How Does Online Reputation Affect Social Media Endorsements and Product Sales? Evidence from Regression Discontinuity Design

Xitong Li*

November 30, 2013

Abstract

Despite the increasing importance of social media marketing, little extant research has explored what factors consumers would take into account in the decision-making of endorsing a product to their peers with established ties via social media. This paper examines if online reputation (restaurants' displayed Yelp ratings), which helps update consumers' perception of product value, is a causal factor that affects consumers' decisions of endorsing via Facebook and purchasing products (the restaurants' vouchers). We build a stylized Bayesian learning model and derive the hypotheses: (1) a higher online reputation leads to increased social media endorsements and voucher sales, but only when the number of review ratings is sufficiently large; (2) these effects are greater for restaurants with more reviews; and (3) these effects are greater for restaurants with a larger variance in the review ratings. Interestingly, the third hypothesis contrasts to the predictions by some established theories (e.g., cue diagnosticity theory). We test the hypotheses using data of Groupon and LivingSocial deals. To identify the causal effects of online reputation, we use a regression discontinuity design by exploiting the institutional feature that displayed Yelp ratings are rounded to the nearest half star. The empirical results largely support our hypotheses. In particular, we find the effect of displayed Yelp ratings on Facebook Likes is greater when the variance of ratings is larger, suggesting that perhaps consumers are risk averse in the decision-making of endorsing a product to their peers. Yet, the effect on voucher sales does not significantly differ with the variance.

* MIT Sloan School of Management, e-mail: xitongli@mit.edu. I thank Erik Brynjolfsson, John Hauser, Stuart Madnick, Catherine Tucker, Duncan Simester, and Juanjuan Zhang for their valuable comments and suggestions. All errors are my own.

1 Introduction

Most online review sites (e.g., Yelp, Amazon) calculate an overall rating score by averaging across all individual ratings of a product. The overall average rating becomes an indicator of online reputation signaling the product quality. Besides that, many review sites prominently display the total number of ratings and make the dispersion of ratings available by showing the numbers of individual ratings at each rating level (often from 1 to 5 star). The central hypothesis underlying such practice is that online reputation together with the number and dispersion of ratings could influence consumers' shopping behaviors.

Prior research (e.g., Chevalier and Mayzlin 2006, Luca 2011) has focused on establishing the casual impact of online reputation on product sales, but neglects consumers' social media endorsements which are also significantly meaningful to firms (Aral et al. 2013). For example, Facebook.com provides the "Like" button allowing users to share and endorse any webpage. The activity that the users have "liked" the webpage is immediately displayed to their friends via Facebook newsfeeds. Recent studies show that consumers' social media activities can increase product awareness (Aral and Walker 2011), drive additional sales (Chen et al. 2011, Li and Wu 2013), and enhance brand loyalty (Rishika et al. 2013). Chompon, an e-commerce platform company, estimates that each Facebook Like is worth $8 for its clients in terms of the immediate next sale.[1] Consumers' social media activities may also have a significant predictive power for firm equity value (Luo et al. 2013). Therefore, engaging with consumers through social media has become "a critical element of any organization's marketing strategy" (Malhotra et al. 2013) and the volume of social media endorsements (e.g., Facebook Likes) is a meaningful and increasingly important indicator to firms' business performance (Miller and Tucker 2013). The importance of social media endorsements is also evident in the fact that there exists a commercial market for buying them.[2]

Consumers' social media endorsements are distinct from product sales, because the motives and costs of endorsing a product to one's peers with established ties via social media are

[1] See http://techcrunch.com/2011/02/18/facebook-shares-are-worth-almost-three-times-more-than-tweets-for-e-commerce/ (accessed on July 9, 2013)

[2] A simple search on Google using the keywords "buy Facebook Likes" provides a list of companies that sell Facebook Likes to paying customers, such as get-likes.com, buylikes.com, and fblikesmart.com.

different from buying it for one’s own consumption. Consumers endorse a product via Facebook, perhaps because the product is interesting (special, unique) and they want to express their preferences for it publicly, or because it is a good deal and they want to inform their friends about it. In either case, consumers may expect to gain “social currency” if their friends appreciate the endorsement (Berger and Schwartz 2011). From the cost perspective, such an endorsement can be done with minimal involvement (i.e., a click on the Facebook “Like” button) and no monetary cost, but consumers may put their self-reputation at risk; endorsing a “bad” product to Facebook friends would probably damage one’s self-reputation (Wojnicki and Godes 2008). Therefore, consumers’ decision-making of endorsing a product is different from their decision-making of purchasing it and deserves to be investigated separately.

Despite the importance and distinction, little extant research has explored what factors consumers would take into account in the decision-making of endorsing a product to their peers. Our study aims to fill this gap in the literature by investigating how online reputation, which helps update consumers’ perception of product value, affects social media endorsements. For the sake of comparison, we also examine the effect of online reputation on product sales. Although psychological theory of consumer choice (Hansen 1976) suggests that the effect of a determinant (herein, online reputation) is often moderated by contextual factors, the moderating role played by the number and variance of individual ratings has not received much attention (Sun 2012). Therefore, we study the moderating effects of the number and variance of ratings, from which we produce important implications about consumers’ endorsing behaviors. Specifically, we seek to answer the following questions in this study:

(1) Does a higher online reputation increase consumers’ social media endorsements and product sales?
(2) How does the number of ratings moderate the effect of online reputation?
(3) How does the variance of ratings moderate the effect of online reputation?

To answer the questions, we, based on the theory of word-of-mouth (WOM) as self-enhancement (Wojnicki and Godes 2008), assume consumers’ propensity to endorse a product via social media is dependent on their expected utility of the product (perception of the product value), so is their propensity to buy. Then, we develop a simple stylized Bayesian learning model and show the structural relationship between a product's review ratings and consumers' posterior expected utility of the product. The analytical results from the stylized model produce testable hypotheses.

Empirically, we examine the situation in which restaurants with review ratings on Yelp.com sell deal vouchers through Groupon.com and LivingSocial.com. Being influenced by the restaurants' displayed Yelp ratings, consumers may endorse the restaurant deals via Facebook and/or buy the vouchers. With a data set collected from Groupon/LivingSocial, Facebook and Yelp, we are able to identify the causal impacts of displayed Yelp ratings on consumers' endorsements via Facebook and voucher sales by using a regression discontinuity (RD) design (Hartmann et al. 2011, Imbens and Lemieux 2008, Lee and Lemieux 2010). In line with the recent econometric literature (Lee and Lemieux 2010), we carefully assess the validity of the RD design implemented in our study using a number of robustness checks. The results show that a restaurant's higher displayed Yelp rating causes to increase consumers' endorsements (i.e., more Facebook Likes) and voucher sales, but only when the number of ratings is sufficiently large. Supporting the assumption that consumers' propensity to endorse a product depends on their expected utility of the product, the empirical findings suggest expected utility (perception of product value) is a key factor that consumers would take into account in the decision-making of endorsing. The magnitudes of the estimated effects are practically significant. For restaurants with at least 20 Yelp reviews, an extra half-star displayed Yelp rating increases the aggregate volumes of Facebook Likes and voucher sales by 26.3% and 17.4%, respectively, controlling for observed (and unobserved) characteristics of restaurant deals. However, these effects decrease significantly and even disappear for restaurants with fewer Yelp reviews.

More importantly, there seems to be no conclusive theoretical prediction for the moderating role of the variance of ratings. On the one hand, some established theories (Basuroy et al. 2006, Feldman and Lynch 1988, Sun 2012) predict that consumers' responses to the average rating would decrease with the variance of ratings. For example, the cue diagnosticity theory (Feldman and Lynch 1988) suggests that consumers would reduce their reliance on the average rating as a quality signal when the variance of ratings is large, because they may find the quality signal is nondiagnostic (Basuroy, et al. 2006). Consequently, consumers would be less responsive to the average rating when the variance of ratings is larger. In a separate study, Sun (2012) develops an analytical model and shows the interaction effect between the average and variance of ratings on product sales is negative.

On the other hand, however, our stylized Bayesian learning model, based on fairly general but different assumptions from the model of Sun (2012), shows that risk aversion could make consumers' posterior expected utility more responsive to the average rating when the variance of ratings is larger. Unlike the conventional Bayesian learning literature (Ching et al. 2011, Roberts and Urban 1988, Zhao et al. 2013), our model does not assume any explicit form for consumers' utility functions. Thus, the theoretical implications of our model hold true with a broad set of utility functions, including that for constant or decreasing absolute risk aversion (CARA / DARA) (Friend and Blume 1975).

Therefore, the competing predictions from our stylized model and alternative theories (Basuroy, et al. 2006, Feldman and Lynch 1988, Sun 2012) raise an interesting empirical question that has important theoretical implications. Consistent with the results of our model, we find the effect of displayed Yelp ratings on Facebook Likes is greater when the variance of ratings is larger, implying that perhaps consumers are risk averse in endorsing restaurant deals via Facebook. Yet, we find the effect on voucher sales does not significantly differ with the variance, possibly because the mechanisms expounded by the competing theories offset each other in terms of purchasing. The fact that the variance of ratings moderates the effect of displayed Yelp ratings on Facebook Likes but not on voucher sales reveals that consumers exhibit different behaviors in endorsing versus purchasing.

Our study contributes to the emerging literature that focuses on demonstrating the importance of consumers' social media activities (Aral, et al. 2013). For example, Rishika, et al. (2013) find consumers' participation in a firm's social media effort leads to an increase in consumer visit frequency. Li and Wu (2013) find consumers' Facebook Likes drive additional product sales. Kosinski et al. (2013) show that Facebook Likes can be used to predict sensitive personal attributes. Despite the emerging literature on the importance of Facebook Likes, little is known about the determinants of consumers' endorsing decisions via social media. Egebark and Ekström (2012) conduct one of the first studies in this research stream by showing that social proximity and number of predecessors have significant impacts on a user's probability to "Like" a Facebook status update. While Egebark and Ekström (2012) study the consumer-side influence on "liking" Facebook status updates, our research examines the seller-side influence on "liking" commercial products. Our study is also distinct from the works by Moe and Schweidel (2012) and Muchnik et al. (2013), because the motives and costs of endorsing a product to one's peers via social media are different from that of writing an online product review (Moe and Schweidel 2012) or voting up a news article (Muchnik, et al. 2013).[3] To the best of our knowledge, we are the first to establish and quantify the causal impact of a seller's online reputation (user-generated review ratings) on consumers' endorsing commercial products to their peers with established social ties. Our findings suggest that consumers take into account their perception of product value when they make the decision of endorsing it to their peers.

Our study also contributes to a large body of literature that examines the impact of review ratings on product sales. The existing literature, however, documents mixed empirical findings. While a considerable number of studies document a higher average rating could increase product sales (Chevalier and Mayzlin 2006, Chintagunta et al. 2010), it has been recognized that the average rating may not necessarily reveal the true product quality (Hu et al. 2009) or influence consumers' purchasing decisions (Eliashberg and Shugan 1997) due to at least two reasons. First, consumers may realize online reviews could be posted by biased consumers (Li and Hitt 2008). Second, firms have incentives to manipulate their online reputation by posting fake reviews (Dellarocas 2006). It is thus not surprising that Liu (2006) finds the valence of movie messages has little explanatory power for movie revenue. Duan et al. (2008) show similar findings and conclude "online user reviews have little persuasive effect on consumer purchase decisions". Therefore, whether a higher online reputation increases product sales (and/or consumers' social media endorsements) is still an open empirical question.

One way that could potentially reconcile the seemingly inconsistent empirical findings about the impact of review ratings is to examine the moderating role played by contextual factors. For example, Zhu and Zhang (2010) find the average rating has an influential impact on

[3] The key distinction lies in the fact that receivers of social media endorsements are consumers' peers with established social ties and thus they may expect to gain social currency or risk their self-reputation from the endorsements, whereas users on online review or news sites are often anonymous and have no established social ties among them.

sales of video games only for less popular games. However, the literature in this research stream is still scant. Our study contributes to this growing literature by reporting that the number of ratings could moderate the effect of average ratings. The variance of ratings is another important moderating factor. For example, Sun (2012) shows the interaction effect between the average and variance of ratings on book sales is negative. Interestingly, we find the effect of displayed Yelp ratings on Facebook Likes is greater when the variance is larger.

Perhaps the two prior studies by Luca (2011) and Anderson and Magruder (2012) are most closely related to our study from the perspective of empirical identification. Our study differs in the following important aspects. First, we seek to establish the causal impact of displayed Yelp ratings on consumers' social media endorsements, whereas the prior studies still focus on the effect on product sales, such as restaurant revenues (Luca 2011) and table reservation availability (Anderson and Magruder 2012). Second, we investigate how such causal impacts are moderated by the number and variance of individual ratings. Last but not least, while Luca (2011) shows the impact of displayed Yelp ratings on restaurant revenues is larger for restaurants with more reviews, Anderson and Magruder (2012) find the opposite: the effect on table reservation availability is smaller for restaurants with more reviews. Our study contributes to the literature by providing new empirical evidence consistent with Luca's findings; we find the effects on Facebook Likes and voucher sales are both larger for restaurants with more reviews.

The rest of this paper is organized as follows. In Section 2, we present a simple stylized Bayesian learning model and derive the hypotheses. We also discuss competing predictions by alternative theories. In Section 3, we describe the research setting and data. In Section 4, we present the identification strategy and estimation specifications. In Sections 5 and 6, we report the empirical results and robustness checks, respectively. Finally, we discuss the implications and conclude the paper in Section 7.

2 Theory

We develop a simple stylized model based on well-established assumptions from the classic Bayesian learning literature (Erdem and Keane 1996, Roberts and Urban 1988) to derive testable hypotheses. The simple stylized model results in a prediction about the moderating role of the variance of ratings which is in contrast to the predictions by some established theories.

2.1 A Simple Stylized Model

When consumers endorse a product to their peers via social media, they communicate not only information but also something about themselves (Berger and Schwartz 2011). Because most people enjoy an enhanced self-image (identity) (Akerlof and Kranton 2000), consumers want their peers to think highly of them and often endeavor to associate themselves with ideal products and brands (Berger and Heath 2007). Based on the theory of word-of-mouth (WOM) as self-enhancement (Wojnicki and Godes 2008), we expect consumers are more likely to endorse good (*vs.* bad) products to their peers. Accordingly, we assume consumers' propensity to endorse a product via social media is dependent on their expected utility of the product (perception of product value), so is their propensity to buy.

Herein, we develop a simple stylized Bayesian learning model to show the structural relationship between a product's review ratings and consumers' posterior expected utility (perception of the product value). Following to the setup of the classic Bayesian learning model by Roberts and Urban (1988), we make the assumptions A1-A4:

- Assumption 1 (A1): A consumer i's prior belief about the value of a product j is X_{ij0}, where $X_{ij0} \sim N(\mu_{ij0}, \sigma_{ij0}^2)$. Herein, σ_{ij0}^2 indicates the information uncertainty in consumer i's prior belief.
- Assumption 2 (A2): Each review rating is an unbiased[4] but imperfect signal of the value of product j, which is normally distributed with mean μ_j and variance σ_j^2. The random disturbance in the signal is normally distributed with zero mean and variance σ_j^2, which reflects "inherent product variability" (Roberts and Urban 1988) and/or "idiosyncratic perceptions" (Erdem and Keane 1996).
- Assumption 3 (A3): Consumers use a Bayesian updating rule to produce their posterior beliefs about the product value.

[4] Selection bias in online reviews (Li and Hitt 2008) may result in the fact that the average rating does not necessarily signal the true product value. We assume each review rating is an unbiased signal, because this study aims to explore how consumers update their posterior beliefs about the product value based upon learning unbiased review ratings.

- Assumption 4 (A4): Consumers are risk averse (and prudent) with a utility function satisfying $u' > 0, u'' < 0, u''' > 0$.

A1-A3 are common assumptions in the Bayesian learning literature (Ching, et al. 2011, Roberts and Urban 1988, Zhao, et al. 2013). Besides the assumption of risk aversion, the Bayesian learning literature often assumes consumers are forward-looking (Ching, et al. 2011). Since this study aims to explore how online reputation would affect consumers' endorsing behaviors, we argue that consumers are less likely forward-looking in this study, because endorsing a product via social media is not trial consumption and would not increase their information sets about the product value.

Note that unlike the conventional Bayesian learning literature (Ching, et al. 2011, Roberts and Urban 1988, Zhao, et al. 2013), A4 does not assume any explicit form for utility function. A4 is a fairly general assumption in that any utility function for either constant or decreasing absolute risk aversion (CARA / DARA)[5] (Friend and Blume 1975) implies A4. In fact, A4 is first introduced by Kimball (1990) in the economics literature as the notion of "prudence" - consumers are risk averse and have a positive precautionary saving motive; consumers' current savings increase with the uncertainty about their future incomes.[6] Subsequently, Eeckhoudt et al. (1995) introduce A4 in a management application.

Suppose there are n_j review ratings about product j. According to A2, we know the mean $\bar{x}_j$ of the n_j review ratings is normally distributed with $E(\bar{x}_j) = \mu_j$, $\sigma^2_{\bar{x}_j} = \frac{\sigma_j^2}{n_j}$, where $\sigma^2_{\bar{x}_j}$ is the variance of the n_j review ratings and indicates the information uncertainty of the review ratings. Based on A1-A3, it can be shown that consumer i's posterior belief about the product value, X_{ij}, is also normally distributed, $X_{ij} \sim N(\mu_{ij}, \sigma^2_{ij})$. As Roberts and Urban (1988) show, the mean and variance (information uncertainty) of consumer i's posterior belief are given by

---

[5] The coefficient of absolute risk aversion is defined as $A(w) = -\frac{u''(w)}{u'(w)}$. Constant absolute risk aversion (CARA) means $A(w)$ is constant and the exponential utility function $u(w) = -e^{-rw}$, where $r > 0$ is unique in exhibiting CARA. Decreasing absolute risk aversion (DARA) means $A'(w) < 0$. Although experimental and empirical evidences are mostly consistent with DARA (Friend and Blume 1975), CARA is often assumed in the Bayesian learning literature (Roberts and Urban 1988) for the sake of mathematical tractability.

[6] As Kimball (1990) explains, "prudence" is meant to suggest one's propensity to prepare and forearm oneself in the face of uncertainty in future income, whereas "risk aversion" simply indicates one dislikes facing uncertainty.

$$\mu_{ij} = \frac{\tau_i \mu_{ij0} + n_j \bar{x}_j}{\tau_i + n_j} \tag{1}$$

$$\sigma_{ij}^2 = (\frac{\tau_i}{\tau_i + n_j})^2 \sigma_{ij0}^2 + (\frac{n_j}{\tau_i + n_j})^2 \sigma_{\bar{x}_j}^2 \tag{2}$$

where τ_i is the relative strength/precision of consumer i's prior belief, $\tau_i = \frac{\sigma_j^2}{\sigma_{ij0}^2}$.

Based on A1-A4, we can prove the following proposition about consumers' posterior expected utility (the proof is given in the appendix).

Proposition 1. Suppose consumers obey the von Neumann–Morgenstern axioms to produce the expected utility for decision-making. Given A1-A4, consumer *i*'s posterior expected utility of product *j* after learning product *j*'s review ratings, $E[U(X_{ij})]$, has the following properties:

(a) $E[U(X_{ij})]$ is increasing and concave w.r.t. the mean $\bar{x}_j$ of product *j*'s review ratings, i.e.,

$$\frac{\partial E[U(X_{ij})]}{\partial \bar{x}_j} > 0, \qquad \frac{\partial^2 E[U(X_{ij})]}{\partial \bar{x}_j{}^2} < 0;$$

(b) $E[U(X_{ij})]$ is decreasing w.r.t. the variance $\sigma_{\bar{x}_j}^2$ of product *j*'s review ratings, i.e.,

$$\frac{\partial E[U(X_{ij})]}{\partial (\sigma_{\bar{x}_j}^2)} < 0;$$

(c) The cross-partial derivative of $E[U(X_{ij})]$ w.r.t. $\bar{x}_j$ and $\sigma_{\bar{x}_j}^2$ is positive, i.e.,

$$\frac{\partial^2 E[U(X_{ij})]}{\partial \bar{x}_j \partial (\sigma_{\bar{x}_j}^2)} > 0.$$

Given that consumers are risk averse (A4), properties (a) and (b) in Proposition 1 are intuitive. Risk-averse consumers would increase their posterior expected utility of a product if the product has a higher average rating. The marginal posterior expected utility induced by the average rating diminishes with a higher average rating. When the variance of ratings is large, risk-averse consumers would decrease their posterior expected utility due to the large information uncertainty of the quality signal provided by the ratings.

By the same token, when the variance is large, it may be expected that risk-averse consumers are less responsive to the average rating because they may reduce their reliance on the review ratings due to the information uncertainty (Basuroy, et al. 2006). Somewhat counterintuitively, property (c) in Proposition 1 shows the opposite: consumers' posterior expected utility is more responsive to the average rating when the variance is larger. Although property (c) seems counterintuitive, it is intuitively understandable. Figure 1 illustrates the intuition. The expected utility is increasing and concave w.r.t. the certainty equivalent which is a function of the average and variance of ratings. When the variance rises, risk-averse consumers reduce their posterior expected utility to a lower level where the marginal expected utility induced by an incremental increase in the average rating (i.e., the first-order derivative of $E[U(X_{ij})]$ w.r.t. $\bar{x}_j$) is greater, because consumers are risk averse and the utility function is concave. It is exactly risk aversion (A4) that results in property (c).

According to property (a) in Proposition 1, a higher average rating increases a consumer's posterior expected utility of the product. However, when the number of reviews n_j goes to zero, Equation (1) shows the weight of the average rating in consumers' posterior beliefs reduces to zero, suggesting that the positive marginal expected utility w.r.t. the average rating may be minimal and empirically undetectable when the number of reviews is too small. After all, if a product only has a few review ratings, consumers may doubt the representativeness of the only few ratings and simply ignore the quality signal of the average rating. On the other hand, when the average rating is calculated based upon a larger sample of reviews, the weight of the average rating, compared to the prior beliefs, increases and the information uncertainty in review ratings reduces. Consequently, the effect of the average rating would increase. By assuming consumers' propensities to endorse and buy a product are dependent on their posterior expected utility, we therefore formulate the following hypotheses.

Hypothesis (H1): *A restaurant's higher online reputation (displayed Yelp rating) increases consumers' social media endorsements and voucher sales, but only when the restaurant has enough reviews.*

Hypothesis (H2): *The effects of a restaurant's online reputation (displayed Yelp rating) on consumers' social media endorsements and voucher sales are greater for restaurants with more reviews.*

According to property (c) in Proposition 1, consumers' posterior expected utility is more responsive to the average rating when the variance of ratings is larger. We hypothesize

Hypothesis (H3A): *The effects of a restaurant's online reputation (displayed Yelp rating) on consumers' social media endorsements and voucher sales are greater for restaurants with a larger variance of ratings*.

2.2 Predictions by Alternative Theories

In regard to the moderating role of the variance of ratings, two established theories in the literature would predict the opposite of H3A. First, the cue diagnosticity theory (Feldman and Lynch 1988) suggests that when the variance of a product's ratings is large, consumers' reliance on the average rating as a specific cue signaling product quality may reduce, because they may find it nondiagnostic and turn to alternative quality signals other than review ratings (Basuroy, et al. 2006). Second, Sun (2012) develops an analytical model, which incorporates consumer preference heterogeneity and mismatch costs, and shows the interaction effect between the average and variance of ratings on product sales is negative. The intuition behind the theory of Sun (2012) lies in that a large variance of ratings can improve consumers' perception of the product quality only if the average rating is low. When the average rating rises, the dominant role played by a large variance would change to signal a high mismatch cost and reduce quality perception. Sun (2012) also provides empirical evidence showing that the average rating of book reviews negatively interacts with the standard deviation of ratings in affecting book sales. Based on the two alternative theories, we hypothesize

Hypothesis (H3B): *The effects of a restaurant's online reputation (displayed Yelp rating) on consumers' social media endorsements and voucher sales are smaller for restaurants with a larger variance of ratings*.

In sum, our simple stylized model based on well-established assumptions from the Bayesian learning literature shows that the effects of online reputation as indicated by the average rating increase with the variance of ratings when consumers are risk averse, while alternative theories (Basuroy, et al. 2006, Feldman and Lynch 1988, Sun 2012) suggest mechanisms for the competing prediction. Therefore, whether the effects of online reputation increase or decrease with the variance of ratings is an empirical question which has important theoretical implications and will be answered in this study.

3 Research Setting and Data

3.1 Setting

We choose the daily-deal businesses as our research setting because of its practical importance and theoretical relevancy. The popularity of using daily deals as a new marketing vehicle has dramatically increased in recent years (Dholakia 2012). As of April 2012, consumers in North America have spent approximately $7 million a day (more than $2.5 billion a year) on daily deals[7] and it is projected to reach $5.5 billion a year by 2016.[8] Many restaurants have been selling deal vouchers through daily-deal sites. While some have attracted thousands of consumers, others have acquired only a few. It is thus important to understand what factors would affect consumers' response to these deals.

Besides the practical importance, the leading daily-deal sites (Groupon, LivingSocial) provide an ideal context for us to identify the causal effects of online reputation on consumers' social media endorsements and product sales.

First, leading daily-deal sites provide the setting where we are able to accurately collect the two outcome variables (i.e., aggregate numbers of consumers' endorsements via Facebook and voucher sales for each deal) so that we could quantify the effects precisely. Figure 2 shows a screenshot of a typical restaurant deal from Groupon. On the deal page, consumers can see the characteristics of the deal, such as restaurant name, discounted voucher price, and the displayed star rating of the restaurant from third-party reputation sites (most likely Yelp.com). Consumers can buy the deal and/or endorse it by clicking the Facebook "Like" button (as circled in Figure 2).

Second, to identify the causal effect of displayed Yelp ratings, it is required to prevent consumers from "interfering" Yelp ratings so that the possible reverse causality is avoided. Based on our inspection, the Yelp star ratings displayed on Groupon deal pages are hard-coded and fixed during the deal promotion. Moreover, since most restaurant deals are only sold for one or two days and the vouchers are often valid for redemption within six months, it is less

---

[7] See http://savvr.com/2012/04/top-10-highest-grossing-daily-deals-of-all-time/, (accessed on March 10, 2013)
[8] See http://streetfightmag.com/2012/09/17/forecast-consumer-daily-deals-spending-to-reach-5-5-billion-by-2016/, (accessed on May 13, 2013)

likely for consumers who buy vouchers to redeem them immediately and post review ratings on Yelp.com when the deal is still on sale. Therefore, displayed Yelp ratings in this setting are largely exogenous.

For data about online reputation, we choose Yelp.com as the data source, because it is perhaps the most well-known and widely-used third-party site providing user-generated reviews about restaurants. Particularly, in most cases where Groupon deals are related to restaurants, their overall Yelp star ratings are prominently displayed on the deal pages (as circled in Figure 2). Thus, consumers are likely to be influenced by the restaurants' Yelp ratings when they look at the restaurant deals, which is supported by the survey conducted by Kimes and Dholakia (2011). Consumers may further go to the restaurants' Yelp profiles through the hyperlinks and check detailed information about the reviews, such as the number and variance of ratings. Correspondingly, we focus on the category of restaurant deals about which Yelp review ratings are most often available.

3.2 Data

We collect the data about restaurant deals from two sources: one is from the dataset provided by Byers et al. (2012) (named as BMZ) and the other is from Yipit.com, an aggregator of daily deals. The BMZ dataset contains a nationwide sample of deals distributed from 19 major cities across the US.[9] The BMZ dataset includes the characteristics (e.g., vendor, discounted voucher price) and accurate voucher sales of each deal. Besides that, the BMZ dataset contains the accurate number of Facebook Likes associated with each deal. In the BMZ dataset, Groupon deals are collected between January 3rd and July 3rd of 2011, and LivingSocial deals are collected between March 21st and July 3rd of 2011. Thus, we turn to Yipit.com and additionally collect LivingSocial deals between January 3rd and March 20th of 2011 (for the same 19 cities). For those LivingSocial deals from Yipit, all relevant deal characteristics are collected but not Facebook Likes. In total, we assemble 3,311 restaurant deals from the 19 US cities between January 3rd and July 3rd of 2011.

[9] The 19 US cities are Atlanta, Boston, Chicago, Dallas, Detroit, Houston, Las Vegas, Los Angeles, Miami, New Orleans, New York, Orlando, Philadelphia, San Diego, San Francisco, San Jose, Seattle, Tallahassee, and Washington DC.

For each restaurant deal, we manually check if the restaurant has a profile on Yelp.com.[10] Since our study aims at identifying the impact of a restaurant's Yelp ratings, we exclude restaurants for which we could not confidently find their Yelp profiles[11] and those with no reviews on their Yelp profiles. For each remaining restaurant's Yelp profile, we use a computer program to automatically extract all individual reviews (including numeric ratings, textual contents, and dates) that are posted before the date of the deal promotion. Ultimately, we assemble a cross-sectional dataset consisting of 2,545 restaurant deals and 129,129 individual review ratings (from 1 to 5 star).

In this study, we have two outcome variables: the number of Facebook Likes (labeled as *Likes*) and number of voucher sales (labeled as *Sales*). *Likes* measure the total number of "Likes" that consumers endorse for a restaurant deal via Facebook. *Sales* measure the total number of vouchers purchased for a restaurant deal.

We collect explanatory and control variables at two aspects. One is about the deal characteristics, including voucher price, discount rate, the number of days that a deal promotion lasts, and a dummy indicating whether it is from Groupon (coded as 1) or LivingSocial. The other is about the restaurant characteristics, including the displayed overall Yelp ratings, number of individual ratings, the mean and variance of individual ratings. We code a proxy variable for a restaurant's business age by calculating the number of days from when the restaurant's earliest Yelp review was posted to the promotion date. Table 1 reports the summary statistics of the key variables in our dataset.

4 Identification

In non-experimental studies, identifying the causal effects of online reputation is a challenging task due to the potential endogeneity problem; online reputation (e.g., the average review rating) is often correlated with unobserved product characteristics that affect consumers' responses. For example, unobserved marketing expenditure is likely correlated with both online

---

[10] We use the restaurant information (e.g., name, street address, zip code, and phone number) to search the restaurants' profiles on Yelp.com.

[11] In order for a restaurant's Yelp profile to be confidently identified, we require (a) the restaurant deal must have a single physical location for redemption, and (b) the restaurant must only have one single Yelp profile.

reputation and product sales. Without reasonably controlling for such unobserved confounders, online reputation may just serve as a predictor of consumers' preference rather than an influencer (Eliashberg and Shugan 1997). To identify the causal effect of displayed Yelp ratings, we need variation in Yelp ratings that is uncorrelated with any deal or restaurant characteristics (e.g., unobserved marketing expenditure). Only the changes in consumers' responses produced by such variation in Yelp ratings could allow us to identify the causal effect.

Fortunately, Yelp's institutional feature of displaying the overall average ratings provides an opportunity for the identification strategy. For a restaurant with multiple review ratings (each ranging from 1 to 5 star), Yelp calculates the average of these ratings and rounds it up or down to the nearest half-star. For example, one restaurant with an average rating of 3.74 is rounded down and displayed as 3.5-star Yelp rating, while the other with an average rating of 3.76 is rounded up and displayed as 4-star. As a result, there is a half-star difference between the displayed overall Yelp ratings of the two restaurants, although their true average ratings are fairly close. The rounded average rating is prominently displayed on the restaurant's Yelp profile (and Groupon's deal page as shown in Figure 2), while the true average rating is not displayed.

For restaurants whose true average ratings fall in a small "window" centered on a threshold (in the above case 3.75), whether one gets rounded up or down is likely to be merely subject to random chance such that they appear to be randomly assigned around the threshold.[12] The only difference between the restaurants on the left and right of a threshold, on average, would be a half-star difference in the displayed Yelp ratings. Therefore, any possible discontinuity in consumers' responses to the restaurant deals (e.g., social media endorsements, voucher sales) could be attributed to the extra half-star displayed Yelp rating. The discontinuity induced by Yelp's displaying rule allows us to identify the causal effect of displayed Yelp ratings by implementing a regression discontinuity (RD) design (Hartmann, et al. 2011, Imbens and Lemieux 2008, Lee and Lemieux 2010).

[12] This is the key identification assumption of regression discontinuity (RD) design, so-called local randomization around the threshold. In this study we conduct a number of robustness checks that assess the validity of the assumption.

Let r_i be the true average Yelp rating of restaurant deal i which may fall in a small (e.g., 0.2-star) bandwidth of a certain threshold c. The value of c can be 1.25, 1.75, 2.25, 2.75, 3.25, 3.75, 4.25 or 4.75. Each of the thresholds with a bandwidth of 0.25 star corresponds to one rating range, such as (3.75±0.25). In total, there are 8 rating ranges between 1 and 5 stars. We pool the data from the 8 rating ranges and use the 0.2-star bandwidth in the main analysis (we also use different bandwidths in robustness checks and get similar results). We use local linear regression (Imbens and Lemieux 2008) as specified in Equation (3) to estimate the causal effects of displayed Yelp rating

$$y_i = \alpha_0 + \beta \times I(r_i \geq c) + \alpha_1 \times (r_i - c) + \alpha_2 \times (r_i - c) \times I(r_i \geq c) + \gamma X_i + \epsilon_i \tag{3}$$

where the outcome variable y_i can be the natural log of deal i's *Likes* or *Sales*. $I(r_i \geq c)$ is the indication function. If $r_i \geq c$, then $I(r_i \geq c) = 1$ and the restaurant's average rating is rounded up to the nearest half-star; otherwise, $I(r_i \geq c) = 0$ and it is rounded down. The displayed Yelp ratings of rounded-up restaurants are, on average, half-star higher than that of the rounded-down restaurants. Because the discontinuity in outcome y_i is likely to be merely induced by the indication function $I(r_i \geq c)$, the coefficient β estimates the causal effect of an extra half-star displayed Yelp rating. In a valid RD design, including control variables is not necessary for estimating the causal effect, because restaurants around a threshold is "locally" randomized (Hartmann, et al. 2011, Imbens and Lemieux 2008, Lee and Lemieux 2010). Still, we include a vector of baseline covariates X_i about deal and restaurant characteristics (determined prior to the deal promotion) as controls to improve precision of the estimation. The full set of baseline controls include city, promotion duration, weekday, log of voucher price, log of number of reviews, whether it is a Groupon or LivingSocial deal, log of restaurant age proxy, and a categorical variable indicating the restaurant's rating range.

To examine the differential impacts of displayed Yelp ratings for restaurants with more or less reviews, we follow the median split method which is commonly used in the current literature (Demers and Lewellen 2003, Efendi et al. 2012, Rishika, et al. 2013). That is, we create a dummy dn_i indicating if the number of restaurant i's reviews is above or equal to the median of the sample.[13] We include the interaction term between dn_i and the indication function $I(r_i \geq c)$. Since the local linear regressions on the left and right of the threshold may have different slopes, we include interaction terms between dn_i and $(r_i - c)$, $(r_i - c) \times I(r_i \geq c)$. Accordingly, we estimate the moderating effect of dn_i using Equation (4) in which the coefficient β_2 identifies the difference between the RD estimates (differential effects) for restaurants with more or less reviews

$$\begin{aligned} y_i = \alpha_0 + \beta_1 \times I(r_i \geq c) + \beta_2 \times I(r_i \geq c) \times dn_i + \beta_3 \times dn_i \\ +\alpha_1 \times (r_i - c) + \alpha_2 \times (r_i - c) \times I(r_i \geq c) \\ +\alpha_3 \times (r_i - c) \times dn_i + \alpha_4 \times (r_i - c) \times I(r_i \geq c) \times dn_i + \gamma X_i + \epsilon_i \end{aligned} \tag{4}$$

Similarly, to examine the differential impacts of displayed Yelp ratings for restaurants with a large or small variance of individual ratings, we create a dummy dv_i indicating if the variance of restaurant i's ratings is above or equal to the median of the sample. We estimate the moderating effect of dv_i using Equation (5) in which β_2 identifies the difference between the RD estimates for restaurants with a large or small variance of ratings. If $\beta_2 > 0$ in Equation (5), then H3A is confirmed; otherwise, if $\beta_2 < 0$ in Equation (5), then H3B is supported.

$$\begin{aligned} y_i = \alpha_0 + \beta_1 \times I(r_i \geq c) + \beta_2 \times I(r_i \geq c) \times dv_i + \beta_3 \times dv_i \\ +\alpha_1 \times (r_i - c) + \alpha_2 \times (r_i - c) \times I(r_i \geq c) \\ +\alpha_3 \times (r_i - c) \times dv_i + \alpha_4 \times (r_i - c) \times I(r_i \geq c) \times dv_i + \gamma X_i + \epsilon_i \end{aligned} \tag{5}$$

5 Results

5.1 Main Effects When Number of Reviews is Sufficiently Large

If a restaurant has only a few reviews on Yelp, consumers may not believe in the displayed overall Yelp rating and may simply ignore it. In order for a restaurant's displayed Yelp rating to have an influential impact, the number of individual reviews needs to be sufficiently large.

[13] Using the dummy variable (the median split method) makes it easy to interpret the moderating effect. Nevertheless, we also use the method of continuous variables: we include the interaction terms with the continuous variables of the number (variance) of ratings to estimate the moderating effects. The results are qualitatively similar.

Thus, we focus on restaurants with at least 20 reviews[14] and estimate the main effects of display Yelp rating on Facebook Likes and voucher sales.

Table 2 presents the OLS estimates with log(*Likes*) as the dependent variable.[15] In Column (1), we only use a set of categorical baseline covariates as controls, including dummies of cities, weekdays, Yelp rating ranges and promotion duration. The significantly positive estimated coefficient of the indication function (the discontinuity) $I(r_i \geq c)$ suggests that crossing a threshold (i.e., an extra half-star displayed Yelp rating) increases consumers' Facebook Likes. In Columns (2) to (5), more baseline covariates are included as controls in the estimation. Including baseline covariates as controls in a valid RD design only helps improve the precision of the estimation but would not reduce bias (if any). As shown in Columns (1)-(5), the point estimates of the discontinuity remain fairly stable and become more precise and significant when additional covariates are included. The stable RD estimates increase our confidence that the RD design is valid; whether a restaurant falls on the left and right of a threshold is "locally" randomized. Table 2 shows consistent evidences that displayed Yelp ratings affects consumers' endorsements via Facebook for restaurants with enough reviews, supporting H1. The magnitude of the estimated effect is also practically significant. Column (5) suggests that for those restaurants with at least 20 reviews, an extra half-star displayed Yelp rating increases the total number of consumers' Facebook Likes by 26.3%.

Table 3 presents the OLS estimates with log(*Sales*) as the dependent variable. In Column (1), we only use a set of categorical covariates as controls and the estimated coefficient of the discontinuity is positive but not significant. When additional covariates are included, the positive coefficient estimates of the discontinuity become more significant. Again, the estimates of the discontinuity in Columns (1)-(5) are fairly stable, enhancing our confidence about the validity of the RD design. Thus, Table 3 shows consistent evidences that displayed Yelp ratings also affects voucher sales, supporting H1. Economically, Column (5) suggests that for those

[14] In our sample, 45.3% (1154) out of the 2545 restaurants have less than 20 Yelp reviews and thus 20 is a substantive cutoff. Qualitative evidence from our interviews confirms that 20 reviews are often enough to make consumers believe the displayed Yelp ratings are meaningful. Other cutoffs (e.g., 15 or 25) provide qualitatively similar results.

[15] There are 77 deals with zero Facebook Likes, accounting for 3% of the sample. To include these deals in the regression, we also use log(*Likes*+0.5) or log(*Likes*+1) as alternative dependent variables and get similar results.

restaurants with at least 20 reviews, an extra half-star displayed Yelp rating increases voucher sales by 17.4%.

For comparison, we also use simple OLS regressions to estimate the effect of displayed Yelp rating on voucher sales. Table A.1 in the Appendix A reports the simple OLS estimates, suggesting that displayed Yelp rating has no effect or even significantly negative effect on voucher sales. The results in Table A.1 are similar to the simple OLS estimates reported by Byers, et al. (2012). The comparison between the estimates in Table 3 and Table A.1 reveals that simple OLS regressions without an appropriate identification strategy may produce misleading results, while results from the RD design are more convincing.

It is worth commenting on the estimates of the key covariates in Tables 2 and 3. The coefficient estimates of number of reviews in both tables are positive and significant, suggesting that restaurants with more reviews are likely to receive more Facebook Likes and voucher sales. This is consistent with prior research (e.g., Liu 2006, Duan et al. 2008) that shows the volume of reviews has a significant predictive power for product sales. More interestingly, while the coefficient estimates of voucher price in Table 3 are all negative (though not significant), the estimates of voucher price in Table 2 are positive and significant. Perhaps a high voucher price is correlated with some unobserved factors (e.g., high quality, specialty) that encourage consumers to endorse the deal via Facebook, but meanwhile it decreases consumers' propensity to buy. The opposite signs of the estimates of voucher price in Tables 2 and 3 reveal that consumers do behave differently in endorsing versus purchasing the deals. The coefficient estimates of restaurant age proxy in both tables are negative and significant, indicating that a younger restaurant is associated with more Facebook Likes and voucher sales. This finding suggests that deal shoppers may favor relatively newer restaurants. Lastly, we find Groupon deals receive more Facebook Likes and voucher sales. This is not surprising, because Groupon as the industry leader has more subscribers than LivingSocial. In general, these findings are consistent with our intuition and enhance our confidence about the credibility of the dataset.

5.2 Moderating Effect of Number of Reviews

To examine the differential impacts of displayed Yelp ratings for restaurants with more or less reviews, we create a dummy dn_i indicating if the number of restaurant i's reviews is above or equal to the sample median. We first estimate the two subsamples separately and then use Equation (4) to estimate if the difference between the RD estimates is significant. The results are reported in Tables 4 and 5.

Column (1) of Table 4 shows the coefficient estimate of the discontinuity is positive and significant, indicating a strong positive effect of displayed Yelp rating on Facebook Likes for restaurants with above-median reviews. By contrast, Column (2) shows the coefficient estimate of the discontinuity is negative but insignificant, indicating that the effect of displayed Yelp rating is minimal for restaurants with below-median reviews. Column (3) confirms that the difference between the RD estimates is positive and significant. Therefore, the results in Table 4 suggest that a higher displayed Yelp rating increases consumers' Facebook Likes only when the restaurants have enough reviews. The effect of displayed Yelp rating on Facebook Likes decreases and even disappears when restaurants have few reviews.

Column (1) of Table 5 shows that coefficient estimate of the discontinuity is positive and significant, indicating a positive effect of displayed Yelp rating on voucher sales for restaurants with above-median reviews. Column (2) shows the coefficient estimate of the discontinuity is negative but insignificant, indicating the effect of displayed Yelp rating is minimal for restaurants with below-median reviews. Column (3) confirms that the difference between the RD estimates is positive and significant. Thus, the results in Table 5 suggest that a higher displayed Yelp rating increases voucher sales only when the restaurants have enough reviews. The results in Tables 4 and 5 support H1 and H2.

5.3 Moderating Effect of Variance of Ratings

To examine the differential impacts of displayed Yelp ratings for restaurants with a large or small variance of ratings, we create a dummy dv_i indicating if the variance of restaurant i's ratings is above or equal to the sample median. We estimate the two subsamples separately and then use Equation (5) to estimate if the difference between the RD estimates is significant. The results are reported in Tables 6 and 7 where unbiased sample variance of ratings is used.

Column (1) of Table 6 shows the coefficient estimate of the discontinuity is positive and significant, indicating a positive effect of displayed Yelp rating on Facebook Likes for restaurants with a large variance of ratings. By contrast, Column (2) shows the coefficient estimate of the discontinuity is negative and insignificant, indicating the effect of displayed Yelp rating is minimal for restaurants with a small variance of ratings. Column (3) confirms that the difference between the RD estimates is positive and significant. Therefore, the results in Table 6 suggest that the effect of displayed Yelp rating on Facebook Likes is greater for restaurants with a larger variance of ratings, supporting H3A.

The coefficient estimate of the discontinuity in Column (1) of Table 7 is positive and smaller than the counterpart estimate in Column (2), indicating that the effect of displayed Yelp rating on voucher sales might be smaller for restaurants with a larger variance of ratings. Yet, neither of the coefficient estimates is precise or significant. Accordingly, Column (3) shows that the difference between the RD estimates is negative but not significant. Despite the insignificance, the negative sign directionally suggests that the effect of displayed Yelp rating on voucher sales might be smaller for restaurants with a larger variance of ratings, consistent with H3B.

To highlight the findings about the moderating effects of number and variance of ratings, we provide RD estimates of the effects of displayed Yelp ratings on Facebook Likes and voucher sales for restaurants with above and below median reviews and variance of ratings in Table 8. The upper panel of Table 8 shows that the effect of displayed Yelp rating on Facebook Likes is largest for restaurants with more reviews and a larger variance of ratings, whereas the effect on voucher sales is largest for restaurants with more reviews but a smaller variance of ratings.

6 Robustness Checks

In this section, we conduct a number of robustness checks to verify if the RD design in our study is valid and the findings are robust.

6.1 Balance Check on Baseline Covariates

If restaurants are truly "locally" randomized around a threshold, we expect all observed baseline covariates of the restaurant deals on the left and right of the threshold would appear to be balanced, just like in a true randomized controlled experiment. In addition to the dummies of cities and weekdays, we collect 17 baseline covariates about deal and restaurant characteristics, such as voucher price, number of reviews, and the true average rating. Table 9 reports the results of balance check and show that all of the 17 covariates of restaurant deals on the left and right of the threshold are balanced. The balance on baseline covariates is consistent with our findings that the RD estimates reported in Tables 2 and 3 are fairly stable when additional covariates are included as controls. The balance check on the observed covariates reinforces the belief that restaurant deals on the left and right of a threshold are comparable, even in terms of unobserved factors. For example, unobserved marketing expenditure is likely correlated with some observed covariates, such as the number of reviews. Table 9 shows that the natural logs of number of reviews for the restaurants above and below the threshold are quite close (3.30 *vs*. 3.37), suggesting that unobserved marketing expenditures between the groups of restaurants are plausibly comparable. Therefore, the balance check enhances our confidence that the RD design in this study is valid and the estimated effects of displayed Yelp ratings can be interpreted as causal.

6.2 Inspection of Possible Review Manipulation

The key identification assumption of a valid RD design is that whether a restaurant's true average rating falls on the left or right of a threshold is "locally" randomized. If some restaurants could precisely manipulate their average ratings (e.g., through posting fake review ratings) to cross a threshold, the identification assumption of the RD design would be invalidated (Hartmann, et al. 2011). Although it is difficult to directly observe restaurants' review manipulation (Mayzlin et al. 2012), prior studies (Anderson and Magruder 2012, Luca 2011) provide both qualitative arguments and empirical evidence that restaurants' incentives to manipulate Yelp ratings is less likely an issue for the RD design used in our study.

Herein, we add two additional arguments. First, Yelp.com has been actively fighting with possible fake reviews by using advanced detection algorithms and punishment policies.[16] Second, in order for the possible review manipulation to invalidate the RD design in our study, the manipulation has to be sufficiently precise such that the true average rating is shifted from the left to the right of a threshold, e.g., from 3.74 to 3.76. Shifting the average rating from 3.74 to 4.24 would not invalidate the RD design in our study, because it is still on the left of a threshold and in the rounded-down group.

---

[16] See http://bits.blogs.nytimes.com/2012/10/18/daily-report-yelp-fights-fake-reviews-with-shaming/ (accessed on June 1, 2013)

In this study we provide additional empirical evidence to further reduce the concern about restaurants' possible manipulation of Yelp review ratings. If some restaurants could precisely manipulate Yelp ratings and shift their average ratings from the left to the right of a threshold, we expect the aggregate distribution of the distance from threshold would be discontinuous at zero and sorting toward the right of the threshold. Figure 3 shows the frequency distribution of the distance from a threshold for all restaurants with distance less than 0.25 (N=2129). The distribution appears symmetric about the threshold; the skewness coefficient is -0.021, far from significantly different from zero (p=0.69). The symmetry of the distribution reduces the concern of restaurants' possible review manipulation. Although there seems to be a peak exactly at the zero point, we provide further evidence (reported in Appendix B) indicating that the peak at zero alone may not necessarily suggest precise review manipulation.

6.3 Placebo Effects on Baseline Covariates

Since baseline covariates are predetermined deal and restaurant characteristics, they would not be affected by displayed Yelp ratings. Thus, we conduct another set of robustness checks to test if any placebo effects of displayed Yelp ratings on baseline covariates could be detected by the RD design. Specifically, we perform the same procedure of RD estimation as we have done with the true outcome variables (Facebook Likes and voucher sales), but instead use a baseline covariate as the dependent variable. If the RD estimates of placebo effect on any baseline covariate are detected as significant, one may raise some concern about the credibility of the estimated effect on the true outcomes. Table 10 reports the RD estimates of placebo effects on four different baseline covariates. None of the RD estimates is significant, suggesting that no placebo effect on the baseline covariates is detected. Therefore, the RD design in our study is valid in the sense that it only allows us to detect the effects on the true outcome variables.

6.4 Different Bandwidths

A narrow bandwidth would make the RD estimates more convincing in terms of "local randomization" and "local linearity", but it often reduces the sample size substantially and leads to an insignificant estimate even if the true effect exists. On the other hand, a wide bandwidth allows more observations in the analysis but may make "local randomization" or "local linear regression" less likely to be valid. Thus, RD estimates may be sensitive to bandwidth selection (Imbens and Lemieux 2008, Lee and Lemieux 2010). To verify if the estimated effect of displayed Yelp ratings is robust, we choose a number of different bandwidths to analyze the data. If the RD estimates are relatively stable with the selection of different bandwidths, the findings would be more credible.

Using different bandwidths from the smallest (0.05-star) to the widest (0.25-star), Table 11 reports the RD estimates of the effect of displayed Yelp rating on Facebook Likes for restaurants with at least 20 reviews. As Columns (2)-(5) show, the RD estimates are all positive and significant when the bandwidth increases from 0.10 to 0.25. Even though the RD estimate in Column (1) is insignificant with the bandwidth of 0.05 (in this case only 253 observations are used in the analysis), the point estimate is still comparable with those in Columns (2)-(5). Considering the small number of observations used in Column (1), the positive effect on Facebook Likes is likely there. Similarly, Table 12 reports the RD estimates of the effect of displayed Yelp rating on voucher sales for restaurants with at least 20 reviews. All the RD estimates in Table 12 are positive and significant. In sum, the results in Table 11 and Table 12 suggest that the findings about the effects of displayed Yelp ratings are robust with the selection of different bandwidths.

6.5 Alternative Measures for Dispersion of Ratings

One of the key findings of this study is that the effect of displayed Yelp ratings on Facebook Likes is greater for restaurants with a larger variance. In Section 5.3, we use the notion of unbiased sample variance for restaurants with more than one review, while it is undefined for restaurants with only one review. Herein, we verify if our findings are robust to a number of alternative measures for dispersion of ratings.

First, since unbiased sample variance is undefined for restaurants with only one review, we use the notion of biased sample variance for those restaurants and define it as zero. In such way, the restaurants with only one review (accounting for 3.2% of the full sample) are included in the analysis. Second, entropy, a concept from information theory (Shannon 2001), is an alternative measure of dispersion and uncertainty in a random variable (Ebrahimi et al. 1999). For a discrete random variable X, each possible value x_i is realized with a probability p_i, then the entropy of X is defined as: $H(X) = -\sum_i p_i \log(p_i)$. Entropy is maximized if p_i is equal across all possible realizations, and it is minimized as zero for a deterministic value. In recent literature, entropy has been used for measuring the dispersion of different opinion groups (Dellarocas et al. 2007) and for mining online product reviews (Zhang and Tran 2008). Third, the Herfindahl–Hirschman index (HHI) in the economic literature is a measure of market share concentration and has been used to capture the consensus in movie critics' reviews by summing up the squares of proportions of pro, con, and mixed opinions (Basuroy, et al. 2006). Since HHI is a measure of opinion consensus, we use the inverse of HHI as an alternative measure of dispersion of review ratings.

Table 13 shows the Pearson correlations between the alternative measures of dispersion of ratings. Since we only additionally define the variance of restaurants with a single review as zero, it is not surprising that the augmented variance is perfectly correlated with unbiased sample variance. On the other hand, entropy and inverse HHI are both positively but not perfectly correlated with unbiased sample variance, suggesting that both are meaningful alternative measures of dispersion. Table 14 reports the estimates of the differential effects of displayed Yelp ratings on Facebook Likes using the three alterative measures and different bandwidths. All the estimates of the interaction term between the discontinuity and large-variance dummy $I(r_i \geq c) \times dv_i$ are positive and significant. The results in Table 14 are consistent with that in Table 6, suggesting that the effect of displayed Yelp rating on Facebook Likes is greater for restaurants with a larger variance of ratings. On the other hand, using these alternative measures of dispersion does not produce any significant differential effects of displayed Yelp ratings on voucher sales, which is also consistent with the results in Table 7. Therefore, we conclude that the dispersion of ratings moderates the effect of displayed Yelp ratings on Facebook Likes, but not on voucher sales.

6.6 Controlling Confounding Factors for Variance of Ratings

The empirical findings that the effect of displayed Yelp ratings on Facebook Likes is greater for restaurants with a larger variance of ratings may result from confounding factors other than variance (or dispersion). For example, voucher price may be associated with the variance of ratings, so perhaps it is voucher price that results in the moderating effects of the variance of ratings, rather than the variance itself. Without controlling for the confounding factors, we cannot conclude that consumers would respond to the variance (dispersion) of ratings.

To reduce this concern, we compare the restaurants with large (dv_i=1) and small (dv_i=0) variances. Table 15 reports the results of the comparison. While restaurants with large and small variances are similar in terms of the number of reviews and restaurant age, the voucher price and true average rating of the two groups are significantly different. Thus, we control these confounding factors by including the interaction terms with them; the results are reported in Table 16. We find that the interaction terms with dv_i (large variance) are all positive and largely significant. Note that in all columns the point estimates of the interaction terms with dv_i (large variance) are fairly stable, suggesting that the estimates are not biased by the confounding factors (voucher price, true average rating). Although the interaction term with dv_i in Column (4) is barely significant (p=0.11), the estimated coefficient of the interaction term with true average rating is close to zero and far from significant, indicating that true average rating is not a significant explanatory factor for the moderating effect of the variance of ratings. Therefore, the results in Table 16 suggest that the moderating effect of the variance on Facebook Likes cannot be explained by the observed confounding factors, such as number of reviews, voucher price, or the true average rating. The results enhance our confidence that it is the large variance (dispersion) of ratings that results in a greater effect of displayed Yelp ratings on Facebook Likes.

7 Conclusion

7.1 Summary of Findings

Little extant research has studied what factors consumers would take into account in the decision-making of endorsing a product to their peers with social ties. We investigate if and how a seller's online reputation affects consumers' social media endorsements and product sales. We derive the testable hypotheses by developing a stylized Bayesian learning model.

Empirically, we examine the situation in which restaurants with review ratings on Yelp sell deal vouchers through Groupon and LivingSocial. We identify the causal impacts of displayed Yelp ratings on consumers' Facebook Likes and voucher sales of restaurant deals. To establish the causal relationships, we implement a RD design and conduct a number of robustness checks to ensure the validity of the RD design. We find a restaurant's higher displayed Yelp rating increases the aggregate number of Facebook Likes and voucher sales, but only for restaurants with enough reviews. The effects of displayed Yelp ratings decrease and even disappear for restaurants with fewer reviews. More interestingly, we find the effect of displayed Yelp ratings on Facebook Likes is greater when the variance of ratings is larger, but the effect on voucher sales does not significantly change with the variance.

7.2 Implications

Our study yields several important implications for theory and practice.

First, social media endorsements as an increasingly important indicator of firms' business performance (Aral, et al. 2013) are distinct from product sales, because the motives and costs of endorsing a product are different from purchasing. Therefore, consumers' decision-making of endorsing via social media deserves to be investigated separately. Our empirical findings suggest that online reputation could affect not only product sales, but also consumers' social media endorsements. Our study is perhaps the first to establish the causal relationship between sellers' online reputation and consumers' social media endorsements for commercial products. Our Bayesian learning model provides a plausible theoretical explanation for the mechanism of the effects of online reputation, that is, through signaling product quality and updating consumers' perception of product value. The results suggest that consumers seem to incorporate their perception of product value into their decision-making of endorsing a product to their peers via social media. The results also show that consumers' social media endorsing behaviors can be predicted well by using a simple Bayesian learning model.

Second, we show that the effects of online reputation are moderated by the number and variance of review ratings. Ignoring the moderating role played by the two contextual factors may lead to misleading results. For example, we find the positive effects of displayed Yelp ratings could only be detected for restaurants with enough reviews, but not for those with few reviews. Our results provide a plausible explanation for the seemingly inconsistent empirical findings about the effects of the valence of online reviews (Chevalier and Mayzlin 2006, Chintagunta, et al. 2010, Duan, et al. 2008, Liu 2006). The findings also offer insights on when and for which restaurants the effects of online reputation would be more salient.

Third and more interestingly, our stylized model based on well-established assumptions from the Bayesian learning literature shows that risk aversion makes consumers' posterior expected utility of a product more responsive to the average rating when the product has a larger variance of ratings, whereas the cue diagnosticity theory (Feldman and Lynch 1988) suggests consumers may reduce their reliance on the average rating and become less responsive to it. Consistent with the prediction of the stylized model, our empirical findings show that the effect of displayed Yelp ratings on consumers' social media endorsements is greater when the variance of ratings is larger, in contrast to the predictions from the alternative theories (Basuroy, et al. 2006, Feldman and Lynch 1988, Sun 2012). The results suggest that perhaps consumers are risk averse in endorsing restaurant deals via Facebook. Yet, we find the effect on voucher sales does not significantly change with the variance. The different moderating roles of the variance of ratings on Facebook Likes and voucher sales reveal that consumers exhibit different behaviors in endorsing versus purchasing products. One possible explanation is that perhaps consumers are relatively less risk averse in purchasing products for their own consumption than they are in endorsing to their peers with social ties and the mechanisms expounded by the competing theories may offset consumers' risk aversion in purchasing.

Fourth, our study reveals that the true causal effect of displayed Yelp ratings is more likely to be detected using a valid RD design, while simple OLS regressions without an appropriate identification strategy may produce misleading results. What's more, beyond the prior studies (Anderson and Magruder 2012, Luca 2011), we provide some new procedures (see Sections 6.1 and 6.2) to inspect if restaurants manipulate the review ratings in a way that may invalidate the RD design. Empirical evidence from our inspection does not support that restaurants in our dataset have precisely manipulated their Yelp ratings in this research setting. The procedures for inspection of possible review manipulation that we use in this study can be applied in other contexts of using the RD design (Anderson and Magruder 2012, Luca 2011).

Last but not least, we show that the average rating, the number and variance of ratings are all important predictors for consumers' response to restaurant deals. Managers (e.g., restaurant owners, daily-deal sites, and even movie studios) may use these simple descriptive statistics of online review ratings in forecasting models for consumers' social media endorsements and product demand.

7.3 Limitations

We recognize that our study has some limitations. First, the modeling assumption (A2) that the random disturbance in the review signal is normally distributed may not reflect the empirical reality of online review ratings; prior empirical findings show that review ratings often follow a binomial distribution (Hu, et al. 2009) and the mean and variance of review ratings are correlated due to the bounding nature of 1 to 5 stars. Although A2 does not perfectly capture the empirical reality of review ratings, we use A2 because of its mathematical tractability and believe it does not compromise the key theoretical implications from the stylized model. The other limitation is that we cannot control unobserved confounding factors for identifying the moderating effects of the variance of ratings, although we find robust results by using a number of alternative measures for dispersion (Section 6.5) and controlling observed confounding factors (Section 6.6).

References

Akerlof, G.A., R.E. Kranton. 2000. Economics and identity. *The Quarterly Journal of Economics* **115**(3) 715-753.

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

Aral, S., C. Dellarocas, D. Godes. 2013. Social Media and Business Transformation: A Framework for Research. *Information Systems Research* **24**(1) 3-13.

Aral, S., D. Walker. 2011. Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Science* **57**(9) 1623-1639.

Basuroy, S., K.K. Desai, D. Talukdar. 2006. An empirical investigation of signaling in the motion picture industry. *Journal of Marketing Research* **43**(2) 287-295.

Berger, J., C. Heath. 2007. Where consumers diverge from others: Identity signaling and product domains. *Journal of Consumer Research* **34**(2) 121-134.

Berger, J., E.M. Schwartz. 2011. What Drives Immediate and Ongoing Word of Mouth? *Journal of Marketing Research* **48**(5) 869-880.

Byers, J.W., M. Mitzenmacher, G. Zervas. 2012. Daily deals: Prediction, social diffusion, and reputational ramifications. *Proceedings of the fifth ACM international conference on Web search and data mining (WSDM'12)* 543-552.

Chen, H., P. De, Y.J. Hu. 2011. IT-Enabled Broadcasting in Social Media: An Empirical Study of Artists' Activities and Music Sales. *Working Paper*.

Chevalier, J.A., D. Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research* **43**(3) 345-354.

Ching, A., T. Erdem, M. Keane. 2011. Learning models: An assessment of progress, challenges and new developments. *Working paper available at SSRN*.

Chintagunta, P.K., S. Gopinath, S. Venkataraman. 2010. The effects of online user reviews on movie box office performance: Accounting for sequential rollout and aggregation across local markets. *Marketing Science* **29**(5) 944-957.

Dellarocas, C. 2006. Strategic manipulation of Internet opinion forums: Implications for consumers and firms. *Management Science* **52**(10) 1577-1593.

Dellarocas, C., X.M. Zhang, N.F. Awad. 2007. Exploring the value of online product reviews in forecasting sales: The case of motion pictures. *Journal of Interactive marketing* **21**(4) 23-45.

Demers, E., K. Lewellen. 2003. The marketing role of IPOs: evidence from internet stocks. *Journal of Financial Economics* **68**(3) 413-437.

Dholakia, U. 2012. How Businesses Fare with Daily Deals as They Gain Experience: A Multi-Time Period Study of Daily Deal Performance. *Working Paper at SSRN*.

Duan, W., B. Gu, A.B. Whinston. 2008. Do online reviews matter? — An empirical investigation of panel data. *Decision Support Systems* **45**(4) 1007-1016.

Ebrahimi, N., E. Maasoumi, E.S. Soofi. 1999. Ordering univariate distributions by entropy and variance. *Journal of Econometrics* **90**(2) 317-336.

Eeckhoudt, L., C. Gollier, H. Schlesinger. 1995. The risk-averse (and prudent) newsboy. *Management Science* **41**(5) 786-794.

Efendi, J., M. Kinney, K. Smith, M. Smith. 2012. Marketing Supply Chain Using B2B Buy-Side E-Commerce Systems: Does Adoption Impact Financial Performance? *Academy of Marketing Studies Journal (forthcoming)*.

Egebark, J., M. Ekström. 2012. Is Liking Contagious? *Unpublished working paper*.

Eliashberg, J., S.M. Shugan. 1997. Film critics: Influencers or predictors? *The Journal of Marketing* **61**(2) 68-78.

Erdem, T., M.P. Keane. 1996. Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science* **15**(1) 1-20.

Feldman, J.M., J.G. Lynch. 1988. Self-generated validity and other effects of measurement on belief, attitude, intention, and behavior. *Journal of Applied Psychology* **73**(3) 421-435.

Friend, I., M.E. Blume. 1975. The demand for risky assets. *The American Economic Review* **65**(5) 900-922.

Hansen, F. 1976. Psychological theories of consumer choice. *Journal of Consumer Research* **3**(3) 117-142.

Hartmann, W., H.S. Nair, S. Narayanan. 2011. Identifying causal marketing mix effects using a regression discontinuity design. *Marketing Science* **30**(6) 1079-1097.

Hu, N., J. Zhang, P.A. Pavlou. 2009. Overcoming the J-shaped distribution of product reviews. *Communications of the ACM* **52**(10) 144-147.

Imbens, G.W., T. Lemieux. 2008. Regression discontinuity designs: A guide to practice. *Journal of Econometrics* **142**(2) 615-635.

Kimball, M.S. 1990. Precautionary Saving in the Small and in the Large. *Econometrica* **58**(1) 53-73.

Kimes, S.E., U.M. Dholakia. 2011. Customer Response to Restaurant Daily Deals. *Available at SSRN 1925932*.

Kosinski, M., D. Stillwell, T. Graepel. 2013. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences* **110**(15) 5802-5805.

Lee, D.S., T. Lemieux. 2010. Regression Discontinuity Designs in Economics. *Journal of Economic Literature* **48**(2) 281-355.

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

Li, X., L. Wu. 2013. Observational Learning and Social-Network Word-of-Mouth: Evidence from Groupon. *Working paper available at SSRN*.

Liu, Y. 2006. Word of Mouth for Movies: Its Dynamics and Impact on Box Office Revenue. *Journal of Marketing* **70**(3) 74-89.

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

Luo, X., J.J. Zhang, W. Duan. 2013. Social Media and Firm Equity Value. *Information Systems Research* **24**(1) 146-163.

Malhotra, A., C. Kubowicz Malhotra, A. See. 2013. How to Create Brand Engagement on Facebook. *MIT Sloan Management Review* **54**(2) 18-20.

Mayzlin, D., Y. Dover, J.A. Chevalier. 2012. Promotional reviews: An empirical investigation of online review manipulation.

Miller, A.R., C. Tucker. 2013. Active Social Media Management: The Case of Health Care. *Information Systems Research* **24**(1) 52-70.

Moe, W.W., D.A. Schweidel. 2012. Online product opinions: Incidence, evaluation, and evolution. *Marketing Science* **31**(3) 372-386.

Muchnik, L., S. Aral, S.J. Taylor. 2013. Social Influence Bias: A Randomized Experiment. *Science* **341**(6146) 647-651.

Rishika, R., A. Kumar, R. Janakiraman, R. Bezawada. 2013. The Effect of Customers' Social Media Participation on Customer Visit Frequency and Profitability: An Empirical Investigation. *Information Systems Research* **24**(1) 108-127.

Roberts, J.H., G.L. Urban. 1988. Modeling multiattribute utility, risk, and belief dynamics for new consumer durable brand choice. *Management Science* **34**(2) 167-185.

Shannon, C.E. 2001. A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and Communications Review* **5**(1) 3-55.

Sun, M. 2012. How Does the Variance of Product Ratings Matter? *Management Science* **58**(4) 696-707.

Wojnicki, A., D. Godes. 2008. Word-of-mouth as self-enhancement. *HBS Marketing Research Paper*.

Zhang, R., T. Tran. 2008. An entropy-based model for discovering the usefulness of online product reviews. *Proceedings of IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)* 759-762.

Zhao, Y., S. Yang, V. Narayan, Y. Zhao. 2013. Modeling consumer learning from online product reviews. *Marketing Science* **32**(1) 153-169.

Zhu, F., X. Zhang. 2010. Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics. *Journal of Marketing* **74**(2) 133-148.

Tables

Table 1. Summary Statistics

	N	Mean	S.D.	Min	Median	Max
Dependent Variables:						
Facebook Likes (*Likes*)	2459	50.35	73.93	0	31	1126
Voucher sales (*Sales*)	2545	910.93	994.77	0	660	26560
Explanatory / Control Variables:						
Voucher price ($)	2545	14.86	30.74	1	12	1500
Original value ($)	2545	30.34	61.67	4	25	3000
Discount rate	2545	50.89	3.03	0	50	83
Is Groupon or LivingSocial deal?	2545	0.84	0.37	0	1	1
Promotion duration (days)	2545	1.72	0.83	0	1	5
Restaurant's displayed Yelp rating	2545	3.62	0.60	1	3.5	5
No. of reviews per restaurant	2545	50.74	80.27	1	24	1186
Restaurant's true average rating	2545	3.61	0.59	1	3.63	5
Variance of a restaurant's ratings	2464	1.19	0.63	0	1.16	8
Proxy of restaurant age (days)	2545	993.64	656.91	1	943	2450

Notes: The notion of unbiased sample variance is used to calculate the variance of ratings, while it is undefined for restaurants with only one review.

Table 2. RD Estimates of Displayed Yelp Effect on Facebook Likes

	(1)	(2)	(3)	(4)	(5)
Discontinuity $I(r_i \geq c)$	0.218*	0.273**	0.270**	0.241**	0.263**
	(0.13)	(0.12)	(0.12)	(0.11)	(0.11)
Distance $(r_i - c)$	-1.32	-1.34	-1.33	-0.721	-0.939
	(0.85)	(0.81)	(0.81)	(0.72)	(0.72)
$(r_i - c) \times I(r_i \geq c)$	0.944	0.680	0.682	0.034	0.264
	(1.10)	(1.06)	(1.06)	(0.95)	(0.95)
log(Number of reviews)		0.470***	0.466***	0.500***	0.543***
		(0.046)	(0.046)	(0.040)	(0.042)
log(Voucher price)			0.038	0.155**	0.160***
			(0.063)	(0.060)	(0.062)
Is a deal from Groupon?				1.35***	1.35***
				(0.067)	(0.067)
log(Restaurant age proxy)					-0.124***
					(0.040)
Promotion duration	Yes	Yes	Yes	Yes	Yes
Dummies of rating ranges	Yes	Yes	Yes	Yes	Yes
Dummies of cities	Yes	Yes	Yes	Yes	Yes
Dummies of weekdays	Yes	Yes	Yes	Yes	Yes
Number of Observations	1017	1017	1017	1017	1017
R^2	0.101	0.195	0.196	0.382	0.389

Notes: Dependent variable is log(*Likes*). OLS estimates of the effect of one extra half-star displayed Yelp rating on the number of Facebook Likes. All regressions use restaurants with at least 20 Yelp reviews and a bandwidth of 0.2 stars. Robust standard errors are reported in parentheses. *p < 0.10, **p < 0.05, ***p < 0.01

Table 3. RD Estimates of Displayed Yelp Effect on Voucher Sales

	(1)	(2)	(3)	(4)	(5)
Discontinuity $I(r_i \geq c)$	0.126	0.157*	0.161*	0.156*	0.174**
	(0.091)	(0.084)	(0.084)	(0.083)	(0.082)
Distance $(r_i - c)$	-1.04*	-1.00*	-0.992*	-0.847	-1.06*
	(0.59)	(0.56)	(0.55)	(0.55)	(0.55)
$(r_i - c) \times I(r_i \geq c)$	0.888	0.679	0.649	0.464	0.708
	(0.81)	(0.77)	(0.77)	(0.76)	(0.75)
log(Number of reviews)		0.379***	0.388***	0.397***	0.448***
		(0.032)	(0.032)	(0.031)	(0.032)
log(Voucher price)			-0.102	-0.074	-0.066
			(0.067)	(0.070)	(0.073)
Is a deal from Groupon?				0.267***	0.276***
				(0.056)	(0.055)
log(Restaurant age proxy)					-0.145***
					(0.034)
Promotion duration	Yes	Yes	Yes	Yes	Yes
Dummies of rating ranges	Yes	Yes	Yes	Yes	Yes
Dummies of cities	Yes	Yes	Yes	Yes	Yes
Dummies of weekdays	Yes	Yes	Yes	Yes	Yes
Number of Observations	1087	1087	1087	1087	1087
R^2	0.137	0.234	0.238	0.250	0.263

Notes: Dependent variable is log(*Sales*). OLS estimates of the effect of one extra half-star displayed Yelp rating on the number of voucher sales. All regressions use restaurants with at least 20 Yelp reviews and a bandwidth of 0.2 stars. Robust standard errors are reported in parentheses. *p < 0.10, **p < 0.05, ***p < 0.01

Table 4. Number of Reviews Moderates Displayed Yelp Effect on Facebook Likes

	(1) # of Reviews >= 24	(2) # of Reviews < 24	(3) Full Sample
Discontinuity $I(r_i \geq c)$	0.260**	-0.072	-0.102
	(0.11)	(0.13)	(0.13)
$I(r_i \geq c) \times dn_i$ (more reviews)			0.315*
			(0.17)
Distance $(r_i - c)$	-1.11	-0.496	0.117
	(0.76)	(0.93)	(0.92)
$(r_i - c) \times I(r_i \geq c)$	0.466	0.975	0.178
	(0.98)	(1.16)	(1.14)
dn_i (more reviews)			-0.363**
			(0.15)
$(r_i - c) \times dn_i$			-1.18
			(1.19)
$(r_i - c) \times I(r_i \geq c) \times dn_i$			0.311
			(1.51)
log(Number of reviews)	0.612***	0.190**	0.456***
	(0.046)	(0.064)	(0.038)
log(Voucher price)	0.152**	0.065	0.125**
	(0.064)	(0.078)	(0.049)
Is a deal from Groupon?	1.39***	1.27***	1.33***
	(0.070)	(0.097)	(0.056)
log(Restaurant age proxy)	-0.132***	-0.076**	-0.102***
	(0.043)	(0.033)	(0.026)
Promotion duration	Yes	Yes	Yes
Dummies of rating ranges	Yes	Yes	Yes
Dummies of cities	Yes	Yes	Yes
Dummies of weekdays	Yes	Yes	Yes
Number of Observations	940	792	1732
R^2	0.401	0.263	0.340

Notes: Dependent variable is log(*Likes*). OLS estimates show the effect of displayed Yelp rating on the number of Facebook Likes is greater for restaurants with more Yelp reviews. The coefficient estimate of $I(r_i \geq c) \times dn_i$ in Column (3) indicates the difference between the RD estimates of the effects of displayed Yelp rating for restaurants with more and less reviews. All regressions use a bandwidth of 0.2 stars. Robust standard errors are reported in parentheses.
*p < 0.10, **p < 0.05, ***p < 0.01

Table 5. Number of Reviews Moderates Displayed Yelp Effect on Voucher Sales

	(1) # of Reviews >= 24	(2) # of Reviews < 24	(3) Full Sample
Discontinuity $I(r_i \geq c)$	0.160*	-0.082	-0.081
	(0.089)	(0.090)	(0.091)
$I(r_i \geq c) \times dn_i$ (more reviews)			0.221*
			(0.128)
Distance $(r_i - c)$	-0.977*	0.091	0.208
	(0.58)	(0.70)	(0.70)
$(r_i - c) \times I(r_i \geq c)$	0.608	-0.364	-0.855
	(0.79)	(0.86)	(0.86)
dn_i (more reviews)			-0.341***
			(0.110)
$(r_i - c) \times dn_i$			0.221*
			(0.13)
$(r_i - c) \times I(r_i \geq c) \times dn_i$			1.43
			(1.18)
log(Number of reviews)	0.482***	0.303***	0.416***
	(0.035)	(0.044)	(0.028)
log(Voucher price)	-0.095	-0.032	-0.060
	(0.076)	(0.060)	(0.051)
Is a deal from Groupon?	0.293***	0.431***	0.336***
	(0.057)	(0.072)	(0.045)
log(Restaurant age proxy)	-0.125***	-0.106***	-0.111***
	(0.036)	(0.024)	(0.020)
Promotion duration	Yes	Yes	Yes
Dummies of rating ranges	Yes	Yes	Yes
Dummies of cities	Yes	Yes	Yes
Dummies of weekdays	Yes	Yes	Yes
Number of Observations	1005	837	1842
R^2	0.263	0.266	0.307

Notes: Dependent variable is log(*Sales*). OLS estimates show the effect of displayed Yelp rating on the number of voucher sales is greater for restaurants with more Yelp reviews. The coefficient estimate of $I(r_i \geq c) \times dn_i$ in Column (3) indicates the difference between the RD estimates of the effects of displayed Yelp rating for restaurants with more and less reviews. All regressions use a bandwidth of 0.2 stars. Robust standard errors are reported in parentheses.

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

Table 6. Variance of Ratings Moderates Displayed Yelp Effect on Facebook Likes

	(1) Large Variance	(2) Small Variance	(3) Full Sample
Discontinuity $I(r_i \geq c)$	0.200*	-0.119	-0.127
	(0.11)	(0.13)	(0.13)
$I(r_i \geq c) \times dv_i$ (large variance)			0.329*
			(0.17)
Distance $(r_i - c)$	-1.46*	0.180	0.183
	(0.78)	(0.92)	(0.90)
$(r_i - c) \times I(r_i \geq c)$	1.55	-0.444	-0.512
	(0.99)	(1.15)	(1.14)
dv_i (large variance)			-0.464***
			(0.14)
$(r_i - c) \times dv_i$			-1.48
			(1.18)
$(r_i - c) \times I(r_i \geq c) \times dv_i$			1.86
			(1.50)
log(Number of reviews)	0.354***	0.403***	0.385***
	(0.037)	(0.036)	(0.026)
log(Voucher price)	0.167***	0.107	0.139***
	(0.063)	(0.075)	(0.048)
Is a deal from Groupon?	1.29***	1.39***	1.34***
	(0.079)	(0.083)	(0.057)
log(Restaurant age proxy)	-0.087**	-0.128***	-0.110***
	(0.036)	(0.040)	(0.026)
Promotion duration	Yes	Yes	Yes
Dummies of rating ranges	Yes	Yes	Yes
Dummies of cities	Yes	Yes	Yes
Dummies of weekdays	Yes	Yes	Yes
Number of Observations	898	834	1732
R^2	0.328	0.387	0.345

Notes: Dependent variable is log(*Likes*). OLS estimates show the effect of displayed Yelp rating on the number of Facebook Likes is greater for restaurants with a large variance of ratings. The coefficient estimate of $I(r_i \geq c) \times dv_i$ in Column (3) indicates the difference between the RD estimates of the effects of displayed Yelp rating for restaurants with large and small variances of ratings. All regressions use a bandwidth of 0.2 stars. Robust standard errors are reported in parentheses. *p < 0.10, **p < 0.05, ***p < 0.01

Table 7. Variance of Ratings Moderates Displayed Yelp Effect on Voucher Sales

	(1) Large Variance	(2) Small Variance	(3) Full Sample
Discontinuity $I(r_i \geq c)$	0.0019	0.026	0.041
	(0.087)	(0.093)	(0.091)
$I(r_i \geq c) \times dv_i$ (large variance)			-0.031
			(0.13)
Distance $(r_i - c)$	-0.762	-0.055	-0.138
	(0.62)	(0.64)	(0.62)
$(r_i - c) \times I(r_i \geq c)$	1.26	-1.14	-1.15
	(0.80)	(0.86)	(0.85)
dv_i (large variance)			-0.250**
			(0.098)
$(r_i - c) \times dv_i$			-0.584
			(0.86)
$(r_i - c) \times I(r_i \geq c) \times dv_i$			2.32**
			(1.16)
log(Number of reviews)	0.323***	0.391***	0.357***
	(0.029)	(0.028)	(0.020)
log(Voucher price)	-0.045	-0.043	-0.048
	(0.077)	(0.056)	(0.050)
Is a deal from Groupon?	0.289***	0.416***	0.348***
	(0.064)	(0.062)	(0.045)
log(Restaurant age proxy)	-0.092	-0.145***	-0.117***
	(0.025)	(0.031)	(0.020)
Promotion duration	Yes	Yes	Yes
Dummies of rating ranges	Yes	Yes	Yes
Dummies of cities	Yes	Yes	Yes
Dummies of weekdays	Yes	Yes	Yes
Number of Observations	954	888	1842
R^2	0.291	0.354	0.310

Notes: Dependent variable is log(*Likes*). OLS estimates show the effect of displayed Yelp rating on the number of voucher sales is smaller for restaurants with a large variance of ratings. The coefficient estimate of $I(r_i \geq c) \times dv_i$ in Column (3) indicates the difference between the RD estimates of the effects of displayed Yelp rating for restaurants with large and small variances of ratings. All regressions use a bandwidth of 0.2 stars. Robust standard errors are reported in parentheses. *p < 0.10, **p < 0.05, ***p < 0.01

Table 8. RD Estimates of Effects of Displayed Yelp Ratings for Different Subsamples

	Above Median Variance	Below Median Variance
Dependent Variables: log(Likes)		
Above Median Reviews	0.273	0.175
	(0.15)	(0.17)
Below Median Reviews	0.041	-0.155
	(0.18)	(0.20)
Dependent Variables: log(Sales)		
Above Median Reviews	0.015	0.179
	(0.12)	(0.13)
Below Median Reviews	-0.059	-0.135
	(0.13)	(0.13)

Notes: RD estimates of the effects of displayed Yelp ratings on Facebook Likes and voucher sales for different subsamples. The upper panel is produced using the model in Column (5) of Table 2. The bottom panel is produced using the model in Column (5) of Table 3. Robust standard errors are reported in parentheses.

Table 9. Balance Check on Baseline Covariates of Deal and Restaurant Characteristics

	Mean		Diff. in	*t*-statistic
	Above Threshold	Below Threshold	Means	
Deal Characteristics				
log(Voucher price)	2.55	2.51	0.037	1.50
Value saved ($)	16.68	14.55	2.14	1.28
Full value ($)	32.71	28.51	4.20	1.26
Discount rate (%)	50.90	50.78	0.12	0.83
Is a deal from Groupon?	0.83	0.85	-0.026	-1.52
Promotion duration (days)	1.73	1.71	0.026	0.67
Restaurant Characteristics				
log(Number of reviews)	3.30	3.37	-0.073	-1.35
log(Restaurant age proxy)	6.65	6.63	0.023	0.53
True average rating	3.61	3.60	0.007	0.29
Variance of ratings	1.19	1.22	-0.027	-1.06
Percent of 5-star ratings	0.246	0.250	-0.003	-0.42
Percent of 4-star ratings	0.371	0.366	0.005	0.73
Percent of 3-star ratings	0.198	0.196	0.002	0.41
Percent of 2-star ratings	0.113	0.111	0.001	0.30
Percent of 1-star ratings	0.072	0.077	-0.005	-1.26
Number of reviews in the past month	2.67	2.71	-0.042	-0.24
Average number of reviews in the past three months	2.57	2.57	0	0

Notes: Balance check compares the baseline covariates of restaurant deals on the left and right of the threshold within a bandwidth of 0.2 star. Dummies of cities and weekdays are also checked but not reported in the table. The balance check confirms that restaurant deals that are above and below a threshold are comparable. *p < 0.10, **p < 0.05, ***p < 0.01

Table 10. RD Estimates of Placebo Effect on Baseline Covariates

	(1) log(Number of reviews)	(2) log(Voucher price)	(3) Is a deal from Groupon?	(4) log(Restaurant age proxy)
Discontinuity $I(r_i \geq c)$	-0.116	0.037	0.017	0.124
	(0.082)	(0.059)	(0.045)	(0.086)
Distance $(r_i - c)$	0.328	0.011	-0.516	-1.47***
	(0.51)	(0.38)	(0.28)	(0.53)
$(r_i - c) \times I(r_i \geq c)$	0.056	-0.143	0.661	1.66**
	(0.69)	(0.51)	(0.38)	(0.74)
log(Number of reviews)		Yes	Yes	Yes
log(Voucher price)	Yes		Yes	Yes
Is a deal from Groupon?	Yes	Yes		Yes
log(Restaurant age proxy)	Yes	Yes	Yes	
Promotion duration	Yes	Yes	Yes	Yes
Dummies of rating ranges	Yes	Yes	Yes	Yes
Dummies of cities	Yes	Yes	Yes	Yes
Dummies of weekdays	Yes	Yes	Yes	Yes
Number of Observations	1090	1090	1090	1090
R^2	0.280	0.139	0.147	0.183

Notes: OLS estimates of the placebo effect of displayed Yelp rating on baseline covariates. None of the placebo tests on covariates is significant. All regressions use restaurants with at least 20 Yelp reviews and a bandwidth of 0.2 stars. Robust standard errors are reported in parentheses. *p < 0.10, **p < 0.05, ***p < 0.01

Table 11. Displayed Yelp Effect on Facebook Likes with Different Bandwidths

	(1) BW=0.05	(2) BW=0.10	(3) BW=0.15	(4) BW=0.20	(5) BW=0.25
Discontinuity $I(r_i \geq c)$	0.228	0.374**	0.313**	0.263**	0.177*
	(0.21)	(0.15)	(0.12)	(0.11)	(0.098)
Distance $(r_i - c)$	-2.93	-2.74	-1.28	-0.939	-0.288
	(6.04)	(2.04)	(1.07)	(0.72)	(0.57)
$(r_i - c) \times I(r_i \geq c)$	6.52	2.28	0.375	0.264	-0.080
	(6.96)	(2.67)	(1.44)	(0.95)	(0.73)
log(Number of reviews)	0.465***	0.587***	0.577***	0.543***	0.512***
	(0.069)	(0.056)	(0.050)	(0.042)	(0.038)
log(Voucher price)	0.016	0.078	0.157**	0.160***	0.134**
	(0.11)	(0.084)	(0.070)	(0.062)	(0.056)
Is a deal from Groupon?	1.31***	1.36***	1.38***	1.35***	1.35***
	(0.12)	(0.091)	(0.074)	(0.067)	(0.063)
log(Restaurant age proxy)	-0.203***	-0.219***	-0.182***	-0.124***	-0.131***
	(0.071)	(0.047)	(0.042)	(0.040)	(0.037)
Promotion duration	Yes	Yes	Yes	Yes	Yes
Dummies of rating ranges	Yes	Yes	Yes	Yes	Yes
Dummies of cities	Yes	Yes	Yes	Yes	Yes
Dummies of weekdays	Yes	Yes	Yes	Yes	Yes
Number of Observations	253	516	764	1017	1250
R^2	0.521	0.455	0.427	0.389	0.365

Notes: Dependent variable is log(*Likes*). OLS estimates of the effect of one extra half-star displayed Yelp rating on the number of Facebook Likes using different bandwidths. All regressions use restaurants with at least 20 Yelp reviews. Robust standard errors are reported in parentheses. *p < 0.10, **p < 0.05, ***p < 0.01

Table 12. Displayed Yelp Effect on Voucher Sales with Different Bandwidths

	(1) BW=0.05	(2) BW=0.10	(3) BW=0.15	(4) BW=0.20	(5) BW=0.25
Discontinuity $I(r_i \geq c)$	0.313**	0.218**	0.157*	0.174**	0.152**
	(0.15)	(0.11)	(0.092)	(0.082)	(0.073)
Distance $(r_i - c)$	-3.51	-0.891	-0.694	-1.06*	-0.976**
	(4.18)	(1.54)	(0.87)	(0.55)	(0.38)
$(r_i - c) \times I(r_i \geq c)$	0.577	-0.618	0.454	0.708	0.709
	(5.08)	(2.13)	(1.16)	(0.75)	(0.52)
log(Number of reviews)	0.482***	0.477***	0.481***	0.448***	0.430***
	(0.053)	(0.044)	(0.036)	(0.032)	(0.028)
log(Voucher price)	-0.154*	-0.128	-0.037	-0.066	-0.067
	(0.084)	(0.098)	(0.084)	(0.073)	(0.067)
Is a deal from Groupon?	0.387***	0.282***	0.298***	0.276***	0.255***
	(0.094)	(0.077)	(0.062)	(0.055)	(0.049)
log(Restaurant age proxy)	-0.128**	-0.173***	-0.174***	-0.145***	-0.152***
	(0.063)	(0.044)	(0.040)	(0.034)	(0.030)
Promotion duration	Yes	Yes	Yes	Yes	Yes
Dummies of rating ranges	Yes	Yes	Yes	Yes	Yes
Dummies of cities	Yes	Yes	Yes	Yes	Yes
Dummies of weekdays	Yes	Yes	Yes	Yes	Yes
Number of Observations	279	560	824	1087	1336
R^2	0.388	0.315	0.288	0.263	0.259

Notes: Dependent variable is log(*Sales*). OLS estimates of the effect of one extra half-star displayed Yelp rating on the number of voucher sales using different bandwidths. All regressions use restaurants with at least 20 Yelp reviews. Robust standard errors are reported in parentheses. *p < 0.10, **p < 0.05, ***p < 0.01

Table 13. Pearson Correlation between Alternative Measures of Dispersion of Ratings

Variable	Mean	S.D.	(1)	(2)	(3)
(1) Unbiased sample variance	1.19	0.63	1.00		
(2) Variance including restaurants with only one review	1.16	0.65	1.00	1.00	
(3) Entropy	1.07	0.34	0.435	0.522	1.00
(4) Inverse HHI	2.82	0.75	0.451	0.525	0.947

Table 14. Displayed Yelp Effect on Facebook Likes Increases When Variance of Ratings is Larger Using Alternative Measurements of Dispersion and Different Bandwidths

	Including restaurants w/ only one review		Entropy		Inverse HHI	
	(1) BW=0.20	(2) BW=0.15	(3) BW=0.20	(4) BW=0.15	(5) BW=0.20	(6) BW=0.15
Discontinuity $I(r_i \geq c)$	-0.184	-0.110	-0.140	-0.135	-0.101	-0.106
	(0.13)	(0.15)	(0.13)	(0.15)	(0.13)	(0.15)
$I(r_i \geq c) \times dv_i$ (large variance)	0.412**	0.352*	0.374**	0.446**	0.306*	0.403**
	(0.17)	(0.19)	(0.17)	(0.20)	(0.17)	(0.20)
Distance $(r_i - c)$	0.600	-0.506	0.834	0.449	0.725	0.266
	(0.91)	(1.38)	(0.90)	(1.51)	(0.90)	(1.50)
$(r_i - c) \times I(r_i \geq c)$	-1.07	0.247	-1.02	-0.155	-0.999	0.301
	(1.14)	(1.84)	(1.13)	(1.96)	(1.14)	(1.95)
dv_i (large variance)	-0.553***	-0.465***	-0.319**	-0.362**	-0.302**	-0.331**
	(0.14)	(0.16)	(0.14)	(0.17)	(0.14)	(0.17)
$(r_i - c) \times dv_i$	-2.14*	-1.13	-2.56**	-2.91	-2.40**	-2.71
	(1.17)	(1.87)	(1.18)	(1.92)	(1.18)	(1.92)
$(r_i - c) \times I(r_i \geq c) \times dv_i$	2.80*	1.57	2.74*	2.33	2.76*	1.71
	(1.49)	(2.42)	(1.51)	(2.50)	(1.51)	(2.49)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Number of Observations	1732	1289	1732	1289	1732	1289
R^2	0.347	0.359	0.339	0.351	0.339	0.351

Notes: Dependent variable is log(*Likes*). OLS estimates show the effect of displayed Yelp rating on the number of Facebook Likes is greater for restaurants with a large variance/dispersion of ratings. The coefficient estimate of $I(r_i \geq c) \times dv_i$ indicates the difference between the RD estimates of the effects of displayed Yelp rating for restaurants with a large and small variance of ratings. Robust standard errors are reported in parentheses. *p < 0.10, **p < 0.05, ***p < 0.01

Table 15. Comparison between Restaurants with Large and Small Variances

	Mean		Diff. in Means	t-statistic
	Large Variance	Small Variance		
log(Number of reviews)	3.38	3.29	0.086	1.58
log(Voucher price)	2.59	2.46	0.13***	5.44
log(Restaurant age proxy)	6.67	6.61	0.058	1.33
True average rating	3.38	3.84	-0.46***	-20.8
Is a deal from Groupon?	0.85	0.83	0.012	0.72
Discount rate	50.82	50.87	-0.050	-0.34

Notes: The comparison indicates that restaurants with large and small variances are different in terms of voucher price and true average rating. *p < 0.10, **p < 0.05, ***p < 0.01

Table 16. Controlling Confounding Factors for Variance of Ratings

	(1)	(2)	(3)	(4)
Discontinuity $I(r_i \geq c)$	-0.127	-0.386	-0.859*	-0.877
	(0.13)	(0.29)	(0.48)	(0.92)
$I(r_i \geq c) \times dv_i$ (large variance)	0.329*	0.325*	0.306*	0.301
	(0.17)	(0.17)	(0.17)	(0.19)
$I(r_i \geq c) \times$ log(Number of reviews)		0.078	0.060	0.060
		(0.076)	(0.078)	(0.078)
$I(r_i \geq c) \times$ log(Voucher price)			0.217	0.217
			(0.18)	(0.18)
$I(r_i \geq c) \times$ True average rating				0.004
				(0.19)
Other controls	Yes	Yes	Yes	Yes
Number of Observations	1732	1732	1732	1732
R^2	0.345	0.346	0.347	0.348

Notes: Dependent variable is log(*Likes*). When additional confounding factors are controlled, the interaction terms with dv_i (large variance) are all positive and largely significant. All regressions use a bandwidth of 0.2 stars. Robust standard errors are reported in parentheses. *p < 0.10, **p < 0.05, ***p < 0.01

Figures

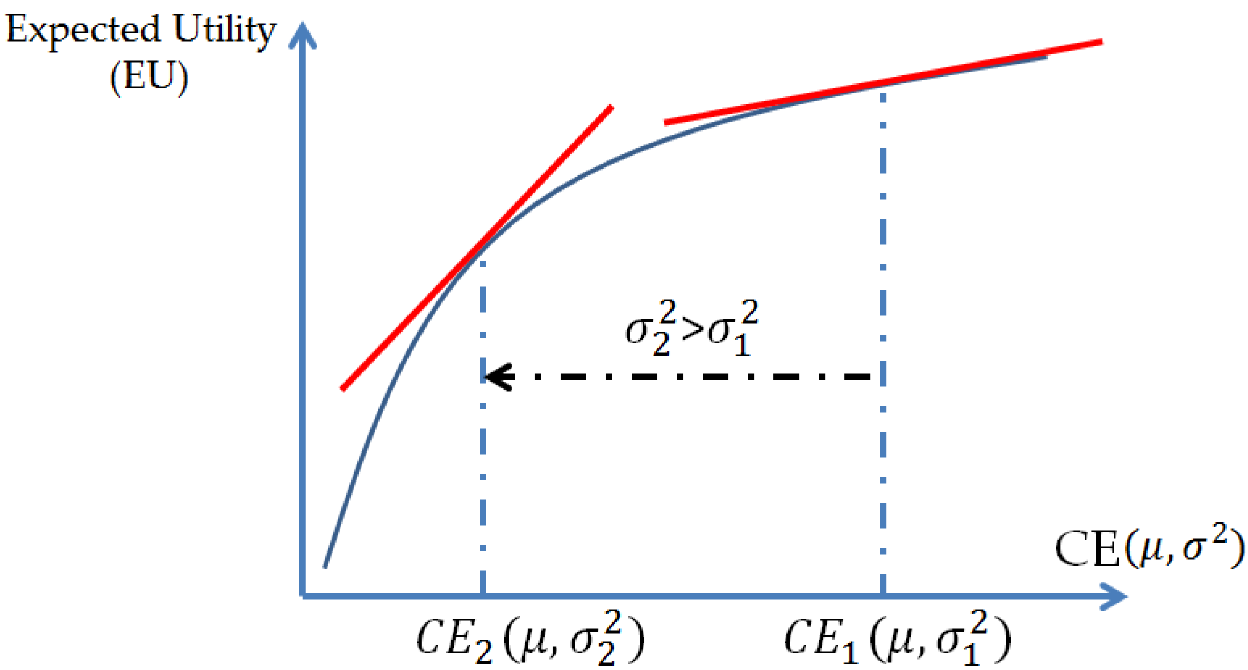

Figure 1. Illustration for the intuition about property (c) in Proposition 1

Notes: The expected utility is increasing and concave w.r.t. the certainty equivalent (CE) which is a function of the average and variance of ratings. A large variance makes the slope of expected utility w.r.t. the average rating steeper so that the expected utility is more responsive to the average rating. Consequently, the cross-partial derivative of the expected utility w.r.t. the average and variance of ratings is positive.

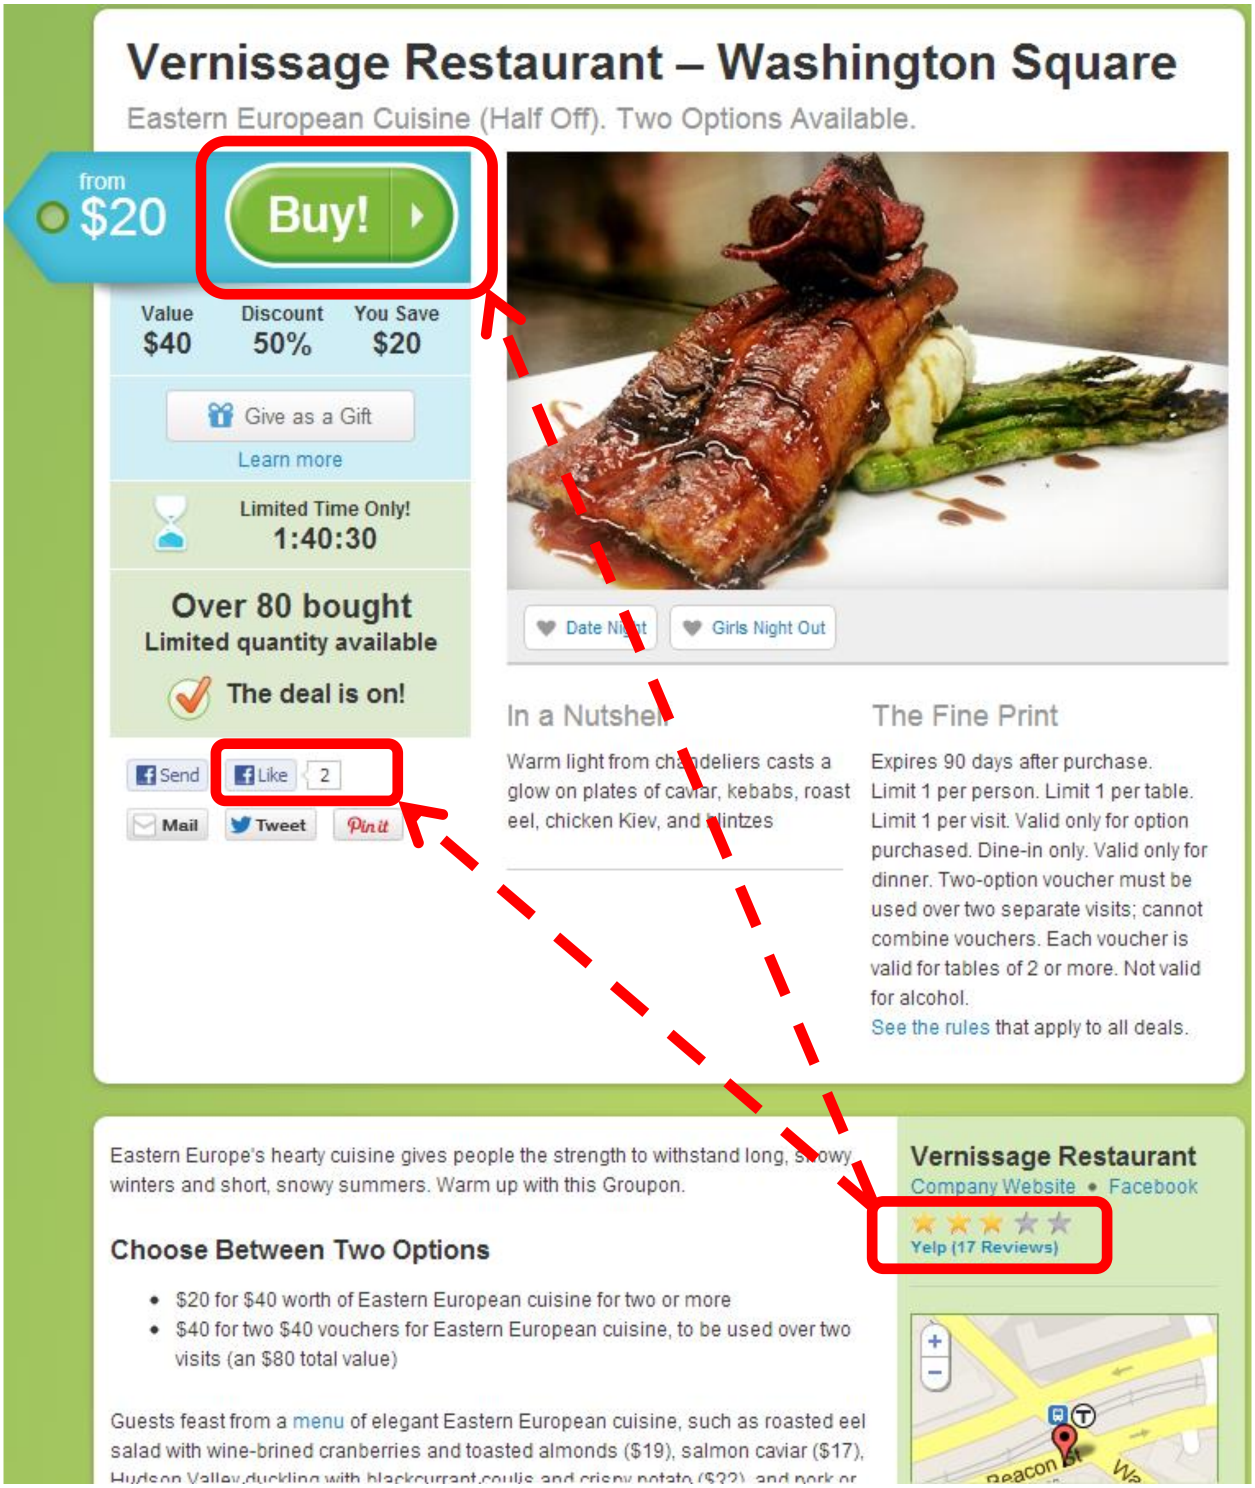

Figure 2. Screenshot of a typical restaurant deal from Groupon.com

Notes: Restaurants' overall Yelp ratings (if any) are often prominently displayed as well as the number of reviews, which could potentially influence consumers to endorse and/or buy the deal.

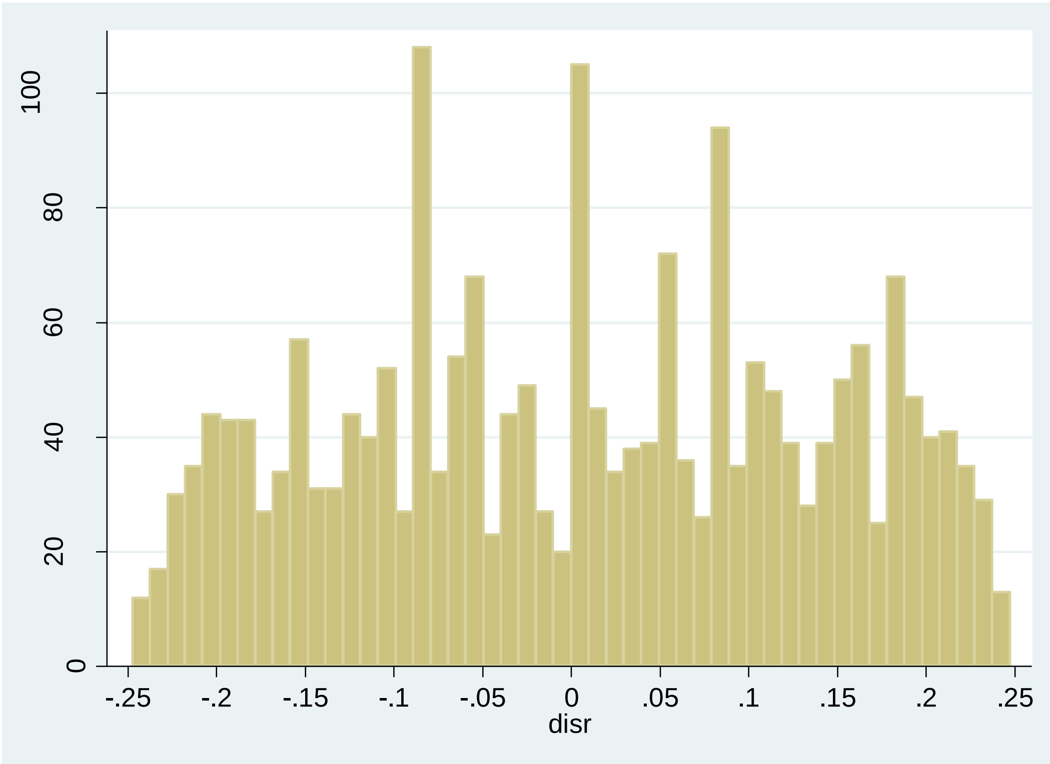

Figure 3. Histogram of the frequency distribution of distance from threshold.

Notes: The histogram plots the frequency distribution of the distance from a threshold for all restaurants with distance<0.25 (N=2129). Note that the distribution appears symmetric about the threshold (skewness coefficient=-0.021, p=0.69), but there is a peak exactly at the zero point.

Appendix A

Table A.1. Simple OLS Estimates of Displayed Yelp Effect on Voucher Sales

	(1) Full Sample	(2) # of Reviews >=20	(3) # of Reviews <20
Displayed average Yelp rating	-0.104	0.012	-0.308***
	(0.064)	(0.083)	(0.10)
Log(Number of reviews)	0.360***	0.429***	0.266***
	(0.018)	(0.028)	(0.051)
Log(Voucher price)	-0.065	-0.067	-0.062
	(0.047)	(0.066)	(0.063)
Is a deal from Groupon?	0.324***	0.260***	0.477***
	(0.042)	(0.050)	(0.077)
Log(Restaurant age proxy)	-0.110***	-0.147***	-0.089***
	(0.019)	(0.030)	(0.025)
Promotion duration	Yes	Yes	Yes
Dummies of rating ranges	Yes	Yes	Yes
Dummies of cities	Yes	Yes	Yes
Dummies of weekdays	Yes	Yes	Yes
Number of Observations	2126	1336	790
R^2	0.299	0.256	0.254

Notes: Dependent variable is log(*Sales*). Simple OLS estimates of the effect of displayed average Yelp rating on number of voucher sales. Robust standard errors are reported in parentheses.
*p < 0.10, **p < 0.05, ***p < 0.01

Appendix B

Herein, we provide further evidence indicating that the peak at zero alone may not necessarily suggest precise review manipulation. This is because when the number of reviews is 4**k* (*k* is an integer), the average ratings may "naturally" fall exactly at a threshold without any manipulation. For example, for restaurants with four reviews, (5,5,4,1) and (5,5,3,2) both result in an average rating of 3.75. In fact, the combinatorial math suggests that for the case of four review ratings, there are 625 (=5^4) different combinations in total, out of which 312 fall at a threshold. That is, even if restaurants are truly randomized and assigned to each of the 625 rating combinations with an equal probability, the aggregate probability that the average rating falls at a threshold is 49.92% (=312/625).[17] By using this simple case, we show that when the number of reviews is 4*k (especially for a small k), the average ratings are very likely to fall at a threshold even without any manipulation.

Figure B.1 shows the frequency distribution of the distance from a threshold for restaurants with number of reviews not as 4*k, called non-4k-type restaurants (N=1631). Note that the peak at the zero point disappears in Figure B.1. Also, the distribution appears symmetric about the threshold (skewness coefficient is -0.017, p=0.78), suggesting that for non-4k-type restaurants (accounting for 76.6%), there seems to be no sorting toward the right of the threshold.

Figure B.2 shows the frequency distribution of the distance from a threshold for restaurants with number of reviews as 4*k, called 4k-type restaurants (N=498). They account for 23.4% of the total restaurants with distance less than 0.25. Consistent with the results from the simple case of four reviews, there is a striking peak at the zero point in Figure B.2. On the other hand, the distribution is also symmetric about the threshold (skewness coefficient is -0.043, p=0.69). Interestingly, given the fact that 4k-type restaurants are likely to fall at a threshold "by nature", restaurants in our dataset are not more likely to be 4k-type, compared to the case in a true randomization (23.4% *vs*. 25%). That is, there seems no evidence that restaurants manipulate the number of reviews to be 4*k so that they could have a higher chance of being rounded up.

In sum, by inspecting the aggregate distribution of the distance from threshold, we find that neither non-4k-type restaurants sort toward the right of a threshold nor restaurants manipulate the number of reviews to be 4*k. The findings suggest that restaurants' incentive to manipulate Yelp ratings is not a concern and enhance our confidence about the validity of the RD design used in our study.

[17] Intuitively, the average of four ratings can only end with 0, 0.25, 0.5, and 0.75. Two of the four endings fall at a threshold, suggesting that the probability is nearly 50%.

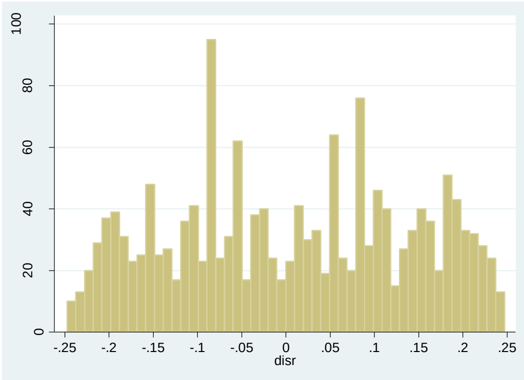

Figure B.1. Histogram of the frequency distribution of distance from threshold, for restaurants with number of reviews not as 4**k* (non-4k-type restaurants).

Notes: The histogram plots the frequency distribution of the distance from a threshold for restaurants with distance<0.25 but the number of reviews not as 4**k* (*N*=1631). Note that there is no peak at the zero point and the distribution appears symmetric about the threshold (skewness coefficient=-0.017, *p*=0.78).

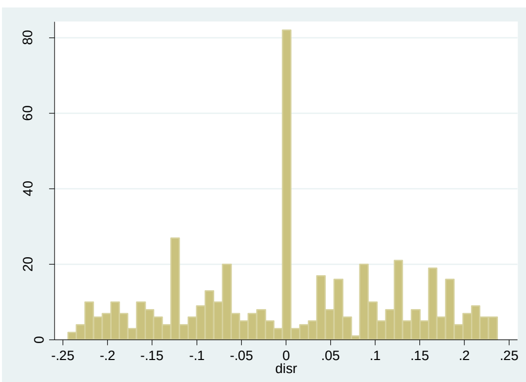

Figure B.2. Histogram of distribution of distance from threshold, only for restaurants with number of reviews as 4**k* (4k-type restaurants).

Notes: The histogram plots the frequency distribution of the distance from a threshold for restaurants with number of reviews as 4**k* (*N*=498). They account for 23.4% of the total restaurants with distance<0.25. Note that there is a striking peak at the zero point and the distribution appears symmetric about the threshold (skewness coefficient=-0.043, *p*=0.69).