Social Learning and Local Consumption Amenities: Evidence from Yelp*

Elliot Anenberg† Chun Kuang‡ Edward Kung§

October 26, 2018

Abstract

Using data from Yelp, we show that consumers learn about restaurant quality from reviews, which means restaurants are more likely to go out of business when receiving poor reviews. Average restaurant quality thus becomes higher in areas with faster learning, which tend to be areas closer to the city center, and areas with younger and more educated populations. To quantify the effect of learning on equilibrium restaurant quality, we estimate a Bayesian learning model of consumer demand and restaurant exit. Simulations show that learning increases average restaurant quality by 0.25 Yelp stars in large markets and by 0.11 stars in small markets. This differential increase in restaurant quality associated with learning accounts for 0.9 percentage points of the house price difference between such areas. Our results have implications for the literature on gentrification and urban revival.

JEL Codes: D83, L15

*The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors.

†Board of Governors of the Federal Reserve System

‡Department of Economics, East Carolina University.

§Department of Economics, UCLA.

1 Introduction

We study the effect of "social learning" on the provision of local consumption amenities. Social learning refers to the ability for individual consumers to easily share their experiences with local goods, and for other consumers to easily access those opinions. In the last decade, social learning has grown exponentially thanks to the digitization of massive user-generated information. For example, Figure 1 shows that the cumulative number of reviews on Yelp increased from less than 10 million in 2008 to over 160 million in 2018.[1]

The increasing richness of the information environment in which consumers and producers operate likely has significant effects on the organization of markets. We hypothesize that social learning improves the average quality of local goods. The mechanism we propose is that social learning allows consumers to learn more quickly about the quality of local goods, making producer profits more sensitive to their true product quality. Faster consumer learning results in an acceleration of exit by low quality producers and enhanced survival and growth of high quality ones. Social learning thus increases long-run average product quality. We expect the intensity of learning, and therefore its effect on average quality, to vary across space. Quality will increase more in areas where learning is faster, such as in urban areas with a larger consumer base leaving reviews.[2]

In this paper, we focus on restaurants as the local good and Yelp as the social learning platform. The main contribution of our paper is to quantify the effect of social learning on average restaurant quality across different types of markets. We begin by documenting some empirical patterns that will motivate key features of our model. We use data from the 10th round of the Yelp Dataset Challenge. These data provide the entire history of reviews for all restaurants listed on Yelp in eight U.S. metropolitan areas, including information on if and when restaurants go out of business. We use variation in restaurant exit patterns and the arrival rate of new reviews to show that (i) consumers use reviews to learn about restaurant quality, (ii) existing consumer reviews have a direct impact on demand, (iii) learning intensity

[1]Yelp is an example of a review platform for local businesses that facilitates social learning. Other examples include TripAdvisor for travel and Rotten Tomatoes for movies.

[2]Social learning should also affect the quality of traded goods. However, the effect is likely to be larger for local goods because information on local goods is sparser to begin with, due to the limited geographic scope of the consumer base. Moreover, learning about traded goods will not generate a differential effect between areas with larger and smaller markets.

varies across markets, and (iv) average restaurant quality increases more in markets with faster learning, through the differential exit rate of high versus low quality restaurants.[3]

In our model, consumers decide whether to visit a restaurant based on the expected quality of dining at the restaurant, the uncertainty in the quality of the restaurant, as well as idiosyncratic factors. Some exogenous fraction of the consumers that visit a restaurant each period leave reviews on Yelp, which are unbiased but noisy signals of average quality. The noise is due to idiosyncratic consumer tastes or shocks to the dining experience. Other consumers rationally update their beliefs about restaurant quality using Bayes' Rule.

Forward-looking restaurants decide whether to exit or stay in business each period. Restaurants earn a constant marginal profit per customer and pay a fixed operating cost (i.e. rent) which can depend on market size. Each restaurant's expected profit therefore depends on both its market size and the consumers' beliefs about its quality, which is an endogenous function of Yelp reviews.

We estimate the model parameters using the Yelp micro data, and then quantify the effect of learning on the long-run average restaurant quality through counterfactual simulations. We consider two counterfactual learning environments: (1) *no* learning, and (2) a pre-Yelp learning environment where the arrival rate of reviews is calibrated to mimic the rate at which Zagat, the predominant restaurant review aggregator prior to Yelp, published restaurant reviews. Relative to the *no* learning counterfactual, social learning improves average restaurant quality by 0.15 Yelp stars in median-sized markets. This effect is not small, as 0.15 Yelp stars is 18% of the standard deviation of restaurant quality and 17% of the standard deviation of consumers' idiosyncratic tastes over restaurants. Moreover, the effect of learning is bigger in larger markets–it is 0.25 Yelp stars in "large" markets (i.e. the 95th percentile in market size). Relative to the pre-Yelp, Zagat learning environment, learning improves average restaurant quality by 0.09 Yelp stars in median-sized markets, 0.14 stars in large markets, and 0.07 stars in small markets.

We close the paper by estimating the effect of the increase in restaurant quality through social learning on consumer welfare. Existing research suggests that restau-

[3]Our finding that reviews affect demand is consistent with Anderson and Magruder (2012) and Luca (2011), who both use a regression discontinuity design to convincingly demonstrate an effect of Yelp reviews on restaurant demand, as measured by either bookings or revenues.

rants are an important component of local amenities, the value of which gets capitalized into house prices.[4] We therefore estimate the hedonic relationship between Yelp stars and house prices using micro data on housing transactions. We control for a rich set of housing and neighborhood characteristics, and to deal with any remaining endogeneity, we implement the Bajari et al. (2012) method of partialling out the effects of time-varying unboservables, which is itself an extension of the repeat sales method (Case and Shiller (1989)). Our estimates suggest that the 0.15 Yelp star increase in average restaurant quality in median-sized markets from social learning is associated with a 1 percent increase in house prices. Applying the hedonic estimate to the differential effects of learning between large and small zipcodes according to our counterfactual simulations, we find that social learning can explain a 0.9 percentage point difference in house prices between large and small zipcodes. This house price effect is equal to 4.2 percent of the difference in quality-adjusted house price differentials between such areas. Social learning therefore appears to be an important—but certainly not the only—mechanism generating higher prices in more urban neighborhoods.

Related Literature

This paper contributes most directly to the literature on how information technology (IT) interacts with cities. Gaspar and Glaeser (1998) first study whether IT would act as a substitute or complement for urban density, in the context of productivity spillovers. Sinai and Waldfogel (2004) pose a similar question in the context of consumption. Would IT act as a substitute for cities, by increasing the access in rural areas to product varieties that would normally only be available in cities? Or would IT act as a complement to cities, by providing information about goods and services that can only be consumed locally and are more densely provided in cities? By showing that social learning leads to increases in restaurant quality, we provide evidence of a specific way in which IT is complementary to the consumption benefits of urban density. Moreover, we show that even *within* cities, IT can have differential effects on larger versus smaller neighborhoods. Our results complement the findings in Anenberg and Kung (2015), who focus on a different channel through which IT interacts with density. Anenberg and Kung (2015) show that IT increases access to

[4]See, for example, Kuang (2017); Couture (2013); Diamond (2016); Glaeser et al. (2018).

food variety in cities by reducing locational uncertainty and facilitating the growth of the food truck industry.

The results of this paper have implications for the broader literature on gentrification and urban revival. Over the past couple of decades, downtown areas of most large American cities have experienced a revival (Couture and Handbury (2017)). A number of explanations for gentrification and urban revival have been proposed, including changing amenity valuation (Glaeser et al. (2001); Couture and Handbury (2017); Baum-Snow and Hartley (2016); Cosman (2017)), increasing value of time and the consequent disutility of commuting (Edlund et al. (2015); Su (2018)), and falling crime (Ellen et al. (2017)). Our paper suggests a mechanism for why amenities may be improving particularly in downtown areas. Although it has long been established that denser, downtown locations provide a greater number and variety of goods (George and Waldfogel (2003); Berry and Waldfogel (2010); Schiff (2015); Cosman (2017)), our paper highlights how social learning would increase quality of these amenities over time.

Our work is also related to several papers that either study the effect of Yelp on restaurants or use Yelp data to study broader economic outcomes. Using a convincing regression discontinuity design, Anderson and Magruder (2012) find that higher Yelp ratings of restaurants significantly decrease restaurant availability, consistent with our empirical finding and model assumption that Yelp reviews affect consumer demand. Luca (2011) uses a similar empirical design to show that high Yelp ratings of restaurants increase consumer demand and restaurant revenues. Luca (2011) also presents empirical evidence generally consistent with the type of Bayesian learning model that we use in this paper. Avery et al. (1999) and Acemoglu et al. (2017) study review platforms from a more theoretical perspective, focusing on the game theoretic basis for reviews and learning. Glaeser et al. (2015), Glaeser et al. (2017), and Kuang (2017) demonstrate that Yelp data could improve measurement of local amenities in both private and public goods markets. Davis et al. (2017) use Yelp data to measure segregation in urban consumption.

The findings in this paper are consistent with the broader literature on learning and product quality. Consistent with our findings, Jin and Leslie (2003, 2009) show that both consumers and restaurants are responsive to posted health grades, and that chain affiliation and regional variation in repeat customers affects responsiveness. Cabral and Hortaçsu (2010) show that sellers on Ebay are more likely to exit when they receive negative feedback, and that subsequent negative feedback is less impactful than the first.

Finally, our paper complements the recent work of Fang (2018), who studies the effect of consumer learning on restaurant quality using a structural demand model with a similar social learning process. Her focus is on quantifying the effects of faster learning on consumer surplus and restaurant revenues, necessitating a more detailed model of consumer demand than we have in our paper. Our focus is on quantifying the effect of faster learning on average restaurant quality, and the interaction of local market size with the speed of learning. We therefore have a dynamic model of restaurant exit, whereas Fang (2018) does not explicitly model restaurant behavior.

2 Data

Our main data source is the 10th round of the Yelp Dataset Challenge. Yelp is the leading online platform for consumers to post reviews about local businesses, with over 160 million consumer reviews for more than 1.5 million businesses across 30 countries.[5] Periodically, Yelp makes a subsample of its business and review data available to researchers as part of the "Yelp Dataset Challenge." To create the subsample, Yelp first selects a number of metropolitan areas, then provides data on all the business listings in that area with at least 3 reviews over its lifetime. Yelp provides information about the business, such as its name, street address, business categories (i.e. restaurant, cafe, groceries, etc), and the entire history of reviews for that business.[6] Yelp also reports an indicator for whether the business is out of business as of the last date in the sample (July 26, 2017). We do not observe the actual exit date of closed restaurants, and we assume that the exit occured in the interval between the date of the last review and July 26, 2017.[7] We also do not observe the actual entry date of each restaurant, and assume that it is the date of first review.[8] In some of

[5]Source: Business Wire and Yelp.

[6]Yelp filters out reviews that appear illegitimate (i.e. made by bots, shills, or competitors). Only reviews that make it past Yelp's review filter are included in the data. This is appropriate because only these reviews are made visible to consumers. See Luca and Zervas (2016) for a further discussion on the process and reliability of Yelp's review filtering algorithm.

[7]In this interval, zero reviews must have been received, else we would observe additional reviews for the restaurant during this interval.

[8]This is a fairly reasonable assumption as the data shows that newly opened restaurants tend to receive a lot of reviews in the first few months of business, likely due to a set of customers who like to try new restaurants.

our analysis, we distinguish between restaurant chains and non-chain independent restaurants. We identify chains using the business name, which is standardized in the Yelp data. If a business name appears at least twice in at least two different metro areas (thus a minimum of 4 appearances), we classify the business as a chain.

For this paper, we focus on restaurant listings (including bars and coffee shops) in the eight U.S. metropolitan areas included in the dataset: Phoenix, AZ, Las Vegas, NV, Charlotte, NC, Cleveland, OH, Pittsburgh, PA, Madison, WI, Urbana, IL, and Akron, OH. We identify a business as a restaurant if the word "restaurant", "bar", or "coffee" is contained in the business's listed categories. Businesses are assigned to metro areas based on their address zipcode and a zipcode-to-CBSA crosswalk provided by the Missouri Census Data Center.

Reviews start appearing in the data as early as late 2004, but we restrict the analysis to the period between January 2012 and July 2017, as restaurant listings in Yelp appear not to be fully representative prior to 2011. After 2011, we can be fairly confident that the Yelp sample presents an accurate picture of the restaurant choices facing consumers. A more complete discussion about the data is provided in the Appendix. Figure 2 shows that there is a high degree of correlation across zipcodes between the number of restaurants in Yelp and the number of restaurants reported in the Census Bureau's ZIP Code Business Patterns files.

Table 1 reports basic statistics for our final sample. Altogether, there are 31,397 businesses and 2,361,282 reviews in our sample. The second panel of Table 1 shows the distribution of star ratings for all reviews in the sample.

3 Empirical Facts

The Effect of Reviews on Restaurant Exit

We now document some empirical facts from the Yelp data that motivate our study. First, we consider the relationship between reviews and restaurant exit. To do this, we run regressions of the following form:

$$exit_{it} = \alpha_1 Stars_{it} + \alpha_2 \#Reviews_{it} + \alpha_3 Stars_{it} \times \#Reviews_{it} + X_{it}\beta + \epsilon_{it} \qquad (1)$$

where $exit_{it}$ is an indicator for whether restaurant i exits in month t, $Stars_{it}$ is the average Yelp stars for reviews cumulative up to month t, $\#Reviews_{it}$ is the number of reviews cumulative up to month t, and X_{it} is a vector of controls.[9]

In Table 2, we report the results for various specifications. In column (1), we do not include any controls. In column (2) we add fixed effects for the metro area, calendar year-month, months since opening, and price category ($-$$$$). In column (3), we estimate the equation using the sample of chain restaurants only, while in column (4), the equation is estimated using the sample of non-chains only.

The results show a statistically and economically significant effect of reviews on restaurant exit. In column (2), a one-star increase in the restaurant's rating reduces the probability of exiting in a month by 0.04 percentage point, which is a 8% reduction over the baseline exit rate of 0.5%. The effect of star rating also appears to be amplified by the total number of reviews. When there are only a few reviews, having a low star rating does not increase exit probability as much as if there are hundreds of reviews. A one standard deviation increase in the number of reviews increases the effect of a one-star increase in rating by an additional 0.08 percentage point. This suggests that consumers use reviews as a noisy signal of restaurant quality, and a greater number of signals leads to a more precise inference about restaurant quality. The coefficient on the number of reviews is positive because if the star rating is low, more reviews increases the probability of exit.

In columns (3) and (4), we consider the effect of reviews on chain and non-chain restaurants separately. The results show that the exit of chains is less sensitive to reviews than the exit of non-chains. This is likely because there is less ex-ante uncertainty about the quality of a chain restaurant than of a non-chain. Potential consumers of chain restaurants are therefore less likely to depend on reviews to learn about quality, which in turn would make chain restaurants' profits and exit decisions less sensitive to reviews. In estimating the learning model we present in Section 4, we use data on non-chains only.

Arrival Rate of Reviews

We are primarily interested in two questions about the arrival rate of reviews. First, do previous reviews affect the arrival rate of new reviews? If they do, under the

[9] We normalize the number of reviews to have mean zero and standard deviation 1. The unnormalized mean is 49.6 and the unnormalized standard deviation is 130. We run a linear probability model because graphical evidence suggests that exit rate is roughly linear in the star rating, and for easy interpretation of the coefficients. The main results are qualitatively robust to using a logit or probit model.

assumption that the arrival rate of reviews is correlated with the number of visitors, this would be additional evidence that reviews affect consumer demand. Second, how do the population characteristics of the restaurant's local neighborhood affect the arrival rate of reviews? We might expect that restaurants in areas closer to the city center, or with younger or more educated populations, may experience a higher arrival rate of reviews.

We run regressions of the form:

$$new\_reviews_{it} = \alpha_1 Stars_{i,t-1} + \alpha_2 \#Reviews_{i,t-1} + \alpha_3 Stars_{i,t-1} \times \#Reviews_{i,t-1} + X_{it}\beta + \epsilon_{it} \quad (2)$$

where $new\_reviews_{it}$ is the number of new reviews that restaurant i receives in month t. Included in X_{it} are the characteristics of the zipcode that i is located in. We include the following zipcode characteristics: log population, share of population between the ages 18 and 34, share of population with a bachelor's degree or higher, share of households with children under the age of 18, and distance to CBD.[10] In addition, we include in X_{it} fixed effects for the metro area, calendar year-month, months since opening, and price category of the restaurant ($-$$$$).

Table 3 reports the results for three specifications. In column (1) we include only the restaurant measures. We add the zipcode characteristics in column (2), and the full vector of controls in column (3). A robust result is that both the star rating and the number of reviews from the previous month positively affect the number of reviews received in the current month. If the number of reviews is positively correlated with the number of visitors, then this evidence suggests that positive reviews increase consumer demand. Moreover, the effect is increasing in the number of existing reviews, which is consistent with a consumer learning model. To provide some evidence that the number of reviews is indeed positively correlated with the number of visitors, we use data on the number of "check-ins" recorded by Yelp for each restaurant in the data over our sample period.[11] The correlation between number of check-ins—a

[10] All zipcode characteristics are from the 2008-2012 ACS 5-year estimates. We do not include any time varying zipcode characteristics.

[11] When a customer is eating at a restaurant, they can choose to "check in" on their smartphone, which can help the customer earn rewards and share their location with their friends on social media. In the data, the number of check-ins is greater than the number of reviews for over 90% of restaurants, suggesting that number of check-ins may be a better proxy for number of visitors than number of reviews. Unfortunately, check-ins are not timestamped in our data like reviews are, which makes it impossible to calculate the number of check-ins for a restaurant by time period.

proxy for number of visitors—–and number of reviews across restaurants is 0.85.

Column (3) shows that zipcodes closer to the city center, with a higher share of young adults, with a higher share of college educated individuals, and with a lower share of households with children, have a higher arrival rate of reviews. These are sensible patterns, as we would expect young adults, college educated individuals, and households without young children to visit restaurants more often, and perhaps be more active in posting reviews to Yelp.

Restaurant Quality and the Arrival Rate of Reviews

A main hypothesis of this paper is that average restaurant quality will be higher in areas with faster learning due to low quality restaurants exiting faster and high quality restaurants surviving longer. We now test these hypotheses directly in the data.

First, we test whether higher quality restaurants are ex-ante more likely to survive in areas where the arrival rate of reviews is higher. We estimate an accelerated failure time model of restaurant survival duration as follows:

$$\log t_i = \alpha_1 quality_i + \alpha_2 arate_i + \alpha_3 quality_i \times arate_i + X_i \beta + \epsilon_i \tag{3}$$

where t_i is the uncensored survival duration (in days) of restaurant i, $quality_i$ is the quality of restaurant i as measured by its average star rating on its last review date in our sample, and $arate_i$ is the "predicted arrival rate" of reviews of the restaurant's zipcode.[12] The "predicted arrival rate of reviews" is equal to the contribution of zipcode characteristics to the expected monthly arrival rate of reviews, as reported in column (3) of Table 3. X_i includes the additional controls: a dummy for whether restaurant i is a chain, a CBSA fixed effect, and a price-category fixed effect. ϵ_i is modeled as an exponential distribution, though the results are robust to different assumptions about ϵ_i. The observed survival duration is censored at July 26, 2017, the time of data release.

Table 4 reports the results from three specifications of model (3). Column (1) is estimated on the full sample of restaurants. Columns (2) and (3) report results estimated from only the sample of chains and non-chains, respectively. The results show

[12] The average star rating as of the last review date is typically associated with many reviews, therefore, is a reasonable proxy for the restaurant's true quality.

that on average, for non-chains, higher quality restaurants survive longer. Restaurants in zipcodes with higher arrival rate exit faster, possibly because such areas tend to be more educated and closer to the CBD, indicating higher rents and operating costs. The most important term is the interaction between the star rating and the zipcode review arrival rate, which is estimated to be positive. The results in column (3) imply that increasing a non-chain restaurant's rating by one star increases its expected survival duration by 21%, and then by an additional 7% if the arrival rate of reviews is one standard deviation higher. The effect of quality and arrival rate on survival only holds for non-chain restaurants. None of the results in column (2), which are estimated with chain restaurants only, are statistically significant. These findings support our hypothesis that higher quality restaurants are more likely to survive than low quality restaurants when the arrival rate of reviews is higher, and that the effect is due to consumer learning.

Finally, we show that the differential survival rate of high versus low quality restaurants leads to higher overall quality of restaurants in areas with higher review arrival rate. This result does not mechanically follow from the results reported in Table 4 because it is possible that restaurants that exit in high arrival rate zipcodes are replaced by low quality restaurants. Panel A of Figure 3 shows a local polynomial fit of the average quality of non-chain restaurants that are open as of July 2017 in a zipcode against the zipcode's predicted arrival rate. Panels B and C of Figure 3 plot, respectively, a local polynomial fit of the *change* in average quality and number of open restaurants from January 2012 to July 2017 against the zipcode's predicted arrival rate. Both the change in average restaurant quality and change in number of restaurants are increasing in the zipcode's predicted arrival rate of reviews, suggesting a possible role for the growth of Yelp in the improvement of local consumption amenities.

4 A Model of Consumer Learning and Restaurant Exit

The results presented in Section 3 show that consumers learn from reviews, reviews affect restaurant survival, and faster learning can lead to higher average restaurant quality. We next develop and estimate a structural model of consumer learning and restaurant behavior that is consistent with these empirical patterns. The model will allow us to quantify the effect of learning on restaurant quality for different types of markets through counterfactual simulations.

Reviews, Quality, and Consumer Demand

Consider a restaurant that enters the market with quality q. In practice, restaurants are heterogeneous in quality, so q for new entrants is assumed to be drawn from a normal distribution with parameters (μ_q, σ_q^2). Each consumer will have their own idiosyncratic perception of quality for a given restaurant, but q represents the utility of dining at the restaurant according to the preferences of the average customer.

Each period, the restaurant faces a mass N of potential customers indexed by i. Each customer decides whether or not to visit the restaurant. The utility of visiting the restaurant is $q_i = q + z_i$, where $z_i \sim N(0, \sigma_z^2)$ is an idiosyncratic taste shifter or shock to the dining experience. q_i is only revealed to the consumer after she visits the restaurant. The utility of not visiting the restaurant is $q_0 + \sigma\epsilon_i$, where ϵ_i is independently distributed across potential customers according to $U[0, 1]$ and σ is a scaling parameter.[13] Unlike z_i, ϵ_i is known by the consumer before deciding to visit the restaurant. We assume that consumers have exponential utility over q_i, leading to a a mean-variance form of expected utility. A consumer will visit the restaurant if:

$$E\left[q_i\right] - \frac{1}{2}\gamma Var\left[q_i\right] \geq q_0 + \sigma\epsilon_i \tag{4}$$

where γ measures the degree of consumer risk aversion. Consumers that visit the restaurant leave reviews at rate λ. The probability of leaving a review is the same regardless of the realized experience q_i. Moreover, customers are honest in their reviews, and simply report q_i. The distribution of review scores for the restaurant is therefore normal with mean q and variance σ_z^2.

Consumers do not know q when deciding whether or not to visit the restaurant. Instead, they update their beliefs about q by applying Bayes' rule to the observed distribution of review scores. Since the consumer's prior beliefs and the signal noise on reviews are both assumed to be normally distributed, consumers' beliefs after

[13]The utility of not visiting the restaurant can be interpreted as the utility of the next best option, such as eating at home, or going to another restaurant. In the Appendix, we show how this specification approximates a more detailed model of spatial competition between restaurants.

applying Bayes' rule will also be normally distributed because of the self-conjugacy of the normal distribution. If Q is the average of the review scores and R is the number of reviews, then the mean and variance of beliefs are:

$$E\left[q_i|Q,R\right]=\frac{\sigma_z^2\mu_q+\sigma_q^2RQ}{\sigma_z^2+\sigma_q^2R} \tag{5}$$

and:

$$Var\left[q_i|Q,R\right]=\frac{\sigma_z^2\sigma_q^2}{\sigma_z^2+\sigma_q^2R}+\sigma_z^2 \tag{6}$$

The number of customers visiting the restaurant with an average review score of Q and a number of reviews of R is therefore:

$$n(Q,R)=\frac{1}{\sigma}\left(E\left[q_i|Q,R\right]-\frac{1}{2}\gamma Var\left[q_i|Q,R\right]-q_0\right)N \tag{7}$$

$n(Q,R)$ is increasing in both Q and RQ. Intuitively, a high review average increases market share and a low review average decreases it, and the effect will be stronger when the number of reviews is higher because the degree of certainty is higher. The direct effect of R is ambiguous. On the one hand, it may increase market share by decreasing the uncertainty in quality that the consumer faces; on the other hand, it can decrease market share if Q is low relative to μ_q.

Arrival Rate of Reviews

We assume that λ is small, so that $\lambda n(Q,R)$ is small relative to $n(Q,R)$. The distribution of the number of new reviews can therefore be approximated by a Poisson process with arrival rate $\lambda n(Q,R)$. In addition, we assume an additive constant to the arrival rate of reviews, λ_0, which represents a stream of reviews by customers who deliberately seek out restaurants to leave reviews for, independent of the restaurant's quality.[14] We again assume that λ_0 is small relative to $n(Q,R)$ so that this stream of reviewers does not measurably impact profits, but may affect the review arrival rate. The probability mass function of the number of new reviews, ΔR, conditional on R

[14]Yelp incentivizes this behavior through its Yelp Elite program, which rewards active reviewers with access to special events and other perks.

and Q is therefore:

$$P(\Delta R|Q,R) = \frac{(\lambda_0 + \lambda n(Q,R))^{\Delta R} e^{-(\lambda_0 + \lambda n(Q,R))}}{(\Delta R)!} \tag{8}$$

which is simply the Poisson distribution.

Restaurant Exit

We assume that the restaurant makes a constant marginal profit per customer, π. Each period, the restaurant must pay a fixed operating cost $c-\xi$ to stay in the market, where ξ is a temporary cost shock each period drawn from a logistic distribution. The restaurant's profit when the reviews are (Q,R) is therefore:

$$\Pi(Q,R) = n(Q,R)\pi - c + \xi \tag{9}$$

The present value of profits associated with exiting the market is normalized to zero.

Restaurants are forward-looking. We assume that they know their true quality, q, and make exit decisions based on their flow profits as well as the expectations of future profits. Since flow profits are determined by consumer beliefs about quality, the true quality q has no direct effect on flow profits once (Q,R) are accounted for. However, q affects the distribution of future reviews, and thus the restaurant's expectations for future profits. For example, a restaurant that has a high q but low Q and low R due to some unlucky reviews will tend to stay in business even though flow profits are low today because the restaurant rationally expects Q to move up to q in the future as more reviews accumulate.

Let $V(Q,R,q)$ be the expected net present value of staying in the market when the reviews are (Q,R) and q is true quality, net of the cost shock ξ. We can write:

$$V(Q,R,q) = n(Q,R)\pi - c + \beta E\left[\max\left\{0,\ V(Q',R',q) + \xi'\right\}\ |\ Q,R,q\right] \tag{10}$$

where the expectation is taken over next period's reviews (Q',R'), and next period's cost shock, ξ'. In the current period, the restaurant will exit if $V(Q,R,q) - \xi \leq 0$. The probability of exit is therefore:

$$e(Q,R,q) = \frac{1}{1 + e^{V(Q,R,q)}} \tag{11}$$

5 Implementation and Estimation

Parameterization

To capture differences between markets that may be unrelated to social learning, we allow the following parameters to also depend on market size: (i) the average quality of new entrants, μ_q, (ii) the constant term in the review arrival rate, λ_0, and (iii) the fixed operating cost c. We write:

$$\mu_q = \mu_{q,cons} + \mu_{q,N} \times N \tag{12}$$
$$\lambda_0 = \lambda_{0,cons} + \lambda_{0,N} \times N \tag{13}$$
$$c = c_{cons} + c_N \times N \tag{14}$$

The arrival rate of new reviews becomes:

$$\begin{aligned}\lambda_0 + \lambda n(Q,R) &= \lambda_{0,cons} + \left(\lambda_{0,N} - \frac{\lambda q_0}{\sigma}\right) N + \frac{\lambda}{\sigma} E\left[q_i|Q,R\right] N - \frac{\lambda\gamma}{2\sigma} Var\left[q_i|Q,R\right] N \\ &= \theta_1 + \theta_2 N + \theta_3 E\left[q_i|Q,R\right] N + \theta_4 Var\left[q_i|Q,R\right] N \end{aligned} \tag{15}$$

and the flow profits become:

$$\begin{aligned} n(Q,R)\pi - c &= -c_{cons} - \left(c_N + \frac{\pi q_0}{\sigma}\right) N + \frac{\pi}{\sigma} E\left[q_i|Q,R\right] N - \frac{\pi\gamma}{2\sigma} Var\left[q_i|Q,R\right] N \\ &= \theta_5 + \theta_6 N + \theta_7 \left(E\left[q_i|Q,R\right] + \frac{\theta_4}{\theta_3} Var\left[q_i|Q,R\right]\right) N \end{aligned} \tag{16}$$

The goal of our estimation will be to estimate the 11 parameters $(\mu_{q,cons}, \mu_{q,N}, \sigma_q, \sigma_z, \theta_{1:7})$.

Data

For estimation, we make a few additional restrictions to the sample described in Section 2. First, we include data on non-chain restaurants only. As is showed in Section 3, the economics of chains and non-chains appear to be quite different. Therefore, we focus on modeling the more information-sensitive restaurant segment. Second, we drop restaurants that ever received more than 30 reviews in a month, which is less than 2% of all observed restaurants. These restaurants may already be famous local fixtures and could face different incentives from other restaurants (some of these receive up to 200 reviews in a month).

An important variable in our model is the market size N. A priori, it is not obvious how this should be measured. Theoretically, it should be an aggregation of the population in the city, weighted by their distance to the restaurant, their propensity to eat at restaurants, and by the amount of competing alternatives.[15] Given that our model predicts a review arrival rate with a linear term in N, we measure N using the zipcode "predicted arrival rate" of reviews as calculated in Section 3. This is a linear combination of the zipcode's population demographics and distance to city-center that is likely to be highly correlated with the restaurant's nearness to the city's restaurant-going population. In unreported results, we employ an alternative measure for N, the number of people living within a 15 minute drive of the restaurant, computed using ACS 2008-2012 zipcode populations and driving times from the Google Distance Matrix API. This measure is motivated by Couture (2013)'s calculation from the NHTS that the average travel time for trips to restaurants is 15 minutes. Our final results are robust to both measures of N.

The final dataset can therefore be described as, for each restaurant i, $D_i = \{\mathbf{q}_{i,t}, e_i, T_i, N_i\}_{t=t_i}^{t=T_i}$ where $\mathbf{q}_{i,t}$ is the set of all review scores accumulated for restaurant i up to time t, e_i is an indicator for whether the restaurant is out of business on July 26, 2017, T_i is the date of last review, N_i is the market size, and t_i is the minimum of January 2012 and the date of the restaurant's first review.

Estimation

Estimation proceeds in three steps. In step 1, we estimate the parameters governing the distribution of entrant quality and idiosyncratic tastes $(\mu_{q,cons}, \mu_{q,N}, \sigma_q^2, \sigma_z^2)$ directly from the distribution of reviews. We then use these estimates to construct consumers' expected value and variance of restaurant quality $E\left[q_{it}|Q_{it}, R_{it}\right]$, $Var\left[q_{it}|Q_{it}, R_{it}\right]$ for each restaurant in each period. In step 2, we estimate the parameters governing the arrival rate of review, $(\theta_1, \theta_2, \theta_3, \theta_4)$, using maximum likelihood, where the likelihood is constructed by plugging the observed arrival rate of reviews into the Poisson probability mass function (8). In step 3, we estimate the parameters governing the exit decisions of the restaurants, $(\theta_5, \theta_6, \theta_7)$ by maximum likelihood, where the likelihood is formed by plugging the observed exits into the exit proba-

[15] Although we do not explicitly model spatial competition, in the Appendix we describe a model of spatial competition in which the number of customers is approximately linear in $E[q_i|Q,R] - \frac{1}{2}\gamma Var[q_i|Q,R]$, as in equation (7).

bilities (11). We provide the technical details of the estimation procedure, including regression equations and likelihood functions, in the Appendix.

Estimation Results

Table 5 reports our parameter estimates. They are sensible in that θ_3 is positive and θ_4 is negative, implying that consumer demand is increasing in expected quality but decreasing in variance. The implied risk aversion parameter is $\gamma = 4.2$, in units of Yelp stars.[16] θ_7 is estimated to be positive, implying that profits are higher, and thus restaurants are more likely to stay in business, when expected quality is higher and variance is lower. θ_6 being negative implies that operating costs (such as rents) are higher in large markets than in small markets. Finally, a positive $\mu_{q,N}$ implies that larger markets have slightly higher quality entrants.

Highlighting the Economic Mechanisms

Figure 4 shows the value function of staying in business, $V(Q, R, q)$, evaluated at baseline estimates, by current star rating Q and number of reviews R. The value function is evaluated at a true quality of $q = 3.75$ and at the median market size N. The value function is increasing in Q due to the direct effect of Q on restaurant flow profits. The effect of R on the value function depends on the value of Q. For high Q, a higher R increases the value function, while for low Q, a higher R decreases the value function. This is because at a low Q, more reviews makes consumers more confident that a restaurant with a low Q is indeed low-quality, and vice versa. However, it is highly unlikely that a restaurant with $q = 3.75$ would ever find itself in the very bottom-right of Figure 4 (very low Q, very high R) because we assume that reviews are unbiased. Note also that the marginal effect of more reviews is declining in R. When there are already a large number of reviews, a couple of more reviews do not change the posterior by much, but when there are few reviews, a small number of new reviews will.

Figure 5 shows the probability of exit by true quality and market size when a restaurant has a current star rating of $Q = 3.7$ and a number of reviews $R = 20$. The

[16] Variance in perceived quality ranges from about 1.5 Yelp stars (when there are zero reviews) to about 0.8 Yelp stars (when there are a large number of reviews—the remaining variance is due to the idiosyncratic variance σ_z^2) so the implied quality-variance tradeoff going from no information to full information is about 1.47 Yelp stars.

figure shows that higher quality restaurants are less likely to exit, even conditional on the current set of reviews, because they know they are high quality and are likely to receive better reviews going forward. Without the dynamic decision process for restaurants in our model, the lines in Figure 5 would be flat. The slope of the exit rate with respect to true quality is increasing in the market size. This is because in larger markets, the arrival rate of reviews is higher. Thus, high quality restaurants expect their review scores to improve more quickly over time. To illustrate this intuition, Figure 6 shows the probability of exit by true quality when the number of reviews is $R = 120$. At this high number of reviews, the average star rating will not subsequently change much in any type of market, and thus the slope differences for markets with different review arrival rates shrink.[17] Level differences due to differences in operating costs still exist.

Consistent with the evidence presented in Section 3 and Table 4, Figure 7 shows that the model predicts high quality restaurants to be more likely to survive in larger markets, and low quality restaurants to be less likely to survive in larger markets. How much of the survival differences by market size in Figure 7 are driven by the learning channel as opposed to differences in operating costs or entrant quality distributions? Figure 8 replicates Figure 7 for a counterfactual set of parameters, in which operating costs and entrant quality distributions are assumed not to vary by market size. Instead, the fixed operating costs and the mean entrant quality is set to that of the median market size. Figure 8 shows that when operating cost differences and entrant quality differences are eliminated, the difference in survival rate of high versus low quality restaurants is still greater in larger markets than in smaller markets. Thus, the learning channel will still generate large differences in the ex-post quality distributions of restaurants by market size, even when cost differences and ex-ante differences are eliminated.

6 The Effect of Social Learning on Restaurant Quality

In this section, we quantify the effects of social learning on long-run restaurant quality by simulating average restaurant quality under counterfactual information environ-

[17]The larger number of reviews also decreases consumer uncertainty, which increases consumer demand all else equal, and contributes to the level differences between Figures 5 and 6.

ments.

Counterfactual Information Environments

First, we consider a *No learning* environment, where the arrival rate of reviews is set to zero. When the arrival rate of reviews is zero, demand will not be sensitive to true quality because it is not possible for consumers to learn about true quality before visiting the restaurant. All differences in restaurant exit behavior across markets will be driven purely by differences in operating costs across markets.

Second, we consider a scenario which likely approximate the pre-Yelp social learning environment. We call this the *Zagat* environment, referencing the restaurant review brand Zagat which was the predominant restaurant review aggregator before Yelp. Before being acquired by Google in 2011, Zagat published regional restaurant ratings in region-specific guidebooks once a year. Zagat ratings are compiled from the reports of hundreds of restaurant-goers who are recruited by Zagat specifically for this purpose, but who are not professional restaurant critics. Zagat ratings are therefore similar to Yelp in the sense that they are generated by consumers, but the coverage is much lower and ratings are published only once a year. Moreover, the information can only be accessed by purchasing the guide, or by accessing it through a library. Taken together, it can be argued that compared with Zagat, Yelp substantially increased the accessibility of restaurant reviews with negligible costs to the consumer.

To calibrate the second counterfactual, we assume that restaurants receive a Zagat rating at a rate that is calibrated to match Zagat's coverage of restaurants before online review platforms became widely used. Based on the 1998 Zagat guide for restaurants in Arizona and New Mexico (Zagat Survey (1997)) roughly 6% of restaurants in the Phoenix area (the largest metro area in our sample) receive a Zagat rating once per year.[18] Since each rating from Zagat reflects reports from potentially hundreds of diners, we set the signal variance, σ_z^2, of the Zagat rating to zero.[19] Thus, we assume that receiving a Zagat rating fully reveals the restaurant's true quality. This assumption likely overstates the amount of information provided by Zagat. In the Appendix, we discuss some other assumptions that we make in the calibration

[18] We describe this calibration, including data sources used, in more detail in the Appendix.

[19] The variance of idiosyncratic taste shocks remains the same, leading to a difference in the signal variance and the taste variance in the counterfactual.

exercise, and why they also likely overstate the amount of information available in the pre-Yelp period.

We fully solve the dynamic programming problem of the restaurant under both counterfactual information environments, so that restaurants have correct expectations about the rate at which they receive future reviews and about the effect of reviews on consumer demand. To simulate long-run restaurant quality under each information regime, we draw an initial 10,000 restaurants from the potential entrant distribution and simulate an entry process to be described in the next subsection. We then simulate the review and exit paths for these restaurants according to our model for 72 periods, i.e. 6 years, which is the length of the sample period from which we estimated the model. If a restaurant exits, it is immediately replaced by another restaurant, drawn again from the entrant quality distribution. We evaluate the average quality of open restaurants after 72 periods. We conduct the simulations for various market sizes, using the 5th, 50th, and 95th percentile of market size distribution in our data.

Accounting for Selective Entry

In our simulations, we assume that exiting restaurants are immediately replaced by a new one. For simulations under our baseline information environment, we can simply draw restaurant entrants using our paramater estimates of μ_q and σ_q^2—i.e. from the estimated distribution of actual entrants. However, when the information environment changes, the distribution of entrants is likely to change also due to selective entry, thus making the parameter estimates of μ_q and σ_q^2 invalid. For example, if there is no learning, then lower quality restaurants should be more likely to enter. Therefore, to account for the effects of selective entry in our counterfactual simulations, we also model an entry decision. In addition, we need to determine the distribution of *potential* entrants that will face the entry decision.

For the entry process, we assume that a potential entrant of quality q enters if:

$$V^{enter}(q) + \xi = -\sum_{t=0}^{5} \beta^t c + \beta^6 V\left(0, 0, q\right) + \xi \geq 0 \tag{17}$$

That is, to enter, the restaurant must pay an entry cost equal to 6 months of operating cost plus an entry cost shock ξ. The restaurant will enter with zero reviews, so the value function associated with being in business at entry can be denoted $V(0, 0, q)$

where V is defined in equation (10). The entry cost shock ξ is i.i.d. logistic, so the probability of entering is:

$$P(enter|q) = \frac{e^{V^{enter}(q)}}{1 + e^{V^{enter}(q)}} \tag{18}$$

V in equation (17) will vary with the information environment, and so the distribution of actual entrants will vary in each counterfactual.

To estimate the quality distribution of potential entrants, $P(q)$, that will face the entry decision just described, we note that we can write:

$$P(q) \propto \frac{P(q|enter)}{P(enter|q)} \tag{19}$$

Note that under our baseline information environment, $P(q|enter)$ is determined directly from our estimates of μ_q and σ_q^2, and $P(enter|q)$ can be computed from baseline parameter estimates using equation (18). Therefore, we can recover $P(q)$ under our baseline information environment. We assume that the distribution of potential entrants, $P(q)$, is fixed across the counterfactual information environments. We now have all the pieces we need to conduct the simulations accounting for selective exit. In each information environment, consumers are assumed to have correct prior beliefs about the mean and standard deviation of the actual entrant quality, though they continue to approximate the distribution as normal.

Simulation Results

Table 6 presents the results of the counterfactual simulations. When there is no learning, long-run average restaurant quality is decreased by 0.15 Yelp stars in median-sized markets, relative to the baseline with learning. The magnitude of this effect is significant. Recall from Table 5 that the standard deviation of the entrant quality distribution is 0.84. Therefore, the increase in restaurant quality attributed to learning in the typical market is equal to 18% of the standard deviation in restaurant quality. The effect of learning on restaurant quality is larger in large markets. Eliminating learning reduces long-run average restaurant quality by 0.25 Yelp stars in the 95th percentile sized markets, and by only 0.11 Yelp stars in the 5th percentile sized markets. The difference in effects between large and small markets (i.e. the difference-in-differences) is 0.14 Yelp stars.

When learning is calibrated to reflect the pre-Yelp, Zagat environment, long-run average restaurant quality is reduced by 0.09 Yelp star in median-sized markets, 0.14 star in large markets, and 0.07 star in small markets, relative to our baseline with social learning. The difference-in-differences between large and small markets is 0.07 Yelp star. The results suggest that learning significantly improves restaurant quality over time, and can accelerate the differences in amenity quality across neighborhoods of different market sizes.

7 Welfare Implications

How important are the differences in Yelp stars reported in Table 6? In this section, we estimate a hedonic regression of house price on housing and neighborhood characteristics, including the average Yelp stars of restaurants in the house's zipcode. Under some strong assumptions about utility-maximizing behavior, the estimate of the effect of restaurant quality on house prices from this regression can be used to infer consumers' marginal willingness to pay for a change in average restaurant quality (Rosen (1974)). Even if those assumptions are relaxed, a positive hedonic coefficient would imply that restaurant quality is capitalized into house prices and therefore valued by the typical home buyer (Bayer et al. (2007)).

The equation we estimate is:

$$lnp_{ijt} = \alpha AvgStars_{jt} + \boldsymbol{X}_{ijt}\beta + \xi_{it} \tag{20}$$

where lnp_{ijt} is the log price of house i in zipcode j in year t. $AvgStars_{jt}$ is the average star rating of open restaurants in $i's$ zipcode j in year t.[20] $\boldsymbol{X}_{ijt}$ is a vector of controls such as the characteristics of the house and neighborhorhood, and time effects to capture volatility in housing values. In some of our specifications, X will include a house fixed effect. ξ_{ijt} reflects omitted or unobserved house or neighborhood characteristics that affect house prices. α is the paramater of interest.

A standard concern with OLS estimation of equation (20) is that improvements in omitted local amenities (e.g. changes in local crime rate, school quality, economic opportunities) are positively correlated with improvements in local restaurant quality,

[20]The average star rating for each restaurant is measured as the average rating at its last review date in our data. We have also tried estimating this equation using the restaurants' inferred quality as implied by the Bayesian learning model, and the results are similar.

leading to upward bias in our estimate of α. If the omitted attributes are time-invariant, then equation (20) can be consistently estimated by including house fixed-effects that absorb any time-invariant heterogeneity (i.e. the repeat sales method). However, it may be too strong an assumption that omitted variables such as school quality are not changing while the variable of interest, restaurant quality, is. We therefore follow the approach suggested by Bajari et al. (2012) (BFKT), which is an extension of the repeat sales idea to control for time-varying unobserved heterogeneity, to consistently estimate α.

The main assumption of BFKT that allows for consistent estimation of equation (20) is that home buyers rationally form expectations about the evolution of omitted housing attributes. Therefore, in a differenced version of equation (20) to be estimated using repeat-sales, the error term ξ_{it} is replaced with the innovation in ξ_{it}, which has both an expected and an unexpected component. The expected component can be written as a function of lagged observed characteristics and lagged price, while the unexpected component is uncorrelated with lagged observables, due to rational expectations. We fully describe the implementation of BFKT in the Appendix, which reduces to a simple two-stage nonlinear least squares estimator.

To estimate equation (20), we obtain data on home prices and house characteristics from housing transactions data collected by CoreLogic. The estimation sample consists of repeat sales between 2009 and 2015, from the eight metro areas present in our Yelp data.[21]

Table 7 presents the results. In column (1), we simply estimate equation (20) using OLS without any controls. The estimate of α implies that a one-star increase in average zipcode restaurant quality is associated with a 38.6 percent increase in house prices. Most likely, this estimate is biased upward because restaurant quality is positively correlated with the error term in the regression. Consistent with this interpretation, when we add the controls in column (2), the estimate of α drops significantly, though still large. The complete list of control variables can be found in the notes to Table 7. For example, the controls include house characteristics like size of the house and fixed effects for the year of sale. In column (3), we add house fixed effects, which can be achieved without a loss of observations since we are estimating on a sample of repeat sales. The estimate of α drops further to 0.096. In column

[21] Even though we have Yelp data through 2017, our estimation sample stops in 2015 because it is the final year of data in our Corelogic sample.

(4), we present our preferred results using the BFKT method. A one-star increase in average restaurant quality is associated with a 6.4 percent increase in house prices. We provide a robustness check on this point estimate using an alternative procedure in the Appendix.

From the estimate in column (4), we see that a 0.15-star increase in average restaurant quality—which is the effect of learning on restaurant quality in a median-size zipcode—is associated with a 1 percent increase in house prices. Focusing on the effect of learning relative to the pre-Yelp learning environment, we find that the 0.09 Yelp star increase in average quality is associated with a 0.58 percent increase in house prices. Learning generates a 0.14 Yelp star differential in restaurant quality between the 5th and the 95th percentile market, which translates to a 0.9 percent price differential. The price of a constant-quality home in the 95th percentile of the market size distribution is 21.7 percent higher than the price in the 5th percentile.[22] Social learning therefore explains 4.2 percent of the cross-sectional gradient in house prices by market size.

Discussion

We now discuss some of the equilibrium responses that our model and results in Sections 4 and 6 abstract from, and we explain why taking them into account should amplify the effects of learning on both the quality of restaurants and the implied welfare consequences. We therefore believe that the welfare effects of learning that we estimate should be interpreted as lower bounds on the effects of learning.

1. Change in the competitive environment

When the learning environment changes and average restaurant quality changes, so too will the competitive environment that each restaurant faces. If learning increases and restaurant quality goes up, then this puts even more competitive pressure on low quality restaurants. This should amplify the effects of learning because there will be a positive feedback effect between learning, restaurant quality, and the ability for low quality restaurants to survive. This process should amplify the effects of

[22]We obtain data from Zillow on median price per square foot in December 2017 for each zipcode in our Yelp sample. We flexibly regress the log price on "predicted arrival rate" of reviews (our zipcode market size variable) at the zipcode level. 21.7 percent is the difference in the predicted value from this regression in the 95th compared to the 5th percentile market size zipcode.

learning on consumer welfare from restaurants, especially if the supply of higher quality restaurants is very elastic.

2. Change in operating costs and market expansion

When the learning environment changes, it changes the expected profits of all restaurants. In particular, it should increase the overall market for restaurants because of the reduction in uncertainty that consumers face. If space is limited, this may put competitive pressure on the rental price for floor space, making it more difficult for all restaurants to survive, but especially the low quality ones. If space is elastic, the market should expand, leading to higher total welfare.

3. Change in demographic composition

Finally, it is possible that learning changes the demographic composition of different neighborhoods. In fact, we have argued that the improvement to consumption amenities driven by social learning can be one factor explaining neighborhood change. If increased learning leads to higher restaurant quality in certain neighborhoods, and individuals who particularly value going out to restaurants move to those neighborhoods, this would lead to even faster learning and a positive feedback effect in such neighborhoods.

4. Improvements to other local amenities

Our model has focused only on the effect of learning on restaurant quality. While restaurants are surely an important component of local amenities, they are not the only one. Faster learning can improve average quality and reduce quality uncertainty for other local goods as well (i.e. entertainment, services), leading to even larger effects on welfare.

5. Improvements to variety

Learning may increase restaurant variety. If consumers have more uncertainty about certain types of cuisines, learning can increase the incentives for high quality restaurants to open with more specialized menus. Since consumers tend to have a taste for variety, this is another way in which learning can increase welfare.

6. Investment in quality

Our model assumes that a restaurant's quality is fixed over time. In practice, restaurants may be able to improve their quality. Learning should increase the returns to investing in quality improvements because learning makes profits more sensitive to true quality. That said, we do not find much empirical evidence for dynamic quality investment. The distribution of changes in average star rating within restaurants is centered around zero, which is more consistent with signal noise than the result of quality investment. Furthermore, the variance of this distribution closely fits the variance predicted by our estimated model, which only allows for variation in average star rating through the variance in consumer tastes, σ_z^2.

8 Conclusion

We show that social learning increases the average quality of restaurants, an important component of local consumption amenities. The effect of social learning is stronger in areas with faster learning, which tend to be areas closer to the city center, and areas with younger and more educated populations. The results suggest that growth in social media and related information technologies may be a contributing factor in gentrification and urban revival.

Understanding the factors driving gentrification and urban revival has important implications for policy on how to address their more controversial consequences, such as rising rental burdens in urban areas. For example, if gentrification is primarily driven by a rise in the value of time and the disutility of commuting, then policy ought to be targeted at improving transportation networks to reduce the commuting cost of living far from the city center. Alternatively, this paper suggests that at least one of the factors driving urban revival is fundamental to the interaction of information technology and density. To the extent that some of the differences in house prices across space reflect differences in amenities, then the improvement in IT may imply a permanent increase in the value of living in denser areas.

In this paper, we focus on a specific channel through which IT contributes to the consumption benefits of cities. We would like for future work to quantify the broader effects that IT is having on urban amenities, but isolating exogenous variation in IT remains a challenge for empirical work.

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar, "Fast and slow learning from reviews," Technical Report, National Bureau of Economic Research 2017.

Altonji, Joseph G, Todd E Elder, and Christopher R Taber, "Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools," *Journal of political economy*, 2005, *113* (1), 151–184.

Anderson, Michael and Jeremy Magruder, "Learning from the crowd: Regression discontinuity estimates of the effects of an online review database," *The Economic Journal*, 2012, *122* (563), 957–989.

Anenberg, Elliot and Edward Kung, "Information technology and product variety in the city: The case of food trucks," *Journal of Urban Economics*, 2015, *90*, 60 – 78.

Avery, Christopher, Paul Resnick, and Richard Zeckhauser, "The market for evaluations," *American Economic Review*, 1999, *89* (3), 564–584.

Bajari, Patrick, Jane Cooley Fruehwirth, Kyoo il Kim, and Christopher Timmins, "A rational expectations approach to hedonic price regressions with time-varying unobserved product attributes: The price of pollution," *American Economic Review*, 2012, *102* (5), 1898–1926.

Baum-Snow, Nathaniel and Daniel A Hartley, "Accounting for central neighborhood change, 1980-2010," 2016.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan, "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy*, 2007, *115* (4), 588–638.

Berry, Steven and Joel Waldfogel, "Product Quality and Market Size," *Journal of Industrial Economics*, 2010, *58* (1), 1–31.

Cabral, Luís and Ali Hortaçsu, "The Dynamics of Seller Reputation: Evidence from Ebay," *Journal if Industrial Economics*, 2010, *58* (1), 54–78.

Case, Karl E. and Robert J. Shiller, "The Efficiency of the Market for Single-Family Homes," *American Economic Review*, 1989, *79* (1), 125–137.

Cosman, Jacob, "Industry Dynamics and the Value of Variety in Nightlife: Evidence from Chicago," Working Paper, Carey Business School, Johns Hopkins University 2017.

Couture, Victor, "Valuing the consumption benefits of urban density," *University of California, Berkeley. Processed*, 2013.

_ **and Jessie Handbury**, "Urban Revival in America, 2000 to 2010," Working Paper 24084, National Bureau of Economic Research November 2017.

Davis, Donald R, Jonathan I Dingel, Joan Monras, and Eduardo Morales, "How Segregated is Urban Consumption?," Technical Report, National Bureau of Economic Research 2017.

Diamond, Rebecca, "The Determinants and Welfare Implications of US WOrkers' Diverging Location Choices by Skill: 1980-2000," *American Economic Review*, 2016, *106* (3), 479–524.

Edlund, Lena, Cecilia Machado, and Maria Micaela Sviatschi, "Bright minds, big rent: gentrification and the rising returns to skill," Technical Report, National Bureau of Economic Research 2015.

Ellen, Ingrid Gould, Keren Mertens Horn, and Davin Reed, "Has Falling Crime Invited Gentrification?," 2017.

Fang, Limin, "Evaluating the Effects of Online Review Platforms on Restaurant Revenues, Consumer Learning and Welfare," Technical Report, University of Toronto 2018.

Gaspar, Jess and Edward L. Glaeser, "Information Technology and the Future of Cities," *Journal of Urban Economics*, 1998, *43* (1), 136–156.

George, Lisa and Joel Waldfogel, "Who Affects Whom in Daily Newspaper Markets?," *Journal of Political Economy*, 2003, *111* (4), 765–784.

Glaeser, Edward L, Hyunjin Kim, and Michael Luca, "Nowcasting the Local Economy: Using Yelp Data to Measure Economic Activity," Technical Report, National Bureau of Economic Research 2017.

_ , _ , and _ , "Measuring Gentrification: Using Yelp Data to Quantify Neighborhood Change," Working Paper 24952, National Bureau of Economic Research August 2018.

_ , Jed Kolko, and Albert Saiz, "Consumer city," *Journal of economic geography*, 2001, *1* (1), 27–50.

_ , Scott Duke Kominers, Michael Luca, and Nikhil Naik, "Big Data and Big Cities: The Promises and Limitations of Improved Measures of Urban Life," Technical Report, National Bureau of Economic Research 2015.

Jin, Ginger Zhe and Phillip Leslie, "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards," *Quarterly Journal of Economics*, 2003, *118* (2), 409–451.

_ and _ , "Reputational Incentives for Restaurant Hygiene," *American Economic Journal: Microeconomics*, 2009, *1* (1), 237–67.

Kuang, Chun, "Does quality matter in local consumption amenities? An empirical investigation with Yelp," *Journal of Urban Economics*, 2017, *100*, 1–18.

Luca, Michael, "Reviews, reputation, and revenue: The case of Yelp. com," *Harvard Business School NOM Unit Working Paper*, 2011, (12-016).

_ and Georgios Zervas, "Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud," *Management Science*, 2016, *62* (12), 3412–3427.

Rosen, Sherwin, "Hedonic prices and implicit markets: product differentiation in pure competition," *Journal of political economy*, 1974, *82* (1), 34–55.

Schiff, Nathan, "Cities and Product Variety," *Journal of Economic Geography*, 2015.

Sinai, Todd and Joel Waldfogel, "Geography and the Internet: Is the Internet a Substitute or a Complement for Cities?," *Journal of Urban Economics*, 2004, *56* (1), 1–24.

Su, Yichen, "The Rising Value of Time and the Origin of Urban Gentrification," *Stanford University*, 2018.

Zagat Survey, *1998 Southwest Restaurants*, New York, New York: Zagat Survey, LLC, 1997.

Table 1: Summary Statistics

Metro Area	# Restaurants	Share Open	# Reviews
Phoenix-Mesa-Scottsdale, AZ	9,866	0.758	823,494
Las Vegas-Henderson-Paradise, NV	6,956	0.742	941,718
Cleveland-Elyria, OH	4,109	0.815	144,982
Charlotte-Concord-Gastonia, NC-SC	3,956	0.795	191,657
Pittsburgh, PA	3,666	0.792	147,814
Madison, WI	1,546	0.779	71,484
Akron, OH	703	0.838	17,170
Champaign-Urbana, IL	595	0.788	22,963
Total	31,397	0.774	2,361,282

	Share of reviews				
Metro Area	1 Star	2 Star	3 Star	4 Star	5 Star
Phoenix-Mesa-Scottsdale, AZ	0.117	0.093	0.120	0.258	0.413
Las Vegas-Henderson-Paradise, NV	0.115	0.093	0.137	0.260	0.395
Cleveland-Elyria, OH	0.111	0.101	0.136	0.283	0.368
Charlotte-Concord-Gastonia, NC-SC	0.111	0.101	0.142	0.289	0.357
Pittsburgh, PA	0.098	0.099	0.150	0.299	0.354
Madison, WI	0.091	0.105	0.150	0.310	0.344
Akron, OH	0.140	0.112	0.130	0.262	0.355
Champaign-Urbana, IL	0.118	0.111	0.154	0.278	0.338
Total	0.113	0.095	0.133	0.267	0.392

Notes: Summarizes data on U.S. restaurants, bars, and coffee shops from the 10th round Yelp Dataset Challenge. “Share Open” refers to the share of businesses listed as open on July 26, 2017 - the date of the data’s publication. Businesses were assigned to metro areas based on their address zipcode and a zipcode-to-CBSA crosswalk provided by the Missouri Census Data Center.

Table 2: The Effect of Reviews on Restaurant Exit

	Dependent Variable: Restaurant Exit			
	(1)	(2)	(3)	(4)
Stars	-0.000232**	-0.000366***	-0.000289**	-0.00210***
	(0.000109)	(0.000117)	(0.000143)	(0.000196)
# Reviews (standardized)	0.00353***	0.00300***	0.00162	0.00143**
	(0.00123)	(0.00113)	(0.00146)	(0.000712)
Stars × # Reviews	-0.000978***	-0.000881***	-0.000542	-0.000516**
	(0.000320)	(0.000304)	(0.000382)	(0.000215)
Additional controls		Y	Y	Y
Chains only			Y	
Non-chains only				Y
Observations	1,412,939	1,324,451	370,375	954,074

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

Notes: Results from linear probability model (2) where dependent variable is an indicator for whether a restaurant exits in a month. Number of reviews is normalized to have mean 0 and standard deviation 1. The unnormalized mean and standard deviation is 49.6 and 130. The average probability of exit in any given month is 0.005. Standard errors are clustered at the zipcode level. Additional controls include fixed effects for: 1) metro area, 2) calendar year-month, 3) months since the restaurant opened (measured by first review), 4) price category.

Table 3: The Arrival Rate of Reviews

Dependent variable: Number of new reviews	(1)	(2)	(3)
Stars_{t-1}	0.378***	0.380***	0.297***
	(0.0222)	(0.0225)	(0.0156)
# Reviews_{t-1}	0.541*	0.514**	0.761***
	(0.291)	(0.261)	(0.278)
$\text{Stars}_{t-1} \times$ # Reviews_{t-1}	0.406***	0.410***	0.358***
	(0.0891)	(0.0820)	(0.0857)
log(Population)		0.0193	-0.00994
		(0.0433)	(0.0197)
Share young (18-34)		-0.468**	0.411***
		(0.236)	(0.137)
College share (Bachelors+)		-0.221	0.762***
		(0.171)	(0.107)
Share households w children		-0.0954	-0.335**
		(0.292)	(0.168)
Distance to CBD (miles)		-0.0150***	-0.0104***
		(0.00477)	(0.00269)
Additonal controls			Y
Observations	1,396,714	1,396,412	1,308,767
R-squared	0.483	0.484	0.533

Robust standard errors in parentheses.

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

Notes: Results from model (2) where dependent variable is the number of reviews restaurant i receives in month t. Number of reviews is normalized to have mean 0 and standard deviation 1, as in Table 2. Other variables are not normalized. The average number of reviews in a month is 1.48. Standard errors are clustered at the zipcode level. Additional controls include fixed effects for: 1) metro area, 2) calendar year-month, 3) months since the restaurant opened (measured by first review), 4) price category.

Table 4: The Effect of Review Arrival Rate on Restaurant Survival Durations

Depvar: Log Survival Duration (days)	(1)	(2)	(3)
Stars (Final)	0.182***	-0.0454	0.211***
	(0.0193)	(0.0637)	(0.0203)
Zip Predicted Arrival Rate	-0.360***	-0.171	-0.304***
	(0.0662)	(0.204)	(0.0717)
Stars (Final) × Pred. Arr. Rate	0.0833***	-0.0578	0.0720***
	(0.0187)	(0.0649)	(0.0201)
Additional controls	Y	Y	Y
Chains only		Y	
Non-chains only			Y
Observations	28,946	6,885	22,061

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

Note: Results from accelerated failure time model described in (3). "Final star rating" is the restaurant's last observed star rating. The zipcode's "predicted arrival rate" of reviews is equal to the contribution of the zipcode's log market population, share young, college share, share households with children, and distance to CBD to the expected arrival rate of new reviews, as reported in column 3 of Table 3. Predicted arrival rate is normalized to have mean 0 and standard deviation 1. The unnormalized mean is 0.83 and the unnormalized standard deviation is 0.16. Additional controls include metro area fixed effects and price category fixed effects.

Table 5: Parameter Estimates

Parameter	Description	Estimate	Std. Err.
$\mu_{q,cons}$	Constant term in mean quality of entrants	3.8277***	0.0420
$\mu_{q,N}$	Effect of market size on mean quality	0.0174*	0.0101
σ_q	Standard deviation of entrant quality	0.8424***	0.0082
σ_z	Standard deviation of idiosyncratic experience (z_i)	0.9057***	0.0046
θ_1	Constant term in arrival rate	0.9562***	0.0382
θ_2	Direct effect of market size on arrival rate	-0.1135	0.0709
θ_3	Effect of market size*expected quality on arrival rate	0.1468***	0.0081
θ_4	Effect of market size*variance of quality on arrival rate	-0.3072***	0.0597
θ_5	Constant term in flow profit function	-0.5423***	0.0006
θ_6	Direct effect of market size on flow profit	-0.0060***	0.0010
θ_7	Effect of market size*expected utility on arrival rate	0.0028***	0.0002

Bootstrapped standard errors in parantheses (50 bootstraps.)

*** p<0.01, ** p<0.05, * p<0.1

Table 6: Simulations Under Counterfactual Information Environments

	Market size		
Information Environment	5th pctile	50th pctile	95th pctile
Long-Run Average Restaurant Quality			
Baseline	3.85	3.90	3.99
Zagat	3.78	3.81	3.85
$\Delta =$	-0.07	-0.09	-0.14
No learning	3.75	3.75	3.73
$\Delta =$	-0.11	-0.15	-0.25

Notes: Simulation results under counterfactual information environments. Long-run is defined to be 72 periods, or 6 years. In the "Zagat" counterfactual, each restaurant receives a Zagat rating at a rate that is calibrated to match Zagat's coverage of restaurants in 1998. Zagat ratings are assumed to perfectly reveal average restaurant quality. In the "No learning" environment, the arrival rate of reviews is zero.

Table 7: Effect of Restaurant Quality on House Prices

	Log Price	Log Price	Log Price	Log Price
Avg Star Rating	0.386	0.238	0.096	0.064
	(0.112)	(0.071)	(0.028)	(0.039)
Controls		X	X	X
House Fixed Effects			X	X
BFKT (2012)				X
N	109689	109689	109689	109689

Avg Star Rating is a measure of average quality of active restaurants (measured in Yelp stars) by zipcode and year. The first three columns show results from OLS regressions. The final column shows results from a 2SNLS estimator following the approach of BFKT (2012) to address the endogeneity of Avg Star Rating. The sample concerns repeat sales of individual homes between 2009-2015 in the CBSAs listed in Table 1. Controls include square feet, lot size, year built, number of bathrooms, three-digit zipcode fixed effects, the zipcode's "predicted arrival rate" of reviews as defined in the main text, and year fixed effects. The housing transaction data come from Corelogic. Standard errors are shown in parenthesis. In the first three columns, standard errors are clustered by zip code. In the final column, standard errors are clustered by census tract following BFKT (2012).

Figure 1: Cumulative Number of Yelp Reviews, millions, 2005-2018

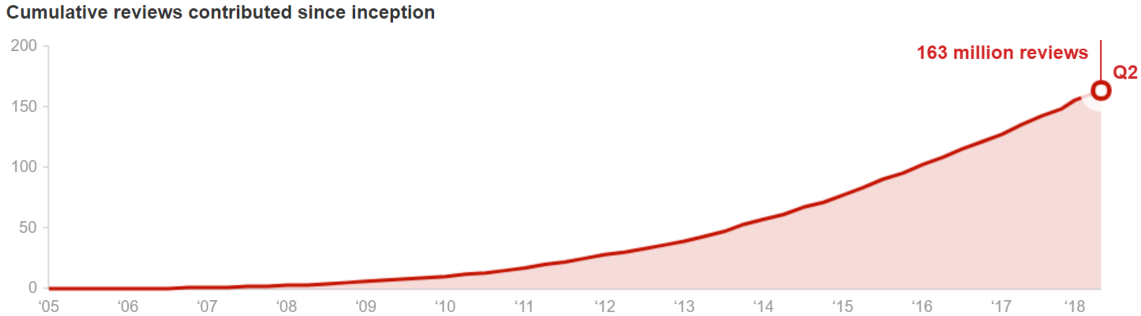

Source: Yelp.

Figure 2: Comparison of Yelp Data to ZIP Code Business Patterns Data

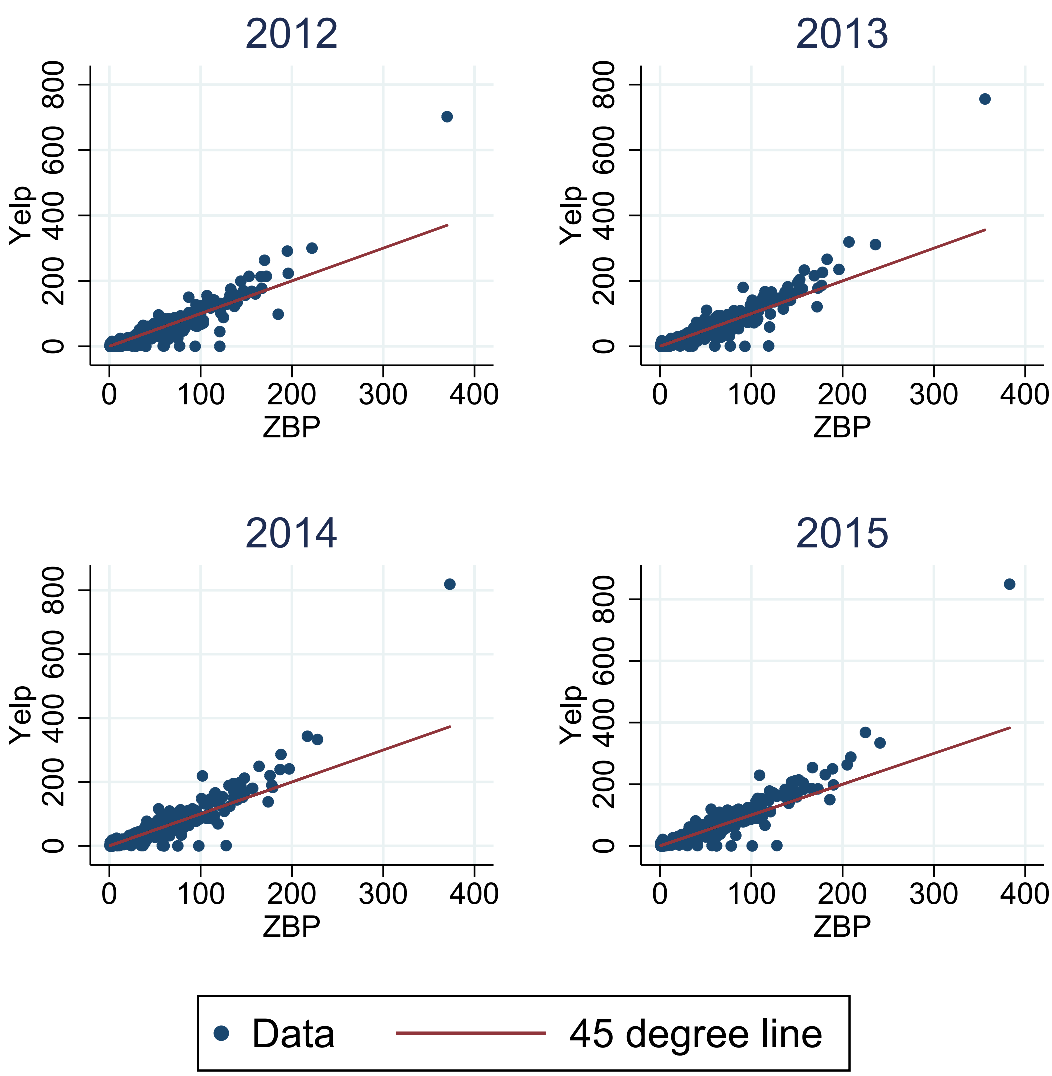

Notes: Each data point represents a zipcode-year. On Y-axis we plot the number of restaurants, bars, and coffee shops in the U.S. Yelp data, taking first review date as entry date and last review date as exit date (for closed restaurants). On the X-axis, we plot the number of eating and drinking places reported in the ZBP (NAICS 2012 code 7224 or 7225).

Figure 3: Restaurant Quality and the Arrival Rate of Reviews

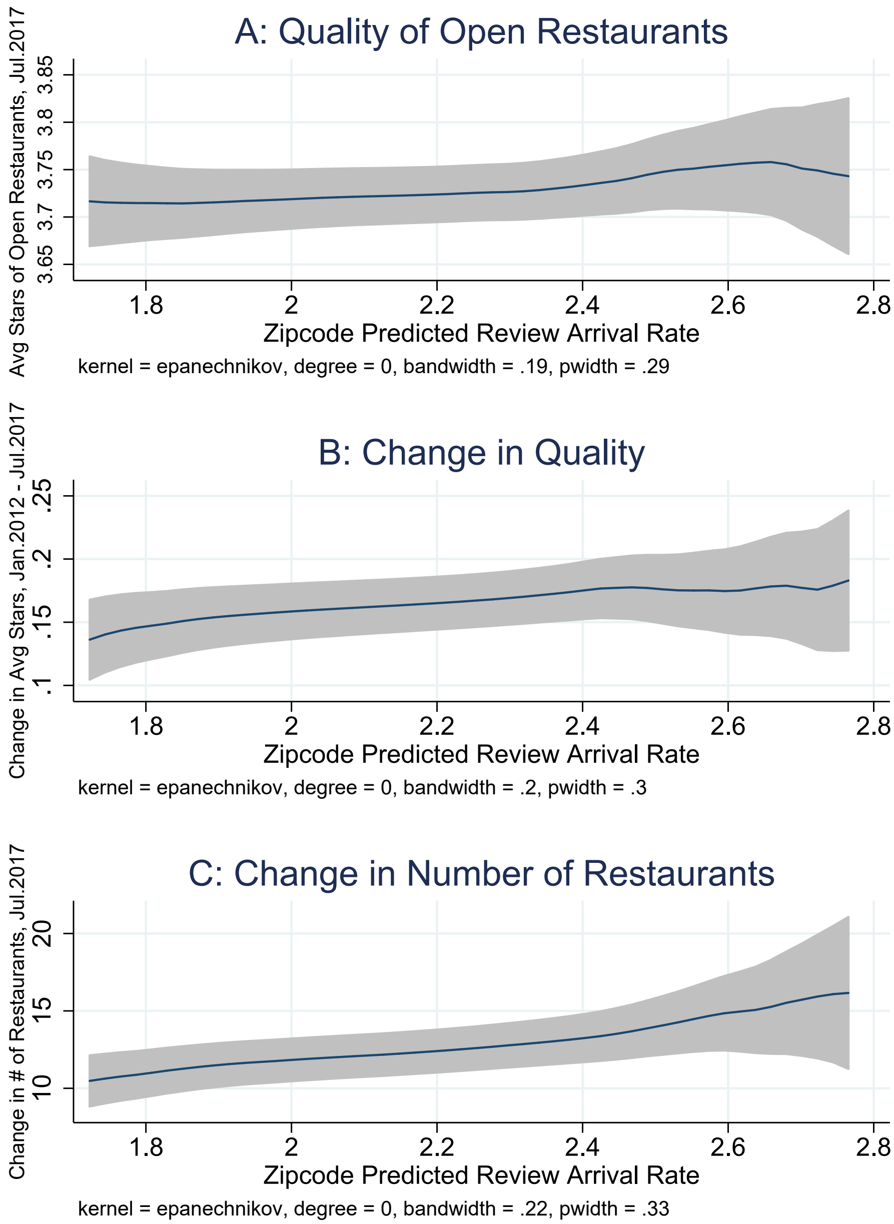

Note: Panel A plots a local polynomial fit of average restaurant quality in a zipcode to the "predicted arrival rate" of reviews in the restaurant's zipcode, for non-chain restaurants that are open as of July 2017. Panel B plots a local polynomial fit of the change in restaurant quality from 2012 to 2017. Panel C plots a local polynomial fit of the change in the number of restaurants from 2012 to 2017. The zipcode's "predicted arrival rate" is equal to the contribution of the zipcode's log population, share young, college share, share households with children, and distance to CBD to the expected arrival rate of new reviews, as reported in column 3 of Table 3. Restaurant quality is measured as the average star rating of the all reviews received as of July 26, 2017 (the date of data release).

Figure 4: Value Function of Staying in Business (true quality=3.75, median market size)

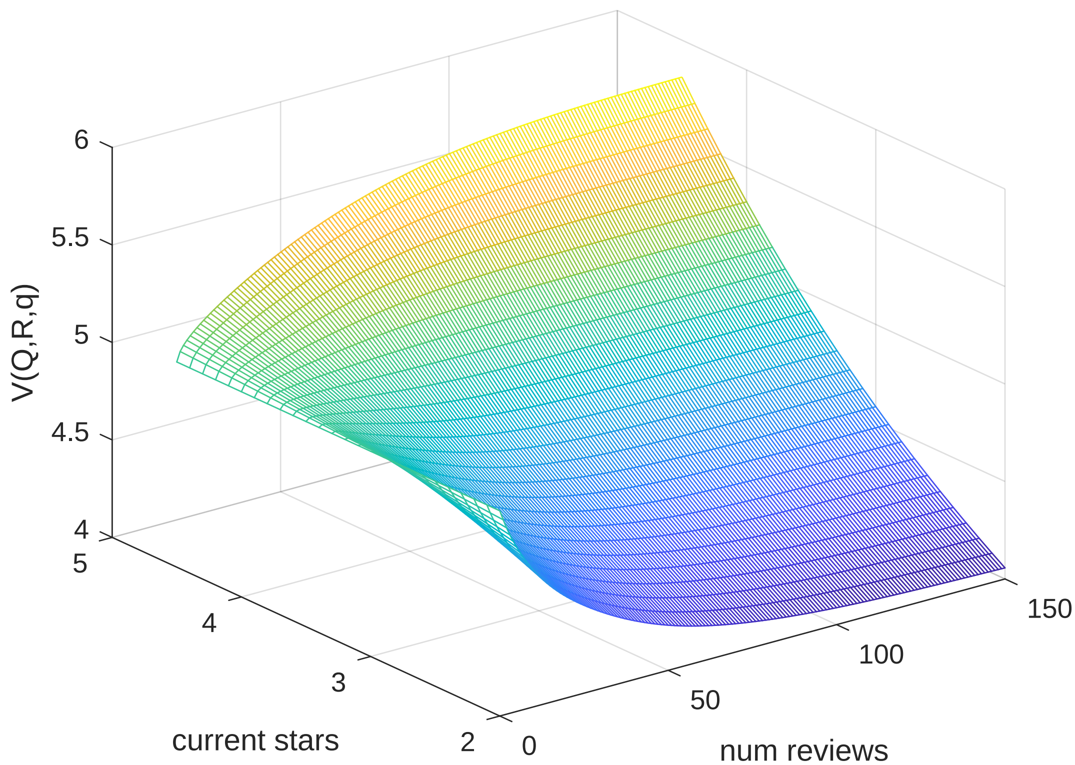

Notes: This figure shows the value function of staying in business, $V(Q, R, q)$ by Q and R for $q = 3.75$ and for a restaurant with median market size, evaluated at baseline parameter estimates given in Table 5.

Figure 5: Probability of Exit by True Quality (stars=3.7, reviews=20)

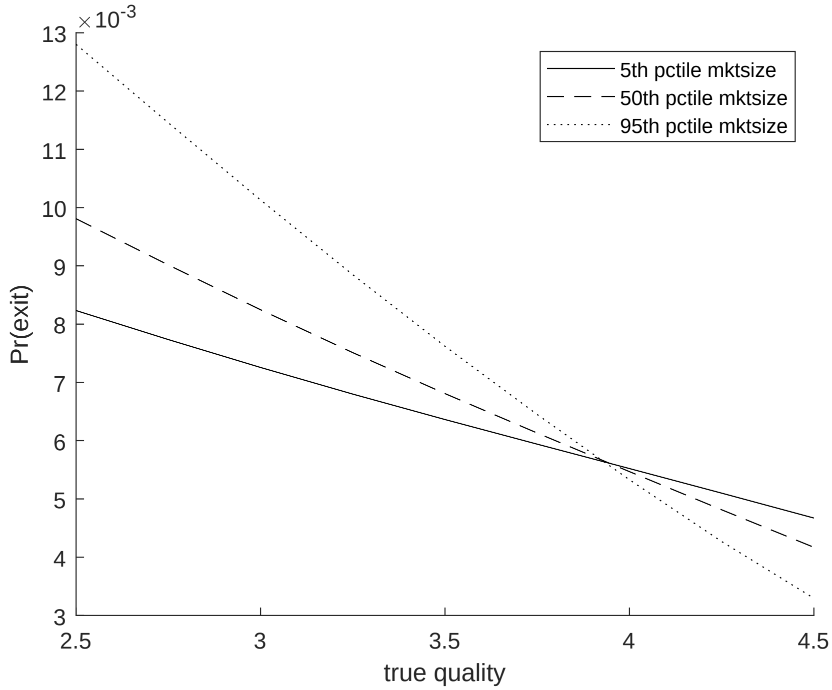

Notes: This figure shows the probability of a restaurant exiting when the current star rating is 3.7 and the number of reviews is 20, evaluated at baseline parameter estimates given in Table 5.

Figure 6: Probability of Exit by True Quality (stars=3.7, reviews=120)

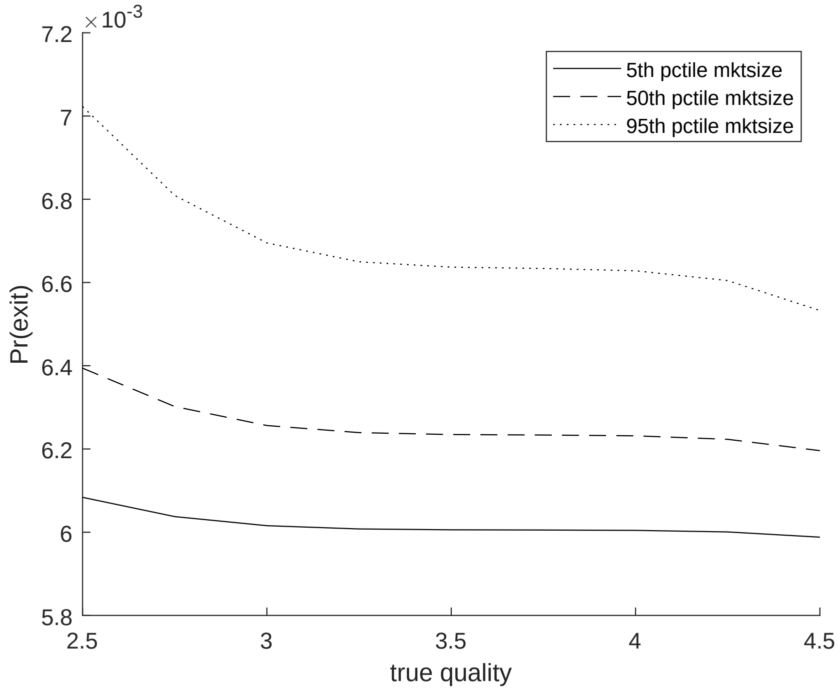

Notes: This figure shows the probability of a restaurant exiting when the current star rating is 3.7 and the number of reviews is 120, evaluated at baseline parameter estimates given in Table 5.

Figure 7: Ex-Ante Probability of Surviving 5 Years

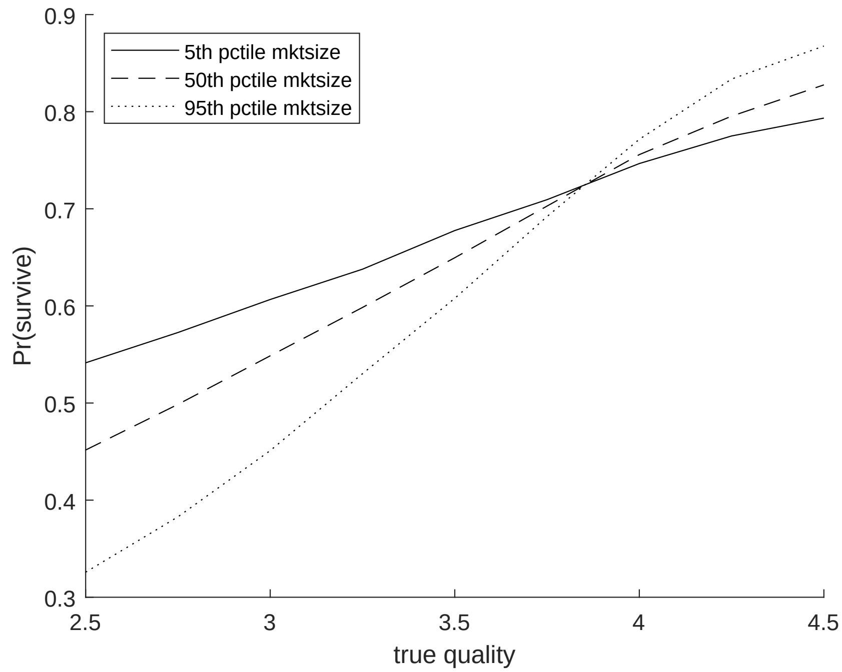

Notes: This figure shows the ex-ante probability of a new entrant surviving for 5 years by true quality, evaluated at baseline parameter estimates. The probability of survival is calculated by simulating the path of reviews and shocks for each entrant 5,000 times.

Figure 8: Ex-Ante Probability of Surviving 5 Years (No cost or ex-ante quality differences by market size)

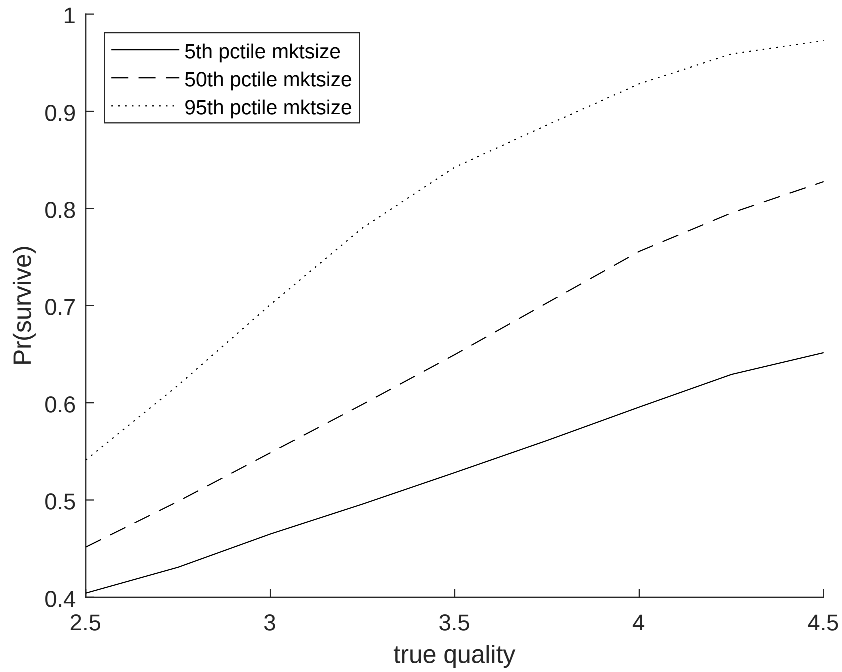

Notes: This figure shows the ex-ante probability of a new entrant surviving for 5 years by true quality, evaluated at counterfactual parameter estimates in which there are no differences in ex-ante restaurant quality or operating costs by market size. Instead, operating costs and ex-ante quality differences are set to that of the median market size at baseline parameters. The probability of survival is calculated by simulating the path of reviews and shocks for each entrant 5,000 times.

A Representativeness of the Yelp Data

Reviews started appearing in the data as early as late-2004/early-2005 for all metro areas in our sample. However, we need to consider the timing of reviews and the representativeness of the sample carefully, as Yelp was only founded in 2004 and has steadily grown in popularity since. Since Yelp data relies primarily on user inputs, data from earlier years may not be representative of the whole restaurant industry in our metro areas.

Figure 9 explores the timing of reviews in our sample. Panel A shows the distribution of review dates for all reviews. The number of reviews increases steadily over time from 2005, reflecting the continued growth and influence of Yelp as an online review platform. Panel B shows the distribution of review dates for only the first reviews of each restaurant. The number of first reviews exhibits an increasing trend from 2005 to 2011, after which it begins to decline. The increase from 2005 to 2011 appears to reflect the growth of Yelp as a platform, while the decline appears to be driven by the natural churn of new restaurants opening for business and existing ones shutting down. This is made clearer in Panels C and D, which show the distribution of first review dates for restaurants that are open or closed, respectively, at the time of data release.[23] Both panels show an increasing trend in the number of first reviews from 2005 to 2011, but the number of first reviews after 2011 stays roughly constant for open restaurants, while the number of first reviews for closed restaurants declines after 2011. This is roughly the pattern one would expect if there were a constant rate of new restaurants opening and old restaurants shutting down, and if closed restaurants are kept in the data after shutting down. Finally, Panels E and F show the distribution of last review dates for open and closed restaurants. As expected, almost all of the open restaurants were last reviewed near in time to the data release (suggesting that Yelp remains an active platform even for well established restaurants), whereas closed restaurants exhibit a more varied distribution in last review dates.

The timing of reviews suggests that the effect of Yelp's growth as a platform on the comprehensiveness of Yelp's restaurant listings is fully seasoned in by 2011. To verify this further, we compare our restaurant sample to data in the Census Bureau's ZIP Code Business Patterns (ZBP) files for 2012-2015. For each zipcode-year, we

[23]Open or closed status is reported in the Yelp dataset. However, the date of closing is not reported. Thus, we only know if the restaurant is open or closed at the time of data release (July 26, 2017).

count the number of restaurants present in our Yelp sample, taking first review date as the entry date and last review date as the exit date for each business (exit date only for closed businesses). We then compare this to the number of eating and drinking places reported in the ZBP, also at the zipcode-year level.[24] Figure 2 in the main text plots the results, and shows that the two datasets are closely correlated. Errors may be due to a variety of reasons, including differences in how establishments are classified,[25] or whether the establishment has employees.[26]

Comparability of Yelp ratings across local markets with different demographics is also important for our analysis, as we are using Yelp ratings to measure quality. If reviewers in different areas apply different standards in evaluating restaurants, it would render Yelp ratings unsuitable for cross-market analysis. To better understand the spatial pattern of restaurant reviews, we explore the distribution of reviews (i) across metro areas for a well-established chain restaurant, namely, McDonald's, and (ii) within metro for all restaurants in zipcodes that vary in their distance to the city center. We present the inter-metro review distributions for McDonald's in Figure 10, and the intra-metro review distributions for all restaurants in each of the four quartiles in terms of their (zipcode) distance to CBD in Figure 11. Zipcodes in the 1st quartile are the closest to the city center, while zipcodes in the 4th quartile are the farthest out. In general, there does not appear to be a spatial pattern in the distribution of reviews across or within metros. Specifically, the distribution of reviews does not appear to be systematically skewed towards lower ratings in denser markets due to the notion that urban consumers might be pickier.

[24]Eating and drinking places are identified in the ZBP using NAICS 2012 codes 7224 "Drinking Places (Alcoholic Beverages)" and 7225 "Restaurants and Other Eating Places."

[25]In Yelp, businesses can be tagged with multiple categories whereas in the ZBP each establishment receives only one NAICS code based on its primary business activity. Thus, a grocery store with an attached cafe may be included in our Yelp restaurant sample but not be classified as an eating or drinking place in the ZBP. Interestingly, this can explain the outlier in the top right corner of each graph in Figure 2. This zipcode just happens to be the Las Vegas strip, and it is possible that there are many eating and drinking places located within hotels that are not counted as separate establishments by the ZBP, but have separate Yelp listings. This leads to an "over-counting" of restaurants in Yelp relative to the ZBP, but in actuality the Yelp sample may represent a more accurate picture of the choices facing consumers than the ZBP data does. See Glaeser et al. (2017) for a further discussion on the comparison of Yelp to ZBP data.

[26]ZBP does not count establishments without payroll.

B Model of Spatial Competition

We now describe a model of spatial competition which gives rise to a consumer demand function which is approximately linear in expected quality and variance, as in equation (7). Let there be $j = 1, \ldots, J$ restaurants in a city. For a restaurant j, let $u_j(R_j, Q_j)$:

$$u_j(R_j, Q_j) = E\left[q_j|R_j, Q_j\right] - \frac{1}{2}\gamma Var\left[q_j|R_j, Q_j\right]$$

The expected utility that consumer i gets from visiting restaurant j is:

$$u_j(R_j, Q_j)\epsilon_{ij}e^{-\kappa d_{ij}}$$

where ϵ_{ij} is Frechet distributed with shape parameter $\alpha = 1$, and d_{ij} is a measure of the commuting cost from consumer i to restaurant j. In addition to the J restaurants, consumers have an outside option with a utility equal to $u_0\epsilon_{i0}$. Now suppose there are $\ell = 1, \ldots, L$ locations, each with n_ℓ consumers. The share of consumers visiting restaurant j from location ℓ is:

$$n_{\ell j} = \frac{u_j(R_j, Q_j)e^{-\kappa d_{\ell j}}}{u_0 + \sum_{k=1}^{J} u_k(R_k, Q_k)e^{-\kappa d_{\ell k}}} n_\ell$$

and the total demand for restaurant j is:

$$\begin{aligned} n_j &= \sum_{\ell=1}^{L} \frac{u_j(R_j, Q_j)e^{-\kappa d_{\ell j}}}{u_0 + \sum_{k=1}^{J} u_k(R_k, Q_k)e^{-\kappa d_{\ell k}}} n_\ell \\ &= u_j(R_j, Q_j) \sum_{\ell=1}^{L} \frac{n_\ell e^{-\kappa d_{\ell j}}}{\Phi_\ell} \end{aligned} \tag{21}$$

where $\Phi_\ell = u_0 + \sum_{k=1}^{J} u_k(R_k, Q_k)e^{-\kappa d_{\ell k}}$ is an aggregate of the utilities for each restaurant for consumers at location ℓ. Expression (21) for consumer demand is approximately linear in $u_j(R_j, Q_j)$ if the contribution of $u_j(R_j, Q_j)$ to $\sum_{\ell=1}^{L} 1/\Phi_\ell$ is small relative to the contribution of other restaurants. This is true in our data when we use our estimated γ from Table 5, and consider zipcodes as the locations and driving time as the cost of commuting.

C Estimation Details

Step 1: Estimate $(\mu_{q,cons}, \mu_{q,N}, \sigma_q^2, \sigma_z^2)$ directly from the distribution of reviews

In step 1, we estimate the parameters governing the distribution of entrant quality and idiosyncratic tastes ($\mu_{q,cons}$, $\mu_{q,N}$, σ_q^2, σ_z^2) directly from the distribution of the first three reviews for each restaurant that opened between January 1, 2012 and July 26, 2017. We use only the first three reviews because there is selection on quality in exit, resulting in biased estimates if we used all the reviews (the bias results from the endogeneity between the unobserved restaurant quality and the number of reviews). However, no restaurant exits before receiving three reviews in our sample, so this sample of reviews has essentially no selection bias.

Let $y_{i,j}$ be the star rating of the jth review for restaurant i. We can write:

$$\begin{aligned} y_{i,j} &= q_i + z_{i,j} \\ &= \mu_{q,cons} + \mu_{q,N} N_i + r_i + z_{i,j} \end{aligned}$$

Here, $r_i = q_i - \mu_{q,cons} - \mu_{q,N} N_i$ is the difference between the restaurant's unobserved true quality and the mean quality for its market size. The average of the first three reviews, $\bar{y}_i$, is:

$$\bar{y}_i = \mu_{q,cons} + \mu_{q,N} N_i + \underbrace{r_i + \frac{1}{3}\sum_{j=1}^{3} z_{i,j}}_{\xi_i}$$

Because there is no selection bias in the first three reviews, $E[\xi_i | N_i] = 0$, and we can estimate $\mu_{q,cons}$ and $\mu_{q,N}$ by regressing $\bar{y}_i$ on N_i. The variance of the residuals is an unbiased estimate of $\sigma_q^2 + \sigma_z^2/3$. The average within-restaurant variance of $y_{i,j}$ is an unbiased estimate of σ_z^2. Subtracting our estimate of $\sigma_z^2/3$ from the variance of the regression residuals therefore gives us an estimate of σ_q^2.

Step 2: Estimate $(\theta_1, \theta_2, \theta_3, \theta_4)$ from the observed arrival rate of reviews

In the second step, we estimate the parameters governing the arrival rate of new reviews. From the first step, we have estimates of μ_q for each restaurant, σ_q^2, and σ_z^2. Coupled with R_{it} and Q_{it} from the data, we compute the mean and variance of the posterior, $E[q_{it} | R_{it}, Q_{it}]$ and $Var[q_{it} | R_{it}, Q_{it}]$, for each restaurant and time period. We then estimate the parameters governing the arrival rate of new reviews, $\theta_{1:4}$, using maximum likelihood estimation. The log likelihood is:

$$LL = \sum_i \sum_{t=t_i}^{T_i} \log P\left(R_{it} - R_{it-1} | Q_{it-1}, R_{it-1}\right) \tag{22}$$

where $P(\cdot|Q, R)$ is the probability mass function of the Poisson distribution as in equation (8).

Step 3: Estimate $(\theta_5, \theta_6, \theta_7)$ *from the observed exits*

In the third step, we use the observed exit patterns to estimate the parameters governing the flow profit function. To do so, we need to first solve the dynamic programming problem which gives us the value function in equation (10). The dynamic programming problem can be solved by backward recursion starting with the value function at some arbitrarily large number of reviews. When the number of reviews is large, the restaurant does not expect its average star rating Q to change. Let $\bar{R}$ be this large value of R, and we approximate $V(Q, \bar{R}, q)$ by solving:

$$V(Q, \bar{R}, q) = n(Q, \bar{R})\pi - c + \beta E\left[\max\left\{0,\ V(Q, \bar{R}, q) + \xi'\right\}\right]$$

for each Q and q. In practice, we choose $\bar{R} = 150$, but the results are not sensitive to this choice. Once the value function at $\bar{R}$ is approximated, we solve for the value functions at smaller R and for each Q and q by backward recursion.

Once the value function is known, we can calculate the probability that a restaurant exits in each period. One complication is that we do not know the true exit period; rather, we know that the exit period is between T_i and T, where T is the last observed period in our data. Moreover, we know that there are no new reviews from period T_i to T. The probability, conditional on true quality q, that exit occurs between T_i and T and that there are no new reviews from T_i to T is equal to:

$$e\left(Q_{i,T_i}, R_{i,T_i}, q\right) \times \frac{1 - \left[\left(1 - e\left(Q_{i,T_i}, R_{i,T_i}, q\right)\right) P\left(0|Q_{i,T_i}, R_{i,T_i}\right)\right]^{T-T_i}}{1 - \left[\left(1 - e\left(Q_{i,T_i}, R_{i,T_i}, q\right)\right) P\left(0|Q_{i,T_i}, R_{i,T_i}\right)\right]} \tag{23}$$

Thus, the likelihood of observing the data for a single restaurant, assuming it has true quality q, is:

$$L_i(q) = \left(\prod_{t=t_i}^{T_i-1} \left[1 - e\left(Q_{it}, R_{it}, q\right)\right]\right) \times \left[1 - e\left(Q_{i,T_i}, R_{i,T_i}, q\right)\right]^{1-e_i} \times \ldots$$
$$\ldots \left[e\left(Q_{i,T_i}, R_{i,T_i}, q\right) \times \frac{1 - \left[\left(1 - e\left(Q_{i,T_i}, R_{i,T_i}, q\right)\right) P\left(0|Q_{i,T_i}, R_{i,T_i}\right)\right]^{T-T_i}}{1 - \left[\left(1 - e\left(Q_{i,T_i}, R_{i,T_i}, q\right)\right) P\left(0|Q_{i,T_i}, R_{i,T_i}\right)\right]}\right]^{e_i} \quad (24)$$

Since true quality is not observed, the log likelihood of all the exit data is:

$$LL = \sum_i \log\left(\int L_i(q) f_i(q)\right) \quad (25)$$

where $f_i(q)$ is the pdf of the normal distribution with mean $\mu_{q,cons} + \mu_{q,N} N_i$ and variance σ_q^2. We estimate the parameters governing the flow profit function, $\theta_{5:7}$, by maximizing this likelihood.

In theory, we could estimate all of the paramaters together in a single step instead of in three separate steps. However, given the large number of paramaters, we choose to estimate the model in stages as described above, which is much lighter computationally. While estimating the model in stages does not affect the consistency of the estimates, it does reduce efficiency. A standard, non-parametric bootstrap is used to calculate the standard errors.

D Zagat Counterfactual

To calibrate the Zagat counterfactual, we obtained the Zagat guide for the southwestern United States (covering Arizona and New Mexico) for 1998 (Zagat Survey (1997)).[27] Zagat reviewed 270 restaurants in the Phoenix metro area in 1998. According to the Census Bureau's Business Patterns data, there were 4,433 restaurants in Phoenix in 1998. Therefore, Zagat's coverage of Phoenix restaurants is 6 percent. Since Zagat publishes its guide roughly once a year and since our model is at a monthly frequency, the Zagat arrival rate of reviews is calibrated as 0.51 Zagat ratings per restaurant per month (i.e. 270/4433/12). We set the arrival rate equal to 0.51

[27]We chose this particular guide because it was available in the public library and because it covers Phoenix, AZ, which is one of the main cities in our estimation sample. We intentionally chose a year from before the 2000's to minimize the potential impact that the growing availability of household internet would have on the learning environment.

for every restaurant in every market in the counterfactual simulation. In the Zagat guide, Zagat also reports that each restaurant review is based on the experiences of potentially hundreds of diners, though they do not report the number of reports they received for each restaurant. We therefore assume that there is no signal variance associated with a Zagat review—it precisely reveals to consumers the average quality of the restaurant.

There are several reasons why we think our calibration of the Zagat environment overstates the amount of information provided to consumers during the pre-Yelp period. First, Zagat does not publish a guide for all cities in our sample. According to advertisements for other Zagat guides available in 1998, there is no guide covering Charlotte, NC, Pittsburgh, PA, or Madison, WI. Yet in the Zagat coutnerfactual, we assume that restaurants in these cities receive Zagat reviews at the same rate as in Phoenix. Second, we assume that a restaurant can receive a Zagat review immediately after opening, but in practice reviews are only published once a year. We expect that requiring Zagat reviews to arrive only at an annual frequency would produce results that are closer to the *No Learning* counterfactual results. We chose not to model lumpiness in the arrival process in part because it creates an additional computational burden as time becomes an additional state variable in the restaurant's value function. Third, it took money and time to access a Zagat guide, but we assume in the counterfactual that Zagat reviews are immediately available to all consumers upon release. Fourth, our calibration of the arrival rate assumes that there is no serial correlation in the year-by-year sampling of restaurants by Zagat. In practice, it seems like restaurants that receive Zagat reviews are more likely to receive reviews the next year, and Zagat appears to oversample restaurants that are already popular and whose quality may be well-known, including many chains.

E BFKT Method

Here we summarize our implementation of the method for estimating hedonic price regressions described in Bajari et al. (2012). For a more complete presentation and further discussion of the method, we refer the reader to Bajari et al. (2012).

Write the hedonic pricing equation as

$$lnp_{it} = \alpha_t + \boldsymbol{X}_{it}\beta + \xi_{it} \tag{26}$$

where i indexes house and t indexes year of sale. α_t denotes a set of time fixed effects and $\boldsymbol{X}$ is a vector of observable house or neighborhood characteristics. ξ are the relevant unobserved characteristics. For expositional simplicity, assume ξ is a scalar, but BFKT show that their method generalizes when we allow for a more general error term.

We assume that ξ evolves according to a first-order Markov process:

$$\xi_{it'} = \gamma(t,t')\xi_{it} + \eta(i,t,t') \tag{27}$$

where $\gamma(t,t')\xi_{it}$ is the expected value of the omitted attribute at time t' conditional on its value at time t. $\gamma(t,t')$ is a paramater allowed to vary flexibly by t and t'. We assume that $E[\eta(i,t,t')|\boldsymbol{I_t}] = 0$ where $\boldsymbol{I_t}$ denotes the information available to home buyers at time t. BFKT provide empirical support for this assumption.

We now rewrite equation (26) for period t' using information from the previous sale of house i at time t to eliminate the unobserved $\xi_{it'}$. Under our assumptions, we can rewrite equation (26) as

$$\begin{aligned} lnp_{it'} &= \alpha_{t'} + \boldsymbol{X}_{it'}\beta + \xi_{it'} \\ &= \alpha_{t'} + \boldsymbol{X}_{it'}\beta + \gamma(t,t')[lnp_{it} - \alpha_t - \boldsymbol{X}_{it}\beta] + \eta(i,t,t') \end{aligned} \tag{28}$$

A regression based on equation (28) rather than (26) may still produce inconsistent estimates because some elements of $\boldsymbol{X}_{it'}$ could be correlated with $\eta(i,t,t')$. For example, the innovation in the unobserved "curb appeal" of the home between t and t' could be correlated with improvement in restaurant quality. Therefore, in a first stage, we instrument for each of the time-varying elements of $\boldsymbol{X}_{it'}$ with its projected value based on $\boldsymbol{X}_{it}$ and other observed variables in the information set of the buyer at time t, $\boldsymbol{I_t}$. The projected value of $\boldsymbol{X}_{t'}$, $E[\boldsymbol{X}_{t'}]$, will be correlated with $\boldsymbol{X_{t'}}$ but not with the error term in equation (28) because of the BFKT assumption that $E[\eta(i,t,t')|\boldsymbol{I_t}] = 0$, thus making the projected value a valid instrument. In practice, following the implementation in BFKT, we include all elements of the information set directly in $\boldsymbol{X}$. So the first stage projection for x^j, one of the J controls in the vector of controls $\boldsymbol{X}$**,** is

$$x^j_{it'} = \pi_0(t,t') + \boldsymbol{X}_{it}\pi_1(t,t') + \nu(i,t,t') \tag{29}$$

where the $\pi's$ are paramaters that are allowed to flexibly vary by t and t'. Estimation then procedures in two stages. In the first stage, we estimate equation (29) for each of the time-varying variables in X by OLS. Then, we plug in the predicted values from this regression into equation (28) and estimate the paramaters of equation (28) using nonlinear least squares. We obtain standard errors corrected for the first stage estimation error using the clustering method that BFKT implement in their application.

In $\boldsymbol{X}$, we include house square feet, lot size, year built, number of bathrooms, three-digit zipcode fixed effects, the zipcode's predicted arrival rate of reviews as defined in the main text, and $AvgStars$ – the characteristic of interest that is defined in Section 7. α_t are year fixed effects. Our choices of controls closely follow the ones used in the application presented in BFKT. As in BFKT, we restrict our sample to homes that are sold exactly twice within our sample period.

Alternative Bounds Estimates

We now present an alternative procedure based on Altonji et al. (2005) (AET) to estimate bounds on the effect of average restaurant quality on house prices. Some differences in this procedure compared to the BFKT procedure are: (i) it does not require a sample of repeat sales, (ii) it uses the full time period for which we have Yelp data, (iii) it uses a different set of control variables, and (iv) it does not impose the rational expectation assumption of BFKT. Reassuringly, we find that our preferred estimate using the BFKT method lies within the bounds suggested by the AET method.

We begin by running a regression of

$$lnp_{it} = \alpha AvgStars_{it} + \boldsymbol{X}_{it}\beta + \epsilon_{it} \tag{30}$$

where lnp_{ij} is the log average house price in zipcode i in month t, $AvgStars_{it}$ is the average quality (in units of Yelp stars) of open restaurants in zipcode i and month t, and $\boldsymbol{X}$ is a vector of controls including CBSA-by-month fixed effects, and relevant zipcode characteristics such as income, share of college educated, and population. Quality-adjusted house prices come from Zillow. $AvgStars$ is computed from our Yelp data. A restaurant's true quality is assumed to be its average star rating as of its last review date. We use data from 2012 to 2017 for the eight metro areas in our

Yelp sample.

Table 8 presents the OLS results. As in 7, we see that the estimate of α declines as we add the control variables, suggestive of omitted variable bias. The estimate in column (2) with all of the controls is 0.09 – somewhat higher than our preferred estimate shown in column (3) of Table 7. The 0.09 estimate may be an overestimate of the true α because some amenities that are correlated with $AvgStars$ are unobserved and cannot be controlled for.

We attempt to place a lower bound on α using the approach of Altonji et al. (2005). The general idea is that because we have deliberately chosen the variables in $\boldsymbol{X}$ to minimize the amount of endogeneity, it is reasonable to assume that:

$$0 \leq \frac{\text{cov}(AvgStars, \epsilon)}{\text{var}(\epsilon)} \leq \frac{\text{cov}(AvgStars, \mathbf{X}\beta)}{\text{var}(\mathbf{X}\beta)} \tag{31}$$

Following Altonji et al. (2005), the procedure to construct bounds on α is as follows. First, choose a candidate parameter for α. Since

$$lnp - \alpha AvgStar = \mathbf{X}\beta + \epsilon$$

with $E[\epsilon|\mathbf{X}] = 0$, we can estimate $\hat{\beta}$ and $\hat{\epsilon}$ under the candidate for α. Then, if $0 \leq \frac{\text{cov}(AvgStar,\hat{\epsilon})}{\text{var}(\hat{\epsilon})} \leq \frac{\text{cov}(AvgStar,\mathbf{X}\hat{\beta})}{\text{var}(\mathbf{X}\hat{\beta})}$, α is in the identified set. Otherwise, it is inconsistent with the model assumptions. Implementing this approach, we find that α is between 0.017 and 0.088. It is reassuring that our preferred estimate of 0.064 in column (3) of Table 7 lies within the bounds suggested by the AET approach.

Table 8: Robustness: Effect of Restaurant Quality on House Prices

	Log Price	Log Price
Avg Star Rating	0.496	0.088
	(0.025)	(0.007)
Controls		X
N	18360	18360

Notes: Regression of log average Zillow house price on average quality of active restaurants (measured in Yelp stars). Observations are at the zipcode-year-month level. The first column includes no controls. The second column includes controls for zipcode average income, log population, share young, college share, share of households with children, distance to CBD, and CBSA-by-year-month fixed effects. Standard errors clustered by CBSA-year-month.

Figure 9: Timing of Reviews

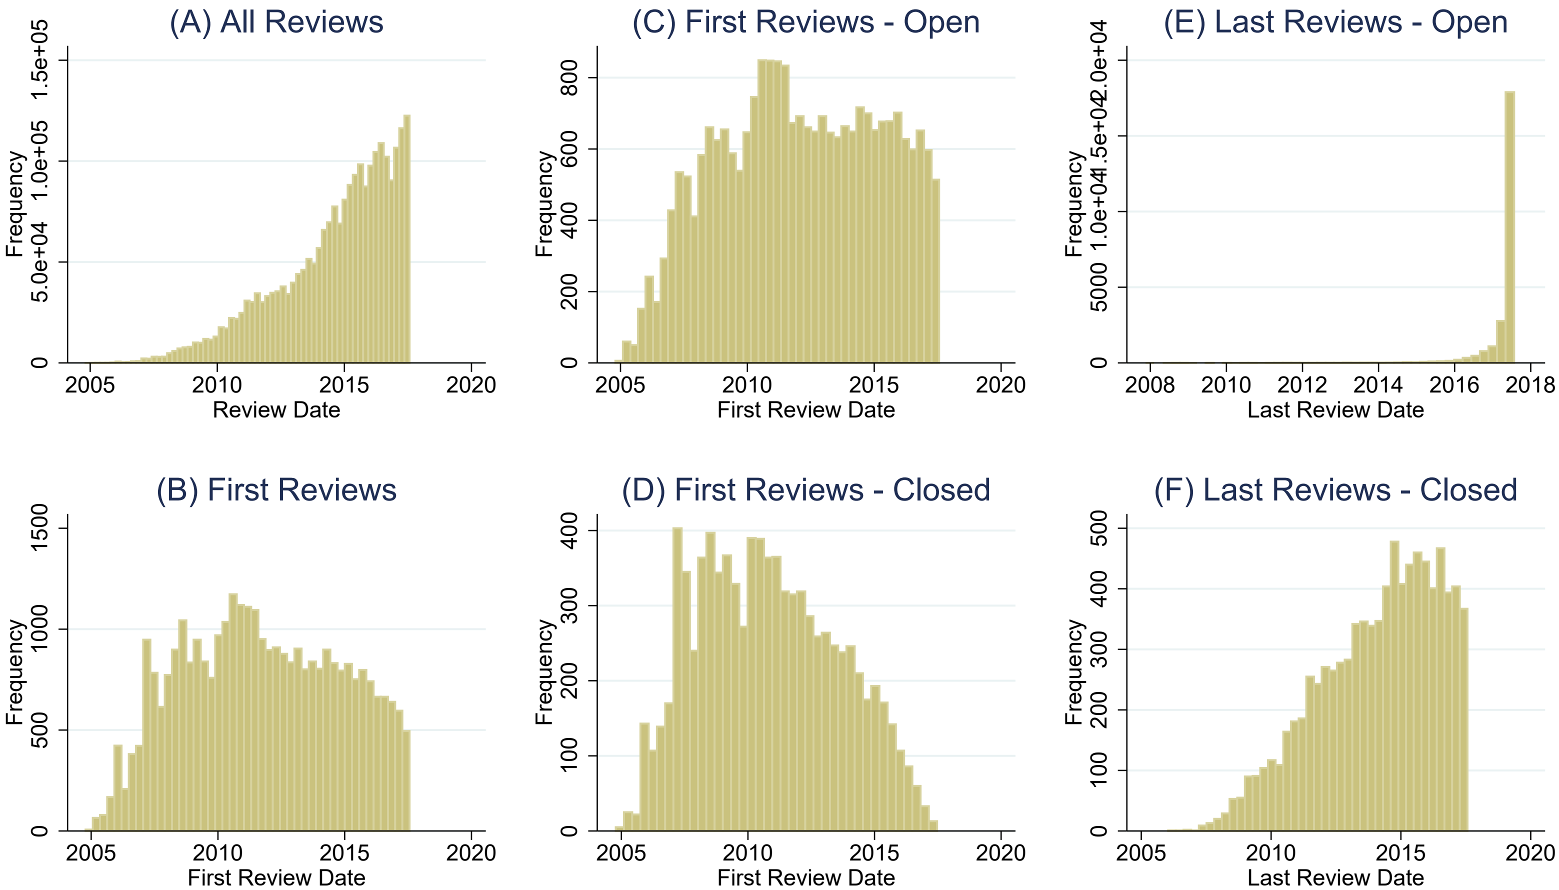

Notes: Summarizes timing of reviews on U.S. restaurants, bars, and coffee shops from the 10th round Yelp Dataset Challenge. Panel A: all reviews. Panel B: all first reviews for businesses. Panel C: first reviews for businesses that are still open at the time of data release. Panel D: first reviews for businesses that are closed at time of data release. Panel E: last reviews for businesses that are still open at the time of data release. Panel F: last reviews for businesses that are closed at time of data release.

Figure 10: Distribution of Reviews for McDonald's across Metro

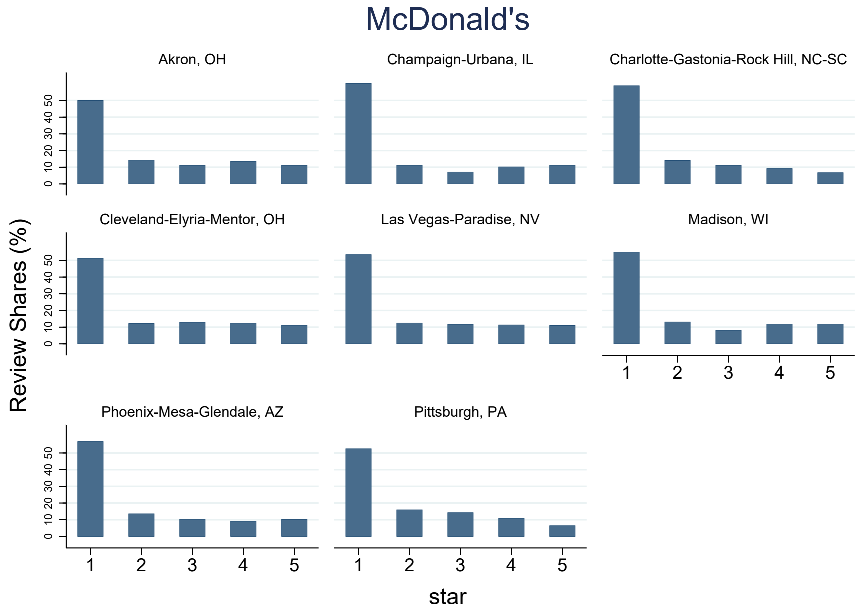

Notes: The graph plots the distributions of reviews for McDonald's separately for each of the eight metro areas in our sample.

Figure 11: Distribution of Reviews for All Restaurants by Quartiles of Distance to CBD

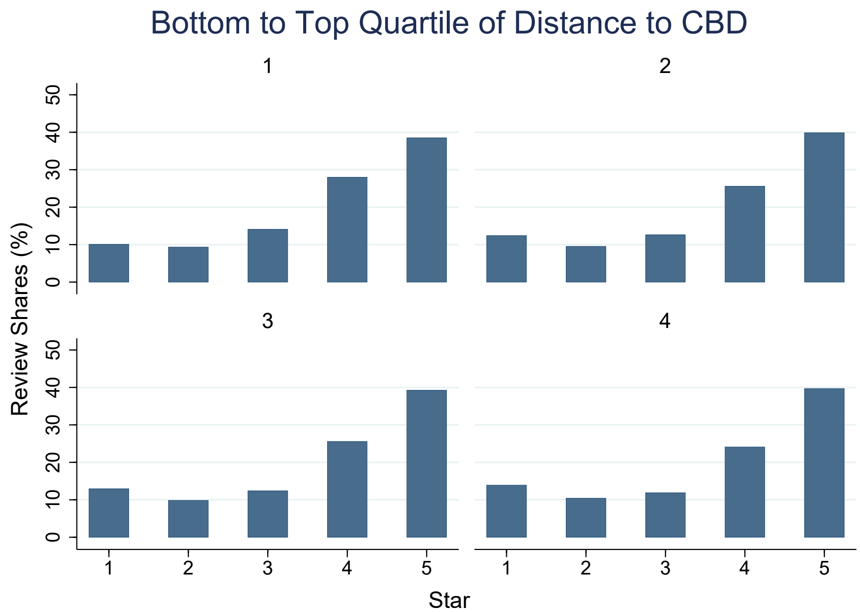

Notes: The graph plots the distributions of reviews for all restaurants in zipcodes by quartiles in terms of distance to CBD. Zipcodes in the 1st quartile are the closest to the city center, while zipcodes in the 4th quartile are the farthest out.