Crowd-Driven Competitive Intelligence: Understanding the Relationship between Local Market Structure and Online Rating Distributions

Dominik Gutt
Paderborn University
dominik.gutt@uni-paderborn.de

Philipp Herrmann
Consultant
philipp_herrmann@hotmail.de

Mohammad S. Rahman
Purdue University
mrahman@purdue.edu

Abstract

Crowdsourced information, such as online ratings, are increasingly viewed as a critical source for understanding local market dynamics. A key aspect of utilizing online ratings to derive competitive market intelligence is to delineate the systematic relationship between local market structure and distributional properties of online ratings. As one of the earliest papers in this stream, combining demographic, population, and restaurant review data from Yelp.com for 372 isolated markets in the U.S., our empirical findings suggest that an increase in competition leads to a broader range of ratings and to a decrease in the average mean rating in a market. These effects are particularly pronounced when the analysis is limited to specific restaurant types where there are fewer opportunities for horizontal differentiation. To gain richer insights into the empirical results, we adopt the classical theoretical lens of an oligopoly where firms vertically differentiate their quality offerings in the presence of heterogeneous consumers and marginal costs that increase quadratically in quality. Moreover, we present evidence in support of both the internal and external validity of Yelp's crowdsourced online ratings, validating the role that online rating distributions can play in helping scholars and managers understand competitive dynamics in local markets.

Keywords: Local Market Structure, Online Ratings, Online Offline Interplay, Competition, Competitive Market Intelligence.

1. Introduction

What is the relationship between local market structure and online rating distributions? While there has been a steady growth of studies focusing on online ratings and other crowdsourced information, there have not been any studies so far that have focused on understanding the influence of local market structure on the distributional properties of online ratings. Although the theoretical work in the industrial organization literature sheds light on the impact of market structure on quality provisioning decisions of firms (Shaked and Sutton 1982, Jones and Mendelson 2011), and online mean ratings are widely believed to reflect those quality provisioning decisions (Chen at al. 2011, Luca 2011), it is less straightforward to determine whether or not online rating distributions differ based on local market structure. Clearly, for every firm operating in a competitive environment, knowing the quality spectrum of a market and its position within this spectrum is valuable information. The practice of gathering this knowledge is usually referred to as Competitive Intelligence (CI), a prominent sub-stream of the nascent Business Intelligence and Analytics (BI & A) activities (Chen et al. 2012). As an example, a restaurant with a mean online rating of two stars (on a scale from 1 to 5 stars) that represents the lowest quality rating in a market would garner very different market power compared to a restaurant with a two-star mean rating in another market in which there are several other restaurants with lower than a two-star ratings.[1] Increasingly, online ratings are viewed as a critical source for understanding local market dynamics as businesses are seeking ways to assess their market position by replacing traditional, imperfect, and cumbersome survey market research techniques with a big data-driven approach. While the traditional survey market research techniques might be capable of assessing one's own service quality (Parsuraman et al. 1988), they offer no insight into the service quality of competitors and into the number of competitors in a market. Moreover, classical ways of assessing the competitiveness of a market such as, the Herfindahl-Hirschman Index (HHI), capture the horizontal competition of a market based on the number of competitors, but they fail to directly take service quality into consideration. Online ratings open the avenues to gather CI combining both the assessment of quality provisioning decisions of firms as well as gauging the competition based on the number of competitors, i.e., the market structure. Naturally, it is important to both practitioners and academics alike to analyze the relationship between market structure and online ratings distribution, especially as it pertains to gaining a deeper understanding of the competitive dynamics of available qualities in a market, e.g., the range of the quality spectrum, the dispersion in the distribution, and the changes in the distributional properties with respect to the market structure. In this regard, online rating distributions represent a promising, but still largely underutilized, source of customer data for CI activities (Kaushik 2010).

Delineating the systematic impact of the local market structure on the distributional properties of online ratings is by no means intuitive. Taking the restaurant market as an example, Figure 1 showcases

[1] Online rating websites usually operate on a rating scale from 1 to 5 stars, where reviewers can post ratings for a specific business – for example a restaurant – which is then averaged and reported as the mean rating.

possible relationships between increasing competition and changes in the distribution of ratings. For example, increasing competition may lead to an increase in the market-level average mean rating if quality dispersion is driven mainly by the upper end of the rating range, as depicted in the bottom left corner. On the other hand, if the dispersion would turn out to be driven mainly by lower ratings in a market, this might also reduce the market-level average mean rating, as shown in the bottom right hand corner. However, if the dispersion in the distribution is primarily driven by both ends of the distribution, as in the top right corner, only the range increases whereas the average mean rating remains the same. One might even imagine a case in which neither the range of ratings nor the average mean rating are changing with increasing competition but merely lead to more restaurants being evenly distributed over the already existing ratings, as in the top left corner. Consequently, we empirically study the relationship between market structure and distributional properties of online ratings as well as provide rich insights into the underlying mechanism that drives the changes in online rating distributions with the market structure.

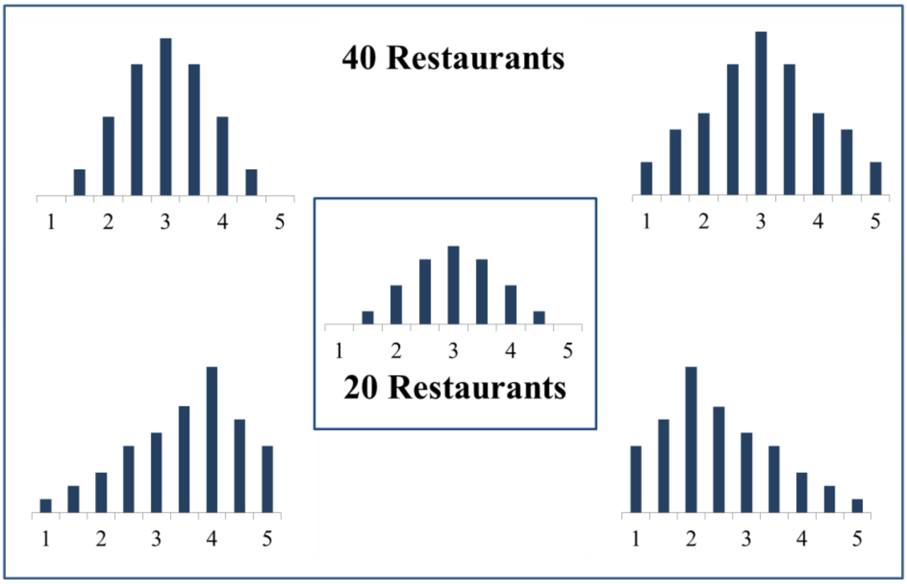

Figure 1: Changes to the Rating Distribution with Increases in Competition

Another way to think about the significance of understanding the impact of market structure on distributional properties of online ratings is to consider gaining insights into the market power of a business. Accordingly, online review management services recommend conducting a competitive analysis by comparing one's own online ratings with those of one's most important competitors (e.g., Shuller 2014). It is reasonable to assume, for example, that a three-star restaurant that faces a lot of competition from restaurants from the lower end of the rating range might price its service differently to a restaurant with the same rating that faces considerable competition from the upper end of the rating range. Therefore, from an internal management perspective, being able to assess the business's competitive position within a market's quality spectrum is important and, more importantly, a central activity of gathering CI. Equally, from an external perspective, it appears that investors are increasingly considering businesses' online ratings as part of their financial decisions (e.g., Huizinga 2015). In this case, considering the online ratings given to a restaurant, one can imagine that prospective investors would seek to assess the position of a restaurant's standing in the competitive environment in the context of a local market. It is conceivable, then, that restaurants situated in a highly contested lower end of the rating distribution may not be deemed as attractive to investors as restaurants in the upper end of the rating range with fewer other businesses offering comparable quality. As such, if the online rating distributions change with the market structure and if a three-star restaurant, for example, is considered in the same way from both an internal and an external perspective irrespective of the market structure, the analyst might misinterpret the competitive position of their business within the quality spectrum and the investor would risk making a flawed decision. While such a flawed decision may raise questions about the appropriateness of crowdsourced data in determining creditworthiness, the real issue may be the failure to understand the dynamics behind rating distributions. Thus, from the standpoint of expanding the appropriate utilization of crowdsourced data, it is also important to investigate the relationship between local market structure and online rating distributions. This becomes even more apparent when observing the differences in the competitive environment of a restaurant with a three-star mean rating in the scenarios shown in Figure 1. A three-star restaurant in a market with 40 restaurants faces a lot more competition from upscale competitors than a three-star restaurant in a market with 20 restaurants as long as the market-level average mean rating increases with competition. By contrast, if the market-level average mean rating decreases with increasing competition, a three-star restaurant in a 40 restaurant market would experience relatively less competition from upscale competitors than in a market with 20 restaurants. Therefore, restaurants with equal mean ratings might have very different market power across markets with differing levels of competitiveness.

As such, a key question for both academics and practitioners is how online ratings can be utilized effectively for deriving competitive market intelligence. As one of the earliest papers in this stream, we pursue the following strategy in our paper: First, we set up an analytical model that serves as the theoretical lens for guiding and better comprehending the main results of our empirical analysis. Second, we present novel evidence from parametric and non-parametric analyses in support of the external and internal validity of Yelp ratings capturing quality provisioning decisions of firms. The external validity is crucial in that if firms want to infer the quality distribution of a local market from the distribution of online ratings, the mean ratings of restaurants need to correctly reflect the qualities in a market. Internal validity of the crowdsourced online rating data is indispensable in order to present a generalizable underlying mechanism for the connection between market structure and online rating distributions. Third, we complement the Yelp data on restaurant ratings and restaurant characteristics for 372 isolated cities in the US with demographics and socio-economic information on these cities. This unique dataset limits the possibility of cross-border city level substitution effects of competition. We then use our dataset to analyze the relationship between market structure and the range of ratings as well as the market-level average ratings. Empirically, we find a decrease in average mean ratings and a broader range of ratings in markets with a greater number of restaurants. Considering the heterogeneity as well as the horizontal differentiation within the restaurant industry and, subsequently, limiting our dataset to, respectively, (1) only Non-Fast-Food Restaurants and (2) only Traditional American Restaurants, we find a strong increase in the absolute magnitude of these effects. While these findings may seem surprising, especially the decrease in average quality in a market with increasing competition, we present extensive analyses that demonstrate the mechanism behind the changes in the distributional properties – higher dispersion in the lower quality spectrum when there are more players in the market.

Crowdsourced Online Ratings and Gathering CI

By analyzing the relationship between online rating distributions and competition, we propose crowdsourced online ratings as a new source of information to gather CI. Not only do we present evidence in support of the internal and external consistency of the dataset itself gathered from Yelp, but in the following we provide a brief discussion on why online ratings also come with some inherent advantages over traditional ways of measuring quality and competition.

First, measuring product or service quality in the field has a long-standing tradition in marketing research and the literature has come up with several survey-based practices (e.g., Parsuraman et al. 1988, Cronin and Taylor 1994). Whereas such survey-based techniques have been frequently used by managers in the past, they primarily serve the purpose of investigating the customer's perception of quality of one's own product or service. In other words, employing survey-based techniques to measure the quality of competitors entails considerable costs and is hardly scalable (Parsuraman et al. 1988), delivering only an incomplete picture of the whole competitive environment a firm is located in and hardly capturing the entirety of competitor firms in the market. Any efforts to increase the scale and scope of the investigation of competitors' qualities often entail third-party involvement incurring substantial costs (e.g., Dahl 2011, Chidi 2013). Consequently, a vital industry branch has emerged for these kind of CI activities, including people sometimes referred to as "competitive spies" who provide various services, ranging from going to a competitor's outlet in order to experience the competitor's service quality to strategically investigating future product launches of rival firms (Chidi 2013). This approach however is costly, sometimes operates at boundaries of the legal and ethical borders, and provides merely sparse quantitative foundation to data-driven decision making.

Second, the industrial organization literature has produced constructs that measure competition as an index of market concentration, based on industry- or market-level boundaries. Empirically, among the most commonly found measures of competition in the literature is the HHI, which is a concentration ratio that captures the distribution of market shares of firms in a certain industry. A major shortcoming of this approach of measuring competition is that it primarily captures horizontal competition and neglects product or service quality as well as the competition pertaining to qualities. Most importantly, its computation requires knowledge on the market shares of the respective firms in the industry and, as a result, the HHI can be typically calculated solely for public firms for which such data is available but not for small or private firms for which this data is mostly inaccessible. Naturally, the HHI is almost never useful in understanding the local market dynamics and, not surprisingly, the literature is vocal about the limited applicability of the HHI (e.g., Ali et al. 2009, Bens et al. 2011, Dedman and Lennox 2009).

By contrast, our approach using crowdsourced online reviews measuring quality and competition sidesteps the major shortcomings of the above mentioned traditional approaches and allows for a big data-driven way to gather CI, especially in local markets. Unlike survey-based techniques, online ratings provide a wealth of information regarding all the relevant competitors at substantially lower costs. While this assessment of quality is based on the assumption that Yelp ratings reflect quality differentiation in a market, we find support for this in existing literature (e.g., Besbes and Scarsini 2015, Chen et al. 2011, Luca 2011) and conduct extensive empirical investigation on this ourselves in sections 4.2 and 4.8. Furthermore, crowdsourced online ratings on Yelp come with industry classifications necessary for the computation of concentration indices. While market shares are not directly available on Yelp, in section 4.7, we calculate the HHI taking the number of ratings a restaurants received as proxies for market share. Our work is the first to leverage online ratings in this way and the results of this analysis are consistent with those of our baseline estimation in section 4.5, suggesting that an HHI using the number of ratings as a substitute for actual market share is able to reflect competition in a market. Thus, using crowdsourced online ratings overcomes weaknesses of traditional quality and competition assessment techniques and, more importantly, broadens the scope of gathering CI in a novel and innovative way.

Contributions

Our paper makes several significant contributions to the literature. First, to the best of our knowledge, we are the first to empirically investigate the relationship between local market structure and online rating distributions. We show that the crowdsourced online reviews can form an internally and externally consistent data source to analyze local market dynamics. Thereby, we contribute to the nascent but growing literature that investigates online customer ratings as a viable source of data to conduct CI activities and market structure analysis (e.g., Zhang et al. 2013, Lee and Bradlow 2011). Our results suggest that competition is an important predictor of the shape of the rating distributions. In particular, an increase in competition leads to an increase in the range of the available quality spectrum in restaurant markets. Our results further indicate that, counterintuitively, the average quality in these markets decreases as the number of competitors increases. Thus, we are the first to employ mean online ratings of businesses in the context of empirically investigating the market competition, as prior studies have mostly focused on the influence of online ratings on product performance. Second, the combination of our carefully collected dataset on isolated markets and Yelp's mean ratings for all restaurants in these isolated markets allows us to observe the competitive environment and quality spectrum in a market as a whole, while ruling out effects of competition from outside these markets. We thus contribute to the literature on the relationship between quality provisioning decisions and market size by considering a sharper set of hypotheses than was typically possible in prior studies (e.g, Olivares and Cachon 2009, Berry and Waldfogel 2010). Furthermore, the richness of our dataset allows us to make claims about the causal relationship between market size and associated online rating distributions. Third, our study shows that Yelp data can be a useful source for capturing local market structure as well as to gauge the quality provisioning decisions of the firms operating the market. Thus, reference disciplines, such as industrial organization, can substantially benefit from the rise of crowdsourced platforms for online ratings. Finally, there is a large economic literature on rationales for cities. Empirical literature in industrial organization has documented the consumption-side benefit of increased product variety in retail (Ellickson 2007), radio (Berry and Waldfogel 1999), restaurants (Berry and Waldfogel 2010), and newspapers (George and Waldfogel 2003) in larger markets. This work implies that consumers face more options and, therefore, achieve greater satisfaction in larger markets and that the welfare benefits of larger markets are driven not only by the number of products but also by the kinds of products available (Berry and Waldfogel 2010). Our work adds to this literature in two ways. First, we provide clear empirical evidence for the increase in the range of the available quality spectrum in larger markets when the marginal cost of increasing quality is non-negligible. Second, we show that there is also a negative side of larger markets for services with non-negligible marginal costs for quality: the average quality in such markets decreases with the number of competitors.

This research has important managerial implications. We demonstrate how business owners might use Yelp to conduct CI, getting an understanding of their own business's quality in relation to the quality of the competitive environment they are operating in. Managers could identify under or over-served quality levels in their respective markets or align their pricing schemes according to their position within that market's quality distribution. This is especially important for restaurant chains that, by their very nature, operate in many different markets. In seeking to provide the same quality throughout all of their restaurants, chains face very different competitive situations in the local markets they serve, depending on the size of that market. Our results help systematize the structural changes in quality distribution of local markets to better adapt to the competitive dynamics of these markets in a chain's decision making and CI activities. This suggests that using online ratings can enable the gathering of CI at a so far unprecedented scale and scope. These insights can also benefit customers who randomly pick a restaurant in an unknown market: in a smaller market, they might be better off choosing a restaurant at random than in a larger market. This signifies the importance of online review websites and other mechanisms to pre-screen the quality of locations to alleviate or prevent this negative side of large markets.

2. Related Literature

Three streams of research are particularly relevant to our study. The first deals with the nascent but important development of BI & A and its sub stream CI. The second analyzes the relationship between local market conditions and the quality provisioning decisions of players in a market. The third examines the effects of different variables on online ratings. In the following, we discuss related work from all of these streams.

There is a growing literature on BI & A, which is often referred to as the techniques, technologies, systems, practices, methodologies, and applications that analyze critical business data (Chen et al. 2012). Naturally, as data is increasingly becoming available to firms through online crowdsourcing, it is of great importance to explore the ways in which these data can be utilized for BI & A purposes. For example, Chen et al. (2012) identify e-commerce and market intelligence as one potential application of BI & A that especially benefits from the emergence of customer-generated web content from social media platforms that enables firms to gather CI by "listening to the voice of the market". Related to this, Lau et al. (2012) design a score card that leverages collective web intelligence to enhance decision making in Mergers and Acquisitions (M & A). They develop a tool that features domain-specific sentiment analysis and business relation mining to operationalize this score card, and present experimental evidence on the usefulness of this tool in decision making, in particular when market information on the target firm is hard to obtain. Brynjolfsson et al. (2011) in turn investigate the impact of data-driven decision making on productivity, financial performance, and market value. Using data collected from a large scale survey, they find that data-driven decision making increases productivity by 5-6%. Furthermore, Zhang et al. (2013) use online reviews of digital cameras from Amazon to compute a product ranking based on text sentiment analysis. They present evidence that the ranking in this comparison network is an important predictor of a product's sales performance. Their study highlights the notion that sales performance not only depends on how the product is perceived in isolation, but also how it is compared against related products in a competitive landscape. In this way, the authors argue that customer reviews constitute a competitive landscape for firms to understand and manipulate and, therefore, the use of customer reviews is crucial for CI. We contribute to this literature by providing evidence suggesting that the easily interpretable mean online rating of a business and the online rating distributions of its local market can be a valuable resource for CI. Moreover, previous studies using online ratings for CI have mainly focused on competitive analysis on the product level (e.g., Zhang et al. 2013, Lee and Bradlow 2011), whilst we enrich this discussion by focusing on the competitive dynamics of local markets.

Several studies have analyzed the relationship between local market structure and the level of quality provided in a market. For example, Cohen and Mazzeo (2010) analyze the retail banking industry and relate the measure of quality to the number of branches in a bank's local network. They find that certain types of financial institutions tend to have fewer branches in the presence of more local competitors. Dick (2007) analyzes the effect of market size on service quality in retail banking and finds that the service quality of banks increases with market size and that dominant banks provide a better quality of service. Olivares and Cachon (2009) study the effect of competition on inventory holdings of General

Motors' dealerships as a measure of service quality provided in a market. They find that, in this industry, competition has a strong positive effect on service quality. More closely related to our study is the work by Berry and Waldfogel (2010). Their study of the newspaper and the restaurant markets finds that while increasing market size leads to higher quality newspapers, larger markets do not offer much additional variety in terms of the number of newspapers, which increases only slowly. For restaurants, they find that the number of available locations rapidly increases with growing market size and that larger markets support a greater number of high quality restaurants. Our study complements this literature in several ways: First, with Berry and Waldfogel (2010) being the only ones to study industries where quality provisioning is primarily tied to marginal costs (e.g., restaurants), our study of online ratings on Yelp enables us to observe the distribution of quality across the entire spectrum (from low to high) in such industries, and thus to analyze the effect of local market structure on the full range of available qualities. Moreover, we are the first to analyze how market structure affects the average quality provided in a market. Finally, our definition of isolated markets allows us to provide a comprehensive identification of the relevant set of competitors for our industry.

The final relevant stream of literature considers different drivers for online rating distributions. For example, Li and Hitt (2008) find that in book ratings, a downward trend over time can be explained by consumers with a high preference for a book tending to post their reviews early. The same selection pattern is also reported by Hu et al. (2009) who find that online ratings on Amazon.com typically follow a J-shape. Another explanation for the disagreement among reviewers is that consumers exhibit a preference for wanting to differ from those who posted earlier reviews (e.g., Moe and Schweidel 2012). A third factor affecting online consumer reviews, according to Muchnik et al. (2013), is social influence through past ratings. In a randomized experiment, the authors investigate how past ratings affect future ratings. They find that negative social influence through past ratings inspires others to correct manipulated ratings, while positive social influence is associated with a 25% increase in final ratings. Another sub-stream of literature investigates whether online reviews reflect the actual quality provisioning decisions of firms. Using online customer reviews for automobiles, Chen et al. (2011) find that customer reviews provide a valid indicator of the objective quality of a car model. Another study by Luca (2011) using Yelp data notes that consumer reviews can be as equally effective at altering demand as, for example, quality disclosures by regulatory bodies (e.g., food hygiene inspections). Using an analytical model, Besbes and Scarsini (2015) investigate whether customers can infer the true quality of a product or service based on the mean rating they observe. They demonstrate that even if the mean rating of a product or service represents the aggregate of subjective opinions of past customers, future customers are still able to infer the underlying true quality of the product if they know that the mean rating is based on the subjective opinions of customers. Most importantly, the relative position of a product or a service compared to its competitors is usually preserved in mean online ratings. Finally, Dellarocas and Narayan (2006) investigate factors that influence the propensity to post an online rating for a recently watched movie. They find that marketing expenditure and disagreement among critics positively affect this propensity and that consumers are more likely to post ratings for both very good and very bad movies. We contribute to this literature by showing how local market structure impacts the whole spectrum of online rating distributions of an offline market. Unlike the extant literature, which is primarily focused on the influence of online ratings on product performance, this study is pioneering the use of online ratings to understand the competitive landscape of a business. Furthermore, we present evidence in support of the external validity of mean online ratings' ability to reflect quality by conducting a series of correlation analyses between Yelp ratings and Zagat ratings. We also provide evidence in support of the internal validity of the underlying data generation mechanism for Yelp and its ratings distributions. Accordingly, we open up an avenue for future research on how local bricks and mortar store owners can employ online ratings to inform long-term offline decisions on quality provision.

3. Theoretical Framework

To gain a richer understanding of the theoretical mechanism underlying the relationship between market structure and distributional properties of online ratings, we adapt the classical lens of an oligopoly. The insights obtained from this analytical framework guides our empirical investigation and provides a platform for better comprehending the results.

Our model considers markets with S potential consumers, j firms that offer vertically differentiated services. Each consumer has a demand of one unit for the service. Consumers always prefer a higher service quality but they differ in their willingness to pay for it. This willingness to pay is captured by the consumer type variable δ_i which is uniformly distributed between zero and one. The utility that consumer i derives from using a service j is $u_i = \delta_i q_j - p_j$ and not using a service results in a utility of $u_i = 0$. On the demand side, therefore, our model resembles the ones presented in Shaked and Sutton (1982, 1983) and Jones and Mendelson (2011).

Firms can differentiate their service offering by choosing different levels of service quality q, and each firm can produce exactly one service quality. All firm decisions are simultaneous. In line with Shaked and Sutton (1982), we assume that marginal production costs increase quadratically when firms increase service quality.[2] In addition, firms have to pay a fixed entry cost of F to enter a market which is independent of q. Each firm faces the following decision problem: First, firms decide whether to enter the market and pay the fixed market entry cost of F at the time of entering the market. Second, firms simultaneously choose their service qualities. Third, firms decide on a price for their own services after having observed the quality decisions made by other firms.

If two or more firms choose the same service quality, the only equilibrium is to price at the unit cost, which would result in negative profits for all firms with similar service quality. In such cases, these

[2] We consider a quadratic increase of marginal costs for reasons of simplicity. All of our results hold for marginal production costs of q^x where $x > 1$.

firms would then prefer not to enter the market. Moreover, firms that make negative profits in equilibrium would also prefer not to enter the market. Thus, in line with Jones and Mendelson (2011), we only consider product equilibria where all firms offer different service qualities and all firms make a profit.

Without loss of generality, we assume that firm 1 provides the highest quality, firm 2 the second highest quality and so on, i.e., $q_1 > q_2 > ... > q_N$. In a market with N firms, the demand d_j for firm j is $d_1 = 1 - \left(\frac{p_2 - p_1}{q_2 - q_1}\right)$ if $j = 1$, $d_j = (\frac{p_{j+1} - p_j}{q_{j+1} - q_j}) + (\frac{p_j - p_{j-1}}{q_j - q_{j-1}})$ if $j \in [2, N-1]$, and $d_N = \left(\frac{p_N - p_{N-1}}{q_N - q_{N-1}}\right) - \frac{p_N}{q_N}$ if $j = N$. Firms maximize their profits given by $\pi_j = d_j S\left(p_j - q_j^2\right) - F$.

In a market that allows for a positive profit for up to one firm, demand and profit for this firm are $d_1 = 1 - \left(\frac{p_1}{q_1}\right)$ and $\pi_1 = (1 - \left(\frac{p_1}{q_1}\right))S(p_1 - q_1^2) - F$ leading to $p_1^* = \frac{1}{2} q_1(1 + q_1) = \frac{2}{9}$, $q_1^* = 1/3$, $d_1^* = \frac{1}{3}$, and $\pi_1^* = \frac{1}{27} S - F$. As there is only one available quality in the market, each consumer experiences a quality of 1/3. Note that only the entry decision depends on the market size S or the market entry costs F; both optimal price and quality are independent of S and F. In particular, the firm only decides to enter if $\pi_1^* \geq 0$ that is $S \geq 27F$.

In a duopoly market that allows for a positive profit for both firms, demand for the first firm is $d_1 = 1 - \left(\frac{p_2 - p_1}{q_2 - q_1}\right)$ and for the second $d_2 = \left(\frac{p_2 - p_1}{q_2 - q_1}\right) - \frac{p_2}{q_2}$. Profits are $\pi_1 = d_1 S(p_1 - q_1^2) - F$ and $\pi_2 = d_2 S(p_1 - q_1^2) - F$ leading to $p_1^* = 0.2267$, $p_2^* = 0.0750$, $q_1^* = 0.4098$, $q_2^* = 0.1994$, $d_1^* = 0.2792$, $d_2^* = 0.3445$, $\pi_1^* = 0.0164S - F$, $\pi_2^* = 0.0121S - F$. Again, each firm's decision to enter the market depends on S and F. In particular, to enter the market a high quality firm needs $S \geq 60.98F$ and a low quality firm needs $S \geq 82.64F$.

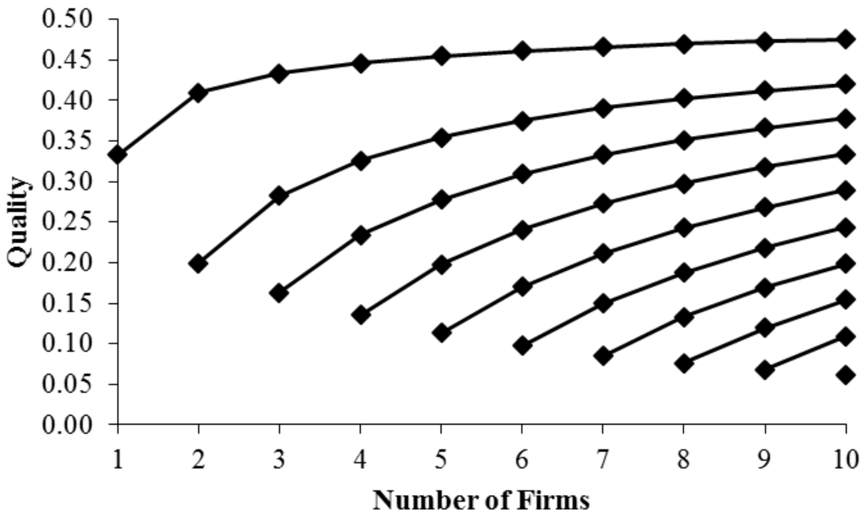

Figure 2: Quality as a Function of the Number of Firms in a Market

Note that the threshold for the high quality firm drops to $S \geq 27F$ if the low quality firm decides not to enter. Comparing the qualities provided in a duopoly market with the quality of a monopoly market shows that firm 1 in the duopoly provides a higher and firm 2 a lower quality than the monopolist in the monopoly market. We can also compute the market-level average quality $\bar{q} = 0.3046$, which is slightly lower than the one found in the monopoly market. It is difficult, however, to derive closed-form solutions for oligopolies with more than two firms. Thus, in line with Jones and Mendelson (2011), we present numerical results for oligopolies with up to ten firms in the following. Figure 2 shows firms' quality choices as a function of the number of firms in a market. The higher the number of firms that enter the market, the more these qualities tend to span the available product space. In other words, the more firms are in a market, the broader the quality spectrum.[3] Note that market size only affects the number of players in a market but has no direct impact on the quality choices of firms, and that the increase in the available quality spectrum decreases with an increasing number of firms in a market.

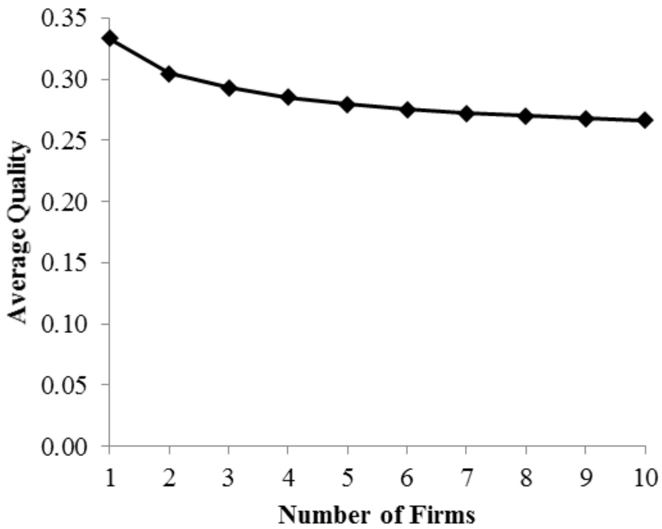

Figure 3: Average Quality as a Function of the Number of Firms in a Market

We can also compute the average of these qualities for a different number of competitors in a market. Figure 3 presents the average in markets with up to 10 firms, showing a decrease in the average quality when the number of firms in a market increases. The magnitude of this decrease is strongest when comparing monopoly and duopoly markets, and decreases with an increasing number of firms in a market. We summarize these observations in our predictions 1 and 2.[4]

Prediction 1: The available range of qualities in a market increases with the number of firms in a market.

[3] Recall that firms decide simultaneously about their quality provisioning once they decide to enter the market, therefore all the different market structures depicted in Figure 2 are independent of each other, e.g., 'firm 2' in a market with total 3 players is not the same 'firm 2' in a market with total 4 players.
[4] In the appendix, we present extensions of the theoretical framework that allows for horizontal differentiation in addition to the vertical differentiation. The key insights remain unchanged.

Prediction 2: The average quality in a market decreases with the number of firms in a market.

In our empirical analyses, we carefully delineate the characteristics of dispersion and shifts in online ratings distributions with changes in the market structure. Interestingly, our empirical results are in line with predictions 1 and 2.

4. Empirical Investigation of Restaurant Ratings on Yelp

Online review websites are a major feature of the current digital era. While in the past decades, review websites experienced tremendous growth, Yelp is probably the most widely recognized platform for recording crowdsourced online ratings for all kinds of businesses. Yelp makes it possible to award online ratings for most businesses, including restaurants, health care, and financial services, amongst others. With a current market value of approximately $843 million (Yahoo Finance 2016), Yelp constitutes a global cornerstone in online consumer conversations. About 75% of Yelp's estimated business value is made up of its advertising products for local businesses, and with an estimated addressable global market value of $130 billion, the company has plenty of growth potential, attracting ever more businesses that are likely to be rated by consumers (Forbes 2015). Unsurprisingly, then, online ratings provided by Yelp have garnered much academic interest in recent years. Luca (2011), for example, finds that a 1-star increase in a restaurant's mean rating increases its revenue by 5-9%, and that customers do use Yelp ratings to infer accurate product quality. Scholars and practitioners are keen to understand the factors that influence these ratings and a significant literature has emerged on the different facets of online review systems (e.g., Dellarocas 2003, Chevalier and Mayzlin 2006). This interest is also in line with the increasing focus on deriving big data-driven competitive intelligence. This underpins our choice of Yelp as data source in our empirical study, and of restaurants in particular, as this is Yelp's hallmark. Furthermore, it allows us to demonstrate how crowdsourced data can be effectively leveraged to generate competitive market intelligence.

In addition to the above reasons, the restaurant industry offers three production features that are crucial to our theoretical model: their marginal production costs typically increase faster than linearly when firms increase service quality because the largest part of a restaurant's quality is linked to its marginal costs (the cost of delivering service quality, the cost of ingredients, and the cost of food preparation). Second, firms vertically differentiate their service offerings, as increasing service quality in restaurants typically increases marginal costs a great deal faster than linearly (Berry and Waldfogel 2010). Essentially, while it is comparably cheap to serve a simple meal, the costs of ingredients and the time needed for both food preparation and service quality substantially increase for higher quality food (e.g., the quality difference between a high-quality French restaurant and a neighborhood pizza outlet[5]) – an

[5] Horizontal differentiation may also play an important role in the restaurant industry. Even the authors of this paper have different tastes and, therefore, prefer different restaurants when choosing from a set of restaurants with comparable qualities. To reduce potential effects of horizontal differentiation, we limit our analysis to restaurants from specific categories in section 4.5.

increase which we expect to be non-linear. Finally, all other factors being equal, we expect that, while most people would prefer higher quality restaurants, their willingness to pay for higher quality is substantially heterogeneous among consumers.

There may be concerns regarding a selection effect concerning the motivations of individuals who post reviews (e.g., Hu et al. 2009) or the issue of fake reviews posted by location owners (e.g., Mayzlin 2006).[6] Nevertheless, these are mitigated as follows: first, the self-selection of those posting reviews should apply to all locations across our sample and be independent of market size. We still address this concern in our robustness checks. Second, we have no reason to believe that fake reviews are more prevalent in some markets than in others. Moreover, online review websites such as Yelp have recently started to identify and tag potential fake reviews as "not recommended" (Mukherjee et al. 2013). Naturally, we do not include these in our analysis. Even though one might assume that Yelp's filter probably fails to identify some fake reviews, as well as possibly misidentifying some genuine reviews as fakes, a recent study has underlined that Yelp's filtering algorithm is overall rather effective in identifying fakes and, thus, represents a conservative measure to exclude suspicious reviews (Luca and Zervas 2016). In sum, all the evidences we have consulted from the extant literature and the comprehensive analyses we have done give us confidence in the Yelp mean rating's ability to reflect the quality of a restaurant in a market. We provide more evidence in support of this in section 4.2.

4.1. Data

In line with Breshnahan and Reiss (1991) and Olivares and Cachon (2009), after having first defined isolated markets, we are able to accurately measure their characteristics as well as the number of firms and the level of competition prevalent in these specific markets. To limit the issue of competitors from neighboring towns that might affect the behavior of the players in our focal locations, we need to find towns that are isolated from each other. We consider a town as isolated if it has no significant neighbors within 20 miles, with 'significant' defined by a population with more than 2,000 inhabitants. Following this criterion and by utilizing a dataset provided by Collard-Wexler (2006), we have identified 372 isolated towns that we have included in our sample. Using a customized web crawler, we collected a broad variety of information on all of these markets from www.city-data.com, in May 2014. In particular, we obtained the following population data for each market: number of inhabitants according to the 2012 Census (POP), and population change since 2000 (POPCHNG). Moreover, we collected the following demographic variables for each market: median age (MEDAGE), median income (MEDINC), median house value (MEDHVAL), median gross rent (MEDRENT), unemployment rate (UNEMP), the

[6] If the motivation to post reviews is altruism – to help other consumer discerns the quality of the restaurant, then subjectivity won't be an issue. If, however, the intention is to express themselves and declare what they are, what they feel, what they think, then the consumers may not express themselves objectively. Besbes and Scarsini (2015), through a rich comprehensive analytical framework, demonstrate that even the in the case of subjective expressions, the mean ratings – average of individual ratings – converge to the underlying objective quality under a very pragmatic set of assumptions.

surrounding land area (LANDAREA), the fraction of people living in urban areas (URBAN), cost of living (COSTOFLIV), percentage of Hispanics (PERCHISP), percentage of American Indians (PERCINDIAN), percentage of Asians (PERCASIAN), and percentage of Blacks (PERCBLACK) in each market.

Table 1: Market-level Summary Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
POP (in 10,000 inhabitants)	371	1.11	1.08	0.26	11.89
POPCHANGE (since 2000)	372	0.04	0.14	-0.51	0.79
MEDAGE	372	36.38	6.19	20.60	58.20
MEDINC (in $10,000)	372	37.03	10.56	18.52	89.73
MEDHVAL (in $10,000)	372	111.78	64.74	38.09	566.80
MEDRENT (in $100)	371	6.05	1.53	3.69	17.99
UNEMP (in % in 2014)	372	7.22	3.13	0.70	20.40
LANDAREA (in 100 square miles)	372	0.26	2.05	0.02	28.74
URBAN (% of total population)	355	0.95	0.10	0.00	1.00
COSTOFLIV (US cost of living index)	372	84.11	10.06	76.10	194.30
PERCHISP	371	0.15	0.20	0.00	0.97
PERCINDIAN	371	0.03	0.13	0.00	1.00
PERCASIAN	366	0.01	0.03	0.00	0.38
PERCBLACK	371	0.09	0.17	0.00	0.83
NUMLOC	372	40.15	34.85	0.00	289.00
NUMRATE	370	249.66	573.92	0.00	5676.00
MKTAVGRATE	369	3.44	0.39	1.90	5.00
MINRATE	369	1.45	0.73	1.00	5.00
MAXRATE	369	4.85	0.36	2.50	5.00
RANGE	369	3.41	0.88	0.00	4.00
LOCPOP (number of locations per 10,000 inhabitants)	371	40.05	24.55	0.00	235.19

Finally, using web crawlers in May 2014, we collected the number of locations (NUMLOC), the location-level mean ratings (the average of all individual ratings for a restaurant), and the total number of posted ratings (NUMRATE) on Yelp.com for all restaurants within these markets. We computed the average of the mean ratings in a market (MKTAVGRATE), the minimum (MINRATE), the maximum (MAXRATE), and the range (RANGE) of these mean ratings for each market. Moreover, we computed the number of locations per 10,000 inhabitants (LOCPOP)[7] as an additional control for competiveness in the market. Table 1 shows the summary statistics of these variables. We extracted the data on market information from www.city-data.com as well as the corresponding data from Yelp.com on exactly the same markets at another point in time in November 2015 in order to construct a panel data set. This data is used in further analyses in section 4.7.

[7] We thank an anonymous reviewer for suggesting this additional control variable.

4.2 Yelp Ratings and Service Qualities in Markets

Although extant literature directly and indirectly advocates that Yelp's mean ratings reflect the intrinsic qualities in a market (e.g., Chen et al. 2011, Luca 2011), it is still beneficial to systematically examine

the internal and external validity of the Yelp data generation process. In other words, it is desirable to demonstrate that the Yelp ratings follow an underlying consistent mechanism, i.e., distributions are similar across cities with similar market structure. Furthermore, the external validity of Yelp ratings provides an important reassurance that the quality differentiation in the market is properly reflected. Employing parametric and non-parametric approaches, we incorporate different strategies to provide evidence confirming the internal and external validity of rating data from Yelp.com. In particular, we show that Yelp ratings are highly correlated with Zagat ratings. Using a binning strategy further enables us to show that cities with a similar number of restaurant locations exhibit very similar rating distributions, indicating that the underlying data generation mechanism is systematic and consistent.

In June-July, 2015, we collected data on curated ratings for 31 major US cities with a total of 5536 restaurants from Zagat.com, a company that looks back on a long-standing history of providing restaurant ratings made by enlisted critics.[8] Conducting a series of correlation analysis and regressions on a matched sample of restaurants with Yelp and Zagat ratings reveals that Yelp and Zagat mean ratings exhibit a significantly strong and positive correlation, underlining that Zagat and Yelp ratings reflect the intrinsic qualities of restaurants in the same way.

Correlations between Yelp ratings and Zagat's food ratings are on average 0.4683 and significant ($p<0.0001$). The strength of the correlations, when considered on a city by city level, ranges from a minimum of $\rho= 0.2858$ to $\rho=0.7727$ and remains significant throughout. For all the 31 cities in our matched sample, the comparatively low value of 0.2858 for San Diego is the only value below 0.3, and the majority of cities – comprising 4278 out of the total 5536 restaurants in our matched sample – are above 0.4. We also conducted separate correlations analyses for different cuisine types such as American, Barbeque, Burger and Chinese cuisine and found correlations to range from $\rho=0.6992$ for Barbeque to $\rho=0.3565$ for Chinese cuisine. In order to account for differences in rating scales (Zagat ratings are given on a scale from 1 to 30 and Yelp ratings on a scale from 1 to 5) we transposed Zagat ratings onto a scale from 1 to 5 and repeated the same correlation analysis. The results remain qualitatively and, by and large, quantitatively the same. To verify that the strong correlations are not driven by a higher number of reviews, we calculated separate correlations for restaurants with different numbers of reviews. Correlations range from $\rho = 0.2151$ for restaurants with fewer than 25 reviews to $\rho = 0.4461$ for restaurants with between 900 and 1100 reviews. However, the group of restaurants with

[8] We utilized undergraduate research assistance to manually match every restaurant between Yelp and Zagat.

fewer than 25 reviews comprises only 210 restaurant out of all 5536 restaurants, this group being the only group of restaurants with a correlation value lower than 0.3.

Additionally, we conducted a series of regressions of Yelp ratings on Zagat's food ratings to show that Zagat's food rating is a strong predictor of Yelp ratings. While Zagat features scores for food, service, and decor, the food score is the strongest predictor for Yelp ratings. Notably, the Yelp ratings predicted by Zagat's ratings show a correlation very similar to the correlation between Yelp and Zagat ratings. In addition to that, the residuals of the regression of Yelp ratings on Zagat ratings are not correlated with Zagat ratings.

Second, we investigate the internal consistency of our Yelp data on isolated markets by setting up bins of cities according to the number of restaurants in a city. If the dataset is internally consistent, online rating distributions should be similar for similar market structures across cities. In contrast, if the group of cities with, say, 20 restaurants shows that the distributions do not follow a consistent pattern – e.g., some are left-skewed, some are right-skewed, and others are normal distributions – this would indicate that the data generation process on Yelp does not follow a consistent underlying mechanism. In other words, a theoretically consistent data generation mechanism would suggest that similar cities would exhibit similar rating distributions instead of a random collection of different distributions. In order to systematically assess the restaurant-level mean rating distributions of cities within a bin, we first categorized cities into bins of 10 restaurants, i.e., we have one bin with cities that have 11 to 20 restaurants, cities of the next bin have 21 to 30 restaurants, and so forth. We designed a bootstrapping program that randomly draws two cities from a bin, compares the cities' distributions, puts the two cities back into the bin and repeats this procedure 2000 times. Analogous to Hu et al. (2009), who used a Kolmogorov-Smirnov Test to analyze the rating distribution of products on Amazon.com, we employ a Kolmogorov-Smirnov Test to determine if distributions of the two cities are equal. If the market-level rating distributions are consistent within the bins, we would expect to find substantial shares of equal distributions per bin.

Table 2: Results for Distribution Tests within City Bins

BINS	Bin 20	Bin 30	Bin 40	Bin 50	Bin 60	Bin 70	Bin 80	Bin 90
% Equal	96.45	96.45	96.05	94.65	86.75	92.7	94.75	95.15
# Cities	79	81	56	33	30	23	13	9

Note: Equal means equal distributions at a 95% confidence level. Bin 20, for example, contains cities with 10 to 20 locations.

The line "%Equal" of Table 2 reports the percentage of the 2000 Kolmogorov-Smirnov Tests performed by the bootstrapping program that yielded an insignificant difference in the distributions (on a 95% confidence level). The percentages of equal distributions within the bins range from 86.75% in bin 60 to 96.45% in bin 20 and bin 30. While the lowest value of 86.75% is quite substantial already, the highest consistency of distributions within a bin becomes apparent for the two bins with the most cities. In sum, we find substantial internal consistency of rating distributions within bins across our whole dataset. We present an example of sets of distributions in Figure 4.

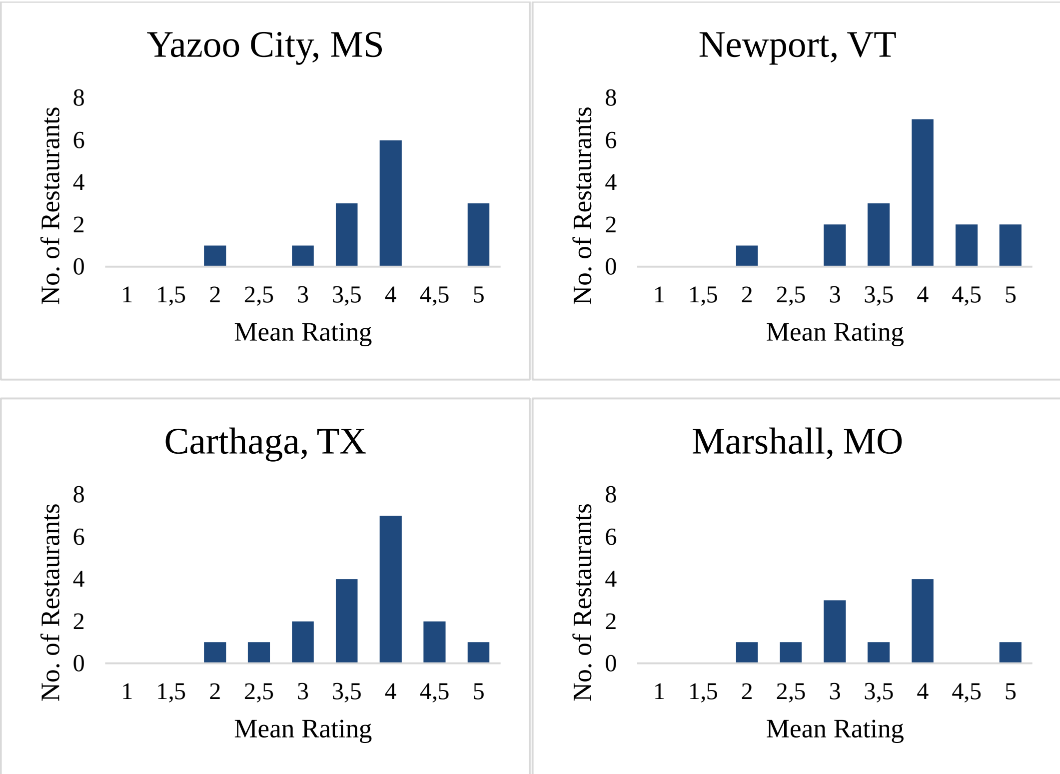

Figure 4: Sample of Distributions from Bin 30

Reassuringly, the results from our systematic analysis above gives us further confidence in the mean ratings of Yelp in effectively representing the quality spectrum in a market. This, coupled with our focus on isolated markets that limits the probability for cross-market substitution effects, provide an ideal foundation for the identification of the effect that the competitiveness of a market has on the distributional properties of online ratings – i.e., the range of mean ratings and the market-average mean rating.

4.3. Empirical Model

To investigate the effect of the local market structure on the distributional properties of online ratings in the restaurant market, we use the following equation:

$$Y_m = \beta_0 + \beta_1 NUMLOC_m + \beta_2 NUMLOC_m * NUMLOC_m + \varsigma X_m + \epsilon_m.$$

where the subscript represents market m, X_m is a vector of market-level control variables, and ϵ_m is the random error term.[9] The dependent variables in this equation are the range of the mean ratings in a market (RANGE) and the average of the mean ratings in a market (MKTAVGRATE), respectively. The

[9] The vector of controls comprises MEDAGE, MEDINC, MEDHVAL, MEDRENT, UNEMP, LANDAREA, URBAN, COSTOFLIV, PERCHISP, PERCINDIAN, PERCASIAN, PERCBLACK and LOCPOP.

key variables of interest in this equation are the number of locations (NUMLOC) and the number of locations squared (NUMLOC*NUMLOC) as measures of competition in a market. Our first theoretical prediction suggests that NUMLOC should positively, and NUMLOC*NUMLOC negatively, affect the range of the mean ratings in a market and that the coefficient on NUMLOC should be significantly larger in absolute magnitude than the coefficient on NUMLOC*NUMLOC. We include NUMLOC*NUMLOC in the empirical analysis because our theoretical model suggests that the positive effect of NUMLOC on RANGE decreases in magnitude with increasing NUMLOC, which should be captured by a negative coefficient for NUMLOC*NUMLOC. Moreover, the convex curve obtained in Figure 3 suggests that NUMLOC should negatively, and NUMLOC*NUMLOC positively, affect the market-level average mean rating, and the negative effect of NUMLOC on MKTAVGRATE decreases in magnitude with increasing NUMLOC. Again, the coefficient on NUMLOC should be significantly larger in absolute magnitude than the coefficient on NUMLOC*NUMLOC.

4.4. Identification

Naturally, it is challenging to establish a causal relationship between, on the one hand, competition measured as the number of locations and, on the other, the range of mean ratings or the market-level average mean rating. Even if we control for a wide range of market-level variables, other unobserved variables may still exist that simultaneously play a role in quality provisioning decisions and the number of locations in a market. For example, a market in which a well-known tourist attraction draws in a substantial number of one-time visitors, the number of restaurants increases that specifically target these one-off customers. However, it is reasonable to assume that these restaurants do not aim for repeat customers and, therefore, provide a comparably lower quality and, consequently, receive less favorable ratings. As a result, we need an exogenous source of variation in the number of locations in a specific market to properly identify the effect of the number of locations on the market-level average mean rating and on the range of mean ratings in a market. In line with Brynjolfsson et al. (2009), in August 2014 we collected data from the US Census ZIP code-level business pattern data (www.census.gov/econ/cbp/) on the historical number of locations in a market ten years ago, in 2004, and use it to further instrument for the current number of locations.[10] As the number of locations in a market should be sticky, at least partially, the number of locations in 2004 (NUMLOC_2004) should be correlated with the number of locations at present (2014). At the same time, we do not expect to find a correlation between the number of locations ten years ago and the current market-level average mean rating, or the range of mean ratings in a market (it is worth noting that, founded in October 2004 based on an email system, Yelp did not start to function properly until late 2005). Moreover, we collected state-level data on the "Small Business Policy Index 2014" (SBPI) from www.sbecouncil.org in December 2015. This index measures the environment in each state as to how policies impact entrepreneurship, investment, and small business

[10] We used the mutually exclusive 2002 NAICS code numbers 722110 (Full-Service Restaurants), 722211 (Limited-Service Restaurants) and 722213 (Snack and Nonalcoholic Beverage Bars) and downloaded the number of businesses with these code numbers for each zip code and, when necessary, aggregated them at a city level.

growth. It is reasonable to expect that the economic state-level environment for small businesses is significantly correlated with the number of locations in a market, as it either promotes or demotes entrepreneurs to found new businesses. In this way, the Index is an appealing source of heterogeneous exogenous variation which we do not expect to be correlated with market-level quality, because the difficulty of setting up a business should not influence the quality provided by it. Therefore, this instrument is likely to fulfill the exclusion criterion. We are also able to separate out alternative effects of competition from the competition effect due to mere increases in the number of locations (NUMLOC), by controlling for the number of locations per 10,000 inhabitants (LOCPOP). This mitigates the concern that cities with an identical number of locations can have different levels of competitiveness due to different population sizes, as a city with 20 restaurants and 10,000 inhabitants might be more competitive than a city with 20 restaurants and 15,000 inhabitants.

4.5. Empirical Results

Table 3 presents the regression estimates of our model with RANGE and MKTAVGRATE respectively, as the dependent variable. The first column of this table shows OLS estimates of our estimation equation with RANGE as the dependent variable. The second column shows 2SLS estimates of the same equation. Column (3) and (4) of this table show OLS and 2SLS estimates with MKTAVGRATE as the dependent variable. For RANGE, as the dependent variable our coefficients of interest, i.e., the β_1s and β_2s, show the signs and significances consistent with the insights from our supplemental theoretical model. In particular, we find a significant positive effect on NUMLOC and a significant negative effect on NUMLOC*NUMLOC for both models.[11] The magnitudes of the coefficients suggest that competition significantly affects the range of available ratings in a market and that this effect decreases gradually with the increasing number of competitors in the market. In particular, for NUMLOC, we have estimated coefficients between 0.02395 (s.e.= 0.0028) for the OLS model and 0.02511 (s.e.= 0.00337) for the 2SLS model. [12] Moreover, for NUMLOC*NUMLOC, the estimated coefficients are -0.00008 (s.e.= 0.00001) for the OLS model and -0.00009 (s.e.= 0.00002) for the 2SLS model. This relates to an increase in the range of mean ratings in a market between 0.02371 (OLS) and 0.02484 (2SLS) if we go from a monopoly to a duopoly. If we compare a market with 20 competitors (the 25th percentile) to a market with 53 competitors (the 75th percentile) the range of mean ratings is between

[11] There might be concerns about skewed estimates, which could be solved by log-transforming our explanatory variable NUMLOC. Therefore, we also estimate an alternative model specification, for all main results, in which NUMLOC and NUMLOC*NUMLOC is replaced by logged values of NUMLOC and we find the results to be qualitatively unchanged. We also ran estimations with logged values of the control variables but the results remain qualitatively unchanged.

[12] Considering the strength of our instruments, a Shea's adjusted partial R^2 of 0.735 for NUMLOC and 0.7326 for NUMLOC*NULOC for the first stage regressions indicates a strong correlation between our instruments and the number of locations in a market (Shea 1997). This is also reflected by the very high F statistics (86.41 for NUMLOC and 47.51 for NUMLOC*NUMLOC) for the joint significance of our instruments, which is well above any critical value for weak instruments (Stock and Yogo 2005). Regarding the correlation of our instruments with the error term, Wooldridge's robust score test of overidentifying shows an insignificant test-statistic implying that our instruments are uncorrelated with the error term. We find qualitatively the same pattern for all following cases. Results of the first stage regressions can be found in the appendix.

0.59763 and 0.61182 smaller in the market with 20 competitors. Comparing the OLS and 2SLS estimates shows only a small difference in magnitude between the models, and neither Wooldridge's score test nor the regression-based test (Wooldridge 1995) can reject the null that NUMLOC and NUMLOC*NUMLOC are exogenous.

Table 3: Regression Results for All Restaurants

	OLS	2SLS IV 2nd Stage	OLS	2SLS IV 2nd Stage
VARIABLES	RANGE	RANGE	MKTAVGRATE	MKTAVGRATE
NUMLOC	**0.02395*** **	**0.02511*** **	**-0.00322** **	**-0.00405** **
	(0.00279)	**(0.00337)**	**(0.00144)**	**(0.00174)**
NUMLOC²	**-0.00008*** **	**-0.00009*** **	**0.00001* **	**0.00001** **
	(0.00001)	**(0.00002)**	**(0.00001)**	**(0.00001)**
MEDAGE	-0.01479**	-0.01437**	0.00766**	0.00842**
	(0.0064)	(0.00642)	(0.00349)	(0.0034)
MEDINC	0.00981*	0.01088*	-0.00558	-0.00879***
	(0.00558)	(0.00582)	(0.00409)	(0.00293)
MEDHVAL	-0.00298***	-0.00305***	0.00079*	0.00089*
	(0.00088)	(0.00089)	(0.00046)	(0.00046)
MEDRENT	-0.0272	-0.03705	0.04191	0.06292**
	(0.0492)	(0.04915)	(0.03048)	(0.02506)
UNEMP	-0.02755	-0.02607	0.01139	0.00575
	(0.02142)	(0.02157)	(0.01018)	(0.00871)
LANDAREA	0.01983	0.01935	-0.00051	-0.00074
	(0.015)	(0.01568)	(0.00405)	(0.00424)
URBAN	1.05271**	1.03653**	0.62253**	0.55356**
	(0.50607)	(0.50634)	(0.26483)	(0.24093)
COSTOFLIV	0.01927**	0.01987**	-0.00401	-0.00528
	(0.00969)	(0.0097)	(0.0037)	(0.00349)
PERCHISP	-0.25407	-0.2516	-0.04015	-0.04192
	(0.24333)	(0.24236)	(0.12107)	(0.12116)
PERCINDIAN	-0.14097	-0.11901	-0.1955	-0.14501
	(0.72958)	(0.72168)	(0.36382)	(0.3534)
PERCASIAN	-1.87797	-1.8273	1.6182	1.50024
	(1.45797)	(1.45021)	(1.18091)	(1.10708)
PERCBLACK	0.12656	0.12178	-0.17455	-0.16452
	(0.25415)	(0.25404)	(0.14806)	(0.14678)
LOCPOP	0.00499***	0.00487**	-0.00082	-0.00116
	(0.00185)	(0.0019)	(0.00102)	(0.00093)
Constant	0.82014	0.75167	2.83458***	3.04406***
	(1.10178)	(1.10674)	(0.47719)	(0.4348)
N	348	347	348	347
R^2	0.31473	0.31493	0.06777	0.0257

*** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

For MKTAVGRATE as the dependent variable, the OLS and the 2SLS estimates of β_1 and β_2 are significant and negative for the former, and positive for the latter. The magnitudes of the β_1 coefficients are -0.00322 (s.e.= 0.00144) (OLS) and -0.00405 (s.e.= 0.00174) (2SLS), respectively. For the β_2 estimates we get 0.00001 (s.e.= 0.00001) (OLS) and 0.00001 (s.e.= 0.00001) (2SLS). Confirming the shape of the theoretically predicted relationship between average quality and the number of competitors depicted in Figure 3, our results suggest that increasing the number of competitors in a market from one to two has the effect of lowering the market-level average mean rating by between -0.00319 or -0.00402 respectively and that the magnitude of this decline decreases with an increasing number of competitors. Considering the combined magnitude of the coefficients, the effect of NUMLOC on MKTAVGRATE is small but still economically significant. For instance, ceteris paribus, the average rating in a market with 20 locations (25th percentile) is between 0.082 and 0.106 higher than in a market with 53 locations (75th percentile). Comparing the OLS and 2SLS estimates shows that the OLS point estimate for NUMLOC is slightly less negative than the 2SLS estimate. However, similar to the models with RANGE as the dependent variable, we cannot reject the null hypothesis that NUMLOC is exogenous at any reasonable level.

In sum, our empirical results show that an increase in the number of locations in a market entails a significant increase in the range of mean ratings at the same time as a significant decrease in the market-level average mean rating. These effects are, however, economically less pronounced than expected. One potential explanation for this result is that the restaurant industry as a whole is too broad to provide a good test for the relationship between market structure and the properties of online rating distributions. For example, the industry comprises anything from fast-food restaurants like Mc Donald's or Burger King to "classical" restaurants like Red Lobsters or Applebees and it is questionable whether these restaurants are competing in the same sphere. If, indeed, they operate in different markets, the question arises whether vertical differentiation, heterogeneous consumer preferences for service quality, and marginal costs that quadratically increase with service quality apply to fast food outlets. More importantly, restaurants can also horizontally differentiate their service offerings. If a small market supports just two restaurants, one is very unlikely to observe, for example, two vertically differentiated Traditional American Restaurants. It is much more likely that the second location is horizontally differentiated by offering Chinese cuisine, for example.[13] These two restaurants from different categories would both be monopolists in their specific cuisine and, therefore, we would not expect to observe a broader range of mean ratings and a lower market-level average mean rating in this context.[14]

[13] This intuition is supported by Neven and Thisse (1989) who find that when the range of qualities a market can accommodate is relatively small as compared to the range of variety it can accommodate, then in equilibrium two firms select a pair of products that are maximally differentiated in the horizontal dimension but minimally differentiated along the vertical dimension.

[14] In fact, we see a strong effect of the number of locations on horizontal differentiation. Regressing the number of locations on the number of cuisines as a measure of horizontal differentiation in a market reveals a strong and significant positive relationship.

4.6. Evidence from Non-Fast-Food Restaurants and Traditional American Restaurants

To account for the potential heterogeneity within the restaurant market, and to limit the impact of horizontal differentiation, we re-estimate our models on a subsample from our original dataset including only restaurants which we characterize as Non-Fast-Food Restaurants, as well as a subsample of restaurants from one specific category, namely Traditional American Restaurants. To build these subsamples, we relied on categorizations provided by Yelp. We assign restaurants that were categorized as "American (traditional)", "American (new)", "Italian", "Steak", and "Seafood" to the subsample of Non-Fast-Food Restaurants, and restaurants that were categorized as "American (traditional)" to the subsample of Traditional American Restaurants.

Tables 4 (Non-Fast-Food Restaurants) and 5 (Traditional American Restaurants) display the estimation results for our models with RANGE and MKTAVGRATE respectively, as the dependent variable. The first columns of these tables show OLS estimates of our estimation equation with RANGE as the dependent variables. The second columns show 2SLS estimates of the same equations. Columns (3) and (4) of these tables show OLS and 2SLS estimates with MKTAVGRATE as the dependent variables.

The estimated β_1 coefficients for the effect of NUMLOC on RANGE for Non-Fast-Food Restaurants are positive (OLS: 0.1712 (s.e. =0.01376), 2SLS: 0.23040 (s.e. =0.02403)), highly significant, and between about 7 times (OLS) and 9 times (2SLS) larger than the estimated coefficients for All Restaurants presented in Table 3. The magnitude of the β_2 coefficients increased by a factor of 23 for OLS and by a factor of 35 for 2SLS. For Traditional American Restaurants, significance levels are qualitatively unchanged and the magnitude of the coefficient doubles again (β_1: OLS: 0.3268 (s.e. =0.03753), 2SLS: 0.5290 (s.e. =0.06878), β_2: OLS: -0.0069 (s.e. =0.00129), 2SLS: -0.01604 (s.e. =0.00322)). For Non-Fast-Food Restaurants and Traditional American Restaurants, the pronounced differences in coefficients from OLS and 2SLS model are also reflected in Wooldridge's score test and the regression-based test (Wooldridge 1995), which rejects the null that NUMLOC and NUMLOC*NUMLOC are exogenous at the 0.01% level. Considering the effect of NUMLOC and NUMLOC*NUMLOC on the market-level average mean rating for Non-Fast-Food Restaurants shows a significant negative relationship for NUMLOC and a significant positive relationship for NUMLOC*NUMLOC. The estimated coefficients of -0.219 (s.e. =0.00754) (OLS) and -0.03478 (s.e. =0.01409) (2SLS) for NUMLOC and 0.00019 (s.e. =0.0011) (OLS) and 0.00041 (s.e. =0.00024) (2SLS) are also around 8 times and 41 times larger than the coefficients for All Restaurants presented in Table 3, which suggests a non-negligible effect of the number of locations in a market on the average mean rating. For example, if we compare a market with 3 Non-Fast-Food Restaurants (the 25th percentile) with a market with 11 Non-Fast-Food Restaurants (the 75th percentile) the market-level average mean rating decreases by between 0.154 and 0.324 or, as the market-level average mean rating for all Non-Fast-Food restaurants is 3.58, by between 4.3% and 9.1%.

Table 4: Regression Results for Non-Fast-Food Restaurants

	OLS	2SLS IV 2nd Stage	OLS	2SLS IV 2nd Stage
VARIABLES	RANGE	RANGE	MKTAVGRATE	MKTAVGRATE
NUMLOC	**0.17120*** **	**0.23040*** **	**-0.02191*** **	**-0.03478** **
	(0.01376)	**(0.02403)**	**(0.00754)**	**(0.01409)**
NUMLOC²	**-0.00198*** **	**-0.00322*** **	**0.00019* **	**0.00041* **
	(0.00023)	**(0.00058)**	**(0.00011)**	**(0.00024)**
MEDAGE	-0.02272***	-0.02047**	0.01202**	0.01077*
	(0.00853)	(0.00931)	(0.00569)	(0.00579)
MEDINC	0.02208***	0.02613***	-0.00541	-0.00588
	(0.00687)	(0.00774)	(0.00454)	(0.00484)
MEDHVAL	-0.00563***	-0.00705***	0.00012	0.00036
	(0.00137)	(0.00178)	(0.00084)	(0.00087)
MEDRENT	-0.08468	-0.10828	0.10080**	0.10858**
	(0.0623)	(0.0685)	(0.044)	(0.04521)
UNEMP	0.02258	0.0154	0.00758	0.00916
	(0.02453)	(0.02521)	(0.01209)	(0.01277)
LANDAREA	0.02085*	0.01960*	0.00828*	0.00963**
	(0.01093)	(0.01187)	(0.005)	(0.00478)
URBAN	0.0918	-0.26739	0.34623	0.44414
	(0.5002)	(0.53704)	(0.27119)	(0.27503)
COSTOFLIV	0.02596**	0.02568**	-0.00904	-0.00936
	(0.01228)	(0.01309)	(0.00843)	(0.00858)
PERCHISP	-0.20991	-0.1136	-0.04303	-0.06412
	(0.27175)	(0.27759)	(0.17292)	(0.17618)
PERCINDIAN	-1.57088*	-1.54933*	-0.24853	-0.31186
	(0.88009)	(0.87054)	(0.59293)	(0.58637)
PERCASIAN	-2.23815**	-1.88718*	0.01726	-0.04891
	(1.10851)	(1.07601)	(0.68994)	(0.72112)
PERCBLACK	0.10469	0.18993	0.07144	0.05843
	(0.36488)	(0.36471)	(0.27056)	(0.26885)
LOCPOP	0.02203**	0.03338*	0.00281	0.00272
	(0.00949)	(0.01738)	(0.00453)	(0.00485)
Constant	-0.72143	0.68606	3.23127***	3.22899***
	(1.31087)	(1.38344)	(0.85576)	(0.87146)
N	326	325	326	325
R^2	0.50124	0.43673	0.05493	0.04513

*** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

For Traditional American Restaurants, the magnitude of the coefficients substantially increases again (NUMLOC: -0.07933 (s.e.=0.0199) (OLS) and -0.18355 (s.e.=0.0483) (2SLS), NUMLOC * NUMLOC: 0.00147 (s.e.=0.00052) (OLS) and 0.00560 (s.e.=0.00189) (2SLS)) meaning that an increase in the number of Traditional American Restaurants from the 25th percentile (1 Traditional American Restaurant) to the 75th percentile (4 Traditional American Restaurants) decreases the market-level average mean rating by between 0.21 and 0.47. Comparing the coefficients from OLS and 2SLS models shows a substantial difference between the coefficients.

Table 5: Regression Results for Traditional American Restaurants

	OLS	2SLS IV 2nd Stage	OLS	2SLS IV 2nd Stage
VARIABLES	RANGE	RANGE	MKTAVGRATE	MKTAVGRATE
NUMLOC	**0.32681*** **	**0.52902*** **	**-0.07933*** **	**-0.18355*** **
	(0.03753)	**(0.06878)**	**(0.0199)**	**(0.0483)**
NUMLOC*NUMLOC	**-0.00693*** **	**-0.01604*** **	**0.00147*** **	**0.00560*** **
	(0.00129)	**(0.00322)**	**(0.00053)**	**(0.00189)**
MEDAGE	-0.01018	-0.00759	0.00856	0.00556
	(0.00893)	(0.01058)	(0.00774)	(0.0085)
MEDINC	0.01358*	0.02007*	-0.01224**	-0.01522**
	(0.00799)	(0.01037)	(0.00536)	(0.00607)
MEDHVAL	-0.00337***	-0.00589***	0.00019	0.00128
	(0.00097)	(0.00219)	(0.00102)	(0.00135)
MEDRENT	-0.03378	-0.04368	0.10055*	0.11752**
	(0.06018)	(0.07931)	(0.05165)	(0.05477)
UNEMP	0.00216	-0.00848	0.01106	0.01548
	(0.02075)	(0.0232)	(0.01785)	(0.01991)
LANDAREA	-0.03014***	-0.02779	0.01931***	0.02062**
	(0.01065)	(0.02022)	(0.0065)	(0.01005)
URBAN	0.69065	0.33922	-0.06325	0.14435
	(0.56189)	(0.60586)	(0.38447)	(0.3953)
COSTOFLIV	0.02087**	0.02459**	-0.01089	-0.01402
	(0.01058)	(0.01159)	(0.00969)	(0.01069)
PERCHISP	0.12595	0.28529	-0.22034	-0.29898
	(0.29411)	(0.30722)	(0.24117)	(0.26734)
PERCINDIAN	-0.54879	-1.42725	-0.39162	-0.12137
	(1.14692)	(1.58031)	(0.76168)	(0.92759)
PERCASIAN	-1.49440*	-0.26267	-0.77341	-1.35607
	(0.79211)	(1.21807)	(0.76126)	(0.82979)
PERCBLACK	0.55898	0.88141**	-0.04455	-0.18767
	(0.4261)	(0.44133)	(0.34587)	(0.36426)
LOCPOP	0.02364	0.07663	0.01362	-0.00075
	(0.02419)	(0.07456)	(0.01159)	(0.02848)
Constant	-1.98898*	-2.66207*	4.23597***	4.63461***
	(1.20777)	(1.38826)	(1.06069)	(1.16084)
N	258	257	258	257
R^2	0.48571	0.39514	0.09934	0.02309

*** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

This is also reflected in the test statistics from Wooldridge's score test and the regression-based test (Wooldridge 1995) which reject the null that NUMLOC and NUMLOC*NUMLOC are exogenous at the 3% significance-level. The strong effect of the number of locations on the range of mean ratings and on the market-level average mean rating for Non-Fast-Food Restaurants and Traditional American Restaurants is also in line with the insights from our theoretical model.

This negative effect of competition on the market-level average mean rating for Non-Fast-Food and Traditional American Restaurants may seem to contradict the findings by Olivares and Cachon (2009), Dick (2007), and Cohen and Mazzeo (2010) who report that quality increases with increasing competition. However, the measured qualities in those studies are substantially affected by fixed costs whereas the quality for restaurants in general, Non-Fast-Food Restaurants and Traditional American Restaurants, in particular (as defined by e.g., service quality, quality of ingredients, and quality of food preparation), should mainly be affected by marginal costs. As discussed in Jones and Mendelson (2011), larger markets entail higher average quality levels when quality improvements are mainly driven by fixed costs, whereas our theoretical framework posits that average quality levels decrease with an increase in the number of competitors in a market in which the quality levels are mainly driven by marginal costs.

4.7 Further Analyses

To further investigate the effect of increasing competition on the rating distribution in a market, we incorporate several additional approaches. In a first approach, in November 2015, we collected data from Yelp.com once again for exactly the same cities in order to construct a panel data set. Consequently, we also collected data for all our control variables included in our baseline model from city-data.com. Thus, we are able to exploit longitudinal variation in the data to verify whether our proposed effect of competitiveness on the range of ratings provided and on the average mean rating holds over time. More importantly, we are able to control for time-invariant unobserved heterogeneity between cities which we are not able to control for in a setting of cross-sectional data. Thus, we first conduct a Hausman specification check for our model with All Restaurants, with only Non-Fast-Food Restaurants, and for Traditional American Restaurants, to investigate whether a structural difference in estimates arises due to time-invariant city-level heterogeneity. The Hausman Test results suggest that a random effects (RE) model specification is appropriate and that unobserved city-level characteristics are not a concern[15]. Tables B1 to B3 in the appendix illustrate the findings of our RE model for All Restaurants. The result of our RE model estimations remain qualitatively the same, compared to the results of the baseline model.

One might argue that one inconvenient assumption underlying the RE estimator is that there is no correlation between city-level observables and time-invariant unobservable characteristics of cities. In order to address this issue, we follow Mundlak (1978) that relaxes this assumption and allows for some of the above mentioned correlation. In particular, we implement group means of the time varying

[15] The only case in which the Hausman Test finds a structural difference in estimators is in the model with Traditional American Restaurants with RANGE as the dependent variable. However, the coefficients for NUMLOC and NUMLOC*NUMLOC are qualitatively equal. Therefore, we report the RE model for RANGE in Table B3.

covariates in the RE model. We find that our estimates for NUMLOC and NUMLOC*NUMLOC remain qualitatively unchanged which confirms that our RE model is properly defined. NUMLOC is positive and significant for RANGE whereas NUMLOC is negative and significant for MKTAVGRATE in all three model specifications. Luca and Zervas (2016) follow the same approach using Yelp data in order to estimate the likelihood of restaurant chains to engage in review fraud. In addition to that, we also model the city-level random effects explicitly using a hierarchical modelling approach, which allows us to assess the variance-covariance structure of the city level characteristics with the residuals. This analysis equally reveals the absence of any substantial correlation (0.06).

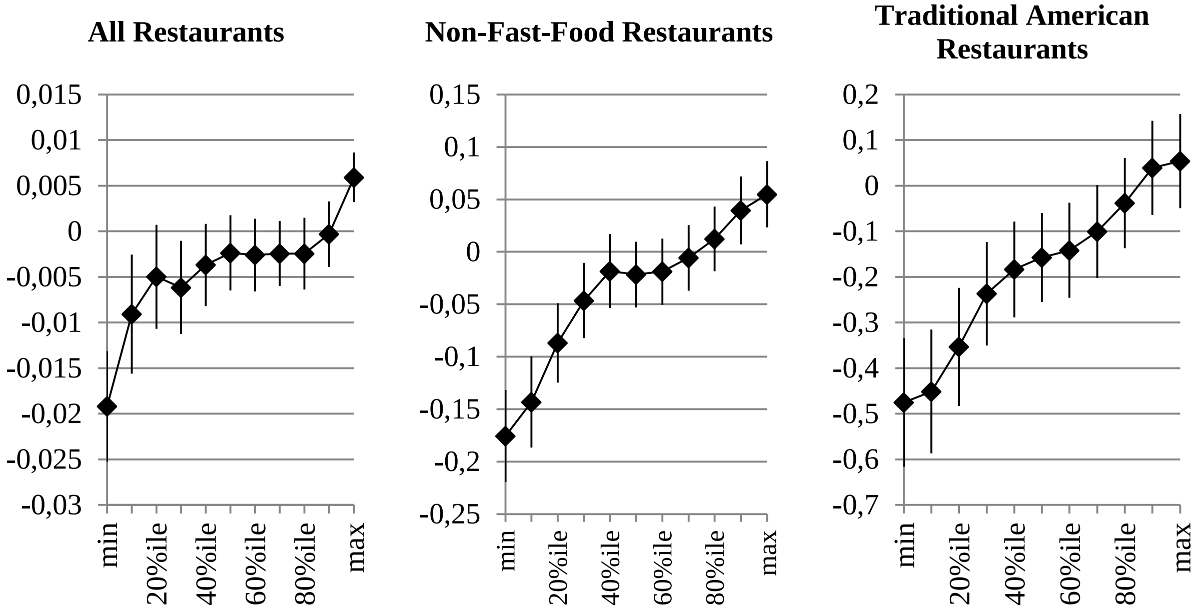

Figure 5: Estimated Coefficients on NUMLOC for Different %iles of the Quality Distribution as Dependent Variables

In another approach, we computed the median, minimum and maximum as well as the 10th percentile, 20th percentile, 30th percentile, 40th percentile, 60th percentile, 70th percentile, 80th percentile, and 90th percentile of the rating distribution for each isolated market and estimated our model for All Restaurants, Non-Fast-Food Restaurants, and Traditional American Restaurants with these percentiles as dependent variables. If the range of the mean rating distribution increases and the market-level average mean rating decreases, we would expect to find a negative effect for the median and for all percentiles left to the median where the absolute magnitude of the negative effect should be strongest for the minimum rating and decreasing for increasing percentiles. For the percentiles above the median, the sign and magnitude of the expected coefficients depend on two opposing effects. The increase in the range of the mean rating should lead to a positive effect while the decrease of the market-level average mean rating should lead to a negative effect, where the former increases with increasing percentiles and the latter remains constant. Therefore, we would expect to find a decreasing absolute magnitude of the negative effect for increasing percentiles, and, if the positive effect is strong enough, a conversion to a positive sign for the highest percentiles. Figure 5 shows the estimates (diamond symbols) as well as the

95% confidence intervals (vertical lines) from 2SLS models with the percentiles on the x-axis as dependent variables for All Restaurants, Non-Fast-Food Restaurants, and Traditional American Restaurants. This figure shows that the estimated coefficients are negative for the lower percentiles, decrease in absolute magnitude for increasing percentiles, and turn positive for the highest percentiles. This finding provides additional evidence that with increasing competition in a market, the range of mean ratings increases and the market-level average mean rating decreases.

Table 6: Regression Results from Ordinal Logistic Regression

	All Restaurants	Non-Fast-Food	Traditional American Restaurants
VARIABLES	QUALITYCOUNT	QUALITYCOUNT	QUALITYCOUNT
NUMLOC	**0.0915*****	**0.390*****	**0.903*****
	(0.00862)	**(0.0578)**	**(0.164)**
MEDAGE	-0.0436**	-0.0361*	-0.0399*
	(0.0196)	(0.0203)	(0.0204)
MEDINC	0.0408**	0.0219	0.0110
	(0.0159)	(0.0175)	(0.0157)
MEDHVAL	-0.00658*	-0.0124**	-0.00463
	(0.00370)	(0.00526)	(0.00312)
MEDRENT	-0.0726	0.0892	0.0364
	(0.156)	(0.149)	(0.136)
UNEMP	-0.0579	0.0405	0.0173
	(0.0456)	(0.0427)	(0.0444)
LANDAREA	-0.0103	0.0168	-0.144**
	(0.0181)	(0.0262)	(0.0700)
URBAN	1.293	1.277	1.023
	(1.386)	(1.053)	(1.153)
COSTOFLIV	0.0435*	0.0610**	0.0296
	(0.0258)	(0.0252)	(0.0247)
PERCHISP	-0.0119	-1.118**	-1.090*
	(0.609)	(0.558)	(0.557)
PERCINDIAN	-0.0617	-0.974	1.960
	(1.997)	(2.127)	(3.315)
PERCASIAN	-4.176	-5.212**	-1.356
	(3.944)	(2.633)	(2.403)
PERCBLACK	0.484	-0.888	-1.001
	(0.726)	(0.678)	(0.825)
N	348	338	283

*** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

Another approach to verifying whether a greater competition leads to a larger variety of available qualities in a market would involve using the discreteness of the star rating system on Yelp.com and counting the number of different qualities available in a market. In this rating system, qualities are measured in full and half stars, ranging from one star for the lowest to five stars for the highest quality rating, which gives us a total of nine potential quality levels for each market. If restaurants differentiate their quality in response to competition, we would expect to see a higher number of different mean ratings in more competitive markets. As the number of different mean ratings in a market is an ordinal variable, we use an ordered logistic regression model to analyze the relationship between this variable and the number of competitors in a market. Columns (1) to (3) in Table 6 show the estimation results for the ordered logistic regression model for the sample of All Restaurants, Non-Fast-Food Restaurants, and Traditional American Restaurants. If the number of locations in a city were to increase by one unit, its ordered log-odds of having a higher number of available quality levels would increase by between 0.09 (s.e.= 0.009) and 0.9 (s.e.= 0.164).

In an additional analysis, we investigate if crowdsourced online ratings can be used to conduct an alternative computation of the HHI that captures competition in a market. This would obviate the missing data problem of the rarely available information on market share. We conduct a simple OLS regression substituting NUMLOC with HHI as an explanatory variable into our model that represents a popular traditional measure of competition. One difficulty in doing this is that the computation of the HHI requires market shares. Since we cannot observe market shares based on our data, we calculate the HHI based on the number of ratings a restaurant received. It is reasonable to assume that the number of ratings a restaurant received has a positive correlation with the market share of a restaurant, as a restaurant that establishes a customer base over time also sees its number of ratings increasing. To the best of our knowledge, we are the first to calculate the HHI in this way based on crowdsourced data. Therefore, this analysis can also deliver evidence in support of the claim that crowdsourced online ratings can be used to obviate a missing data problem by conducting an alternative computation of the HHI. While ultimately, inferring whether our computation of the HHI based on the number of reviews is equivalent to the calculation based on real market shares is not possible in this case, we can however check whether the results using the alternative calculation of the HHI are consistent with the results in our baseline estimation in 4.5. Additionally, we investigate how our operationalization of competition, NUMLOC, compares to traditional measures of competition, in order to provide evidence in support of the claim that our construct adequately captures competition. For each city, we calculated the HHI by summing up the total squared percentage market shares of ratings for each restaurant.[16] Thus, similar to the traditional computation, the HHI can take up values between 0 and 10,000, where 0 indicates high competition among a lot of competitors and 10,000 indicates a perfect monopoly. Reassuringly, we find that NUMLOC is statistically significantly and substantially negatively correlated to our computation of the HHI (Pearson's r=-0.4473). This means that a higher HHI, where 10,000 equals a pure monopoly, is associated with a lower value for NUMLOC. Thus, we expect coefficients for HHI and HHI*HHI with opposite signs compared to those for NUMLOC and NUMLOC*NUMLOC in our baseline model.

[16] For instance, if there are two restaurants in the market with each 200 individual ratings, then the HHI is calculated as $50(\%)^2 + 50(\%)^2 = 2500$.

Indeed, the results for this analysis largely remain qualitatively unchanged to our baseline model, as depicted in Table B6.

4.8. Robustness Checks

Our results provide direct evidence for the effect of local market conditions on online consumer ratings. In the following, we demonstrate that these results withstand a wide range of robustness checks.[17]

First, one might argue that our results are affected by the truncation of our key variables of interest. The average of mean ratings in a market MKTAVGRATE can only take values between one and five and the range of mean ratings values between zero and four. We consider this specific characteristic of our variables by estimating Tobit versions of all our models. The resulting coefficients are all larger in absolute magnitude compared to the respective coefficients from the linear regression models. We are confident, therefore, that the results presented in Tables 3 to 5 can be interpreted as a lower boundary for the real effects rather than being driven by the truncation of our key variables of interest.

Second, Yelp may not cover all restaurants in a specific location and may, therefore, represent an inappropriate measure of local market structures. To address this concern, we collected data on the number of restaurants for each of our isolated markets on Yellowpages.com and compared this number with the number of locations collected from Yelp.com. Although the resulting average number of restaurants on Yellowpages.com (46.0) is significantly larger than the average number on Yelp (40.1),[18] we find that these two measures are highly correlated ($\rho = 0.959$) and that our results do not change if we substitute the number of locations collected on Yelp.com with the number of locations collected on Yellowpages.com. These results show that our results are not biased by an inappropriate measure of local market structures.

Third, one might argue that the increase in the range of mean ratings is attributable to an increase only in high or low qualities, e.g., the lowest (highest) quality in a market remains constant, while the highest (lowest) quality increases (decreases). We address this issue by re-estimating our models with the minimum (MINRATE) and maximum (MAXRATE) mean ratings in a market as dependent variables. An increase in the number of locations leads to a decreasing minimum mean rating and to an increasing maximum mean rating for all types of restaurants. This result reaffirms our finding that the range of available qualities increases with market size.

Another concern might be that our results are driven by restaurants with only a few very positive or very negative ratings which do not represent the true quality of a location. For example, the owner (or their friends) of a newly opened location posting a five star rating would almost certainly increase the observable range of mean ratings. The same applies if the competitors of a newly opened location in

[17] The results of these robustness checks are available upon request.
[18] One potential explanation is that due to crowd-driven nature of Yelp data generation process, it reflects restaurant closures much more promptly than Yellow Pages. The Yellow Pages listings were collected in September 2014.

turn post very negative ratings. We address this issue by limiting our dataset to restaurants which received at least 2 or at least 5 ratings respectively. We find a significant positive effect of the number of locations on the range of mean ratings for all types of restaurants and all models. The magnitude of the estimated coefficients slightly decreases with an increasing threshold. This decrease indicates that a small part of the variation in the range of mean ratings is attributable to locations with very few ratings. Nevertheless, the positive effect of the number of locations on the range of mean ratings remains economically and statistically significant even if we exclude these locations. We observe the same pattern for the effect of NUMLOC on the market-level average mean rating for all models and all subsets of our data. The magnitude of these effects also decreases slightly with an increasing threshold. Thus again, a smaller part of the decrease in the market-level average mean rating with more competitors is attributable to locations with very few ratings. Still, even after excluding these locations from our sample, the reported effects remain economically and statistically significant.

Similarly to our binning strategy introduced earlier, we also test whether our empirical baseline results hold if we compare cities from different bins. We again utilize our bootstrapping program that randomly draws two cities with replacement, this time each from a different bin of city sizes, and tests for differences in the average rating and in the range of ratings. Running the program for 2000 rounds, we find that cities from bin 20 have a higher average rating in at least 50.9% of cases (when compared to cities from bin 30) and at most in 59.8% percent of cases (when compared to cities from bin 40). Similarly, the range of ratings of cities from bin 20 is smaller in at least 52.75% of cases (when compared to cities from bin 30) and at most in 67.7% percent of cases (when compared to cities from bin 40).[19] Results are robust if we use inter quartile range as opposed to the normal range. The difference in ranges is more pronounced than the difference in average ratings, which mirrors the main results of our basic model.

In spite of the robust empirical evidence supporting the Yelp's mean ratings, we still would like to address the concern about the effect of the rater's motivation on review distributions, and on mean ratings as an indicator of quality. As outlined in our literature section, a sub-stream of literature investigated social influences on online ratings (e.g., Li and Hitt 2008, Muchnik et al. 2013). For example, Dai et al. (2014) found that, unlike Hu et al's (2009) findings on the distribution of Amazon ratings, rating distributions on Yelp.com are not plagued by the J-shaped distribution which normally indicates a tendency to rate when opinions are extreme. Nevertheless, in order to exclude this concern and account for social influence on ratings, we implement a de-biasing approach proposed by Dai et al. (2014). Using Yelp data as well, Dai et al. find that a social influence bias increases with the number of ratings a restaurant obtains which might cause the simple mean rating to disguise accurate quality. The

[19] It is worth nothing that, in the absence of statistical power enjoyed by parametric tests because of distributional assumptions, the observed level of consistency is reassuring. After all, the cost of fewer assumptions in nonparametric tests is that they are generally less powerful than their parametric counterparts (i.e., when the alternative is true, they may be less likely to reject H0).

authors develop and implement an aggregation mechanism in order to construct an optimal mean that accounts for social influence. The construction of the optimal mean basically relies on a weighted average based on the number of reviews a restaurant has. We apply this reweighting strategy developed by Dai et al. to the mean ratings of the restaurants in our basic model and find that the conclusions remain qualitatively unchanged. We conclude that social influence hidden behind the mean of Yelp's ratings is not a concern in our case. In addition to that, we run model specifications of our basic model using only restaurants with similar numbers of ratings (only restaurants with less than 10 ratings, less than 25 ratings, and 10 to 20 ratings), and the results remain qualitatively unchanged.

While some evidence supports the claim that social influence on ratings are not a concern in our setting, one might argue that reviews might deviate from an accurate quality measurement in that they follow a certain trend over time, similar to the observations that Li and Hitt (2008) report for online ratings for books. Our identification of the effect of NUMLOC on MKTAVGRATE and RANGE might be plagued by this time trend if this trend is correlated with NUMLOC. In order to exclude these concerns, we collected time series data on all individual customer-level reviews for all restaurants in our original sample in December 2015. First, we conduct a Run Test which tests whether a time series follows a random order or whether it exhibits some trend. A "run" is defined as the sequences of consecutive observations above or below the sample median and the number of runs relative to the total number of observations indicates whether a variable follows a random order or a trend (Madansky 1988). According to the Run Test, only a small proportion of restaurants (<15%) follows a time trend. Second, we correlate the results of the Run Test with NUMLOC and it turns out that the number of locations in a given restaurant market is not correlated with a time trend. For further robustness, we estimated a logit model with a dummy variable as regressand that indicates whether or not a time series follows a trend, and indicator variables for the city bins from section 4.2 as regressors, including the same control variables as in our main model. None of the city bins has a significant influence on whether or not the ratings of a restaurant follow a time trend.

5. Conclusion and Discussion

Online rating websites for offline locations have seen tremendous growth during the past few years. The substantial amount of crowdsourced data created this way, in turn, represents a potentially appealing resource for firms to conduct CI in their local markets. It is all the more surprising, then, that the impact of local market structure on online rating distributions on these websites, and the usability of online ratings to conduct CI, has not yet been studied. This paper attempts to empirically fill this gap in the literature in two ways: First, by presenting evidence that suggests that crowdsourced ratings, by and large, follow an internally and externally consistent data generation process and, therefore, can be used for CI purposes. Second, and more importantly, by analyzing how changes in local market structure influence the distributional properties of online ratings, we demonstrate an effective way of deriving insights related to local competitive dynamics. To supplement the empirical analysis, we build a theoretical model of an oligopoly where firms vertically differentiate their quality offerings. This model sheds light on how an increasing number of competitors in a market leads to a broader range of qualities offered, at the same time as to a decrease in the average qualities.

Combining demographics, population data, and review data for 372 isolated markets in the U.S., we find evidence that increasing competition leads to a broader range of mean ratings and to a decrease in market-level average mean ratings. These effects are particularly pronounced when we limit our analysis to specific restaurant types where there are no significant opportunities for horizontal differentiation. In addition, we present extensive empirical examination of whether or not Yelp mean ratings can internally and externally consistently embody the quality provisioning decisions of firms. Moreover, the empirical evidence presented on the internal and external consistency of the rating data lends support to the theoretical findings of Besbes and Scarsini (2015) that, even the in the case of ratings seen as subjective expressions, the mean ratings – i.e., the average of individual ratings of a restaurant– converge to the underlying objective quality under a very pragmatic set of assumptions and that the position of a product or a service relative to competitors is usually preserved in online ratings. Furthermore, even if there was a difference in the rating behavior of diners in bigger or smaller cities (e.g., diners in bigger cities have a tendency to rate more critically than those in small cities, which could be coined as a kind of *framing* effect) which can be partially held responsible for the shape of the observed online rating distribution, the significance of our results remains unaffected. This is because a key argument of this paper is that the market power of a restaurant that received a mean rating of three stars and which is situated in a small market should be considered quite differently to a restaurant that received the same online rating but is situated in a larger market. Most prominently, they face different competitive environments because the three-star restaurant from a smaller market faces less competition from restaurants with lower ratings than its respective counterpart from the larger market. Consequently, even if this was partially caused by different rating attitudes (or *framing* of the raters), our argument that the restaurants with equal ratings from markets of different sizes should not be treated equally, still holds. Future research should be cognizant of our findings, especially if mean ratings of firms from different locations are to be used. Moreover, literature on BI & A can build on the evidence presented here, in particular, the fact that Yelp's mean ratings form an internally and externally consistent data source for conducting CI activities.

By pioneering a way to substitute traditional practices in gathering CI, this paper is poised to alter industry practices in a non-trivial manner. We show that, in CI gathering practice, Yelp ratings (or other similar websites) can be used to identify the quality spectrum that is currently available in a market and the changes of the quality spectrum that will likely occur when market size changes. Generally, business owners may map their own quality provided on the quality spectrum of the market. Specifically, business owners can assess if a specific market is underserved and use this opportunity for strategic benefits. We consider these implications to be especially valuable for chain restaurants which by nature operate in many different markets. Seeking to provide similar quality to their customers at each of their restaurants, chains might face very different competitive dynamics in markets with different market structure, as highlighted in our work. Thus, all else being equal, we argue that in markets with 20 restaurants, for example, "classical" restaurant chains like Applebees or Red Lobster face relatively more competition from the upper end of the rating range than in a market with 40 restaurants, where competition from the lower end of the rating range will be more pronounced, as the market-level average mean rating decreases. Furthermore, as investors are increasingly employing online ratings to assess the creditworthiness of businesses, the findings of our work help make better decisions as restaurants with similar ratings have to be assessed differently if they are located in different market structures. With the increasing emphasis on data-driven decision making, the importance of using publicly available sources to understand critical market factors or investment opportunities, as demonstrated in this study, will undoubtedly continue to grow.

The results presented in this paper have important implications for everyday life too. In general, people who search for high quality restaurants (or other high quality locations from industries in which increases in quality involve marginal costs) are better off in larger markets with more competitors, as the likelihood of having a high class location (and all other specific quality levels) is higher in these markets. However, people who randomly pick a restaurant to eat are worse off in these markets as the average quality is lower than in smaller markets with fewer competitors. In other words, the consumption benefits of agglomeration in larger markets that have been identified in prior work (e.g., Berry and Waldfogel 2010) come at the expense of lower average qualities. This result emphasizes the importance of online review websites and other mechanisms to pre-screen the quality of locations in large markets to alleviate or prevent this cost.

Appendix A: Theoretical Framework Extensions

Implications of horizontal competition

While we employ an analytical model of vertical differentiation in our work to investigate service quality, one might argue that additional differentiation along the horizontal dimension might also have an influence on service quality which might be appropriate to consider in the context of restaurants. Here, vertical differentiation (denoted by the subscript v) refers to the fact that the same product, say traditional American food, can be ranked by quality, and different firms produce different qualities. On the other hand, horizontal differentiation (denoted by the subscript z) refers to a situation in which products or services cannot be distinguished by objective quality but by features such as color or cultural background, which are only related to taste. Even though both services are of equally high quality, consumers, for example, might value French cuisine higher than traditional American cuisine just because of their taste. To illustrate these two types of differentiation and to depict the consequences that an introduction of horizontal differentiation has into our model we can model a two times two market structure with the two different qualities h and l and with two different taste-related features, say blue and yellow (indicated by the subscripts b and y). Here, online ratings equal quality, so $R = q$, therefore in the pure vertical model with two qualities q_h and q_l, this means that $R_h > R_l$ and that the average rating in the vertical model $\bar{R}_v$ is $R_h > \bar{R}_v > R_l$. On the other hand, for a pure horizontal model with only one quality level, for example high quality, we have ratings equal to $R_z = q_{hb}$ and $R_z = q_{hy}$. Note that the objective quality q_h of both restaurants is identical and that both ratings are high and only differ in the taste-related dimension b and y, so that $R_z = q_{hb} = q_{hy}$. Hence, the average quality in the pure horizontal model equals q_h. Combining both the pure vertical and the pure horizontal model to a two times two market with four restaurants, yields a market average rating $\bar{R}_{vz} = \frac{q_{hb}+q_{hy}+q_{lb}+q_{ly}}{4}$. It is then straightforward to show that the market average rating for the model that includes both vertical and horizontal competition, $\bar{R}_{vz}$ is equivalent to the rating average in the pure vertical case with two restaurants, so that $\bar{R}_v = \bar{R}_{vz}$.

Although the qualities and, therefore, the prices do not differ between the pure vertical model with two firms and the vertical-horizontal model with four firms, the firms' decision to enter the market, dependent on S and F, change. Specifically, recall from section 3 that in the case of vertical competition of two firms, a high quality firms needs $\geq 60.98F$, and a low quality firms needs $S \geq 82.64F$, to enter the market. Whereas the factors 60.98 and 82.64 do not change in the vertical-horizontal model because it is solely calculated on the basis of price and quality (which do not differ), the customer base has to choose between the two offered product variants b and y. This implies that in order to accommodate for 4 firms, a market in the vertical-horizontal model needs twice as many consumers as a market in the pure vertical model with 2 firms. By contrast, the lowest quality firm in a market with pure vertical differentiation and 4 firms would need $S \geq 625F$ to enter the market, which is well below two times

82.64. This is in line with the assumption that a market with a total of four firms is much more likely to exhibit two vertically differentiated qualities and two different horizontal variants of a product, than it is to exhibit four firms which differentiate only vertically.

This is in line with a rather extensive literature which points to the idea that under mild assumptions, each model of vertical differentiation equals a model of horizontal differentiation and vice versa (e.g., Cremer and Thisse 1991, Anglin 1992). Introducing the influence of taste on ratings - as in Perloff and Salop (1985) - we assume that people have taste that a restaurant can or cannot meet. Therefore, we model the influence of taste on ratings as a continuous match/mismatch parameter ε that takes on values between -1 and 1. If a consumer prefers French food over Traditional American Food (both of the same quality), the parameter ε is positive if the consumer rates a French restaurant, and negative if the consumer rates a traditional American restaurant. Then, a restaurant's rating is $R = q + \varepsilon$. This is similar to Perloff and Salop's (1985) model of "spurious differentiation". Applied to our context this means that consumers give higher or lower ratings to restaurants unrelated to their actual quality. If each epsilon is drawn from an i.i.d. distribution with zero mean which is uncorrelated to the vertical dimension (as in Perloff and Salop 1985), i.e., quality, then the match/mismatch ε is mean preserving. Therefore, in the presence of horizontal differentiation, the objective market average quality measured by online ratings does not differ compared to the model of pure vertical differentiation since $\bar{R}_v = \bar{R}_{vz}$. These insights are related to Vogel (2008) who finds that in a vertical-horizontal model with endogenous locations choice, a firm's price, market share, and profit are only influenced through the marginal costs (which is captured in our baseline analytical model of vertical differentiation) of its competitors but they are independent of the competitors location.

Appendix B: Additional Results

Table B1: RE-Model for All Restaurants

VARIABLES	RANGE	MKTAVGRATE
NUMLOC	**0.0209*** **	**-0.000513**
	(0.00228)	**(0.00135)**
NUMLOC*NUMLOC	**-0.0000713*** **	**0.00000183**
	(0.0000111)	**(0.00000485)**
MEDAGE	-0.00665	0.00575**
	(0.00509)	(0.0027)
MEDINC	0.00659	-0.0034
	(0.00475)	(0.00218)
MEDHVAL	-0.00177**	0.000214
	(0.000712)	(0.00037)
MEDRENT	-0.00171	0.0241
	(0.0394)	(0.0201)
UNEMP	-0.0107	0.00213
	(0.0158)	(0.00593)
LANDAREA	0.0269***	0.00243
	(0.00568)	(0.00286)
URBAN	0.719	0.363**
	(0.525)	(0.183)
COSTOFLIV	0.00973	-0.00521*
	(0.00773)	(0.00298)
PERCHISP	-0.16	-0.0893
	(0.214)	(0.115)
PERCINDIAN	-0.169	-0.0563
	(0.186)	(0.131)
PERCASIAN	0.384	0.0672
	(0.514)	(0.28)
PERCBLACK	-0.0515	-0.121
	(0.218)	(0.134)
LOCPOP	0.000849	0.000452
	(0.00121)	(0.000578)
Constant	1.589*	3.293***
	(0.947)	-0.363
N	650	650
R^2	0.2514	0.0374

*** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

Table B2: RE Model for Non-Fast-Food Restaurants

VARIABLES	RANGE	MKTAVGRATE
NUMLOC	**0.155*** **	**-0.00989**
	(0.0131)	**(0.00634)**
NUMLOC*NUMLOC	**-0.00176*** **	**0.0000627**
	(0.000267)	**(0.0000927)**
MEDAGE	-0.0128	0.0104**
	(0.00813)	(0.00426)
MEDINC	0.0136***	-0.000548
	(0.00514)	(0.0034)
MEDHVAL	-0.00267***	0.000106
	(0.00102)	(0.000568)
MEDRENT	-0.0476	0.0144
	(0.0524)	(0.035)
UNEMP	-0.0116	0.0227**
	(0.0168)	(0.0089)
LANDAREA	0.0019	0.0100***
	(0.00704)	(0.00336)
URBAN	-0.215	0.166
	(0.528)	(0.249)
COSTOFLIV	0.0086	-0.00397
	(0.0106)	(0.00607)
PERCHISP	-0.0494	-0.184
	(0.215)	(0.131)
PERCINDIAN	0.376	0.0212
	(0.475)	(0.209)
PERCASIAN	-1.139	-0.105
	(1.227)	(0.394)
PERCBLACK	0.42	-0.154
	(0.266)	(0.235)
LOCPOP	0.00899	0.00327
	(0.00831)	(0.00244)
Constant	0.892	3.194***
	(1.236)	(0.627)
N	606	606
R^2	0.4366	0.0442

*** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

Table B3: RE-Model for Traditional American Restaurants

VARIABLES	RANGE	MKTAVGRATE
NUMLOC	**0.286*** **	**-0.0382*** **
	(0.0289)	**(0.0137)**
NUMLOC*NUMLOC	**-0.00582*** **	**0.000571**
	(0.001)	**(0.000351)**
MEDAGE	-0.000835	0.00884
	(0.0081)	(0.00565)
MEDINC	0.00852	-0.0034
	(0.0062)	(0.0039)
MEDHVAL	-0.00218**	0.0000158
	(0.000888)	(0.000723)
MEDRENT	0.0083	0.0374
	(0.0585)	(0.0389)
UNEMP	0.00952	0.0113
	(0.0194)	(0.0109)
LANDAREA	-0.0154***	-0.00681
	(0.00596)	(0.00435)
URBAN	0.631	-0.36
	(0.558)	(0.383)
COSTOFLIV	0.00563	-0.00679
	(0.00992)	(0.00702)
PERCHISP	0.159	-0.205
	(0.262)	(0.178)
PERCINDIAN	0.107	0.430***
	(0.385)	(0.143)
PERCASIAN	-1.629**	-0.229
	(0.788)	(0.614)
PERCBLACK	0.224	0.113
	(0.336)	(0.324)
LOCPOP	0.00976	0.0131*
	(0.0168)	(0.00774)
Constant	-1.016	4.071***
	(1.186)	(0.816)
N	468	468
R^2	0.4581	0.0790

*** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

Table B4: First Stage Results NUMLOC

	All Restaurants	Non-Fast-Food	Traditional American
VARIABLES	NUMLOC	NUMLOC	NUMLOC
SBPI	0.14674	-0.01842	-0.02416
	(0.15434)	(0.03958)	(0.0222)
SBPI_2	-0.00124	0.00005	0.00016
	(0.00104)	(0.00028)	(0.00015)
RESTAURANTS_ 2004	0.96595***	0.19279***	0.06851***
	(0.10089)	(0.03112)	(0.01512)
RESTAURANTS_ 2004_2	0.00417***	0.00091***	0.00054***
	(0.00075)	(0.00031)	(0.00013)
MEDAGE	-0.43214***	-0.13937***	-0.08294***
	(0.14217)	(0.0402)	(0.02076)
MEDINC	0.24153*	0.07684**	0.04000***
	(0.12426)	(0.0337)	(0.01536)
MEDHOUSEVAL	-0.02466	-0.00161	-0.00348
	(0.02784)	(0.00956)	(0.00446)
MEDGROSSRENT	1.97746*	0.60503**	0.25245*
	(1.0646)	(0.29716)	(0.14823)
UNEMP	-0.19641	-0.06495	0.02016
	(0.25951)	(0.06981)	(0.03981)
LANDAREA	0.55925	0.3145	0.14692*
	(0.35593)	(0.19432)	(0.07627)
URBAN	12.17981**	2.68827	0.37425
	(5.70082)	(1.64856)	(0.97198)
COSTOFLIV	-0.11526	-0.08413	-0.07667**
	(0.17906)	(0.05345)	(0.03417)
PERCHISP	8.32536*	-0.42992	-0.24717
	(4.46817)	(1.12269)	(0.56242)
PECINDIAN	-60.08611***	-16.26813***	-7.90879***
	(21.77671)	(4.85503)	(2.4357)
PERCASIAN	-18.00341	-5.74589	0.54678
	(14.98295)	(3.77079)	(1.68236)
PERCBLACK	5.76524	0.85842	-0.4696
	(4.19768)	(0.95306)	(0.62178)
LOCPOP	0.36179***	0.58500***	0.67360***
	(0.10586)	(0.10135)	(0.13932)
Constant	-12.72159	2.15521	5.89541*
	(17.07363)	(5.10191)	(3.057)
N	347	337	257
F	86.41	53.72	61.66
Shea's Partial R^2	0.735	0.3661	0.2294

*** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

Table B5: First Stage Results NUMLOC*NUMLOC

	All Restaurants	Non-Fast-Food	Traditional American
VARIABLES	NUMLOC*NUMLOC	NUMLOC*NUMLOC	NUMLOC*NUMLOC
SBPI	23.63801	2.98225	-0.27993
	(29.02915)	(3.12247)	(0.70115)
SBPI_2	-0.16897	-0.01974	0.00124
	(0.19812)	(0.02158)	(0.00468)
RESTAURANTS_2004	-126.86641***	-5.9371	-1.0723
	(23.30112)	(3.67123)	(0.65292)
RESTAURANTS_2004_2	3.66366***	0.17055***	0.03384***
	(0.23698)	(0.04727)	(0.00683)
MEDAGE	-61.26741**	-6.60486**	-1.56712*
	(27.9855)	(3.12082)	(0.83077)
MEDINC	60.71543**	5.69849**	1.31041*
	(27.47271)	(2.85211)	(0.67342)
MEDHOUSEVAL	-9.09616	-0.5693	-0.23486
	(7.7512)	(1.02909)	(0.23187)
MEDGROSSRENT	416.37783	17.49309	4.08475
	(268.96583)	(25.57426)	(5.97235)
UNEMP	-5.83667	-3.99822	0.00362
	(44.47982)	(5.20918)	(1.47858)
LANDAREA	108.38677**	14.86539	3.61109**
	(52.0161)	(11.46396)	(1.47909)
URBAN	925.97088	48.22038	5.6452
	(1,184.06)	(150.68833)	(45.12357)
COSTOFLIV	-38.03361	-6.19147	-1.31497
	(36.39131)	(4.42546)	(1.11428)
PERCHISP	345.48108	12.84117	4.63769
	(934.19518)	(80.80727)	(18.19959)
PECINDIAN	-12571.50079**	-857.50883**	-281.75026**
	(5,596.96)	(351.93396)	(116.30942)
PERCASIAN	-2,830.06	-137.9865	122.70954
	(3,084.83)	(334.87801)	(76.97132)
PERCBLACK	596.09246	86.4501	26.47613
	(675.39077)	(63.25912)	(25.93303)
LOCPOP	80.06738**	36.25587***	19.76030**
	(33.41052)	(12.64692)	(8.16025)
Constant	-884.56928	205.38153	72.28463
	(3,477.18)	(439.22862)	(120.71337)
N	347	337	257
F	47.51	12.34	18.53
Shea's Partial R^2	0.7326	0.2395	0.1383

*** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

Table B6: OLS Results with HHI as a Substitute for NUMLOC

	All Restaurants		Non-Fast-Food		Traditional American	
VARIABLES	RANGE	MKTAVG-RATE	RANGE	MKTAVG-RATE	RANGE	MKTAVG-RATE
HHI	**-5.18870*** **	**0.91868** **	**-3.26420*** **	**1.22477** **	**-0.57531**	**1.49908** **
	(0.89129)	**(0.46297)**	**(0.83879)**	**(0.48971)**	**(1.12867)**	**(0.61338)**
HHI*HHI	**1.40680***	**-0.69528**	**0.24334**	**-0.7662**	**-1.49163***	**-0.72919**
	(0.83438)	**(0.54366)**	**(0.66907)**	**(0.49528)**	**(0.83212)**	**(0.53952)**
MEDAGE	-0.01031	0.00723*	-0.02335***	0.01211**	-0.01199	0.00877
	(0.0063)	(0.00404)	(0.00861)	(0.0057)	(0.01004)	(0.008)
MEDINC	0.00708	-0.00525*	0.01561**	-0.00515	0.01165	-0.01183**
	(0.00513)	(0.00316)	(0.00616)	(0.00458)	(0.00769)	(0.00551)
MEDHOUSEVAL	-0.00121	0.00065	-0.00323***	0	-0.00227**	0.00018
	(0.00085)	(0.00056)	(0.00115)	(0.00082)	(0.00106)	(0.00102)
MEDGROSSRENT	-0.03566	0.0417	-0.0196	0.09613**	0.0249	0.08522
	(0.03991)	(0.02792)	(0.05493)	(0.04324)	(0.06385)	(0.05328)
UNEMP	-0.00675	0.01058	-0.00822	0.01122	-0.00629	0.01222
	(0.01518)	(0.00907)	(0.01946)	(0.01201)	(0.01955)	(0.01847)
LANDAREA	0.01484*	-0.00019	0.03514***	0.00531	-0.01828	0.01496**
	(0.00825)	(0.01023)	(0.0105)	(0.00635)	(0.01516)	(0.00682)
URBAN	0.66409	0.63890**	0.24275	0.31365	0.24972	-0.02376
	(0.43138)	(0.25144)	(0.54324)	(0.27901)	(0.57414)	(0.42294)
COSTOFLIV	0.01381	-0.00377	0.0163	-0.00901	0.01104	-0.01047
	(0.01026)	(0.00555)	(0.01056)	(0.0087)	(0.00906)	(0.00996)
PERCHISP	-0.09908	-0.0493	0.07279	-0.08623	0.31528	-0.22197
	(0.19013)	(0.11291)	(0.22928)	(0.17662)	(0.29563)	(0.25508)
PECINDIAN	-0.27189	-0.20483	-1.85919***	-0.21751	-2.32076*	-0.03854
	(0.7464)	(0.43051)	(0.70434)	(0.61094)	(1.30515)	(0.75862)
PERCASIAN	-0.82986	1.52454*	-1.31876	-0.09779	-0.57723	-0.69956
	(0.88508)	(0.79712)	(1.37513)	(0.77062)	(0.84299)	(0.83419)
PERCBLACK	0.11736	-0.19449	0.38898	0.01421	0.45293	-0.00063
	(0.24139)	(0.14429)	(0.36328)	(0.27447)	(0.42579)	(0.35805)
LOC_POP	0.00251*	-0.00069	0.01950***	0.00172	0.04151**	0.00924
	(0.00147)	(0.00106)	(0.00566)	(0.00431)	(0.01617)	(0.01012)
CONSTANT	2.74007***	2.62446***	2.20874*	2.80842***	1.04769	3.38256***
	(1.01219)	(0.57252)	(1.25919)	(0.86393)	(1.10814)	(1.04676)
N	348	348	326	326	258	258
R^2	0.45908	0.0711	0.58468	0.05944	0.50789	0.10184

*** p<0.01, ** p<0.05, * p<0.1, Robust standard errors in parentheses

References

Ali, A., Klasa, S., and Yeung, E. .2009. The limitations of industry concentration measures constructed with Compustat data: Implications for finance research. *Review of Financial Studies*, 22(10), 3839-3871.

Anglin, P. M. 1992. The relationship between models of horizontal and vertical differentiation. *Bulletin of economic research*, 44(1), 1-20.

Bens, D. A., Berger, P. G., and Monahan, S. J. .2011. Discretionary disclosure in financial reporting: An examination comparing internal firm data to externally reported segment data. *The Accounting Review*, 86(2), 417-449.

Berry, S. T., and J. Waldfogel. 1999. Public radio in the United States: does it correct market failure or cannibalize commercial stations?. *Journal of Public Economics* 71(2), 189-211.

Berry, S., and J. Waldfogel. 2003. Social Inefficiency in Radio Broadcasting. *RAND Journal of Economics* 30(3), 397-420.

Berry, S., and J. Waldfogel. 2010. Product Quality and Market Size. *Journal of Industrial Economics* LVIII(1), 1-31.

Besbes, O., and M. Scarsini. 2015. On information distortions in online ratings. Columbia Business School Research Paper, (13-36).

Breshnahan, T. F., and P. C. Reiss. 1991. Entry and Competition in Concentrated Markets. *Journal of Political Economy* 99(5), 977-1009.

Brynjolfsson, E., Y. Hu, and M. S. Rahman. 2009. Battle of the Retail Channels: How Product Selection and Geography Drive Cross-Channel Competition. *Management Science* 55(11), 1755-1765.

Brynjolfsson, E., L. M. Hitt, and H. H. Kim. 2011. Strength in Numbers: How Does Data-Driven Decisionmaking Affect Firm Performance?. Available at SSRN: http://papers.ssrn.com/sol3/Papers.cfm?abstract_id=1819486

Chen, Y., S. Fay, and Q. Wang. 2011. The Role of Marketing in Social Media: How Online Consumer Reviews Evolve. *Journal of Interactive Marketing* 25(2), 85-94

Chen, H., R. H. Chiang, and V. C. Storey. 2012. Business Intelligence and Analytics: From Big Data to Big Impact. *MIS quarterly*, 36(4), 1165-1188.

Chevalier, J. D., and D. Mayzlin. 2006. The Effect of Word of Mouth on Sales: Online Book Reviews. *Journal of Marketing Research* 43(3), 345-354.

Chidi, G. .2013. Confessions of a Corporate Spy. Available Online: http://www.inc.com/magazine/201302/george-chidi/confessions-of-a-corporate-spy.html.

Cohen, A., and M. J. Mazzeo. 2010. Investment Strategies and Market Structure: An Empirical Analysis of Bank Branching Decisions. *Journal of Financial Service Research* 38(1), 1-21.

Collard-Wexler, A. 2006. Isolated Town Data. Available Online: http://pages.stern.nyu.edu/~acollard/Data.html.

Cremer, H., and J. F. Thisse. 1991. Location models of horizontal differentiation: a special case of vertical differentiation models. *The Journal of Industrial Economics*, 383-390.

Cronin Jr, J. J., and Taylor, S. A. .1994. SERVPERF versus SERVQUAL: reconciling performance-based and perceptions-minus-expectations measurement of service quality. *The Journal of Marketing*, 125-131.

Dai, W., Jin, G. Z., Lee, J., and M. Luca (2014). Optimal Aggregation of Consumer Ratings: An Application to Yelp. com. *Harvard Business School NOM Unit Working Paper*, (13-042).

Dahl, D. .2011. 10 Tips on How to Research Your Competition. Available Online: http://www.inc.com/guides/201105/10-tips-on-how-to-research-your-competition.html

Dedman, E., and Lennox, C. .2009. Perceived competition, profitability and the withholding of information about sales and the cost of sales. *Journal of Accounting and Economics*, 48(2), 210-230.

Dellarocas, C. 2003. The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms. *Management Science* 49(10), 1407-1424.

Dellarocas, C., and R. Narayan. 2006. A Statistical Measure of a Population's Propensity to Engage in Post-Purchase Online Word-of-Mouth. *Statistical Science* 21(2), 277-285.

Dick, A. A. 2007. Market Size, Service Quality, and Competition in Banking. *Journal of Money, Credit and Banking* 39(1), 49-81.

Ellickson, P. 2007. Does Sutton Apply to Supermarkets? *RAND Journal of Economics* 38(1), 43-59.

Forbes. Yelp Is Worth $30 Per Share. Available online: http://www.forbes.com/sites/greatspeculations/2015/10/02/yelp-is-worth-30-per-share/#6387258f70dc.

Forman, C., A. Ghose, and A. Goldfarb. Competition Between Local and Electronic Markets: How the Benefit of Buying Online Depends on Where You Live. *Management Science* 55(1), 47-57.

Gallino, S., and A. Moreno. 2014. Integration of Online and Offline Channels in Retail: The Impact of Sharing Reliable Inventory Availability Information. *Management Science* 60(6), 1434-1451.

George, L., and Waldfogel, J. 2003. Who Affects Whom in Daily Newspaper Markets? *Journal of Political Economy* 111(4), 765-784.

Hu, N., P. Pavlou, and J. Zhang. 2009. Overcoming the J-Shaped Distribution of Product Reviews. *Communication of the ACM* 52(10), 144-147.

Huizinga, M. 2015. Could Great Yelp Ratings Make a Better Loan. Available Online: http://www.bankingexchange.com/blogs-3/unconventional-wisdom/item/5831-could-great-yelp-ratings-make-a-better-loan.

Jones, R., and H. Mendelson. 2011. Information Goods vs. Industrial Goods: Cost Structure and Competition. *Management Science* 57(1), 164-176.

Joo, M., K. C. Wilbur, B Cowgill, and Y. Zhu. 2014. Television Advertising and Online Search. *Management Science* 60(1), 56-73.

Kaushik, A. 2010. The Definitive Guide to (8) Competitive Intelligence Data Sources! Available Online: http://www.kaushik.net/avinash/competitive-intelligence-data-sources-best-practices/.

Kuruzovich, J., S. Viswanathan, R. Agarwal, S. Gossain, and S. Weitzman. 2008. Marketspace or Marketplace? Online Information Search and Channel Outcomes in Auto Retailing. *Management Science* 19(2), 182-201.

Lau, R. Y., S. S. Liao, K. F. Wong, and D. K. Chiu. 2012. Web 2.0 environmental scanning and adaptive decision support for business mergers and acquisitions. *MIS Quarterly* 36(4), 1239-1268.

Lee, T. Y., and E. T. Bradlow. 2011. Automated marketing research using online customer reviews. *Journal of Marketing Research* 48(5), 881-894.

Li, X., and L. M. Hitt. 2008. Self-selection and information role of online product reviews. *Information Systems Research* 19(4), 456–474.

Luca, M. 2011.Reviews, reputation, and revenue: The case of Yelp.com. *Harvard Business School Working Paper* 12-016. Available Online: http://www.hbs.edu/faculty/Pages/download.aspx?name=12-016.pdf.

Luca, M. and G. Zervas . 2016. Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud. *Management Science* (Articles in Advance).

Madansky, A. 1988. Prescriptions for working statisticians. Springer New York.

Mayzlin, D. 2006. Promotional Chat on the Internet. *Marketing Science* 25(2), 155-163.

Mazzeo, J. M. 2002. Product Choice and Oligopoly Market Structure. *RAND Journal of Economics* 33(2), 221-242.

Moe, W. W., and D. A. Schweidel. 2012. Online Product Opinions: Incidence, Evaluation, and Evolution. *Marketing Science* 31(3), 372-386.

Muchnik, L., S. Aral, and S. J. Taylor. 2013. Social Influence Bias: A Randomized Experiment. *Science* 341(6146), 647-61.

Mukherjee, A., V. Venkataraman, B. Liu, and N. Glance. What Yelp Fake Review Filter Might be Doing. *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media.*

Mundlak, Y. .1978. On the pooling of time series and cross section data. *Econometrica: Journal of the Econometric Society*, 69-85.

Neven, D., and J. F. Thisse. (1989). On quality and variety competition (No. CORE Discussion Paper (8920)). UCL.

Olivares, M., and G. P. Cachon. 2009. Competing Retailers and Inventory: An Empirical Investigation of General Motors' Dealerships in Isolated U.S. Markets. *Management Science* 55(9), 1586-1604.

Parasuraman, A., Zeithaml, V. A., & Berry, L. L. .1988. Servqual. Journal of retailing, 64(1), 12-40.

Perloff, J. M., and S. C. Salop. 1985. Equilibrium with product differentiation. *The Review of Economic Studies*, 52(1), 107-120.

Shaked A., and J. Sutton. 1982. Relaxing Price Competition Through Product Differentiation. *Review of Economic Studies* 49(1), 3-13

Shaked A., and J. Sutton. 1983. Natural Oligopolies. *Econometrica* 51(5), 3-13.

Shea, J. 1997. Instrument Relevance in Multivariate Linear Regression Models: A Simple Measure. *Review of Economics and Statistics* 79(2), 348-352.

Shuller, C. 2014. 5 Easy Ways to Conduct Your Market and Competitive Analysis Using Online Reviews. Available Online: http://www.reviewtrackers.com/5-easy-ways-conduct-market-competitive-analysis-online-reviews/.

Stock, J. H., and M. Yogo. 2005. Testing for Weak Instruments in Linear IV regression. in D. W. K. Andrews and J. H. Stock, eds. *Identification and Inference for Econometric Models: Essays in Honor of Thomas J. Rothenberg*, 80-108, New York: Cambridge University Press.

Sun, M. 2011. How Does the Variance of Product Ratings Matter? *Management Science* 58(4), 696-707.

Vogel, J. 2008. Spatial competition with heterogeneous firms. *Journal of Political Economy*, 116(3), 423-466.

Wooldridge, J. M. 1995. Score Diagnostics for Linear Models Estimated by Two Stages Least Squares. In: G. S. Maddala, P.C.B. Phillips, and T. N. Srinivasan (Eds.), *Advances in Econometrics and Quantitative Economics: Essays in Honor of Professor C. R. Rao* Blackwell, Cambridge, MA.

Wooldridge, J. M. .2008. Introductory Econometrics: A Modern Approach. 4 uppl. Mason, Ohio: South-Western Cengage Learning.

Yahoo Finance. Yelp Inc. Available online: http://www.bankingexchange.com/blogs-3/unconventional-wisdom/item/5831-could-great-yelp-ratings-make-a-better-loan [last accessed: February 11, 2016].

Zentner, A., M. Smith, and C. Kaya. 2013. How Video Rental Patterns Change as Consumers Move Online. *Management Science* 59(11), 2622-2634.

Zhang, Z., C. Guo, and P. Goes. 2013. Product comparison networks for competitive analysis of online word-of-mouth. *ACM Transactions on Management Information Systems (TMIS)*, 3(4), 20.