercept and a slope. However, it can be easily shown that only the difference between the intercepts for the linear function of A_j^p and that of A_j can be identified because the common part of the intercepts for both equations can be taken out of the learning component

η_{jt} , and once taken out, it will be absorbed by restaurant cuisine fixed effects. Therefore, I formulate the linear function of the initial prior $A_j^p = \theta_r^p R_j$ without an intercept, but the function of the true mean quality $A_j = \theta_0 + \theta_r R_j$ does have an intercept. Once A_j^p and A_j are identified for all rating classes, the intercept and slope parameters θ_0 , θ_r^p , and θ_r can be easily identified.

E Derivation of the Likelihood Function

In this section, I derive the closed form of the log-likelihood function. For cleaner presentation, I rewrite the model by excluding all control variables:

$$z_{jt} = \eta_{jt} + \xi_{jt} \quad (44)$$

$$\eta_{jt} = (1 - \beta_{jt-1})\eta_{jt-1} + \beta_{jt-1}A + \nu_{jt}, \nu_{jt} \sim N(0, \sigma_{\nu_{jt}}^2) \quad (45)$$

$$\xi_{jt} = \rho\xi_{jt-1} + \varepsilon_{jt}, \varepsilon_{jt} \sim N(0, \sigma_{\varepsilon}^2) \quad (46)$$

where z_{jt} can simply be seen as the residual by excluding all the effects from the linear controls. Rearrange the above variables to eliminate ξ_{jt} , we have

$$z_{jt} = \rho z_{jt-1} + \eta_{jt} - \rho\eta_{jt-1} + \varepsilon_{jt} \quad (47)$$

$$\eta_{jt} = (1 - \beta_{jt-1})\eta_{jt-1} + \beta_{jt-1}A + \nu_{jt} \quad (48)$$

Let $\mathbf{z}_j = (z_{j0}, z_{j1}, \dots, z_{jT_j})$ denote the time series revenue data for restaurant j and $p(\mathbf{z}_j|\mathbf{x}, \boldsymbol{\theta})$ be the likelihood for observing \mathbf{z}_j given the structural parameters $\boldsymbol{\theta}$ and independent variables \mathbf{x} . Then the log likelihood function of observing the revenue streams for all restaurants is

$$\mathcal{L} = \sum_{j=1}^n \log(p(\mathbf{z}_j|\mathbf{x}, \boldsymbol{\theta})) \quad (49)$$

where n is the total number of restaurants in the sample.

To obtain this log likelihood function, we need to compute $\log(p(\mathbf{z}_j|\mathbf{x}, \boldsymbol{\theta}))$ for each restaurant. Note that due to the serial correlation in ξ_{jt} , we will not be able to count the likelihood associated with the first observation z_{j0} . Given that η_{jt} is unobserved, we need to integrate it out. The most straightforward approach is to integrate out η_{jt} over the joint distribution of $(\eta_{j1}, \dots, \eta_{jT_j})$; however, it is computationally cumbersome. Instead, we can do this sequentially:

$$\begin{aligned} \mathcal{L} &= \log(p(\mathbf{z}_j|\mathbf{x}, \boldsymbol{\theta})) \\ &= \log\left(\int p(z_{j1:T_j}, \eta_{j1:T_j}|\mathbf{x}, \boldsymbol{\theta})d\eta_{j1:T_j}\right) \\ &= \log\left(\prod_{t=1}^{T_j} p(z_{jt}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1})\right) \\ &= \log\left(\prod_{t=1}^{T_j} \int p(z_{jt}, \eta_{jt}, \eta_{jt-1}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1})d\eta_{jt}d\eta_{jt-1}\right) \\ &= \log\left(\prod_{t=1}^{T_j} \int p(z_{jt}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1}, \eta_{jt}, \eta_{jt-1})p(\eta_{jt}, \eta_{jt-1}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1})d\eta_{jt}d\eta_{jt-1}\right) \\ &= \log\left(\prod_{t=1}^{T_j} \int p(z_{jt}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1}, \eta_{jt}, \eta_{jt-1})p(\eta_{jt}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1}, \eta_{jt-1})p(\eta_{jt-1}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1})d\eta_{jt-1}d\eta_{jt}\right) \\ &= \sum_{t=1}^{T_j} \log\left(\int p(z_{jt}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1}, \eta_{jt}, \eta_{jt-1})p(\eta_{jt}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1}, \eta_{jt-1})p(\eta_{jt-1}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1})d\eta_{jt-1}d\eta_{jt}\right) \quad (50) \end{aligned}$$

where $p(\cdot)$ denotes density function. The first line of this equation shows that the integrated log likelihood can be obtained by integrating over the entire path of $\eta_{j1:T_j}$. The initial period's η_{j0} is the initial prior, a parameter to be estimated. Therefore, we do not take integration over the initial period's likelihood. The second line and those that follow show that this integrated log likelihood function can be written as the sum of each period's log likelihood conditional on past observations. In particular, $p(z_{jt}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1})$ can be attained from integrating the product of $p(z_{jt}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1}, \eta_{jt})$, $p(\eta_{jt}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1}, \eta_{jt-1})$ and $p(\eta_{jt-1}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1})$ over η_{jt} and η_{jt-1} . The first two probability functions are known; the only unknown element is $p(\eta_{jt-1}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1})$, but it can be obtained through the following equations:

$$p(\eta_{jt-1}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1}) = \int \frac{p(z_{jt-1}, \eta_{jt-1}, \eta_{jt-2}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-2})}{p(z_{jt-1}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-2})} d\eta_{jt-2} \quad (51)$$

This equation implies an iterative process, where the likelihood of z_{jt} given all past observations, $p(z_{jt}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1})$, can be calculated recursively. Once we know the first period's likelihood, $p(z_{j1}|\mathbf{x}, \boldsymbol{\theta}, z_{j0})$, we can obtain $p(\eta_{j1}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:1})$ from equation 51. With that we can calculate $p(z_{j2}|\mathbf{x}, \boldsymbol{\theta}, z_{j1})$ using the last line in equation 50. Once we continue to do this iteratively, we can obtain $p(z_{jt}|\mathbf{x}, \boldsymbol{\theta}, z_{j0:t-1})$ for all t . Given that all errors in this model are Gaussian, a closed-form of these probability functions exist. All of the above conditional probabilities come from normal distributions, and their mean and variances can be written as follows:

$$\hat{\eta}_{jt|t-1} = (1 - \beta_{jt-1})\hat{\eta}_{jt-1|t-1} + \beta_{jt-1}A \quad (52)$$

$$\hat{\sigma}_{\eta_{jt|t-1}}^2 = (1 - \beta_{jt-1})^2 \hat{\sigma}_{\eta_{jt-1|t-1}}^2 + \sigma_{v_{jt}}^2 \quad (53)$$

$$\tilde{z}_{jt} = z_{jt} - \rho z_{jt-1} - \hat{\eta}_{jt|t-1} + \rho \hat{\eta}_{jt-1|t-1} \quad (54)$$

$$S_{jt} = \sigma_e^2 + \hat{\sigma}_{\eta_{jt|t-1}}^2 + \rho(\rho - 2(1 - \beta_{jt-1}))\hat{\sigma}_{\eta_{jt-1|t-1}}^2 \quad (55)$$

$$\hat{\eta}_{jt|t} = \hat{\eta}_{jt|t-1} + \left(\hat{\sigma}_{\eta_{jt|t-1}}^2 - \rho(1 - \beta_{jt-1})\hat{\sigma}_{\eta_{jt-1|t-1}}^2 \right) S_{jt}^{-1} \tilde{z}_{jt} \quad (56)$$

$$\hat{\sigma}_{\eta_{jt|t}}^2 = \left[\sigma_e^2 \hat{\sigma}_{\eta_{jt|t-1}}^2 + \rho^2 \hat{\sigma}_{\eta_{jt-1|t-1}}^2 \hat{\sigma}_{\eta_{jt|t-1}}^2 - \rho^2 (1 - \beta_{jt-1})^2 \hat{\sigma}_{\eta_{jt-1|t-1}}^4 \right] S_{jt}^{-1} \quad (57)$$

$$\hat{z}_{jt|t-1} = \rho z_{jt-1} + \hat{\eta}_{jt|t-1} - \rho \hat{\eta}_{jt-1|t-1} \quad (58)$$

$$\hat{\sigma}_{z_{jt|t-1}}^2 = S_{jt} \quad (59)$$

where $\hat{\eta}_{jt|t-1}$ is the mean of η_{jt} conditional on past data $z_{j0:t-1}$. $\hat{\eta}_{jt|t}$ is similarly defined. $\hat{\sigma}_{\eta_{jt|t-1}}^2$ is the variance of $\eta_{jt|t-1}$ given past data $z_{j0:t-1}$, and $\hat{\sigma}_{\eta_{jt|t}}^2$ is also similarly defined. \tilde{z}_{jt} is the residual. S_{jt} is an intermediate factor that facilitates the updating of $\hat{\eta}_{jt|t}$ and $\hat{\sigma}_{\eta_{jt|t}}^2$. $\hat{z}_{jt|t-1}$ is the mean of z_{jt} conditional on past observations $z_{j0:t-1}$, and $\hat{\sigma}_{z_{jt|t-1}}^2$ is the variance.

Once we have $p(z_j|\mathbf{x}, \boldsymbol{\theta})$ for each restaurant, we can obtain the total log likelihood for all restaurants \mathcal{L} as shown in equation 49, and the structural parameters can be estimated by maximizing \mathcal{L} :

$$\boldsymbol{\theta}^* = \underset{\boldsymbol{\theta}}{\operatorname{argmax}} \sum_{j=1}^n \log(p(z_j|\mathbf{x}, \boldsymbol{\theta}))$$

F Technical Issues Regarding Welfare Calculation

As shown in Section 5.2, to quantify consumer welfare, I need to calculate the true mean utility Δ_{jt}^T , the counterfactual expected mean utility Δ_{jt}^c , revenue share s_{jt}^c , revenues of the outside option Rev_{0mt}^c and the counterfactual revenues for all restaurants. I explain how to obtain these counterfactual variables in this section.

First, we need to know the true mean utility difference $\Delta_{jt}^T - \Delta_{0mt}$ for all young independent restaurants listed online. This difference can be predicted given that we have uncovered all the structural parameters and each restaurant's true quality in the structural estimation. Specifically, it can be obtained from

$$\Delta_{jt}^T - \Delta_{0mt} = \mathbf{X}_{mt} \hat{\boldsymbol{\theta}}_x - \hat{\alpha} \log(p_j) + \sum_{k=1}^5 \theta_k \tau^k + \hat{A}_j + \hat{\xi}_{jt} \quad (60)$$

where $\hat{\xi}_{jt}^c$ represents the sum of the estimated aggregate demand shocks ξ_{jt} and experience signal shocks δ_l .

To extract $\hat{\xi}_{jt}^c$, I substitute equation 42 into equation 15 and rewrite the revenue function as follows:

$$z_{jt} = \mathbf{X}_{mt} \hat{\boldsymbol{\theta}}_x - \hat{\alpha} \log(p_{jt}) + \sum_{k=1}^5 \theta_k \tau^k + \frac{\hat{A}_j^p}{1 + \hat{\sigma}_A^2 N_{jt-1}} + \frac{\hat{\sigma}_A^2 N_{jt-1} \hat{A}_j}{1 + \hat{\sigma}_A^2 N_{jt-1}} + \frac{\hat{\sigma}_A^2}{1 + \hat{\sigma}_A^2 N_{jt-1}} \left(\sum_{l=1}^{N_{jt-1}} \hat{\delta}_l \right) + \hat{\xi}_{jt} \quad (61)$$

Recall from equation 42 that η_{jt} has a deterministic part $\frac{\hat{A}_j^p}{1 + \hat{\sigma}_A^2 N_{jt-1}} + \frac{\hat{\sigma}_A^2 N_{jt-1} \hat{A}_j}{1 + \hat{\sigma}_A^2 N_{jt-1}}$ and a noisy part $\frac{\hat{\sigma}_A^2}{1 + \hat{\sigma}_A^2 N_{jt-1}} (\sum_{l=1}^{N_{jt-1}} \hat{\delta}_l)$. This noisy part can be lumped into the aggregate demand shocks $\hat{\xi}_{jt}$ to form the aggregate residual $\hat{\xi}_{jt}^c$:

$$\begin{aligned} \hat{\xi}_{jt}^c &= \frac{\hat{\sigma}_A^2}{1 + \hat{\sigma}_A^2 N_{jt-1}} \left(\sum_{l=1}^{N_{jt-1}} \hat{\delta}_l \right) + \hat{\xi}_{jt} \\ &= z_{jt} - \mathbf{X}_{mt} \hat{\boldsymbol{\theta}}_x + \hat{\alpha} \log(p_{jt}) - \sum_{k=1}^5 \theta_k \tau^k - \frac{\hat{A}_j^p}{1 + \hat{\sigma}_A^2 N_{jt-1}} - \frac{\hat{\sigma}_A^2 N_{jt-1} \hat{A}_j}{1 + \hat{\sigma}_A^2 N_{jt-1}} \end{aligned} \quad (62)$$

To construct $\Delta_{jt}^c - \Delta_{0mt}$, I use a similar equation as equation 60:

$$\Delta_{jt}^c - \Delta_{0mt} = \mathbf{X}_{mt} \hat{\boldsymbol{\theta}}_x - \hat{\alpha} \log(p_{jt}) + \sum_{k=1}^5 \theta_k \tau^k + \frac{\hat{A}_j^p}{1 + \hat{\sigma}_A^2 N_{jt-1}^c} + \frac{\hat{\sigma}_A^2 N_{jt-1}^c \hat{A}_j}{1 + \hat{\sigma}_A^2 N_{jt-1}^c} + \hat{\xi}_{jt}^c \quad (63)$$

where $N_{jt-1}^c \equiv \sum_{l=1}^{t-1} n_{jl}^c$, and $n_{jl}^c = (\kappa + \lambda_y \log(Y_{ml}) + \lambda_g \log(T_{ml})) \frac{Rev_{jl}^c}{\gamma y_{mt-1}}$; the component $\frac{\hat{A}_j^p}{1 + \hat{\sigma}_A^2 N_{jt-1}^c} + \frac{\hat{\sigma}_A^2 N_{jt-1}^c \hat{A}_j}{1 + \hat{\sigma}_A^2 N_{jt-1}^c}$ represents the deterministic part of the expected quality η_{jt}^c in the counterfactual. The noisy part of η_{jt}^c should be $\frac{\hat{\sigma}_A^2}{1 + \hat{\sigma}_A^2 N_{jt-1}^c} (\sum_{l=1}^{N_{jt-1}^c} \hat{\delta}_l)$, which differs from the noisy part of η_{jt} in the total number of signals N_{jt-1} . To retain the shocks in terms of experience signals that firms encounter, I use the estimate of this component from the real world to approximate its value in the counterfactual; that is, I set this component in the counterfactual to $\frac{\hat{\sigma}_A^2}{1 + \hat{\sigma}_A^2 N_{jt-1}^c} (\sum_{l=1}^{N_{jt-1}} \hat{\delta}_l)$ ³⁷. Furthermore, I also retain the aggregate demand shocks from the factual world $\hat{\xi}_{jt}$. By this construction, the sum of these two types of shocks $\hat{\xi}_{jt}^c$ are added to the equation for determining $\Delta_{jt}^c - \Delta_{0mt}$.

To obtain N_{jt-1}^c , we need the counterfactual revenue stream Rev_{jt}^c . To get Rev_{jt}^c , we first need to calculate the revenue share s_{jt}^c , and then use that to multiply the market size. s_{jt}^c can be attained using the following equation:

$$s_{jt}^c = \frac{\exp(\Delta_{jt}^c - \Delta_{0mt})}{1 + \sum_{k \in JO_{mt}} \exp(\Delta_{jt}^c - \Delta_{0mt}) + \frac{\sum_{l \in Oth_{mt}} Rev_{lt}}{Rev_{0mt}}} \quad (64)$$

where JO_{mt} is the set of young independent restaurants listed online; Oth_{mt} is the set of all other full-service restaurants in market m at time t in addition to the young independent restaurants listed online. Note that each of their mean utilities can be approximated by $(\log(Rev_{lt}) - \log(Rev_{0mt}))$. Taking the exponential of the individual mean utility and summing them up gives us $\frac{\sum_{l \in Oth_{mt}} Rev_{lt}}{Rev_{0mt}}$.

³⁷This approximation is innocuous since a closer examination of this component shows that it is the average of all experience signals up to time $t - 1$. This component in the counterfactual differs from that in the factual world only in terms of the precision or variance, but not in terms of the sign, negative or positive. Therefore, firms that encountered good shocks in terms of experience signals in the real world will also experience good shocks in the counterfactual.

Once the shares are computed, the counterfactual revenue stream of a restaurant can be easily obtained by applying equation 32, and the revenue of the outside option in the counterfactual Rev_{0mt}^c can be computed in the same way.

Now we have attained Rev_{0mt}^c , $\Delta_{jt}^T - \Delta_{0mt}$, $\Delta_{jt}^c - \Delta_{0mt}$, and s_{jt}^c , and we know $\Delta_{jt} - \Delta_{0mt}$ and s_{jt} from the data; we can calculate $CS_{mt} - CS_{mt}^c$ for all markets.

G Tables

Table 19: Restaurant Category Classification

Category	Description
1	breakfast, brunch, sandwiches, cafe, and vegan
2	Asian, including Chinese, Japanese, Korean and Indian
3	bars, gastropubs, traditional and new American
4	Tex-Mex, cajun, Caribbean and other southern cuisines
5	European, including French, Spanish, German, Greek, and other Mediterranean cuisines

Table 20: Restaurant Prices

Category	\$	\$\$	\$\$\$	\$\$\$\$
Dollar Amount	\$10	\$20	\$45	\$70

Table 21: Effects of Yelp Exposure on Survival Rates

	(1) action	(2) action	(3) action	(4) action
Log Yelp (θ_y)	-0.000172 (0.00440)	-0.00602** (0.00276)	-0.000726 (0.00115)	-0.00426** (0.00204)
Log Yelp×Rating (θ_{yr})	0.000114 (0.00108)	0.00155** (0.000690)	-0.0000190 (0.000336)	0.00123** (0.000557)
Log Yelp×Chain Dummy (θ_{yc})	-0.00566 (0.00486)	0.000379 (0.00125)	-0.00231 (0.00158)	0.000342 (0.000956)
Log Yelp×Chain Dummy×Rating (θ_{ycr})	0.00131 (0.00131)	-0.000212 (0.000336)	0.000625 (0.000441)	-0.000251 (0.000312)
Controls	X	X	X	X
Time FE	Yes	Yes	Yes	Yes
Time FE×Chain Dummy	Yes	Yes	Yes	Yes
Restaurant FE	Yes	Yes	Yes	Yes
Age Group	Old	Young	Old	Young
Rating	Google Nov, 2016	Google Nov, 2016	Yelp Nov, 2016	Yelp November 2016 Rating
N	121,083	177,373	131,079	177,373

Standard errors in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Estimates for Controls are shown in Section F of the Online Appendix.