Advertising strategy in the presence of reviews: an empirical analysis*

Brett Hollenbeck
Anderson School of Management, University of California Los Angeles

Sridhar Moorthy
Rotman School of Management, University of Toronto

Davide Proserpio
Marshall School of Business, University of Southern California

November 20, 2017

*This paper was presented at the QME Conference in June 2017. We thank members of the audience, and the official discussant, Tobias Klein, for their comments and suggestions. Sridhar Moorthy also thanks the Social Sciences and Humanities Research Council of Canada for their generous support of this research vide Grant #435–2017–0121.

Abstract

Over the last fifteen years, one of the major developments online has been the growth and proliferation of review websites such as TripAdvisor. The ready availability of independent information from past users poses interesting questions for marketing strategy. What role does advertising play in the new environment? How should firms adjust their advertising strategy to the presence of reviews? In this paper we address these questions in the context of the hotel industry. Using a data set of TripAdvisor hotel reviews and another describing hotels' advertising expenditures, we show, first, that overall ad spending decreased from 2002 to 2015, suggesting that online reviews have had the effect of displacing advertising. Second, there is a negative causal relationship between TripAdvisor ratings and advertising spending in the cross-section: hotels with higher ratings spend less. This suggests that user ratings and advertising are substitutes, not complements. Third, this relationship is stronger for independent hotels than for chains, and stronger in competitive markets than in noncompetitive markets. The former suggests that a strong brand name provides some immunity to reviews, and the latter suggests that when ratings are pivotal, the advertising response might be particularly strong. Finally, we show that the relationship between user ratings and advertising has strengthened over time, as websites such as TripAdvisor have become more influential. This provides further confirmation that the effect of online ratings on advertising operates through the demand side, and not the supply side. Hotels seem to react to reviews if and only if consumers react to them.

1 Introduction

Over the last fifteen years or so, one of the major developments online has been the growth and proliferation of review websites such as TripAdvisor and Yelp.[1] According to the Pew Research Center, in 2016, 82% of U.S. adults read reviews occasionally or regularly before purchasing a product for the first time; 40% did so almost always.[2] The ready availability of independent experiential information from past users poses interesting questions for marketing strategy. What role does advertising play in the new environment? How should firms adjust their advertising strategy to the presence of online reviews? In this paper we report on these questions in the context of the hotel industry.

Hotels are experience goods (Nelson 1970). Consumers do not observe key aspects of hotel quality before purchase, but they do after purchase. Word-of-mouth has therefore always been an important source of information for hotels. What has changed in recent years is the scale and scope of word-of-mouth. Whereas in an earlier era consumers would have shared their experiences with a small group of friends and family, now they are effectively broadcasting their opinions to the entire world by posting reviews online.

There are good theoretical reasons for why firms should react to these developments. For advertising strategy ought to differ depending on the amount of independent information available to consumers. Consider, for example, a hypothetical scenario in which consumers have no independent sources of information. In this scenario, advertising has free reign to work its magic in the myriad ways discussed in the advertising literature: creating and maintaining awareness, providing product information (Grossman and Shapiro 1984), signaling quality (Milgrom and Roberts 1986, Nelson 1974), persuasion (Chioveanu 2008, Sutton 1991).[3] On the other hand, in the opposite scenario in which all consumers have access to, and use, independent sources of information, the scope for what advertising can do is nec-

[1]In fact, these are among the most visited websites on the Internet. According to `alexa.com`, TripAdvisor ranked 257 and Yelp 238, in global traffic on April 27, 2017.

[2]See: `http://www.pewinternet.org/2016/12/19/online-reviews/`

[3]For a review of the empirical and theoretical literature on advertising see Bagwell (2007).

essarily much more limited—perhaps non-existent even. Firms with good reviews might see little point in advertising because reviews do the same job for free; firms with bad reviews see little point in advertising because it would be hard to overcome the effect of the negative reviews. Reality lies somewhere in between these extremes. First, not all consumers read reviews despite the growing numbers who do. Advertising may target the ones who don't. Second, even among those who read reviews, advertising might still have a role to play. For instance, it could direct the search process, or it could play a role at the back end of the search process. Depending on which of these narratives prevails, the time-series relationship between reviews and advertising could be one of substitutes or one of complements (Lei and Moorthy 2017).

The hotel industry is an ideal setting to study the relationship between online reviews and advertising, for several reasons. First, it was one of the earliest adopters of online reviews. A large corpus of reviews has accumulated, showing good variation, both in the cross-section and in the time-series. Second, because hotels are experience goods and serve people from many, widely-dispersed locations, online word-of-mouth is especially important relative to offline word-of-mouth. Finally, the industry is large and important in its own right. Hotels generated $196 billion in sales and employed 2 million individuals in 2012 according to the U.S. Economic Census. The industry also advertises heavily, spending $2.1 billion in 2015.

Our empirical analysis is based primarily on two data sets, one comprising all TripAdvisor hotel reviews from 2002 onwards, and the other describing the advertising strategies of the hotels featured in those reviews. The review data set contains all U.S. hotels listed on TripAdvisor, and the advertising data set includes information about each hotel's monthly advertising spend, disaggregated by media—TV, newspapers, magazines, radio, Internet (display), and outdoor. The first contribution of this paper is simply the collection and matching of these two data sets, allowing for the first time-series cross-sectional study of the empirical relationship between online reviews and firms' advertising strategies. This is noteworthy because our data are not just a sample of a particular market in a particular time period, but rather the entire experience of an industry over virtually the entire period online reviews have existed.

Second, the nature of the TripAdvisor data allows us to make a causal claim that user ratings affect advertising strategy via the demand side, and not the supply side. Ordinarily, this would be difficult to do: user ratings signal underlying quality, and underlying quality, arguably, affects both the demand-side and the cost-side. However, TripAdvisor's summary ratings involve an exogenous rounding-off procedure that produces essentially random variation in displayed ratings around particular thresholds. Individual ratings are averaged and either rounded up or rounded down to the nearest .5 or whole number. This provides us with a ready-made regression-discontinuity design (RDD) to identify the demand-side effect of online user ratings on advertising strategy.[4]

Our main results are the following. First, ad spending by hotels falls from 2002 to 2015. This suggests a broad industry-wide substitution of online ratings for ad spending over time. Second, this decline in advertising is greater for chains than for independent hotels. While chains start out in 2002 with a higher (per-property) spending level, by 2015 they are spending about the same as independent hotels. Third, there is a negative relationship between online ratings and advertising spending in the cross-section: higher-rated hotels spend less than lower-rated hotels. This effect obtains both at the extensive margin (decision to advertise) and at the intensive margin (advertising level given advertising). User ratings and advertising spending are therefore substitutes in the cross-section, as well as in the time-series. Fourth, this relationship is stronger for independent hotels than for chains, consistent with prior research showing that online reviews have larger effects on independent hotels' sales than on chain hotels' sales (Hollenbeck 2017). Fifth, when we examine how the ratings-advertising relationship operates in markets differing in "ratings competitiveness"—highly competitive markets being those where the standard deviation of ratings is small—we find that highly competitive markets show a stronger negative relationship between ratings and

[4]For other applications of regression-discontinuity designs on user reviews see Anderson and Magruder (2012), Lei (2017), and Luca (2016).

ad spending. Finally, comparing the relationship between online ratings and advertising spending in the early years of TripAdvisor (2002-2005) with more recent years (2012-2015), we find that the relationship has strengthened over the years.

In Section 6 we discuss the significance of these results and what they imply about how user reviews and advertising interact in the Internet era.

2 Background

As noted earlier, survey evidence shows that a large number of people consult online reviews before making purchase decisions. It is only logical, then, that online reviews should affect sales. Indeed, this is what the literature finds: Chevalier and Mayzlin (2006) and Sun (2012) for books, Anderson and Magruder (2012) and Luca (2016) for restaurants, Jin and Kato (2006) and Cabral and Hortaçsu (2010) for eBay auctions, and Lewis and Zervas (2016) for hotels. Luca (2016) shows, in addition, that the effect is particularly strong for independent restaurants (compared to chain restaurants). In a similar vein, Hollenbeck (2017) shows that online reviews have significantly reduced the revenue premium enjoyed by hotel chains over independent hotels.

In contrast to this large literature showing the effect of online reviews on demand, there is relatively little work on how online reviews affect firms' actions. Melnik and Alm (2002), Resnick et al. (2006) and Houser and Wooders (2006) find that prices realized in eBay auctions increase in seller reputation. Lewis and Zervas (2016), using TripAdvisor data, show that hotel prices increase in response to online ratings. Lei (2017) analyzes Yelp restaurant ratings data from 2014 and finds, like us, that ratings had a negative effect on advertising spending in the cross-section.

On the theoretical side, the seminal papers are the two by Yubo Chen and Jinhong Xie: Chen and Xie (2005, 2008). The first is a duopoly model with horizontal and vertical differentiation; the second is a monopoly matching model in two dimensions. Different assumptions characterize the two papers. In the first, reviews are assumed to provide accurate information about product quality, while advertising can mislead those who don't read reviews; in the second, reviews provide information comprehensively (i.e., on both product dimensions), but may or may not be accurate, while advertising-supplied information is accurate, but may or may not be comprehensive. These different characterizations are justified by different interpretations of reviews and advertising. In the first paper, reviews are interpreted as third-party reviews, and both reviews and advertising are assumed to influence product quality perceptions, whereas in the second, reviews are interpreted as user reviews, and both reviews and advertising provide information on horizontal attributes. Despite these differences, the central concern of both papers is what happens to advertising strategy as reviews become available. The main result of the first paper is that when reviews can be incorporated into ads, but the horizontal differentiation is so strong that prices don't change, then both the low-quality firm and the high-quality firm reduce their advertising expenditures. The second paper argues that information provided in ads may increase or decrease as reviews become available.

Finally, Lei and Moorthy (2017) examine the cross-sectional relationship between user ratings and advertising spending in a duopoly model where firms are vertically differentiated by their review endowments. They distinguish between consumers who search reviews on their own, in an open-ended manner, versus consumers who search, brand-by-brand, if and only if they are exposed to a firm's advertising. Advertising's role, then, is to activate an expansion market. Two types of equilibria result: in one the firm with stronger reviews behaves like a fat cat and eschews the expansion market; in the other, there is a mixed strategy equilibrium with both firms randomizing between pursuing the expansion market and not pursuing it—the only difference being that the firm with stronger reviews pursues the expansion market with higher probability. With the former, there is a substitution effect between ratings and advertising spend; with the latter, there is a complementarity effect.

3 Data

To study the effect of online ratings on hotels' advertising spending empirically we examine data from three sources: Kantar Media, TripAdvisor, and Smith Travel Research (STR).

Kantar Media. Advertising data from Kantar Media (and its previous incarnations, TNS Media Intelligence and LNA) have been the basis for a number of studies in advertising (e.g., Caves and Greene 1996, Shum 2004, Kim and McAlister 2011, Honka et al. 2017). Kantar's data cover all major media: TV, radio, magazines, newspapers, Internet display, and outdoor.[5] Their methodology is essentially a bottom-up approach combining direct monitoring of ads and information supplied by media outlets.[6]

We obtained the monthly advertising expenditures of 15,039 distinct U.S.-based hotel entities over a period of nearly 15 years, from January 2002 to December 2015. Expenditures, by media, are available for each hotel brand and "product," brand being a higher-level aggregation than product. Generally, product refers to a specific hotel property, but more generally it refers to a specific "advertised product." For example, for Best Western Hotels, brand is "Best Western Hotels," and there are over 100 advertised products, including "Best Western Hotels: Bethelehem PA," "Best Western Hotels: Miami Beach," and "Best Western Hotels & Minnesota State Tourism: Combo."[7] A brand's ad expenditure is the sum of its expenditures over all its products.

[5]To be more specific, 18 media categories are identified: Network TV, Spot TV, Spanish Language Network TV, Cable TV, Syndication, Magazines, Sunday Magazines, Local Magazines, Hispanic Magazines, B-to-B Magazines, National Newspapers, Newspapers, Hispanic Newspapers, Network Radio, National Spot Radio, Local Radio, US Internet Display, and Outdoor.

[6]For more details see `http://stradegy.kantarmediana.com/Stradegy/Help/Methodology.aspx?.pl=Methodology`.

[7]As this example indicates, product generally indicates a specific location, but may also indicate the type of ad. For example, we see products such as "Online," "Mobile App," or "Corporate Promotion" in our data. For independent hotels, generally, there is only one product, which may or may not include the specific location where the hotel is located.

TripAdvisor. Our hotel user ratings data come from TripAdvisor. Launched in 2000, TripAdvisor is now one of the most popular review platforms on the Internet.[8] In an average month, it has about 350 million unique visitors worldwide.[9] In addition, TripAdvisor's summary ratings are displayed on other travel websites such as Hotels.com and Expedia.com, making its reach even larger. As of May 2016, TripAdvisor had over 500 million customer reviews on over 6 million accommodations, restaurants, and attractions.

For each hotel in its database, TripAdvisor displays a summary rating that ranges from 1 to 5 stars (actually, "filled-in bubbles"). Summary ratings are average ratings, averaged over all the reviews received by a hotel. However, it is not the exact average. TripAdvisor rounds the average rating up or down to the nearest half- or full-star. Thus, average ratings between 1 and 1.24 get rounded down to 1, average ratings between 1.25 and 1.74 get rounded to 1.5, average ratings between 1.75 and 2.24 get rounded to 2.0, and so on. As noted earlier, this rounding-off procedure is key to our identification strategy. It provides the basis for our regression-discontinuity analyses.

We collected hotel information and review history for all of the U.S.-based accommodations (hotels, B&Bs, motels, hostels, and inns) listed on TripAdvisor—a total of 91,783 properties. For every property, we obtained detailed consumer-facing information such as the TripAdvisor unique identifier, property name, and location (state, city, ZIP code, and street address). For every review, we have its unique identifier, full text, publication date, and star-rating.

STR. Finally, from STR, a company that tracks the hotel industry, we obtained a list of all the hotel chains in the U.S., and information about their characteristics: class (economy, midscale, upper midscale, upscale, upper upscale, and luxury), and number of properties in

[8]Indeed, it is one of the top-50 most visited websites on the Internet (`http://www.comscore.com/Insights/Rankings/Revised-Top-50-Digital-Media-Properties-for-October-and-November-2016`).

[9]See: `https://www.tripadvisor.nl/pages/factsheet.html`.

the chain. In addition, for a large fraction of U.S. hotel properties, we obtained basic "census" data, namely, hotel name, location, price, class category, and capacity (# of rooms).[10]

Linking the data sets

The first challenge in using these data sets is matching hotels between TripAdvisor review data and Kantar advertising spending data. While we have detailed information about hotels in the TripAdvisor data (name and address, for example), Kantar data are organized by brand and product.

To link the two data sets, we followed a multi-step process. The first step was to match hotel name and location (when available in the Kantar data) across the two data sets using "fuzzy string-matching." These algorithmic matches were then turned over to Amazon Mechanical Turk workers to check manually whether the linked hotels were indeed correctly matched. We hired U.S.-based workers with a HIT approval rate greater than 95% for this purpose. For every potential match, we solicited two independent opinions from two different workers. We accepted the matches marked correct by both workers, and manually checked the ones where was disagreement. This process resulted in 4,563 Kantar entities (30% of the Kantar sample), corresponding to 4,409 TripAdvisor hotels.[11] These 4,409 hotels have a total of 479,884 reviews from January 2002 to December 2015.[12]

Finally, in order to get more information about hotel characteristics (such as chain versus independent, hotel class, etc.), we linked the 4,409 hotels to the STR hotel census using name and address. This yielded 3,211 matches (73% of 4,409). Some of our analyses are based on the 4,409-hotel data set and some on the 3,211-hotel data set.

[10] The STR hotel census contains information about 63,502 properties, which is about 69% of the properties listed on TripAdvisor.

[11] Out of the approximately 15,039 entities in the Kantar data, 4,104 (about 27%) are hotel chains, and the rest are independent hotels. We find a similar proportion in the linked data set: 1,113 chains and 3,450 independent hotels.

[12] Out of the 4,409 matched hotels, only 4,334 have at least one review.

Table 1: Summary statistics

	2002	2015
Hotels	2,131	4,239
Fraction of months with advertising	0.17	0.11
Fraction of hotels reviewed	0.26	1.00
Ratings and reviews		
Avg. hotel rating	3.92	4.09
Reviews per hotel	1.72	554.94
Average monthly advertising expenditure(thousands of $)		
Internet	0.09	0.15
Newspapers	0.83	0.20
Magazines	0.88	0.28
Outdoor	0.16	0.09
Radio	0.02	0.03
Television	0.04	0.03
Total	2.04	0.77

4 Descriptive evidence

Summary statistics on TripAdvisor user ratings and advertising spending are reported in Table 1 in two columns, one for the year 2002 (the start of our sample) and one for the year 2015 (the end of our sample). Comparing the two columns shows how user ratings and advertising spending have changed over a fourteen year period.

First, both average user ratings and number of reviews have increased over time. Ratings increased by about 0.2 stars, while the number of reviews grew exponentially, from about 2 reviews per hotel in 2002 to over 550 per hotel in 2015. Figure 1 (a) illustrates the growth in reviews graphically.

Second, advertising spending decreased over the same period. Advertising spending per property decreased by about 62% from 2002 to 2015 (Figure 1 (b)), from about $2,000/month in 2002 to about $770/month in 2015.[13] This decrease was particularly pronounced for print

[13] These are nominal spending numbers, not inflation-adjusted numbers. If we were to plot the latter, the decrease would be even steeper.

advertising; ad dollars in newspapers and magazines decreased by about 75%. Outdoor advertising spending follows next with a decrease of about 43%, while TV and radio advertising spending hardly change. The only ad medium to show an increase is Internet display; it increases by 67% from 2002 to 2015. Figure 2 shows monthly average advertising spending by different types of hotels: different hotel classes, and independent versus chain hotels. Luxury hotels advertise the most, spending almost as much as all the other hotel tiers combined; they also decline the most. Comparing chains to independent hotels, ad spending decreased more for the former.

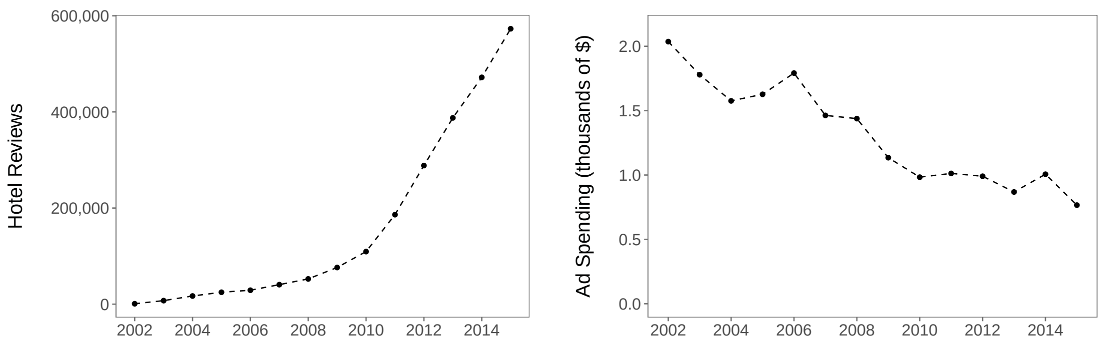

(a) (b)

Figure 1: Number of reviews and ad spending over time

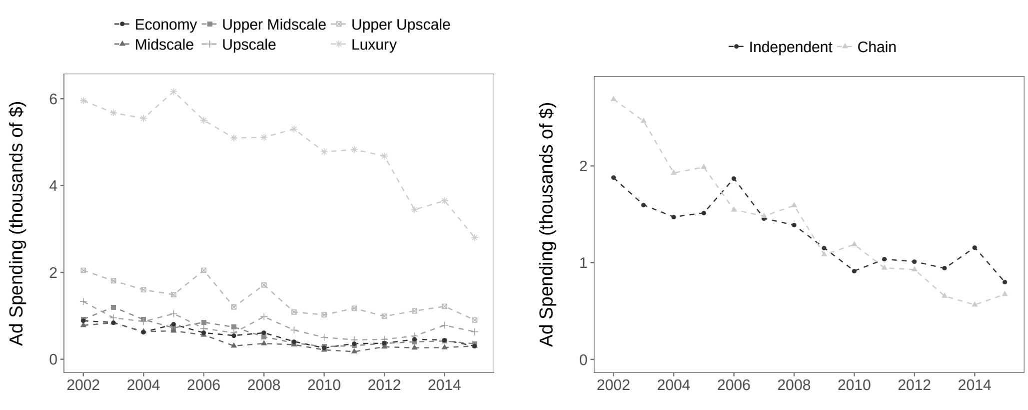

(a) (b)

Figure 2: Ad spending by hotel class and independent versus chain hotels

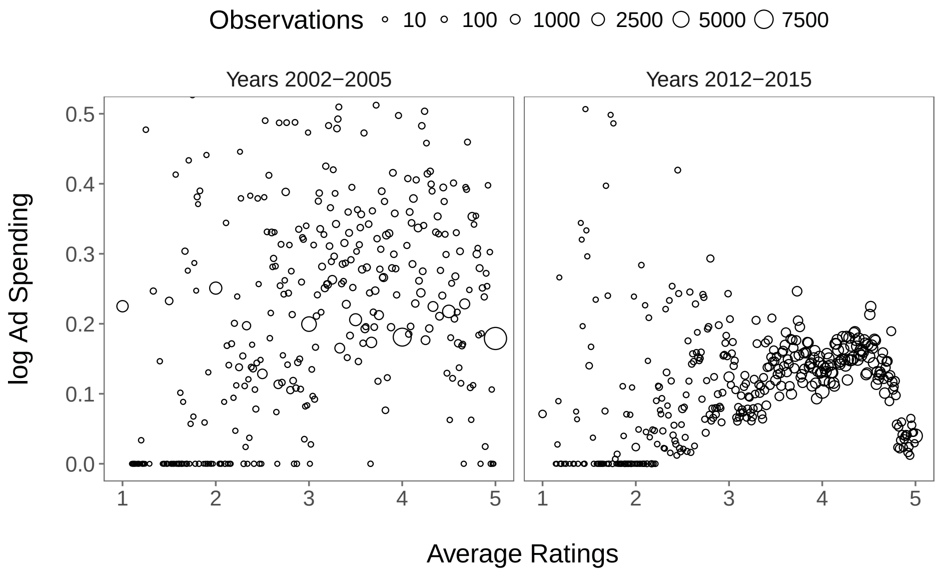

Figure 3: Relationship between advertising and hotel ratings: 2002-2005 versus 2012-2015.

Finally, the relationship between advertising and ratings itself is plotted in Figure 3 for two time periods, the years 2002-2005 and the years 2012-2015. We make two observations. First, comparing the two panels, the relationship between ratings and ad spending is very noisy in the early years, whereas it is fairly well-defined in the later years. Since TripAdvisor's influence has steadily increased over the years, this suggests that hotels started reacting to TripAdvisor reviews only when they had to. In 2002-2005, not too many people were visiting TripAdvisor, so hotels did not feel the need to adjust their ad spending to their user ratings. By contrast, by 2012-2015, large numbers of people are reading TripAdvisor reviews, and in the face of this widespread dissemination of independent quality information, hotels felt they had to fine-tune their ad spending to their ratings.

Second, looking at the recent data, hotels at the extremes of the rating distribution spend less on advertising than those in the middle of the distribution. This replicates other results of a similar nature reported in the literature (e.g., Horstmann and Moorthy (2003) for restaurants and Dhar and Moorthy (2017) for movies). One reason why this might be so is suggested by taking a closer look at the rating distribution. As Figure 4 shows, user ratings for hotels are tightly clustered around 4.0-4.5. When competing hotels have similar ratings, advertising rivalry ought to intensify, as each hotel tries to gain an edge via advertising what it could accomplish via ratings. In equilibrium, all hotels spend more. Under these circumstances, any change in user ratings that allows a hotel to pull away from the pack could produce big changes in its advertising expenditure. In Sections 5.4 and 5.5 we verify these conjectures.

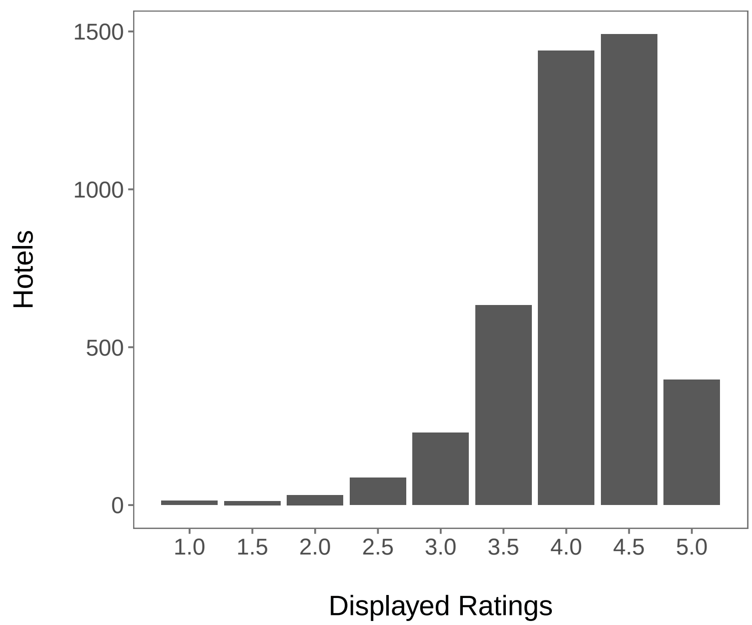

Figure 4: Distribution of user ratings for all hotels

5 Empirical analysis

5.1 Econometric framework

In this section, we estimate a demand-side causal effect of online ratings on advertising spending using a regression discontinuity design. The design exploits a feature of the TripAdvisor platform which generates plausibly exogenous variation in displayed online ratings.

As noted earlier, TripAdvisor's displayed summary ratings are average ratings rounded to the nearest half- or full-star. Thus, for example, a hotel with a real average rating of 3.74 is shown as a 3.5-star hotel, while a hotel with a real average rating of 3.75 stars is shown as having 4 stars. This rounding mechanism creates variation in the hotel ratings displayed to consumers that is effectively exogenous to the hotel's true quality in rounding-threshold neighborhoods. To the extent that hotel marginal costs depend on hotel quality, they depend on true quality, not the discontinuous changes in displayed ratings produced by TripAdvisor's rounding rule. Therefore, any changes in a firm's advertising strategy correlated with crossing those rounding thresholds represent a causal demand-side effect, not a supply-side effect.

Implicit in this identification strategy are two assumptions. First, average ratings are a continuous function of underlying quality. In practice, this means that we need to focus our estimation on a narrow band of observations around each threshold. Our large data set allows us to live with this restriction, while still retaining sufficient statistical power. Second, conditional on the underlying quality, variation in which firms are above the threshold and which are below is essentially random.

The last assumption would be problematic if firms could manipulate ratings (Mayzlin et al. 2014). The concern would be that hotels might write fake "user reviews" awarding high ratings to themselves and giving low ratings to their competitors. (Of course, if manipulation was an equal-opportunity game, all firms would end up where they started.) If firms could successfully pull off this manoeuver around the rounding thresholds, it would make sense for them to reduce their ad spending, but the resulting negative correlation would be spurious. TripAdvisor, of course, assures its readers that this is not a problem:

> We dedicate significant time and resources ensuring that the content on TripAdvisor reflects the real experiences of real travelers. We have quality assurance specialists who have brought a wide range of professional experience to enhance our prevention methods and our team spends thousands of hours every year ensuring the integrity of content on TripAdvisor. We also use automated tools that help flag questionable content for review, and our large and passionate community of millions of travelers keeps an eye out on our site as well.[14]

[14] `https://www.tripadvisor.ca/vpages/review_mod_fraud_detect.html`

This sounds plausible: after all, TripAdvisor's business model relies on its reputation for integrity. Still, to allay any residual concerns about review manipulation, we perform a series of statistical tests. These are described in detail in the Appendix; here we provide a brief synopsis.

The primary check is whether the selection variable—in this case, average rating—is continuous in a neighborhood of the rounding thresholds. If there is upward manipulation of ratings we should expect to see relatively few firms with ratings just below the threshold and a clump of firms with ratings just above the threshold; if there is both upward and downward manipulation of ratings, we should see clumps both below and above the threshold. Figure 6 in the Appendix shows that the distribution of average ratings is essentially uniform, with neither dip nor bump below the rounding cutoffs and no bulge above the rounding cutoffs.[15] The visual evidence is confirmed with the standard McCrary test of continuity for RD designs (McCrary 2008). As secondary checks, we follow Anderson and Magruder's (2012) suggestion to check whether variables plausibly associated with fake reviews and fake reviewers are continuous around the rounding thresholds. These tests also turn up no evidence of review manipulation.

Having assured ourselves that TripAdvisor's summary ratings are accurate reflections of consumers' actual experiences, we proceed to estimating our RDD regression. The specification is as follows:

$$\begin{aligned} \log \text{Ad Spending}_{it} &= \beta_1 \text{Above Threshold}_{it} + \beta_2 \text{Avg Ratings}_{it} \\ &+ \beta_3 \text{Above Threshold}_{it} \times \text{Avg Ratings}_{it} + \epsilon_{it}, \end{aligned} \tag{1}$$

The dependent variable is the logarithm of advertising spending of hotel i in the period $[t, t + 6 \text{ months}]$ and Avg Ratings_{it} is the average rating of hotel i at time t. Ad spending is summed over the following 6 months because ads tend to be purchased far ahead of when

[15] While round numbers such as the rounding cutoffs are more common, the rounding cutoffs themselves are no more likely than other non-rounding round numbers such as 4.15 and 4.35.

they are delivered; also, by aggregating over 6 months we reduce noise in the variable, even at the cost of being overly conservative in finding an effect of ratings on advertising. Above Threshold_{it}, whose coefficient is the main object of interest, is an indicator of whether the average rating of hotel i at time t falls above the rounding threshold (for it to be rounded up to the nearest half- or full-star). The inclusion of separate slopes for average ratings above and below the threshold allows for a more flexible specification; it increases our confidence that β_1 actually represents the difference in advertising spending between hotels that differ by a half-star in their displayed ratings. All specifications include year-month fixed effects. Finally, as far as the range of i is concerned, we only use hotels that are within 0.05 stars of each rounding threshold. In the Appendix we carry out a number of tests to check the robustness of our RDD implementation.

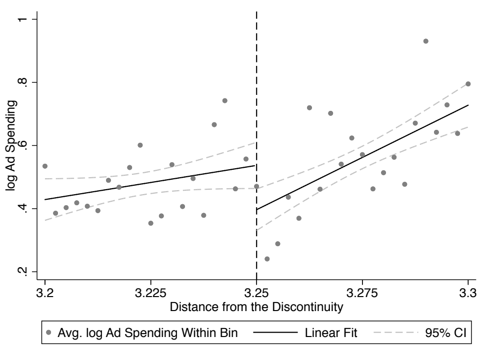

(a)

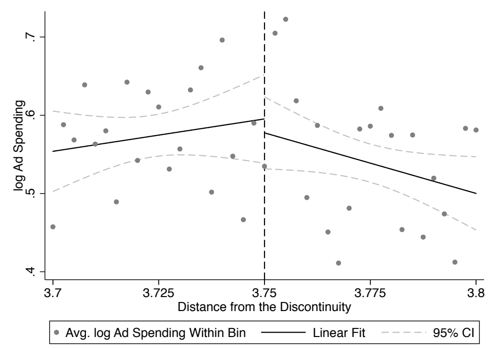

(b)

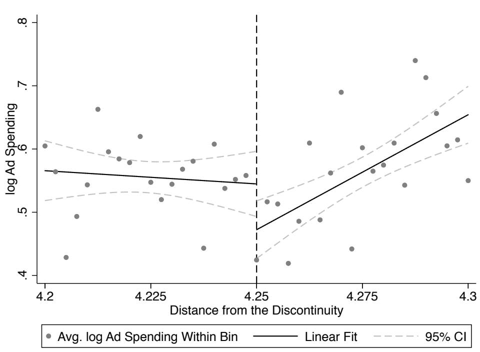

(c)

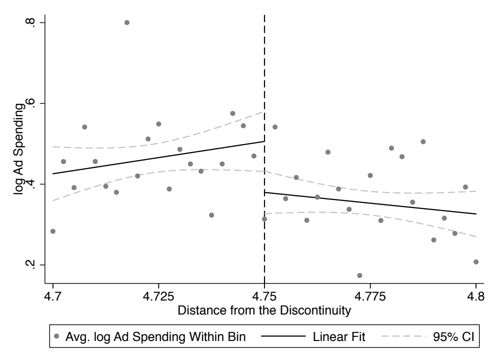

(d)

Figure 5: Log of total ad spending in the following six months by average TripAdvisor ratings

Table 2: How online user ratings affect ad spending: RDD estimates

	3.25	3.75	4.25	4.75
Above Threshold	−0.144**	−0.028	−0.088**	−0.123**
	(0.050)	(0.038)	(0.034)	(0.047)
Avg. Ratings	2.415	1.028	0.177	1.599
	(1.257)	(0.940)	(0.878)	(1.244)
Above Threshold × Avg. Ratings	4.125*	−2.381	3.372**	−2.740
	(1.712)	(1.243)	(1.162)	(1.545)
Month FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
N	8263	17090	19354	8169
R^2	0.050	0.013	0.022	0.029

Note: The dependent variable in each column is log of ad spending in the following 6 months. All columns use a bandwidth of 0.05 stars around the rounding cutoff in the column header. Only firms with 20 or more reviews are included. Robust standard errors in parentheses.

Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

5.2 How online ratings affect advertising spending

We start our analysis by presenting evidence that hotels do show sensitivity to their TripAdvisor summary ratings when choosing their ad spending level, and they do so in a way that suggests that they view the ratings and advertising are substitutes, not complements. In Figure 5 we show the relationship between logarithm of total ad spending and hotel ratings for different rounding thresholds: 3.25, 3.75, 4.25, and 4.75.[16] Recall that crossing a threshold increases the displayed rounded average rating by half a star. Thus, Figure 5 shows that for all but the 3.75 threshold, an increase in displayed ratings around the threshold results in a significant reduction in the amount of ad spending.

Table 2 presents the RDD estimates. Three main results stand out. First, there is indeed a causal effect between displayed online ratings and hotels' advertising spending in the cross-section. Despite the large amount of noise in the hotel ad spending data, we are still able to consistently detect a statistically significant effect. Second, the effect of ratings on ad spending is negative: as ratings go up, ad spending goes down. This is so even for

[16] At levels below 3 stars there are not enough observations to estimate an effect of being above or below the threshold. Fewer than 5% of hotels have a rating below 3 stars in 2015.

Table 3: Ad spending effects on the intensive and extensive margins

	Intensive Margin	Extensive Margin
Above Threshold	−0.230***	−0.020**
	(0.049)	(0.007)
Avg. Ratings	2.840*	0.240
	(1.236)	(0.173)
Above Threshold × Avg. Ratings	1.630	0.076
	(1.648)	(0.229)
Month FE	Yes	Yes
Year FE	Yes	Yes
Observations	17514	57334
R^2	0.033	0.200

Note: The dependent variable in columns 1 is a dummy variable for whether or not the firm advertised in the following 6 months; in column 2 it is the log of ad spending in the following 6 months for firms that advertised in the preceding year. Both columns use a pooled RDD with a bandwidth of 0.05 stars. Only firms with 20 or more reviews are included. Robust standard errors in parentheses.

Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

the threshold 3.75 where the coefficient is insignificant. Because of the consistent negative effect, in future analyses we pool all the thresholds together.

Because many firms do not advertise in a given period, we next examine whether the effect of ratings on advertising operates on the extensive or the intensive margin. In other words, are firms responding to their ratings by varying whether they advertise at all, or do they respond by varying the quantity of advertising? To test this, we repeat the RD analysis, first with a dummy dependent variable which equals 1 if the firm advertises in the following 6 months, and second on the subset of firms who had a positive ad spending level in the past year. Results of these analyses are presented in Table 3. We see a significant response in both margins. Firms are less likely to advertise if they are above the rounding threshold and, among the firms who advertise, the amount they spend is significantly lower for firms above the threshold.

Next, we look at whether the effect of ratings on ad spending varies by media. We use the same specification as Equation 1 but the dependent variable now is log of total ad spending within a particular medium over the following 6-month period. Results are presented in Table 4. As the table shows, our earlier finding is replicated within each media as well. Print, with a reduction in ad spending of 6% is most responsive, followed by Internet display and outdoor (4% and 2%, respectively); radio and television are the least responsive (1.5% and 1.3%, respectively).

Table 4: Ad spending by media type

	Internet	Outdoor	Print	Radio	TV
Above Threshold	−0.038***	−0.019*	−0.060***	−0.015***	−0.013**
	(0.008)	(0.008)	(0.018)	(0.004)	(0.005)
Avg. Ratings	0.360	0.617**	0.783	0.190	0.212
	(0.214)	(0.195)	(0.463)	(0.111)	(0.123)
Above Threshold × Avg. Ratings	0.369	−0.462	0.380	0.166	−0.096
	(0.283)	(0.260)	(0.610)	(0.149)	(0.159)
Month FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
N	52876	52876	52876	52876	52876
R^2	0.0070	0.0030	0.023	0.0033	0.0024

Note: The dependent variable is log of ad spending in the following 6 months in particular media. All columns use a pooled RDD with a bandwidth of 0.05 stars. Only firms with 20 or more reviews are included. Robust standard errors in parentheses.

Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

5.3 Hotel type effects

The previous section demonstrated that displayed online ratings have a significant effect on advertising spending, and that this effect operates in all media, and on both the extensive and intensive margins. In this section we explore whether the effect varies by whether the hotel is independent or part of a chain.

Table 5 presents the results. For chain hotels, the relationship is negative, but not statistically significant; for independent hotels, the relationship is a negative and highly significant ($p < .001$). In fact, the coefficient for independent hotels is not only statistically

Table 5: Ad spending by hotel type

	All hotels	Chain	Independent
Above Threshold	−0.089***	−0.051	−0.101***
	(0.020)	(0.038)	(0.024)
Avg. Ratings	1.087*	0.158	1.406*
	(0.520)	(0.969)	(0.615)
Above Threshold × Avg. Ratings	0.865	1.133	0.862
	(0.686)	(1.280)	(0.812)
Month FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
N	52876	15805	37071
R^2	0.014	0.031	0.013

Note: The dependent variable is log of ad spending in the following 6 months for the hotels covered by the column heading. All columns use a pooled RDD with a bandwidth of 0.05 stars. Only firms with 20 or more reviews are included. Robust standard errors in parenthesis.

Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

significant, it is also economically meaningful; to wit, a half-star increase in TripAdvisor ratings reduces independent hotels' ad spending by about 10 percent. This suggests that the ad responses of independent hotels drives much of the overall result. Why might independent hotels respond more than chain hotels? We think this is because chains have stronger brands than independent hotels. The chains' national branding and customer loyalty insulates them from reviews—both on the positive side and on the negative side. When demand is insensitive to reviews, these hotels don't see a reason to respond.

5.4 Market competitiveness effects

Here we examine how the competitiveness of a market, defined by how similarly-rated competitors are, affects the relationship between ratings and ad spending. We use the STR definition of a market as a Metropolitan Statistical Area, and then calculate the standard deviation of average ratings in each market each year. In markets with relatively little differentiation in hotel ratings, even a small difference in ratings might be pivotal; so in

Table 6: Ad spending effects by market competitiveness

	High ratings std. dev.	Low ratings std. dev.
Above Threshold	−0.034	−0.16***
	(0.043)	(0.042)
Avg. Ratings	−2.91*	3.78***
	(1.181)	(1.037)
Above Threshold × Avg. Ratings	4.64**	−3.17*
	(1.527)	(1.373)
Month FE	Yes	Yes
Year FE	Yes	Yes
N	9238	13023
R^2	0.004	0.003

Note: The left (right) column refers to markets where the standard deviation of average ratings is higher (lower) than the median standard deviation. The dependent variable in each column is log of ad spending in the following 6 months. All columns use a pooled RDD with a bandwidth of 0.05 stars. Only firms with 20 or more reviews are included. Robust standard errors in parentheses.

Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

such markets we should see a bigger ad response than in markets where ratings are already well-differentiated to begin with.

Table 7 compares results between markets with above-the-median differentiation in user ratings and markets with below-the-median differentiation in user ratings. The results are as predicted: ratings have a strong negative effect on ad spending in less-differentiated markets, but not in well-differentiated markets.

5.5 Early versus late-period effects

In Figure 3 we saw that the ratings-advertising relationship in 2002-2005 was mostly noise, whereas in 2012-2015 the noise had substantially died down, and an inverted-U relationship was plain to see. We think this difference is driven by TripAdvisor's growing influence. In 2002, we suspect, few people visited TripAdvisor; by 2015, however, TripAdvisor had become one of the top-50 most visited websites (see footnote 8).

Table 7 verifies this intuition. It shows the results of two RD regressions, one based on 2002-2005 data and another based on 2012-2015 data. As can be seen, changes in user ratings

Table 7: Effect of ratings on ad spending, early versus late

	2002-2005	2012-2015
Above Threshold	−0.206	−0.106***
	(0.168)	(0.027)
Avg. Ratings	6.132	1.166
	(3.950)	(0.705)
Above Threshold × Avg. Ratings	−0.058	0.275
	(5.337)	(0.930)
Month FE	Yes	Yes
Year FE	Yes	Yes
N	1725	26510
R^2	0.032	0.0016

Note: The dependent variable is log of ad spending in the following 6 months during the period of the column heading. All columns use a pooled RDD with a bandwidth of 0.05 stars. Only firms with 20 or more reviews are included. Robust standard errors in parentheses.
Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

do not result in changes in ad spends in the 2002-2005 period, but they do in the 2012-2015 period. In fact, as expected, the negative relationship in the later period is even stronger than in our earlier regressions where we aggregated across all periods (Table 5, column 1).

This analysis makes an important point. It is not the presence of reviews *per se* that induces firms to alter their advertising level, but rather, the belief that people may be attending to them.

6 Discussion

Our results tell an interesting story about the role of advertising in the Internet era. As more and more people get information about products and services from other users, rather than from the companies themselves, advertising's role must necessarily change. To understand what those changes might be, it is useful to frame the discussion in terms of Nelson's (1970) search goods-experience goods framework. Search goods are those whose quality can be observed before purchase; experience goods are those whose quality can be observed only after purchase. We noted earlier that hotels are experience goods. However, they become search goods in the presence of reviews.

What role does advertising play for these newly-coined search goods? As Nelson (1974) noted, for search goods, advertising can convey information directly. Indeed, that seems to be what happens in some categories. For example, Dhar and Moorthy (2017) note that movies that get good ratings from critics, proudly tout those ratings in their ads; similarly, mutual funds which get high ratings from Morningstar feature those ratings in their ads (Krouse and Grind 2017). What about firms that get poor reviews? They will obviously not want to advertise that fact. What are their alternatives? One option is to stay silent. Another option might be to use advertising to change the subject. For instance, they can advertise other search attributes on which they are strong, or, changing the subject entirely, not even talk about the product, i.e., pursue image advertising. The case for advertising to complement ratings is strongest when firms with good ratings advertise their ratings and firms with relatively poor ratings stay silent. Indeed, this is what happens for limited-release movies (Dhar and Moorthy 2017). However, this is not a common scenario: typically, even firms with poor ratings have advertising options. In that case, the correlation between ratings and advertising spending must necessarily become less positive.

In practice, even firms with good reviews may not be able to use user ratings in their advertising. Both of our examples above featured professional reviewers, not users. This may not be a coincidence. Ad claims using user ratings are unlikely to be credible (Dean and Biswas 2001). Also, review websites may not allow their user ratings to be used in advertising. This raises the question, can advertising have a role when it can't incorporate review information? We believe it can. First, as noted above, advertising can convey information about other search attributes not covered in the reviews, and, of course, image advertising is always an option. For example, a movie might advertise its stars or its genre. Second, and perhaps more important, the brand awareness/familiarity generated by advertising can itself be used as an accessory to the consumer's search process. For instance, it can direct the user's search through reviews, as Hoyer and Brown's (1990) laboratory experiments suggest, or it can play a role at the back end of the search process after reviews have been consulted.

Let us review our evidence in light of this framework. First, the fact that ad spending by hotels declined over a thirteen-year time-period in which (inflation-adjusted) ad prices were either stable or declining suggests that demand-side forces are behind the advertising decline, not supply-side forces. Inevitably, it points to a diminished role for advertising in the presence of user reviews. This effect has been predicted for quite some time—even predating the Internet (Archibald et al. 1983). More recently it has led some observers to predict the "twilight of brands" (Surowiecki 2014) and others to predict "the death of advertising" (DeMers 2016).

Our results provide a more nuanced view of these predictions. While advertising might have a more diminished role in the presence of user reviews, what role it still has, has perhaps become even more strategically important for the firms involved. Firms with different user ratings must use advertising differently because advertising's role differs for different review endowments. Simply put, advertising has a more limited role to play for well-reviewed firms because free advertising is a good substitute for paid advertising. By contrast, firms with less stellar reviews need paid advertising all the more, either to direct the consumer's search through reviews in a way that might prove most favorable to them—for example, by trying to induce brand-based search instead of open-ended search—or to target consumers who don't read reviews. Finally, in markets where ratings are largely homogeneous, we expect a big rivalry in advertising spending as each firm tries to get an edge via brand recognition what it could not get through user reviews. Under these circumstances, a small difference in ratings might prove pivotal, inducing a large effect on ad spending for the firm that breaks free. Our results in Section 5.4 support this view. With respect to the role of brands, while the Internet might have contributed toward leveling the playing field between established brands and little-known brands, our results in Section 5.3 show that established brands—such as those of the chain hotels—continue to enjoy some immunity from reviews.

Finally, the observation that the relationship between user ratings and advertising spending has strengthened over time is testimony to the increasing influence of TripAdvisor, and further confirmation that user ratings affect ad spending through the demand-side and not the supply side. Essentially, hotels react to online reviews if and only if consumers do.

7 Conclusion

This paper investigates the relationship between online user reviews and advertising spending in the hotel industry using a 14-year panel of TripAdvisor hotel reviews matched to advertising data from Kantar Media. Our results suggest that hotels treat TripAdvisor ratings as a substitute for advertising as well as treat advertising as a substitute for TripAdvisor ratings. The former is observed in the secular decline in ad spending over the last thirteen years, just as online reviews have gained in popularity. The latter effect manifests in the cross-sectional relationship between ad spending and user ratings; firms with lower ratings advertise more than firms with higher ratings.

Beneath the broad substitution relationship, there are several interesting nuances. Independent hotels are more responsive to reviews than chains, and hotels in more competitive markets are more responsive than hotels in less competitive markets. Over time, all hotels have become more responsive, just as TripAdvisor's influence has grown, This tells us that it is not the presence of reviews *per se* that triggers a reaction, but rather the recognition that consumers are reacting to them.

The negative relationship between user ratings and advertising spending seen here is, we believe, quite robust. It should apply in many categories, not just hotels; indeed, Lei (2017) finds it in restaurants. However, it is unlikely to be universal, as Dhar and Moorthy's (2017) work with movies suggests. This is only to be expected; given the many roles of advertising, anything else would be a surprise.

References

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

Archibald, Robert B, Clyde A Haulman, Carlisle E Moody Jr. 1983. Quality, price, advertising, and published quality ratings. *Journal of Consumer Research* **9**(4) 347–356.

Bagwell, Kyle. 2007. The economic analysis of advertising. *Handbook of industrial organization* **3** 1701–1844.

Cabral, Luís, Ali Hortaçsu. 2010. The dynamics of seller reputation: Evidence from ebay*. *The Journal of Industrial Economics* **58**(1) 54–78.

Caves, Richard E., David P. Greene. 1996. Brands' quality levels, prices, and advertising outlays: empirical evidence on signals and information costs. *International Journal of Industrial Organization* **14**(1) 29–52.

Chen, Yubo, Jinhong Xie. 2005. Third-party product review and firm marketing strategy. *Marketing Science* **24**(2) 218–240.

Chen, Yubo, Jinhong Xie. 2008. Online consumer review: Word-of-mouth as a new element of marketing communication mix. *Management science* **54**(3) 477–491.

Chevalier, Judith A, Dina Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of marketing research* **43**(3) 345–354.

Chioveanu, Ioana. 2008. Advertising, brand loyalty and pricing. *Games and Economic Behavior* **64**(1) 68–80.

Dean, Dwane Hal, Abhijit Biswas. 2001. Third-party organization endorsement of products: An advertising cue affecting consumer prepurchase evaluation of goods and services. *Journal of Advertising* **30**(4) 41–57.

DeMers, Jayson. 2016. Is traditional advertising dead? *Forbes* March 22, 2016.

Dhar, Tirtha, Sridhar Moorthy. 2017. On the marketing of experience goods: the case of movies. Working Paper, University of Toronto.

Hollenbeck, Brett. 2017. Online Reputation Mechanisms and the Decreasing Value of Brands. Working Paper, UCLA.

Honka, Elisabeth, Ali Hortaçsu, Maria Ana Vitorino. 2017. Advertising, consumer awareness, and choice: Evidence from the US banking industry. *The RAND Journal of Economics* **48**(3) 611–646.

Horstmann, Ignatius J., Sridhar Moorthy. 2003. Advertising spending and quality for services: The role of capacity. *Quantitative Marketing and Economics* **1**(3) 337–365.

Houser, Daniel, John Wooders. 2006. Reputation in auctions: Theory, and evidence from eBay. *Journal of Economics & Management Strategy* **15**(2) 353–369.

Hoyer, Wayne D., Steven P. Brown. 1990. Effects of brand awareness on choice for a common, repeat-purchase product. *Journal of consumer research* 141–148.

Jin, Ginger Zhe, Andrew Kato. 2006. Price, quality, and reputation: Evidence from an online field experiment. *The RAND Journal of Economics* **37**(4) 983–1004.

Kim, MinChung, Leigh M. McAlister. 2011. Stock market reaction to unexpected growth in marketing expenditure: Negative for sales force, contingent on spending level for advertising. *Journal of Marketing* **75**(4) 68–85.

Krouse, Sarah, Kirsten Grind. 2017. Some fund ads misrepresent morningstar ratings. *Wall Street Journal* November 19, 2017.

Lei, Ying. 2017. Local restaurants' advertising response to a better online rating. Working Paper, Peking University.

Lei, Ying, Sridhar Moorthy. 2017. Advertising strategy in the presence of reviews: A theoretical analysis Working paper, University of Toronto.

Lewis, Gregory, Georgios Zervas. 2016. The welfare impact of consumer reviews: A case study of the hotel industry .

Luca, Michael. 2016. Reviews, reputation, and revenue: The case of yelp. com. Working Paper, Harvard Business School.

Mayzlin, Dina, Yaniv Dover, Judith Chevalier. 2014. Promotional reviews: An empirical investigation of online review manipulation. *The American Economic Review* **104**(8) 2421–2455.

McCrary, Justin. 2008. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of econometrics* **142**(2) 698–714.

Melnik, Mikhail I, James Alm. 2002. Does a seller's eCommerce reputation matter? evidence from eBay auctions. *The Journal of Industrial Economics* **50**(3) 337–349.

Milgrom, Paul, John Roberts. 1986. Price and advertising signals of product quality. *Journal of political economy* **94**(4) 796–821.

Nelson, Phillip. 1970. Information and consumer behavior. *Journal of political economy* **78**(2) 311–329.

Nelson, Phillip. 1974. Advertising as information. *Journal of political economy* **82**(4) 729–754.

Proserpio, Davide, Georgios Zervas. 2017. Online reputation management: Estimating the impact of management responses on consumer reviews. *Forthcoming Marketing Science* .

Resnick, Paul, Richard Zeckhauser, John Swanson, Kate Lockwood. 2006. The value of reputation on eBay: A controlled experiment. *Experimental Economics* **9**(2) 79–101.

Shum, Matthew. 2004. Does advertising overcome brand loyalty? Evidence from the breakfast-cereals market. *Journal of Economics & Management Strategy* **13**(2) 241–272.

Sun, Monic. 2012. How does the variance of product ratings matter? *Management Science* **58**(4) 696–707.

Surowiecki, James. 2014. Twilight of the brands. *The New Yorker* February 17, 2014.

Sutton, John. 1991. *Sunk costs and market structure: Price competition, advertising, and the evolution of concentration*. MIT press.

Appendix

Here we describe (i) the various tests we did to assure ourselves that TripAdvisor ratings were being manipulated by firms, and (ii) whether alternative implementations of the RD design affect any of our conclusions.

1. Tests to check for manipulation of TripAdvisor ratings

Several tests are possible to check for ratings manipulation and fake reviews. The most direct approach is to test whether the density of average ratings—the subject of the manipulation, if any—is continuous around the rounding thresholds. Other tests, following Anderson and Magruder (2012), check whether various variables that might be associated with fake reviews and fake reviewers are continuous around the rounding thresholds.

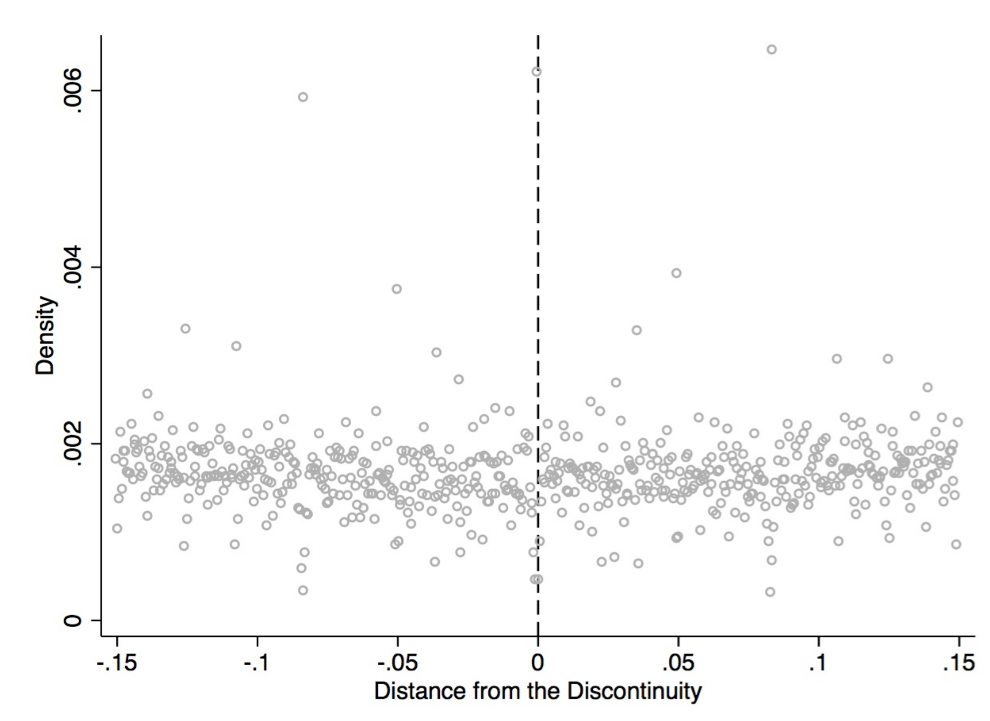

Figure 6: McCrary test: distribution of average ratings near rounding thresholds

Figure 6 illustrates the direct approach. The density of the running variable, the hotels' average rating, seems constant around the rounding thresholds. There are no jumps, either up or down, either below or above the thresholds. A formal McCrary (2008) test confirms what we see visually. The RDD estimates using the density of the running variable and two different bandwidths (0.05 and 0.1 stars) are reported in Table 8. The coefficient of interest is insignificant in both column 1 and 2, suggesting that the running variable is not being manipulated.

Table 8: McCrary test.

	BW=0.05	BW=0.1
Above Threshold	−0.0008	0.0001
	(0.0010)	(0.0003)
Avg. Ratings	−0.0304	−0.0036
	(0.0247)	(0.0032)
Above Threshold × Avg. Ratings	0.0893	0.0039
	(0.0489)	(0.0043)
Month FE	Yes	Yes
Year FE	Yes	Yes
N	251	501
R^2	0.66	0.33

Note: The dependent variable is density of the running variable computed using a bin size of 0.0004 stars. All columns use pooled RDDs. Only firms with 20 or more reviews are included. Robust standard errors in parentheses.

Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Second, we check for the continuity of various variables associated with fake reviews/reviewers. One set of variables is based on the idea that fake reviewers will have fewer reviews overall than genuine reviewers. Another, based on Mayzlin et al.'s (2014) observation that positive fake reviews tend to be 5-stars, checks whether the fraction of 5-stars reviews is significantly different above and below the thresholds. Finally, Proserpio and Zervas (2017) show that hotel management responding to reviews is an effective strategy to increase hotel ratings. We, therefore, check whether hotels use review responses to move their ratings above the threshold. The results of these analyses are in Table 9. Again we find no evidence of manipulation.

In short, neither the direct test, nor the indirect tests, show any evidence of ratings manipulation.

Table 9: Test for review manipulation: reviews and reviewers characteristics

	Fraction of 5-stars Reviews	Average Reviews per Reviewer	Reviewers with 0 Reviews	Reviewers with < 6 Reviews	Fraction of Responses
Above Threshold	0.003	−0.016	0.000	0.001	−0.002
	(0.003)	(0.018)	(0.003)	(0.002)	(0.004)
Avg. Ratings	0.183*	1.047*	−0.124	−0.019	0.138
	(0.084)	(0.455)	(0.067)	(0.054)	(0.090)
Above Threshold × Avg. Ratings	−0.002	−1.061	0.189*	0.026	−0.354**
	(0.113)	(0.615)	(0.091)	(0.074)	(0.117)
Month FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
N	57334	46482	47234	47234	57334
R^2	0.027	0.31	0.066	0.25	0.24

Note: The dependent variable in column 1 is fraction of 5-stars reviews; in column 2 it is average reviews per reviewer prior to the current review; in columns 3 and 4 it is, respectively, fraction of reviewers with no prior reviews and fraction of reviewers with less than 6 prior reviews. All columns use pooled RDDs with a bandwidth of 0.05 stars. Only firms with 20 or more reviews are included. Robust standard errors in parentheses.

Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

2. Robustness of the RD specification

Quadratic polynomial estimator. All our analysis in the text was based on a local linear estimator. In this section we test the sensitivity of our results to a different functional form, a quadratic polynomial estimator. We do so by adding to Equation 1 a quadratic term for the cumulative average rating variable, and allow for a separate effect of the quadratic term above and below the discontinuity:

$$\begin{aligned} \log \text{Ad Spending}_{it} &= \beta_1 \text{Above Threshold}_{it} + \beta_2 \text{Avg Ratings}_{it} + \beta_3 \text{Avg Ratings}^2_{it} \\ &+ \beta_4 \text{Above Threshold}_{it} \times \text{Avg Ratings}_{it} \\ &+ \beta_5 \text{Above Threshold}_{it} \times \text{Avg Ratings}^2_{it} + \epsilon_{it}. \end{aligned} \tag{2}$$

We report the estimates of this model in Table 10. The results are qualitatively similar to those using a local linear estimator, suggesting that our results are not sensitive to the functional form used for the estimation.

Table 10: Sensitivity analysis: quadratic polynomial estimator

	All Firms
Above Threshold	−0.079*
	(0.031)
Avg. Ratings	−0.759
	(2.239)
Avg Ratings^2	−28.171
	(40.591)
Above Threshold × Avg. Ratings	4.879
	(2.782)
Above Threshold × Avg Ratings^2	−17.404
	(52.242)
Month FE	Yes
Year FE	Yes
N	52876
R^2	0.00042

Note: The dependent variable is log of ad spending over the next 6 months. All columns use pooled RDDs with a bandwidth of 0.05 stars. Only firms with 20 or more reviews are included. Robust standard errors in parentheses.

Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Different bandwidth. We test the sensitivity of our results to alternative choices of bandwidth. We report the results for all hotels and three different bandwidth in Table 11. The coefficient of interest, Above Threshold, is statistically significant for every bandwidth we tested, however the effect of ratings it is stronger for smaller bandwidth choices.

Placebo test. As a final check, we perform a placebo test. We estimate the regression discontinuity on placebo thresholds, i.e., thresholds where there is not a jump in the displayed average ratings. We report the results for the pooled RDD for placebo thresholds just below the real discontinuity, i.e., 3.2, 3.6, 4.2, and 4.7, in Table 12. The coefficient of interest, Above Threshold, is statistically insignificant for all specifications. These results give further credibility to our RDD analysis.

Table 11: Sensitivity analysis: different bandwitdhs

	BW=0.025	BW=0.05	BW=0.1
Above Threshold	−0.124***	−0.089***	−0.038*
	(0.030)	(0.020)	(0.017)
Avg. Ratings	3.867*	1.087*	1.195*
	(1.589)	(0.520)	(0.519)
Above Threshold × Avg. Ratings	−0.915	0.865	−1.194*
	(1.974)	(0.686)	(0.519)
Month FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
N	26157	52876	104488
R^2	0.017	0.014	0.014

Note: The dependent variable is log of ad spending over the next 6 months. All columns use pooled RDDs with a bandwidth of 0.05 stars. Only firms with 20 or more reviews are included. Robust standard errors in parentheses.

Significance levels: * p<0.05, ** p<0.01, *** p<0.001.

Table 12: Sensitivity analysis: placebo test

	All Firms
Above Threshold	−0.011
	(0.021)
Avg. Ratings	−0.293
	(0.539)
Above Threshold × Avg. Ratings	0.007
	(0.702)
Month FE	Yes
Year FE	Yes
N	52796
R^2	0.012

Note: The dependent variable in each column is the log of ad spending over the next 6 months. All columns use a pooled RDD around placebo cutoffs (3.2, 3.6, 4.2, 4.7 stars), with a bandwidth of 0.05 stars. Only firms with 20 or more reviews are included. Robust standard errors in parentheses.

Significance levels: * p<0.05, ** p<0.01, *** p<0.001.