

## Does Yelp Affect Restaurant Demand?

Michael L. Anderson and Jeremy Magruder

Internet review forums increasingly supplement expert opinion and social networks in informing consumers about product quality. We estimate the effect of positive Yelp.com ratings on restaurant reservation availability and find that an extra half-star rating causes restaurants to sell out 19 percentage points (49%) more frequently, with larger impacts when alternate information is more scarce.

### Also in this issue

#### Immigration Reform 2013: Implications for California Agriculture

Philip L. Martin.....5

#### Some Impacts of Recession on California's Nursery and Floral Industry

Hoy Carman.....9

### In the Next Issue...

#### China's Growing Role in Agricultural Trade

by Colin A. Carter  
and Sandro Steinbach

**B**eliefs on product quality play an important role in shaping consumer demand. For many goods, consumers face *ex ante* uncertainty regarding the quality of the good and rely on imperfect signals to infer quality. Traditionally, expert opinion and social learning have helped consumers resolve these information asymmetries.

For an expert's take, consumers may consult *Consumer Reports* when buying an automobile or household appliance or they may read reviews by professional critics when selecting a movie or choosing among dining options. Alternatively, consumers may confer with peers who own the automobile or who have eaten at the restaurant. In recent years, however, online sites that cheaply aggregate consumer reviews have recently expanded and have begun supplementing both of the traditional mechanisms. But are these sites playing an important role in determining consumer demand?

Despite the theoretical potential of digital word-of-mouth to influence consumer choices, it is difficult to estimate its impact on purchasing decisions. Products that receive positive reviews are ones that appeal to consumers (i.e., they have high unobserved quality), and these products would likely experience high sales even in the absence of positive reviews. In a recent paper published in the *Economic Journal*, we leverage a feature

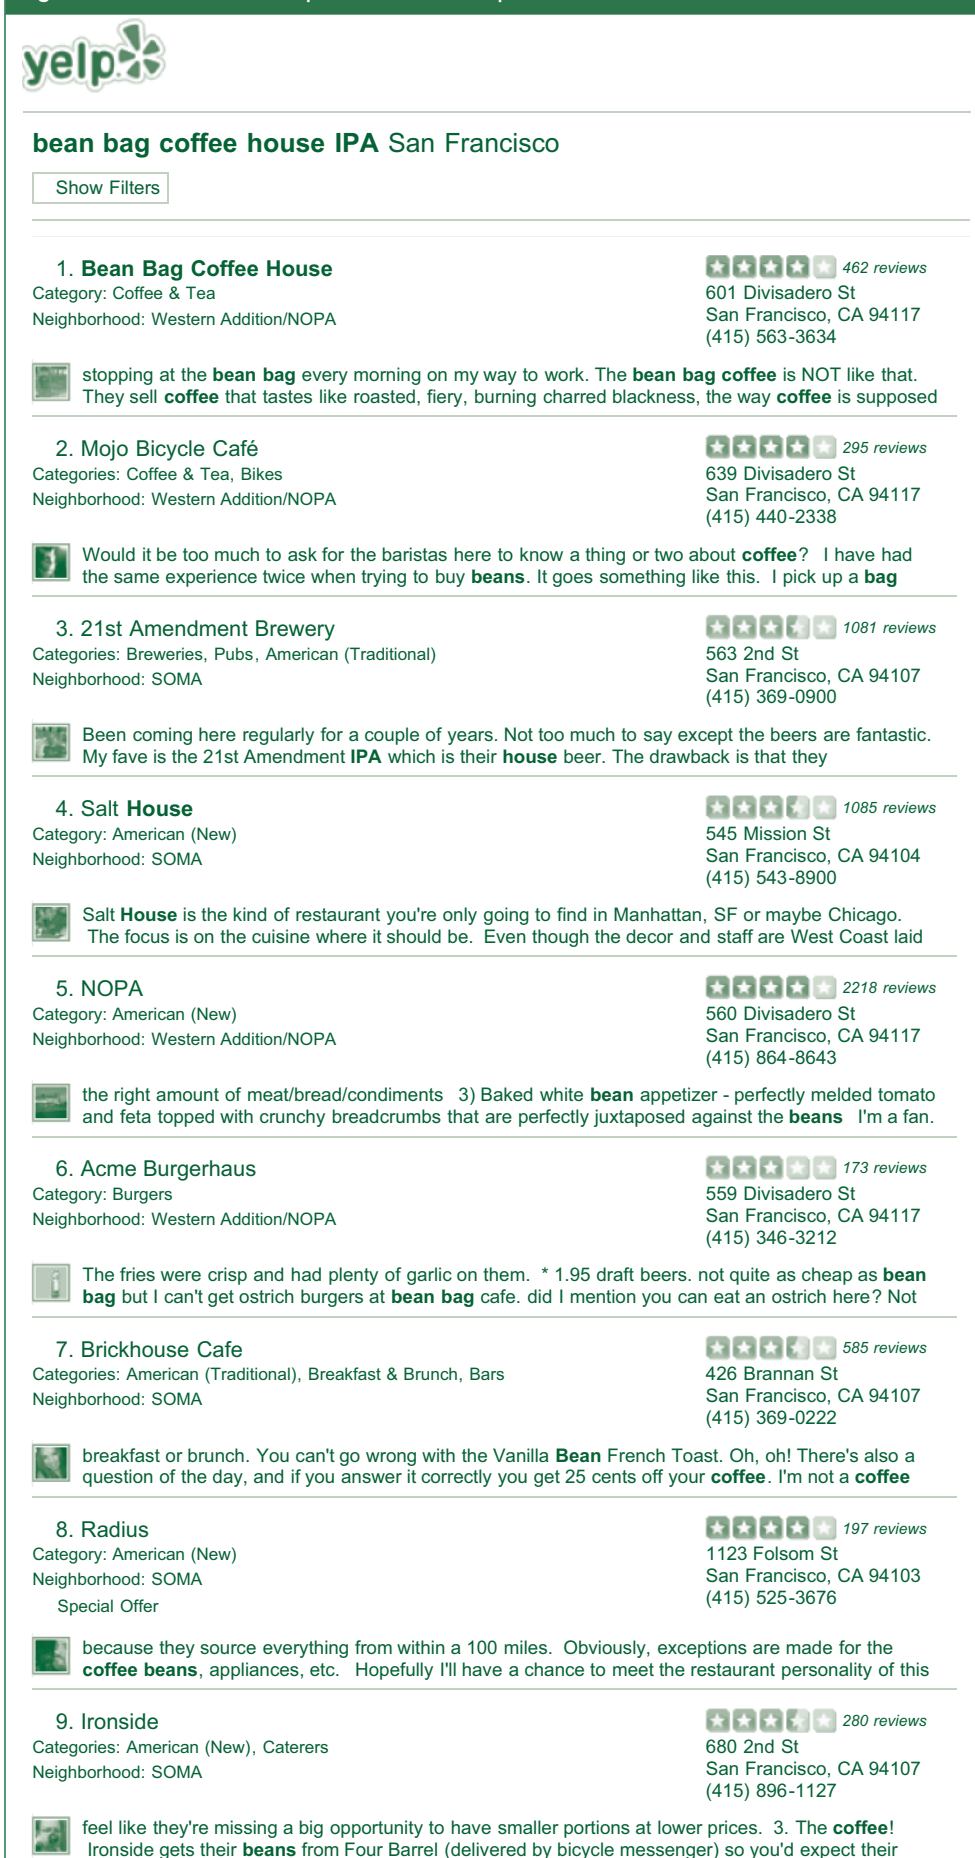
of the display system at Yelp.com—a popular site that allows users to leave reviews of local businesses—to estimate the effect of positive Yelp ratings on restaurant customer flows.

Yelp reviewers assign businesses ratings from one to five stars in whole-star increments. When a user searches Yelp.com, Yelp presents a list of businesses that meet the search criteria or fall within the category of interest. Figure 1 on the next page reproduces an excerpt from a sample search on Yelp.com. Businesses are sorted according to relevance and rating, and for each business the average rating is prominently displayed, rounded to the nearest half star. The number of stars in the average rating is easily visible, particularly because the color of the stars changes at whole star thresholds.

We downloaded the entire history of reviews from Yelp.com for each restaurant in San Francisco, CA and recorded the date of the review, the rating assigned (1–5), and the reviewer's unique user identifier. We then reconstructed the average rating and total number of reviews for each restaurant at every point in time and matched these data with reservation availability data from a large online reservation website.

As Figure 1 demonstrates, Yelp aggregates all reviews for a given business and displays the average rating prominently. However, when Yelp computes

Figure 1. Results of a Sample Search on Yelp.com




the average rating, they round off to the nearest half-star. Two restaurants that have similar average ratings can thus appear very differently on Yelp. For example, a restaurant with an average rating of 3.24 displays a 3-star average rating, while a restaurant with an average rating of 3.26 displays a 3.5-star average rating. In actuality, the true underlying quality of these two restaurants is similar on average, allowing us to identify the effect of the Yelp rating on customer demand while controlling for unobserved quality. If Yelp reviews have significant impacts on consumer demand, then we should observe a sharp increase in reservations at each major rounding threshold (e.g., 3.25 stars and 3.75 stars).

In the paper we use a technique known as regression discontinuity to estimate the effect of Yelp. Here, we present several figures that graphically summarize the results from the regression discontinuity estimator. Figure 2 plots mean 7:00 p.m. reservation availability by Yelp rating. Panel A focuses on the window where restaurants have either 3 or 3.5 stars, and Panel B focuses on the window where restaurants have either 3.5 or 4 stars. There are clear jumps in the mean availability at both the 3.5 and 4 star thresholds.

Moving from 3 to 3.5 stars—which occurs when a restaurant’s rating crosses 3.25 stars—reduces the likelihood of availability from about 85% to about 60%. Moving from 3.5 to 4 stars—which occurs when a restaurant’s rating crosses 3.75 stars—reduces the likelihood of availability further to below 40%.

Interestingly, for the most part, it appears that a step function is a good approximation to the overall relationship between Yelp ratings and restaurant availability. That is, restaurant availability appears to respond primarily to the displayed rating, and not the underlying average review score (which presumably measures the restaurant’s true quality from the perspective of

Figures 2a and 2b. Average Reservation Availability at 7 p.m. by Restaurant Yelp Rating



consumers). Overall, we find that a half-star increase in Yelp ratings decreases reservation availability by 19 percentage points during peak dining hours.

If Yelp is providing information about new restaurants, that information should be most valuable among restaurants that are unfamiliar to patrons. We divide restaurants into familiar/unfamiliar groupings along two dimensions. First, restaurants with fewer than 500 reviews are likely to be less frequented and less well known than those with more than 500 reviews. For restaurants with fewer than 500 reviews, an extra half-star on Yelp reduces reservation availability by 20 to 30 percentage points depending on the reservation time. In contrast, for restaurants with more than 500 reviews, for whom there is likely less hidden information about quality, there is no discontinuous change at any threshold associated with additional Yelp stars.

A second test for whether the Yelp effect is due to solving information problems groups restaurants according to whether there are external sources of quality information. Here, we note that quality information is easily available for restaurants which have a Michelin star or those which appear in the *San Francisco Chronicle's* annual Top 100 Restaurants listing. In contrast, crowd-sourced information may be more

important for restaurants excluded from these prestigious rankings. We again find that an extra half-star on Yelp reduces reservation availability by 20 to 30 percentage points at all three times for restaurants without external recognition but that the Yelp ranking does not similarly advantage restaurants which have been externally accredited. These results support the hypothesis that Yelp is most valuable when there is less external information about restaurants, though other differences between the two groups of restaurants may also play some role.

The high return to positive Yelp ratings naturally creates an incentive for restaurants to manipulate their own ratings by leaving false reviews. Manipulation is feasible in this context because Yelp is crowd-sourced—any restaurateur can, in principle, leave himself a 5-star review. Furthermore, the significant increases in business at Yelp thresholds create a strong incentive for restaurants to attempt to manipulate their ratings to fall above a threshold.

Is it possible that the increases in demand that we observe in Figure 2 at Yelp rounding thresholds are the result of specific restaurants strategically manipulating their ratings so that they fall right above the rounding thresholds? If so, this would invalidate our research design, because restaurants

above the rounding threshold would not be directly comparable to restaurants below the rounding threshold.

However, if specific restaurants manipulate their reviews to fall right above the thresholds, then some of restaurants above the thresholds have “true” Yelp ratings (i.e., the ratings they would receive absent manipulation) that are lower than their observed Yelp ratings.

To generate a significant drop in reservation availability at the threshold, these restaurants must sell out virtually all the time, despite the fact that they receive low ratings from true Yelp reviewers. It seems *ex ante* surprising that a restaurant that receives poor reviews would be extremely crowded, though it is theoretically possible.

Using a short theoretical model, we show that although restaurants face incentives to manipulate Yelp ratings, it does not make sense for them to try to stay right above the Yelp rounding threshold. The intuition is simple: given that a random stream of reviews will change each restaurant’s average rating over any time period, a restaurant which is just above a threshold has a very similar likelihood of just missing that threshold after new reviews come in as a restaurant which is just below the threshold. Both restaurants therefore

face similar incentives to try and push their Yelp scores into safer territory.

We also present a variety of empirical tests that consistently show no evidence of any manipulation behavior that would cause restaurants to cluster just above the thresholds. For example, restaurants leaving fake reviews for themselves should have more 5-star reviews and fewer reviews per reviewer. We find no evidence that restaurants with these types of characteristics cluster just above the thresholds.

Two questions emerge when considering the effects of Yelp ratings on restaurant demand. First, do the effects represent the transmission of information on restaurant quality or do they represent a marketing effect generated by Yelp's ranking system? Second, what changes in customer flows and profits are consistent with the observed changes in reservation availability?

Our estimates may not represent a pure effect of information regarding restaurant quality if the order in which Yelp lists restaurants on its website (e.g., in Figure 1) is a function of a restaurant's displayed average rating rather than its true average rating. In that case, restaurants just above a Yelp threshold would be significantly more likely to be seen by consumers browsing Yelp than restaurants just below a Yelp threshold. However, we find that the order in which Yelp lists restaurants is not affected by the displayed average rating (after controlling for the restaurants' true underlying average ratings). We thus conclude that increased information about restaurant quality causes higher-rated restaurants to have lower availability, rather than any effect of increased visibility.

To gauge what changes in customer flows could be consistent with our result that an extra half-star on Yelp causes a 19 percentage point decrease in reservation availability, we performed a series of simple statistical calibrations. First, we recorded the capacity

of each restaurant in a subsample of 73 restaurants. Next, we assumed that a restaurant has no reservation availability if the number of seats reserved for a given evening reaches its capacity. Finally, we examined the average customer flows that would be consistent with reservation availability rates of 58% (the average rate above the Yelp thresholds) and 39% (the average below the Yelp thresholds) under different assumptions about the distribution of arriving customers. Our calibrations suggest that the median restaurant might experience an increase in customer flows of 6% or more if its reservation availability drops from 58% to 39%.

Modest changes in customer flows, however, can have a significant impact on profits in an industry with high fixed costs and high margins. A typical mid-to-high-end restaurant with \$20,000 per week in sales and a margin of 68% on food and beverage sales, earns approximately \$2,000 per week in pre-tax profit. In comparison, a 6% increase in customer flow translates into an additional revenue gain of \$816 per week in pre-tax profit.

Of course, the increase in profit will be lower if the restaurant is capacity-constrained or if it has to expand staffing levels to maintain service. Nevertheless, the calibrations suggest that a typical restaurant could experience substantial gains in profit when crossing a Yelp threshold.

In summary, the effects we estimate are large, and they indicate a valuable use of crowd-sourced information: because Yelp collects and aggregates the experiences of a large number of patrons, Yelp provides a convenient forum to solve asymmetric information problems about the quality of unfamiliar restaurants. Tightening the link between restaurant quality and restaurant patronage may well have positive benefits for society.

Crowd-sourced quality information may improve the average quality

of consumed meals via two mechanisms. First, it can redirect consumers to higher quality restaurants. Second, it can induce lower quality restaurants to shut down or improve their quality in response to changes in customer demand. We provide direct evidence of the first mechanism, but we cannot speak to the second mechanism. With the rapid spread of Yelp and other similar crowd-sourcing websites, this suggests that market evolution may be an important avenue of future research.

Suggested Citation:

Anderson, M.L., and J. Magruder. 2013. "Does Yelp Affect Restaurant Demand?" *ARE Update* 16(5):1-4. University of California Giannini Foundation of Agricultural Economics.

*Michael Anderson and Jeremy Magruder are both assistant professors in the Department of Agricultural and Resource Economics at UC Berkeley. They can be contacted by e-mail at [mmanderson@berkeley.edu](mailto:mmanderson@berkeley.edu) and [jmagruder@berkeley.edu](mailto:jmagruder@berkeley.edu), respectively.*

**For additional information, the authors recommend:**

Anderson, M.L. and Magruder, J. 2012. "Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database," *The Economic Journal*, 122(563): 957-989.