Promotional Reviews: An Empirical Investigation of Online Review Manipulation

Dina Mayzlin, Yaniv Dover, and Judith Chevalier[1]
Yale University

June 2012

[1]The authors contributed equally, and their names are listed in reverse alphabetical order. We thank the seminar participants at the Yale Center for Customer Insights Conference, the Marketing Science Conference in Houston, and the 9th ZEW Conference on the Economics of Information and Communication Technologies, University of Pennsylvania marketing seminar, NYU Information Systems seminar, Dartmouth marketing seminar, Washington University St. Louis marketing seminar, University of Houston marketing seminar, and Stanford marketing seminar. Correspondence: 135 Prospect St., P.O. Box 208200, New Haven, CT 06520, e-mail: dina.mayzlin@yale.edu, yaniv.dover@yale.edu, judith.chevalier@yale.edu.

Abstract

Online reviews have been shown to impact consumer behavior. However, the authenticity of online user reviews remains a concern because, on many sites, firms can manufacture positive reviews for their own products and negative reviews for their rivals. In this paper, we marry the diverse literature on economic subterfuge with the literature on organizational form. We undertake an empirical analysis of promotional reviews, examining both the extent to which fakery occurs and the market conditions that encourage or discourage promotional reviewing activity. Specifically, we examine hotel reviews, exploiting the organizational differences between two travel websites: Expedia.com, and Tripadvisor.com. While anyone can post a review on Tripadvisor, a consumer could only post a review of a hotel on Expedia.com if the consumer actually booked at least one night at the hotel through the website. We examine differences in the distribution of reviews for a given hotel between Tripadvisor and Expedia. We show in a simple model that the net gains from promotional reviewing are likely to be highest for independent hotels that are owned by single-unit owners and lowest for branded chain hotels that are owned by multi-unit owners. Our methodology thus isolates hotels with a disproportionate incentive to engage in promotional reviewing activity. We show that hotels with a high incentive to fake have a greater share of five star (positive) reviews on Tripadvisor relative to Expedia. Furthermore, we show that the hotel neighbors of hotels with a high incentive to fake have more one star (negative) reviews on Tripadvisor relative to Expedia.

PRELIMINARY AND INCOMPLETE- DO NOT CITE WITHOUT PERMISSION.

1 Introduction

User-generated online reviews have become an important resource for consumers making purchase decisions; an extensive and growing literature documents the influence of online user reviews on the quantity and price of transactions.[1] In theory, online reviews should create producer and consumer surplus by improving the quality of the match between consumers and products. However, one important impediment to the improvement in match quality is the possible existence of fake or "promotional" online reviews. Specifically, reviewers with a material interest in the consumer's purchase decision may post reviews that are designed to influence consumers and to resemble the reviews of disinterested consumers. While there is a substantial economic literature on misrepresentation (reviewed below), the specific context of advertising disguised as user reviews has not been extensively studied.

The presence of undetectable (or difficult to detect) fake reviews may have at least two deleterious effects on consumer and producer surplus. First, consumers who are fooled by the promotional reviews may make suboptimal choices. Second, the presence or potential presence of biased reviews may lead consumers to mistrust reviews. This in turn forces consumers to disregard or underweight helpful information posted by disinterested reviewers. For these reasons, the Federal Trade Commission in the United States recently updated its guidelines governing endorsements and testimonials to also include online reviews. According to the guidelines a user must disclose the existence of a material connection between himself and the manufacturer.[2] To the best of our knowledge, there has not been wide-scale enforcement of these laws in the United States, although the FTC did investigate (but did not fine) Ann Taylor LOFT for breaking the law in giving bloggers gift cards for coverage of its fashion show (see Zmuda (2010)). Relatedly, in February 2012, the UK Advertising Standards Authority ruled that TripAdvisor must not claim that it offers "honest, real, or trusted" reviews from "real travelers". The Advertising Standards Authority, in its decision, held that TripAdvisor's claims implied that "consumers could be assured that all review content on the

[1] Much of the earliest work focused on the effect of Ebay reputation feedback scores on prices and quantity sold; for example, Resnick and Zeckhauser (2002), Melnik and Alm (2002)), and Resnick et al. (2006). Later work examined the role of consumer reviews on product purchases online; for example, Chevalier and Mayzlin (2006), Anderson and Magruder (2012)), Berger et al. (2010), Chintagunta et al. (2010).

[2] The guidelines provide the following example, "An online message board designated for discussions of new music download technology is frequented by MP3 player enthusiasts...Unbeknownst to the message board community, an employee of a leading playback device manufacturer has been posting messages on the discussion board promoting the manufacturer's product. Knowledge of this poster's employment likely would affect the weight or credibility of her endorsement. Therefore, the poster should clearly and conspicuously disclose her relationship to the manufacturer to members and readers of the message board" (http://www.ftc.gov/os/2009/10/091005endorsementguidesfnnotice.pdf)

TripAdvisor site was genuine, and when we understood that might not be the case, we concluded that the claims were misleading." (www.asa.org/ASA-action/Adjudications).

In order to examine the potential importance of these issues, we undertake an empirical analysis of the extent to which promotional reviewing activity occurs and the firm characteristics and market conditions that result in an increase or decrease in promotional reviewing activity. The first challenge to any such exercise is that detecting promotional reviews is difficult. After all, promotional reviews are designed to mimic unbiased reviews. For example, inferring that a review is fake because it conveys an extreme opinion is flawed; presumably, individuals who had an extremely positive or negative experience with a product may be particularly inclined to post reviews. In this paper, we empirically exploit a key difference in website business models. In particular, some websites accept reviews from anyone who chooses to post a review while other websites only allow reviews to be posted by consumers who have actually purchased a product through the website (or treat "unverified" reviews differently from those posted by verified buyers). If posting a review requires making an actual purchase, the cost of posting disingenuous reviews is greatly increased. We examine differences in the distribution of reviews for a given product between websites where faking is difficult and websites where faking is easy.

Specifically, in this paper, we examine hotel reviews, exploiting the organizational differences between Expedia.com, and Tripadvisor.com. Tripadvisor is a popular website that collects and publishes consumer reviews of hotels, restaurants, attractions and other travel-related services. Anyone can post a review on Tripadvisor. Expedia.com is a website through which travel is booked; consumers are also encouraged to post reviews on the site, but, a consumer can only post a review if the consumer actually booked at least one night at the hotel through the website in the six months prior to the review post. Thus, the cost of posting a fake review on Expedia.com is quite high relative to the cost of posting a fake review on Tripadvisor. Further, since the reviewer had to undertake a credit card transaction on Expedia.com, the reviewer is not anonymous to the website host and thus, the potential for detection might also be higher. [3]

We present a simple analytical model that examines the equilibrium levels of manipulation of two horizontally-differentiated competitors who are trying to convince a consumer to purchase their product. The model demonstrates that the amount of potential reputational risk determines the

[3] As discussed above, TripAdvisor has been criticized for not managing the fraudulent reviewing problem. TripAdvisor recently announced the appointment of a new Director of Content Integrity. Even in the presence of substantial content verification activity on TripAdvisor's part, our study design takes as a starting point the potential for fraud in TripAdvisor's business model relative to Expedia.

amount of manipulation in equilibrium. We marry the insights from this model to the literature on organizational form and organizational incentive structures. Based on the modelas well as on the previous literature we examine the following hypotheses: 1) independent hotels are more likely to engage in review manipulation (post more fake positive reviews for themselves and more fake negative reviews for their next-door competitors) than branded chain hotels, 2) small owners are more likely to engage in review manipulation than large owner hotels, and 3) hotels with a small management company are more likely to engage in review manipulation than hotels that use a large management company.

Our main empirical analysis is akin to a differences in differences approach (although, unconventionally, neither of the differences is in the time dimension). Specifically, we examine differences in the reviews posted at Tripadvisor and Expedia for different types of hotels. For example, consider calculating for each hotel at each website the ratio of five star (the highest) reviews to total reviews. We ask whether the difference in this ratio for Tripadvisor vs. Expedia is higher for independent vs. branded chain hotels, whether the difference is higher for hotels that are owned by large owners vs. small owners, and whether the difference is higher for hotels that use large management companies vs. small management companies. Either difference alone would be problematic. Tripadvisor and Expedia reviews could differ due to differing populations at the site. Independent versus chain hotels could have different distributions of true quality, for example. However, our approach isolates whether the two hotel types' reviewing patterns are significantly different across the two sites. Similarly, we examine the ratio of one star (the lowest) reviews to total reviews for hotels that are close geographic neighbors of independent vs. chain hotels, hotels with small owners vs. large owners, and hotels with large management companies versus small management companies. That is, we measure whether the neighbor of independent hotels fare worse on Tripadvisor than on Expedia, for example.

The results are largely consistent with our hypotheses. That is, we find that hotel characteristics (such as ownership, affiliation and management structure) affect the amount of review manipulation. We find that there is relatively more positive manipulation than negative manipulation, even though the order of magnitude of the two is similar. We also find that the total amount of review manipulation, while economically significant, is relatively modest: we estimate that an independent hotel owned by a small owner will generate 7 more fake positive reviews (out of 114) and 4 more fake negative reviews than a chain hotel with a large owner.

The paper proceeds as follows. In Section 2 we discuss the previous literature. In Section 3 we present a simple analytical model and hypotheses. In Section 4 we describe the data and present summary statistics. In Section 5 we present our methodology and results, which includes main results as well as robustness checks. In Section 6 we conclude and also discuss limitations of the paper.

2 Previous Literature

Broadly speaking, our paper is informed by the literature on firm's strategic communication, which includes research on advertising and persuasion. In advertising models the sender is the firm, and the receiver is the consumer who tries to learn about the product's quality before making a purchase decision. In these models the firm signals the quality of its product through the amount of resources invested into advertising (see Nelson (1974), Milgrom and Roberts (1986), Kihlstrom and Riordan (1984), Bagwell and Ramey (1994), Horstmann and Moorthy (2003)) or the advertising content (Anand and Shachar (2009), Anderson and Renault (2006), Mayzlin and Shin (2011)). In models of persuasion, the receiver can influence the receiver's decision by optimally choosing the information structure (Crawford and Sobel (1982) and Chakraborty and Harbaugh (2010) show this in the case where the sender has private information, while Kamenica and Gentzkow (2011) show this result in the case of symmetric information). One common thread between all these papers is that in all of them the sender's identity and incentives are common-knowledge. That is, the receiver knows that the message is coming from a biased party, and hence is able to to take that into account when making her decision. In contrast, in our paper there is uncertainty surrounding the sender's true identity and incentives. That is, the consumer who reads a user review on Tripadvisor does not know if the review was written by an unbiased customer or by a biased source.

The models that are most closely related to the current research are Mayzlin (2006) and Dellarocas (2006). Mayzlin (2006) presents a model of "promotional" chat where competing firms, as well as unbiased informed consumers post messages about product quality online. Consumers are not able to distinguish between unbiased and biased word of mouth, and try to infer product quality based on online word of mouth. Mayzlin (2006) derives conditions under which online reviews are persuasive in equilibrium: online word of mouth influences consumer choice. She also demonstrates that producers of lower quality products will expend more resources on promotional reviews. Compared to a system with no firm manipulation, promotional chat results in welfare loss due to distortions in

consumer choices that arise due to manipulation. The welfare loss from promotional chat is lower the higher the participation by unbiased consumers in online fora. Dellarocas (2006) also examines the same issue. He finds that there exists an equilibrium where the high quality product invests more resources into review manipulation, which implies that promotional chat results in welfare increase for the consumer. Dellarocas (2006) additionally notes that the social cost of online manipulation can be reduced by developing technologies that increase the unit cost of manipulation and that encourage higher participation of honest consumers.

While the literature has not extensively studied biased reviewing, the potential for biased reviews affecting consumer responses to user reviews has been recognized. Perhaps the most intuitive form of biased review is the situation in which a producer posts positive reviews for its own product. In a well-documented incident, in February 2004, an error at Amazon.com's Canadian site caused Amazon to mistakenly reveal book reviewer identities. It was apparent that a number of these reviews were written by the books' own publishers and authors (see Harmon (2004)).[4] Other forms of biased reviews are also possible. For example, rival firms may benefit from posting negative reviews of each other's products. In assessing the potential reward for such activity, it is important to assess whether products are indeed sufficient substitutes to benefit from negative reviewing activity. For example, Chevalier and Mayzlin (2006) argue that two books on the same subject may well be complements, rather than substitutes, and thus, it is not at all clear that disingenuous negative reviews for other firm's products would be helpful in the book market. Consistent with this argument, Chevalier and Mayzlin (2006) find that consumer purchasing behavior responds less intensively to positive reviews (which consumers may estimate are frequently fake) than to negative reviews (which consumers may assess to be more frequently unbiased). However, there are certainly other situations in which two products are obviously substitutes; for example, in this paper, we hypothesize that two hotels in the same location are substitutes.[5]

A burgeoning computer science literature has attempted to empirically examine the issue of fakery by creating textual algorithms to detect fakery. Since the entire goal of a fake reviewer is to mimic a real reviewer; identifying textual markers of fakery is difficult. For example, the popular

[4]Similarly, in 2009 in New York, the cosmetic surgery company Lifestyle Lift agreed to pay $300,000 to settle claims regarding fake online reviews about itself. In addition, a web site called fiverr.com which hosts posts by users advertising services for $5 (e.g.: "I will drop off your dry-cleaning for $5") hosts a number of ads by people offering to write positive or negative hotel reviews for $5.

[5]In theory, a similar logic applies to the potential for biased reviews of complementary products (although this possibility has not, to our knowledge, been discussed in the literature). For example, the owner of a breakfast restaurant located next door to a hotel might gain from posting a disingenuous positive review of the hotel.

press has widely cited the methodology described in Ott et al. (2011) in identifying fake reviews. The researchers hired individuals on the Amazon Mechanical Turk site to write persuasive fake hotel reviews. They then analyzed the differences between the fake 5-star reviews and "truthful" 5-star reviews on Tripadvisor to calibrate their psycholinguistic analysis. However, it is possible that the markers of fakery that the researchers identify are not representative of differently-authored fake reviews. For example, the authors find that truthful reviews are more specific about "spatial configurations" than are the fake reviews. However, the authors specifically hired fakers who had not visited the hotel. We can not, of course, infer from this finding that fake reviews on Tripadvisor authored by a hotel employee would in fact be less specific about "spatial configurations" than true reviews. Since we are concerned with fake reviewers with an economic incentive to mimic truthful reviewers, we are skeptical that textual analysis can provide durable mechanisms for detecting fake reviews.[6] Some other examples of papers that use textual analysis to determine review fakery are Jindal and Liu (2007), Hu et al. (2012), and Mukherjee and Glance (2012).

Kornish (2009) uses a different approach to detect review manipulation. She looks for evidence of "double voting" in user reviews. That is, one strategy for review manipulation is to post a fake positive review for one's product and to vote this view as "helpful." That is, Kornish (2009) uses a correlation between review sentiment and usefulness votes as an indicator of manipulation. This approach is vulnerable to the critique that there may be other (innocent) reasons for such correlation, such as confirmatory bias: if most people who visit a product's page are positively inclined towards the product, more positive reviews will be marked as useful since these reviews confirm the initial belief.

Previous literature has not examined the extent to which the design of websites that publish consumer reviews can discourage or encourage manipulation. In this paper, we exploit those differences in design by examining Expedia versus Tripadvisor. The literature also has not empirically tested whether manipulation is more pronounced in empirical settings where it will be more beneficial to the producer. Using data on organizational form, quality, and competition, we examine the relationship between online manipulation and market factors which may increase or decrease the incentive to engage in online manipulation. We will detail our methodology below; however, it is important to understand that our methodology does not rely on identifying any particular review as unbiased (real) or promotional (fake).

[6] One can think of the issue here as being similar to the familiar "arms race" between spammers and spam filters.

Since review manipulation involves rule-breaking (most sites ask reviewers to pledge that they are not incentivised to write the review, and that the review represents an honest opinion of the product), this paper also relates to the economics literature on cheating. The most closely related papers in that stream are Duggan and Levitt (2002), Jacob and Levitt (2003), and Dellavigna and Ferrara (2010). In all three papers the authors do not observe rule-breaking or cheating ("throwing" sumo wrestling matches, teachers cheating on student achievement tests, or companies trading arms in embargoed countries) directly. Instead, the authors infer that rule-breaking occurs indirectly. That is, Duggan and Levitt (2002) document a consistent pattern of outcomes in matches that are important for one of the players, Jacob and Levitt (2003) infer cheating from consistent patterns test answers, and Dellavigna and Ferrara (2010). In all of these papers we see that cheaters respond to incentives. Importantly for our paper, Dellavigna and Ferrara (2010) show that a decrease in reputation costs of illegal trades results in more illegal trading. Our empirical methodology is similar to this previous work. First, we also do not observe review manipulation directly and must infer it from patterns in the data. Second, we hypothesize and show that the rate of manipulation is affected by differences in reputation costs for players in different conditions. The innovation in our work is that by using two different platforms with dramatically different costs of cheating we are able to have a benchmark.

Of course, for review manipulation to make economic sense, online reviews must play a role in consumer decision-making. Substantial previous research establishes that online reviews effect consumer purchase behavior (see, for example, Chevalier and Mayzlin, 2006). There is less evidence specific to the travel context. Vermeulen and Seegers (2009) measure the impact of online hotel reviews on consumer decision-making in an experimental setting with 168 subjects. They show that online reviews increase consumer awareness of lesser-known hotels and positive reviews improve attitudes towards hotels. Similarly, Ye and Gu (2009) use data from a major online travel agency in China to demonstrate a correlation between traveler reviews and online sales.

Finally, our research is related to the literature on ownership incentives. Our research design depends on smaller owner operators having sharper incentives to bear costs to post reviews and on larger hotel entities recognizing the potential for negative spillovers from being caught undertaking fraudulent activities for the entity's other properties. In this sense, our research is related to an extensive literature on differences in incentives between company-owned and franchised units of service industry chains (see, for example, Blair and Lafontaine (2005)). However, our unusually

rich dataset allows us to exploit the fact that ownership patterns in the hotel industry are actually quite complicated. For example, as discussed previously, a hotel can be franchised to a quite large franchisee company; that franchisee company is less incentivized to engage in fraudulent activity than a small franchisee. In our paper, we advance the literature on ownership by utilizing data on these complex ownership structures.

3 A Simple Model and Hypotheses

We propose a very simple and stylized model to fix ideas. The game consists of two competing firms, A and B, and a continuum of consumers. The time line of the game is the following:

1. ***Stage I***: Nature draws the true quality of each firm (q_A and q_B). We assume that the firms' true quality is not observable to any of the game's players.[7] The prior belief on the firm qualities are: $q_A \sim Normal(q_0, \sigma_q^2)$ and $q_B \sim Normal(q_0, \sigma_q^2)$. Here, the two firms a priori are identically distributed, but the model can be easily generalized to the case where the prior means are not equal. Unless otherwise noted, we assume that all other parameters of the model are common knowledge.

2. ***Stage II***: The firms set prices (p_A and p_B), which are observed by all the players.

3. ***Stage III***: Each firm can surreptitiously (and simultaneously) manufacture positive reviews for itself and negative reviews for its competitor. The reviews are posted by a third party platform that does not verify the reviewers' identity. That is, consumers can not differentiate between real and manufactured (or biased) user reviews. We denote by $e_{i,i}$ the effort that firm i invests into positive self-promotion (manufactured positive reviews), and by $e_{i,j}$ the effort that firm i invests into negative reviews for firm j. While the actual firms' efforts are not observed by the consumers, consumers do observe the user ratings for both firms. Hence we can think of the set of user ratings (which consists of real and fake reviews) providing a *signal* to the consumer on the firm's true quality. In particular, the signals arising from user

[7] The case where only firms, but not the consumers, observe each other's true quality yields similar results, but is considerably more complicated.

ratings are the following:

$$s_A = q_A + e_{A,A} - e_{B,A} + \varepsilon_A \tag{1}$$

$$s_B = q_B + e_{B,B} - e_{A,B} + \varepsilon_B \tag{2}$$

That is, the signal generated from user reviews on firm A's quality consists of the true quality (q_A), the positive self-promotion effort by firm A ($e_{A,A}$), the negative effort by its competitor ($e_{B,A}$), as well as a noise term (ε_A) that reflects random shocks experienced by unbiased reviews: $\varepsilon_i \sim Normal(0, \sigma_\varepsilon^2)$. We also assume that the noise terms are independent across firms.

4. We model the manipulation effort as costly to the firm. We can think of this cost as the reputation-related risks associated with this kind of promotion. That is, if the firm is caught doing this kind of activity, it will suffer damage to its reputation, where the damage may differ if the firm is caught doing self-promotion or generating negative review for its competitors. The chance of getting caught is increasing (at an increasing rate) in the intensity of the promotional activity: the cost is convex in the manipulation effort. Hence we assume that $\frac{\partial C(e_{i,i},e_{i,j})}{\partial e_{i,i}} > 0$, $\frac{\partial C(e_{i,i},e_{i,j})}{\partial e_{i,j}} > 0$, $\frac{\partial^2 C(e_{i,i},e_{i,j})}{\partial^2 e_{i,i}} > 0$, and $\frac{\partial^2 C(e_{i,i},e_{i,j})}{\partial^2 e_{i,j}} > 0$. The following assumed simple functional form satisfies these conditions: $C(e_{i,i}, e_{i,j}) = \frac{\delta}{2}(e_{i,i})^2 + \frac{\gamma}{2}(e_{i,j})^2$. Here δ signifies the damage caused to the firm if it caught doing self-promotion, and γ the damage if it is posting negative reviews for its competitor.

5. ***Stage IV***: Finally, the consumer chooses the product that maximized her utility. We assume that the products are horizontally differentiated. We use a simple Hotelling model of differentiation to model consumer choice, where firm A is located at $x = 0$, firm B is located at $x = 1$, and the consumer at location x chooses A if

$$E[q_A|s_A] - tx - p_A \geq E[q_B|s_B] - tx - p_B \tag{3}$$

We assume that consumers are uniformly distributed on the interval $[0, 1]$. Since consumers do not observe the true quality directly, their expected utility from A and B is inferred from the signals generated from user reviews.

We next solve for the firms' optimal actions by backward induction. We start with the consumer's inference in stage 4. After observing the signal s_A and s_B, the consumers' posterior beliefs on the firms' qualities are:

$$E[q_A|s_A] = (1-\mu_s)q_0 + \mu_s(s_A - \hat{e}^*_{A,A} + \hat{e}^*_{B,A}) \tag{4}$$

$$E[q_B|s_B] = (1-\mu_s)q_0 + \mu_s(s_B - \hat{e}^*_{B,B} + \hat{e}^*_{A,B}) \tag{5}$$

where $\mu_s = \frac{\sigma_q^2}{\sigma_\varepsilon^2+\sigma_q^2}$ $(0 < \mu_s < 1)$ is the optimal weight that the consumer puts on the firms' reviews, and $\hat{e}^*_{A,A}$ and $\hat{e}^*_{B,A}$ are the inferred equilibrium effort levels since the consumer does not observe the firms' manipulation activity directly.

Assuming market coverage, the consumer who is indifferent between the two products is located at point $\hat{x}$, where

$$\hat{x} = \frac{1}{2} + \frac{E[q_A|s_A] - E[q_B|s_B] + p_B - p_A}{2t} \tag{6}$$

Hence, the market shares of firms A and B are $\hat{x}$ and $1-\hat{x}$, respectively. This implies the following profit functions for firms A and B, respectively in stage 3:

$$\Pi^*_{A,Stage\,3} = \underset{e_{A,A},e_{A,B}}{max} \left(p_A E_{q_A,q_B,\varepsilon_A,\varepsilon_B} \left[\frac{1}{2} + \frac{E[q_A|s_A] - E[q_B|s_B] + p_B - p_A}{2t} \right] - \delta_A \frac{e^2_{A,A}}{2} - \gamma_A \frac{e^2_{A,B}}{2} \right) \tag{7}$$

$$\Pi^*_{B,Stage\,3} = \underset{e_{B,B},e_{B,A}}{max} \left(p_B E_{q_A,q_B,\varepsilon_A,\varepsilon_B} \left[\frac{1}{2} + \frac{E[q_B|s_B] - E[q_A|s_A] + p_A - p_B}{2t} \right] - \delta_B \frac{e^2_{B,B}}{2} - \gamma_B \frac{e^2_{B,A}}{2} \right) \tag{8}$$

Substituting (4) and (5) into (7) and (8), and taking the expectation, we can re-write the firm's maximization problem as the following:

$$\Pi^*_{A,Stage\,3} = \max_{e_{A,A},e_{A,B}} \left(p_A \left[\frac{1}{2} + \frac{\mu_s(e_{A,A} + e_{A,B} - \hat{e}^*_{A,A} - \hat{e}^*_{A,B} + c_A) + p_B - p_A}{2t} \right] - \delta_A \frac{e^2_{A,A}}{2} - \gamma_A \frac{e^2_{A,B}}{2} \right) \tag{9}$$

$$\Pi^*_{B,Stage\,3} = \max_{e_{B,B},e_{B,A}} \left(p_B \left[\frac{1}{2} - \frac{\mu_s(e_{B,B} + e_{B,A} - \hat{e}^*_{B,B} - \hat{e}^*_{B,A} + c_B) + p_A - p_B}{2t} \right] - \delta_B \frac{e^2_{B,B}}{2} - \gamma_B \frac{e^2_{B,A}}{2} \right) \tag{10}$$

where $c_A = -e_{B,A} - e_{B,B} + \hat{e}^*_{B,A} + \hat{e}^*_{B,B}$ and $c_B = -e_{A,B} - e_{A,A} + \hat{e}^*_{A,B} + \hat{e}^*_{A,A}$. Proposition 1 below summarizes the optimal manipulation levels for the firms as well as a key comparative static result:

Proposition 1. *In stage 3 (after the firms have committed to prices p_A and p_B),[8] the optimal promotional levels are the following:*

$$e^*_{A,A} = \frac{p_A \mu_s}{2\delta_A t};\; e^*_{A,B} = \frac{p_A \mu_s}{2\gamma_A t} \tag{11}$$

$$e^*_{B,B} = \frac{p_B \mu_s}{2\delta_B t};\; e^*_{B,A} = \frac{p_B \mu_s}{2\gamma_B t} \tag{12}$$

The Corollary below summarizes several key results that we will use in our empirical analysis:

Corollary 1. *The following results are implied by Proposition 1:*

1) A decrease in the reputational costs of manipulation increases the intensity of this activity: $\frac{\partial e^*_{B,A}}{\partial \gamma_B} > 0$, $\frac{\partial e^*_{B,A}}{\partial \gamma_B} > 0$, $\frac{\partial e^*_{B,A}}{\partial \gamma_B} > 0$, $\frac{\partial e^*_{B,A}}{\partial \gamma_B} > 0$.

2) Firms engage in negative manipulation of reviews of their competitors: $e^*_{A,B} > 0$ and $e^*_{B,A} > 0$, *and this activity increases as the costs of manipulation decrease. Hence, a firm that is located close to a competitor will have more negative reviews than a firm has no close competitors (which will have no fake negative reviews), and the number of fake negative reviews is greater if the competitor has lower costs of manipulation.*

[8] The equilibrium promotional levels here represent a partial equilibrium since they take the prices as given. In the Appendix, we solve for the full equilibrium of the game by endogenizing the prices: solving for the equilibrium prices as function of δ, γ and t. We show that the key comparative static - that the firm decreases the amount of review manipulation as the costs of promotion increase remains true in the full equilibrium as well.

Finally, we turn to the effect that review manipulation has on consumer choice. In the basic model consumer can invert the firm's problem and perfectly discounts the amount of manipulation. That is, in equilibrium, $e^*_{A,A} = \widehat{e}^*_{A,A}$, $e^*_{A,B} = \widehat{e}^*_{A,B}$, $e^*_{B,B} = \widehat{e}^*_{B,B}$, and $e^*_{B,A} = \widehat{e}^*_{B,A}$. Since fake reviews are perfectly discounted, the consumer would make the same choices in the current setting where fake reviews are possible and in one where fake reviews are not possible. Despite the fact that fake reviews do not affect consumer choices in equilibrium, firms prefer to post reviews. That is, if the firm chooses not to engage in manipulation, the consumer who expects fake reviews will think that the firm is terrible.

Next we consider a realistic extension of the model which changes the observability assumption. That is, suppose that the consumer does not observe the costs of each firm but forms an expectation on the costs based on prior beliefs. We believe that this assumption is more realistic for our empirical setting. We can show that this results in an outcome where a firm with lower manipulation cost has a higher share and the firm with higher manipulation cost has a lower share compared to the case where review manipulation is not possible. That is, this Proposition shows that manipulation of reviews may create distortions in choices under imperfect observability.

Proposition 2. *Assume for simplicity that $\delta = \gamma$. Suppose that the consumer does not observe the firms' costs of manipulation. That is, with probability α the firm has high cost of manipulation: $\delta = \delta_H$, and with probability $1 - \alpha$ the firm has low cost of manipulation: $\delta = \delta_L$. Consider the case where both types pool on price – consumers can not infer the firm's cost of manipulation from the price. Here $e^*_{L,i,i} = e^*_{L,i,j} = \frac{p_A \mu_s}{2\delta_L t}$, $e^*_{H,i,i} = e^*_{H,i,j} = \frac{p_A \mu_s}{2\delta_H t}$,* and $\widehat{e}^*_{i,i} = \widehat{e}^*_{i,j} = \frac{p_A \mu_s}{2(\alpha\delta_L + (1-\alpha)\delta_H)t}$. Here the consumer under-estimates the amount of manipulation for low-cost type of firm and over-estimates the amount of manipulation for high-cost firm: $e^*_{L,i,i} > \widehat{e}^*_{i,i} > e^*_{H,i,i}$ and $e^*_{L,i,j} > \widehat{e}^*_{i,j} > e^*_{H,i,j}$. This results in a higher share for low-cost firm and a lower share for high-cost firm compared to the case with no manipulation.

Based on the results of this simple model, we formulate the following hypotheses:

1. **Hypothesis 1**: A firm with lower potential reputational costs associated with review manipulation will create more fake reviews.

2. **Hypothesis 2**: A firm that is located close to a competitor will have more fake negative reviews than a firm with no close neighbors.

3. **Hypothesis 3**: A firm that is located close to competitor with low potential reputational costs will have more fake negative reviews than a firm that is located next to a competitor with high costs.

4 Data

User generated internet content has been particularly important in the travel sector. In particular, TripAdvisor-branded websites have more than 50 million unique monthly visitors and contain over 60 million reviews. While our study uses the US site, TripAdvisor branded sites operate in 30 countries. As Scott and Orlikowski (2012) point out, by comparison, the travel publisher Frommer's sells about 2.5 million travel guidebooks each year.

Our data derive from multiple sources. First, we identified the 25th to 75th largest US cities (by population) to include in our sample. Our goal was to use cities that were large enough to "fit" many hotels, but not so large and dense that competition patterns among the hotels would be difficult to determine. We then "scraped" data on all hotels in these cities from Tripadvisor and Expedia. Some hotels will not be listed on one or the other site and some hotels will not have reviews on one or the other site (typically, Expedia). At each site, we obtained the text and star values of all user reviews, the identity of the reviewer (as displayed by the site), and the date of the review. We also obtain data from STR, a market research firm that provides data to the hotel industry (www.str.com). To match the data from STR to our Expedia and Tripadvisor data, we use name and address matching. Our data consist of 3082 hotels matched between Tripadvisor, Expedia, and STR. Our biggest hotel city is Atlanta with 160 properties, and our smallest is Toledo, with 10 properties. Of the 3082 hotels matched across sites, 2931 have reviews on both sites.

Table 1 provides summary statistics for review characteristics, using hotels as the unit of observation, for the set of hotels that have reviews on both sites. Unsurprisingly, given the lack of posting restrictions, there are more reviews on Tripadvisor than on Expedia. On average, our hotels have nearly three times the number of reviews on Tripadvisor as on Expedia. Also, the summary statistics reveal that on average, Tripadvisor reviewers are more critical than Expedia reviews. The average Tripadvisor star rating is 3.50 versus 3.95 for Expedia. Based on these summary statistics, it appears that hotel reviewers are more critical than reviewers in other contexts. For example, numerous studies document that eBay feedback is overwhelmingly positive. Similarly, Chevalier and Mayzlin (2006) report average reviews of 4.14 out of 5 at Amazon and 4.45 at barnesandnoble.com

Table 1: User Reviews at Tripadvisor and Expedia

	Mean	Standard deviation	Minimum	Maximum
Number of Tripadvisor reviews	119.6	172	1	1675
Number of Expedia reviews	42.2	63.2	1	906
Average Tripadvisor star rating	3.52	0.75	1	5
Average Expedia star rating	3.95	0.74	1	5
Share of Tripadvisor 1 star reviews	0.14	0.17	0	1
Share of Expedia 1 star reviews	0.07	0.14	0	1
Share of Tripadvisor 5 star reviews	0.31	0.19	0	1
Share of Expedia 5 star reviews	0.44	0.26	0	1
Total number of hotels	2931			

for a sample of 2387 books.

Review characteristics are similar if we use reviews, rather than hotels as the unit of observation. Our dataset consists of 352,854 TripAdvisor reviews and 123,893 Expedia reviews. Of all reviews, 8.1% of TripAdvisor reviews are 1s and 38.0% of Trip Advisor reviews are 5s. For Expedia, 4.7% of all review are 1s while 48.5% of all reviews are 5s. Note that these numbers differ from the numbers in the table because hotels with more reviews tend to have better reviews. Thus, the average share of all reviews that are 1s is lower than the mean share of 1 star reviews for hotels.

We use the STR categorizations to identify the hotel category (economy, midscale, upper-midscale, upscale, upper upscale and luxury) and we used data from STR on the year that the hotel property was built to construct the hotel age. We also use STR to obtain the hotel location; we assign each hotel a latitude and longitude designator and use these to calculate distances between hotels of various types. Most importantly, we use STR data to construct the various measures of organizational form that we use for each hotel in the data set. A hotel can be an independent,

a franchised unit of a chain, or a company-owned unit of a chain. In general, franchising is the primary organizational form for the largest hotel chains in the US. For example, Choice Hotels, Marriott Hotels, and Starwood Hotels are all made up of more than 99% franchised units. Within the broad category of franchised units, there are a wide variety of organizational forms. STR provides us with information about each hotel's owner. The hotel owner (franchisee) can be an individual owner-operator or a large company. For example, Archon Hospitality owns 41 hotels in our focus cities. In Memphis, for example, Archon owns two Hampton Inns (an economy brand of Hilton), a Hyatt, and a Fairfield Inn (an economy brand of Marriott). Typically, the individual hotel owner (franchisee) is the residual claimant for the hotel's profits, although the franchise contract generally requires the owner to pay a share of revenues to the parent brand. Owners often, though not always, subcontract day to day management of the hotel to a management company. Typically, the management company charges a few of 3 to 5 percent of revenue, although agreements which involve some sharing of gross operating profits have become more common in recent years.[9] In some cases, the parent brand operates a management company. For example, Marriott provides management services for approximately half of the franchisee-owned hotels under the Marriott nameplate. Like owners, management companies can manage multiple hotels under different nameplates. For example, Crossroads Hospitality manages 29 properties in our data set. In Atlanta, they manage a Hyatt, a Residence Inn (Marriott's longer term stay-focused brand), a Doubletree, and a Hampton Inn (both Hilton brands). As discussed above, our model suggests that hotels with a relationship to a large company– either parent brand, or owning entity, and possibly management company – have a higher cost of posting promotional reviews and have a lower potential benefit from posting promotional reviews than do hotels that operate independently of such entities. While a consumer can clearly observe whether a hotel is a member of a branded chain, the ownership and management structure of the hotel are more difficult to infer for the consumer.

In constructing variables, we focus both on the characteristics of the hotel and characteristics of the hotel's neighbors. Table 2 provides summary measures of the hotel's own characteristics. First, we construct a dummy for whether the hotel is an independent or part of a branded chain, using the characterizations reported in STR: 18% of hotels in our sample are independent. The top 5 parent companies of branded chain hotels in our sample are: Marriott, Hilton, Choice Hotels, Intercontinental, and Best Western. Second, we construct a dummy for whether the hotel is owned

[9] See O'Fallon and Rutherford (2010).

Table 2: Hotel Affiliation, Ownership and Management and Structure

Hotel Status	Share of All Hotels With Reviews	Share of Independent Hotels	Share of Chain Affiliated Hotels
Independent	0.17	1.00	0.00
Marriott Corporation Affiliate	0.14	0.00	0.17
Hilton Worldwide Affiliate	0.12	0.00	0.15
Choice Hotels Int'l Affiliate	0.11	0.00	0.13
Intercontinental Hotels Grp Affiliate	0.08	0.00	0.10
Best Western Company Affiliate	0.04	0.00	0.04
Multi-unit owner	0.31	0.16	0.34
Multi-unit management company	0.52	0.35	0.55
Multi-unit owner AND multi-unit management company	0.26	0.12	0.24
Total Hotels in Sample = 2931			

by a multi-unit ownership entity identified by STR. For example, non-independent hotels that are not owned by a franchisee but owned by the parent chain will be characterized as owned by a multi-unit ownership entity, but so will hotels that are owned by a large multi-unit franchisee. Furthermore, while independent hotels do not have a parent brand, they are in some cases operated by large multi-unit owners. In our sample, 15% of independent hotels and 33% of branded chain hotels are owned by a multi-unit owners. As discussed above, these larger groups will be more concerned about the reputational spillovers of being caught undertaking promotional reviewing activity. Third, for some specifications, we will also examine hotels operated by large multi-unit management companies, which is the case for 32% of independent hotels and for 54% of branded chain hotels.

We then characterize the neighbors of the hotels in our data. The summary statistics for these measures are in Table 3. That is, for each hotel in our data, we first construct a dummy variable that takes the value of one if that hotel has a neighbor hotel within 0.5km. As the summary statistics show, 76% of the hotels in our data have a neighbor. We next construct a dummy that takes the value of one if a hotel has a neighbor hotel that is an independent. Obviously, this set of ones is a subset of the previous measure; 31% of the hotels in our data have an independent neighbor. We also construct a dummy for whether the hotel has a neighbor that is owned by a multi-unit owner. Again, the set of hotels that have a one for this measure are a subset of the hotels that

Table 3: Hotel Characteristics of Neighbor Hotels Within 0.5 km Radius

Hotel Status	Share of All Hotels With Reviews	Share of Independent Hotels	Share of Chain Affiliated Hotels
Hotel has a neighbor	0.76	0.72	0.77
Hotel has an independent neighbor	0.31	0.27	0.50
Hotel has a multi-unit owner neighbor	0.49	0.52	0.49
Hotel has a multi-unit management entity neighbor	0.59	0.58	0.59

have a neighbor. However, as discussed above, this set is not a proper subset of hotels that have a non-independent neighbor; some independent hotels are owned by multi-unit owners and many non-independent hotels are franchised to a small owner-operator. In our data 48% of the hotels have a neighbor owned by a multi-unit owner company. For some specifications, we also examine the management structure of neighbor hotels. We construct a variable that takes the value of one if a hotel has a neighbor hotel operated by a multi-unit management entity, which is the case for 58% of hotels in our sample. For our robustness specifications, we construct measures of hotel relatedness. A hotel is totally unrelated to another hotel if it is not a brand of the same parent (so, a Courtyard Marriott and Marriott are related), if it is not owned by the same ownership entity, and if it is not managed by the same management company. We construct a dummy variable that equals one if a hotel has a neighbor that is totally unrelated, which is the case for 54% of the hotels. Again, this variable will equal one for a subset of the hotels that have a neighbor of any sort.

5 Methodology and Results

As Section 4 describes, we collect reviews from two sites, Tripadvisor and Expedia. There is a key difference between these two sites which we utilize in order to help us identify the presence of review manipulation: while anybody can post a review on Tripadvisor, only those users who

purchased the hotel stay on Expedia in the past six months can post a review for the hotel.[10] This implies that it is far less costly for a hotel to post fake reviews on Tripadvisor versus posting fake reviews on Expedia; we expect that there would be far more review manipulation on Tripadvisor than on Expedia. In other words, a comparison of the difference in the distribution of reviews for the same hotel could potentially help us identify the presence of review manipulation. However, we can not infer promotional activity from a straightforward comparison of reviews for hotels overall on Tripadvisor and Expedia since the population of reviewers using Tripadvisor and Expedia may differ; the websites differ in characteristics other than reviewer identity verification.

Here we take a differences in differences approach (although, unconventionally, neither of our differences are in the time dimension): for each hotel, we examine the difference in review distribution across Expedia and Tripadvisor and across different competitive/ownership conditions. We use the results of Section 3 to argue that the incentives to post fake reviews will differ across different competitive/ownership conditions. That is, we hypothesize that hotels with greater incentive to manipulate reviews will post more fake positive reviews for themselves and more fake negative reviews for their hotel neighbors on Tripadvisor, and we expect to see these effects in the difference in the distributions of reviews on Tripadvisor and Expedia.

Consider the estimating equation:

$$\frac{NStarReviews_{ij}^{TA}}{Total\ Reviews_{ij}^{TA}} - \frac{NStarReviews_{ij}^{Exp}}{Total\ Reviews_{ij}^{Exp}} = X_{ij}B_1 + Own_{ij}B_2 + NeighOwn_{ij}B_3 \ + \sum \gamma_j \ + \varepsilon_{ij} \tag{13}$$

This specification estimates correlates of the difference between the share of reviews on TA that are N star and the share of reviews on Expedia that are N star for hotel i in city j. Our primary interest will be in the most extreme reviews, 1-star and 5-star. X_{ij} contains controls for hotel characteristics; these hotel characteristics should only matter to the extent that Tripadvisor and Expedia customers value them differentially. Specifically, we include the hotel's "official" star categorization common to Tripadvisor and Expedia, dummies for the six categorizations of hotel type provided by STR (economy, midscale, luxury, etc), and hotel age. Own_{ij} contains the own-

[10]Before a user posts a review on Tripadvisor, she has to click on a box that certifies that she has "no personal or business affiliation with this establishment, and have not been offered any incentive or payment originating the establishment to write this review." In contrast, before a user posts a review on Expedia, she must log in to the site, and Expedia verifies that the user actually purchased the hotel within the required time period.

hotel organizational and ownership characteristics. In our primary specifications, these include the indicator variable for independent and the indicator variable for membership in a large ownership entity. $NeighOwn_{ij}$ contains the variables measuring the presence and characteristics of other hotels within 0.5km. Specifically, we include an indicator variable for the presence of a neighbor hotel, an indicator variable for the presence of an independent neighbor hotel, and an indicator variable for the presence of a neighbor hotel owned by a large ownership entity. The variables γ_j are indicator variables for city fixed effects.

We start by examining the effects of own-hotel organizational and ownership characteristics (Own_{ij}) on the incentive to manipulate reviews. We argue that branded chain hotels have a higher reputational cost of review manipulation compared to independent hotels since if any single chain hotel is caught posting fake reviews, all the hotels in the chain will suffer damage to their reputation. Similarly, we argue that a multi-unit hotel owner has a higher reputational cost of manipulation since all the hotel owner's properties will suffer if the hotel is caught faking reviews. Finally, the same argument can be made for a multi-unit management company. In all of these cases, an entity that is associated with more properties has more to lose from being caught manipulating reviews: the negative reputational spillovers are higher. Hence, using hypothesis 1 from Section 3, we claim that 1) independent hotels have a higher incentive to post fake positive reviews (have a higher share of 5-star reviews on Tripadvisor versus Expedia) than branded chain hotels, 2) small owners have a higher incentive to post fake positive reviews than multi-unit owner hotels, 3) hotels with a small management company have a higher incentive to post fake positive reviews than hotels that use multi-unit management company. Finally, an alternative explanation for independent hotels having a higher share of positive reviews on Tripadvisor is that the Tripadvisor population likes independent hotels more than the Expedia population. While we can not rule out this alternative explanation, the same critique does not apply to the small owner variables since the ownership structure is not easily observable by customers or reviewers. That is, since neither the identity of the ownership entity (e.g.: Crossroads Hospitality) nor how many units it owns is observable to the reviewers, it is unlikely that reviewers on the different sites would exhibit different preferences for hotels that are owned by multi-unit entities versus single-unit entities. Similarly, we can argue that the size of the management company should not affect the relative preference for the hotel across the two sites.

Finally, we turn to the effect of $NeighOwn_{ij}$variables on review manipulation. Using hypothesis

2 from Section 3, we claim a hotel with a next-door neighbor will have more fake negative reviews (have a higher share of 1-star reviews on Tripadvisor than on Expedia) than a hotel with no next-door neighbor. In addition, using hypothesis 3 from Section 3, we claim that the following next-door neighbor characteristics will result in an increase in fake negative reviews: 1) having a neighbor that is independent, 2) having a neighbor that has a small owner, and 3) having a neighbor that is managed by a small management company.

5.1 Main Results

In this Section we present the estimation results of the basic differences in differences approach to identify review manipulation. Table 4 presents the results of the estimation of Equation 13). Heteroskedasticity robust standard errors are used throughout. We first turn to the specification where the dependent variable is the difference in the share of 5-star reviews. That is, the dependent variable is:

$$\frac{5StarReviews_{ij}^{TA}}{Total\ Reviews_{ij}^{TA}} - \frac{5StarReviews_{ij}^{Exp}}{Total\ Reviews_{ij}^{Exp}},$$

This is our measure of possible positive review manipulation. Consistent with our hypothesis that independent hotels optimally post more positive fake reviews, we see that independent hotels have 2.8 percentage points higher difference in the share of 5-star reviews across the two sites than branded chain hotels. Since hotels on Tripadvisor have on average a 31% share of 5-star reviews, the magnitude of the effect is large. As we mentioned before, while this result is consistent with manipulation, we can not rule out the possibility that reviewers on Tripadvisor tend to prefer independent hotels over branded chain hotels to a bigger extent than Expedia customers.

More interestingly, it is more difficult to believe that there is a strong disparity across sites in preferences for hotels with multi-unit owners, a hotel characteristic that is virtually unobservable to the consumer. Consistent with our hypothesis that multi-unit owners will find review manipulation more costly, and therefore engage in less review manipulation, we find that hotels that are owned by a multi-unit owner have 3.2 percentage point smaller difference in the share of 5-star reviews across the two sites. This translates to about four fewer 5-star reviews on Tripadvisor if we assume that the share of Expedia reviews stays the same across these two conditions and that the hotel has a total of 114 reviews on Tripadvisor, the site average. While we include neighbor effects in this specification, we do not have strong hypotheses on the effect of neighbor characteristics on the difference in the share of 5-star reviews across the two sites, since there is no apparent incentive

for a neighboring hotel to practice positive manipulation on the focal hotel. Indeed, in the 5-star specification, none of the estimated neighbor effects are significant.

We next consider to the specification where the dependent variable is the difference in the share of 1-star reviews. Our dependent variable is thus:
$\frac{1StarReviews_{ij}^{TA}}{Total\,Reviews_{ij}^{TA}} - \frac{1StarReviews_{ij}^{Exp}}{Total\,Reviews_{ij}^{Exp}}$,

This is our measure of negative review manipulation. Unlike the previous specification, here, we do not expect to see any effects of the hotel's organizational structure on its own share of 1 star reviews since a hotel is not expected to negatively manipulate its own ratings. Instead, our hypotheses concern the effects of the presence of neighbor hotels on negative review manipulation. The results are in Column 2 of Table 4. Our coefficient estimates suggest that the presence of any neighbor increases the difference in the 1-star share across the two sites, even though the effect is not significant. However, the presence of an independent hotel within 0.5km results in an increase of 1.8 percentage point in the difference in the share of 1-star reviews across the two sites relative to a non-independent neighbor. Our point estimates imply that having an independent neighbor versus having no neighbor results in a 2.8 percentage point increase in 1 star reviews (0.89 percentage points for having any neighbor plus 1.88 for the neighbor being independent). These are large estimated effects given that the average share of 1-star reviews is 15% for a hotel on Tripadvisor. Again, we hypothesize that multi-unit owners bear a higher cost of review manipulation and thus will engage in less review manipulation. Our results show that having a hotel with a multi-unit owner within 0.5km results in 2 percentage point decrease in the difference in the share of 1-star reviews across the two sites, relative to having a neighbor that is a single-unit owner.

What do the results in Table 4 suggest about the extent of manipulation of reviews overall on an open platform such as Tripadvisor? As we discuss above, the amount of manipulation depends on the exact hotel characteristics. As an example, let's consider the difference in positive manipulation under two extreme cases: a) a branded chain hotel that is owned by a multi-unit owner (the case with the lowest predicted and estimated amount of manipulation) and b) an independent hotel that is owned by a small owner (the case with the greatest predicted and estimated amount of manipulation). Our estimates suggest that, assuming the TripAdvisor average of 114 total reviews, we would expect about 7 more positive reviews in case b versus case a. Similarly, we can perform a comparison for the case of negative manipulation by neighbors. Consider case c) being located next door to a branded chain hotel that is owned by a multi-unit owner and (the case with the

Table 4: Estimation Results of Equation 13

		Difference in share of 5 star reviews	Difference in share of 1 star reviews
X_{ij}	Site rating	-0.0140 **	-0.0095 *
		(0.0067)	(0.0054)
	Hotel age	0.0003	0.0005 ***
		(0.0002)	(0.0001)
	Hotel tier controls?	Yes	Yes
Own_{ij}	Hotel is Independent	0.0280***	0.0113
		(0.0102)	(0.0096)
	Multi-unit owner	-0.0322 ***	-0.0028
		(0.0084)	(0.0047)
$NeighOwn_{ij}$	Has a neighbor	-0.0155	0.0091
		(0.0119)	(0.0104)
	Has independent neighbor	-0.0037	0.0191**
		(0.0097)	(0.0079)
	Has multi-unit owner neighbor	-0.0081	-0.0204 ***
		(0.0095)	(0.0073)
γ_j	City-level fixed effects?	YES	YES
	Num. of observations	2931	2931
	R-squared	0.11	0.09

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

Heteroskedasticity robust standard errors in parentheses. All neighbor effects calculated for 0.5km radius.

smallest predicted and estimated amount of manipulation) and case d) being located next door to an independent hotel that is owned by a small owner (the case with the greatest predicted and estimated amount of manipulation). Our estimates suggest that there would be a total of 4 more fake negative reviews in case d versus case c.

While it appears that the total amount of negative manipulation is lower than the amount of positive manipulation, it is useful to note that, given the overall average star rankings on Tripadvisor is above 3, an incremental 1 star review will change the average stars more than an incremental 5 star reviews. Unfortunately, it is impossible for us, given these data, to measure the effect that these ratings changes will have on sales. While Chevalier and Mayzlin (2006) show that 1 star reviews hurt book sales more than 5 star reviews help book sales, those findings do not necessarily apply to this context. Chevalier and Mayzlin note that two competing books on the same subject may indeed be net complements, rather than net substitutes. Authors and publishers, then, may gain from posting fake positive reviews of their own books, but will not necessarily benefit from posting negative reviews of rivals books. Thus, in the contexts of books, 1 star reviews may be more credible than 5 star reviews. We have seen that, in the case of hotels, where two hotels proximate to each other are clearly substitutes, one cannot infer that a 1 star review should be treated by customers as more credible than a 5 star review.

Finally, note that while it appears that the total amount of manipulation is economically significant in that we would expect it to distort choices, the amount of manipulation is small enough so that it should not destroy the informational value of Tripadvisor reviews. That is, we can speculate that while firms engage in review manipulation, and this sometimes distorts consumer choices, consumers still find reviews informative and persuasive. This is of course consistent with the observed popularity of Tripadvisor.

Our preceding analysis is predicated on the hypothesis that promotional reviewers have an incentive to imitate real reviewers as completely as possible. This is in contrast to the computer science literature, described above, that attempts to find textual markets or fake reviews. Nonetheless, we do separately examine one category of “suspicious” reviews. These are reviews that are posted by one-time contributors to Tripadvisor. The least expensive way for a hotel to generate a user review is to create a fictitious profile on Tripadvisor (which only requires an email address), and following the creation of this profile, to post a review. This is, of course, not the only way that the hotel can create reviews. Another option is for a hotel to pay a user with an existing review history to

post a fake review; another possibility is to create a review history in order to camouflage a fake review. Here, we examine "suspicious" reviews– the review for a hotel is the first and only review that the user ever posted. In our sample, 26% of all Tripadvisor reviews are posted by one-time reviewers. These reviews are more likely to be extreme compared to the entire Tripadvisor sample: 24% of one-time reviews are 1-star versus 15% in the entire Tripadvisor sample, and 39% of one-time reviewers are 5-star versus 31% in the entire Tripadvisor sample. Of course, the extremeness of one-time reviews does not in and of itself suggest that one-time reviews are more likely to be fake; users who otherwise do not make a habit of reviewing may be moved to do so by an an unusual experience with a hotel.

In Table 5 we present the results of the following three specifications. In the first column, we present the results of a specification where the dependent variable is the share of one-time contributor user reviews on Tripadvisor. Thus, our dependent variable is:

$$\frac{one\!-\!time\ Reviews_{ij}^{TA}}{Total\ Reviews_{ij}^{TA}}.$$

This captures the incidence of these suspicious reviews and includes potential positive as well as negative manipulation. Consistent with our earlier results, we find that an independent hotel has a 9 percentage point increase in the share of these reviews, which is a very large effect since the average share of one-time reviews amongst all hotels is 26%. Also consistent with our previous results, our point estimates suggest that a multi-unit owner has 1.6 percentage point decrease in the share of these reviews, and neighboring multi-unit hotel results results in a 1.9% decrease in the share. There is one variable in our specification that does not have the anticipated sign. The presence of any neighbor is negatively associated with "suspicious" reviews; our model would predict that this association would be positive.

The other two specifications in Table 5 address the valence of these reviews. For these specifications, the dependent variable is:

$$\frac{one\!-\!time\ NStarReviews_{ij}^{TA}}{Total\ Reviews_{ij}^{TA}} - \frac{NStarReviews_{ij}^{Exp}}{Total\ Reviews_{ij}^{Exp}}.$$

That is, we look at the difference between the share of N-star "suspicious" reviews on TripAdvisor and the overall share of N-Star reviews on Expedia. Column 2 shows the case where N=5. The effect of hotel independence is positive, as predicted, but not significantly different from zero. Multi-unit owner has a statistically significant 2.4 percentage point lower difference in the share of 5-star reviews across the two sites, which is consistent with our hypotheses and earlier results. The neighbor effects are not statistically significant, as they weren't in the specifications that used all

TA 5-star reviews. Column 3 shows the case where N=1. Here, we find that the presence of an independent hotel next door increases the difference in the share of 1-star reviews across the two sites by a statistically significant 2.1 percentage points, while having a hotel owned by a multi-unit owner next door decreases the difference in the share of 1-star reviews by 2.1 percentage points. The presence of a neighbor has an estimated positive effect, as predicted, but as in our previous specifications, is not statistically significant. Overall, these results confirm our prior results that manipulation of reviews takes place in a way that is consistent with predicted hotel incentives.

5.2 Robustness Checks

In this Section, we undertake a number of further checks that the results are robust to a variety of reasonable specifications. First, we consider additional variables concerning hotel structure. Specifically, we include a variable that equals one if the hotel is managed by a multi-unit management company. As we explain in Section 5, the management company is not residual claimant to hotel profitability the way that the owner is, but nonetheless, obviously has a stake in hotel success. Thus, we expect that a multi-unit management company would have a lower incentive to post fake reviews than a single-unit manager (which in many cases is the owner). We also include the neighbor analog of this variable, a variable that takes the value of one if the hotel has a neighbor that is managed by a multi-unit management company. In the first column in Table 6, we use the share difference in 5 star reviews as the dependent variable. We see that indeed a hotel that is managed by a multi-unit management company has a statistically significant 1.9 percentage point decrease in the difference of the share of 5-star reviews between the two sites which we interpret as a decrease in positive manipulation. Notably, the inclusion of this variable does not alter our previous results; independent hotels have more 5-star reviews on TripAdvisor relative to Expedia and hotels with multi-unit owners have fewer. There are, as before, no significant neighbor effects for 5-star reviews. Column 1 of Table 7 , repeats this same specification for 1-star reviews. Here, as before, we have no predictions for the own hotel characteristics (although we do find here that, with the inclusion of the large management dummy, the large owner dummy becomes statistically significantly different from zero). We do have predictions for neighbor characteristics. As before, we find that having an independent hotel neighbor significantly predicts more one star reviews, and that having a large owner chain neighbor predicts fewer one star reviews (although this effect is now not statistically significant). A large management chain is a negative but not statistically significant predictor of

Table 5: Results for Tripadvisor one-time contributor reviewers

		Share of one-time contributor user reviews	Difference in share of 5 stars reviews	Difference in share of 1 stars reviews
X_{ij}	Site rating	-0.0094 (0.0044)	-0.0080 (0.0078)	-0.0150 ** (0.0074)
	Hotel quality-tier controls?	YES	YES	YES
	Hotel age	0.0006 *** (0.0001)	0.0003 (0.0002)	0.0007 *** (0.0002)
Own_{ij}	Hotel is Independent dummy	0.0916 *** (0.0083)	0.0128 (0.0123)	-0.0137 (0.0121)
	Multi-unit owner	-.0160 *** (0.0054)	-0.0245 ** (0.0118)	0.0013 (0.0085)
$NeighOwn_{ij}$	Has a hotel neighbor dummy	-0.0164 * (0.0084)	-0.0124 (0.0157)	0.0128 (0.0144)
	Has independent hotel neighbor dummy	0.0023 (0.0066)	-0.0005 (0.0128)	0.0214 * (0.0118)
	Has a neighbor that is multi-unit owner dummy	-0.0192 *** (0.0065)	-0.0111 (0.0130)	-0.0212 * (0.0110)
γ_j	City-level fixed effects?	YES	YES	YES
	Num. of observations	3063	2874	2874
	R-squared	0.29	0.06	0.08

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

Heteroskedasticity robust standard errors in parentheses. All neighbor effects calculated for 0.5km radius.

one star reviews. Contrary to our hypothesis, the hotel neighbor dummy becomes negative in this specification, although also not statistically significant. Altogether, there is suggestive evidence that larger management companies are associated with less review manipulation.

In the second column in Table 6 and Table 7, we present specifications that include a new neighbor variable: a dummy that is one if the neighbor has absolutely no affiliation with the focal hotel (either via the same parent brand, the same owner company, or the same management company). We do not find that this variable has a significant effect (although we would not expect it to in the 5-star review specifications). Also, the other coefficients are not significantly impacted by including this variable.

In the right column of Table 6 and Table 7, we present 5-star and 1-star specifications that also include hotel chain fixed effects for the ten largest hotel brands. Inclusion of these chain fixed effects allows TripAdvisor and Expedia patrons to have a very general form of different preferences. They can have not only different preferences for hotel quality tiers and hotel age (all included in the controls in our base specifications), but also can have different preferences for different individual hotel brands. These specifications produce results very similar to the base specifications discussed in Table 5. The only change that inclusion of this variable causes compared to the earlier results is that the independent own hotel dummy in the 5-star specification is no longer statistically significant.

We also examine the relationship between our results and the results that would obtain by substituting data from Expedia for data from Orbitz. Until recently, Orbitz, like Expedia, only accepted reviews from individuals who had booked their stay at orbit.com. Starting in late 2010, Orbitz allowed others to submit hotel reviews, but reviews from verified customers are identified as "Verified" and are given higher weight in calculating the Orbitz Reviewer Score for each property. In our robustness results, we use only verified reviews from Orbitz. Thus, these reviews are analogous to Expedia reviews. Summary statistics are shown in Table 8. Orbitz is less attractive to us as a review site than Expedia. There are 104 hotels that have reviews at TripAdvisor and Expedia but no reviews at Orbitz. For hotels with reviews at both Orbitz and TripAdvisor, the hotels have only about three-quarters the number of Orbitz reviews as Expedia for the hotels in our sample. However, if our results are driven by important (and subtle) differences between the customer pools at Expedia and TripAdvisor, robustness of our results for Orbitz may be valuable. [

We do not use unverified reviews from Orbitz because there are very few of them. For our hotels, we have a total of 87716 verified Orbitz reviews and only 692 unverified reviews. The

Table 6: Specifications with 5-star reviews as dependent variable

		Difference in share of 5-star reviews	Difference in share of 5-star reviews	Difference in share of 5-star reviews
X_{ij}	Site rating	-0.0125 *	-0.0137 **	-0.0137 ***
		(0.0068)	(0.0067)	(0.0068)
	Hotel quality-tier controls?	YES	YES	YES
	Hotel age	0.0003	0.0003	0.00004
		(0.0002)	(0.0002)	(0.0002)
Own_{ij}	Hotel is Independent dummy	0.0255 **	0.0290 **	0.0097
		(0.0102)	(0.0103)	(0.0120)
	Hotel is part of a multi-unit owner dummy	-0.0264 ***	-0.0315 ***	-0.0199 **
		(0.0086)	(0.0083)	(0.0086)
	Hotel is managed by a multi-unit management company dummy	-0.0198 **	–	–
		(0.0091)		
	Hotel chain specific dummy	–	–	YES
$NeighOwn_{ij}$	Has a hotel neighbor dummy	-0.0123	-0.0107	-0.0143
		(0.0139)	(0.0133)	(0.0118)
	Has independent hotel neighbor dummy	-0.0041	-0.0019	-0.0061
		(0.0097)	(0.0098)	(0.0096)
	Has multi-unit owner hotel neighbor dummy	-0.0026	-0.0032	-0.0061
		(0.0113)	(0.0104)	(0.0095)
	Has a hotel neighbor managed by a multi-unit management company dummy	-0.0080	–	–
		(0.0136)		
	Has a hotel neighbor with no-affiliation dummy	–	-0.0124	–
			(0.0118)	
γ_j	City-level fixed effects?	YES	YES	YES
	Num. of observations	2931	2931	2931
	R-squared	0.11	0.11	0.13

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

Heteroskedasticity robust standard errors in parentheses. All neighbor effects calculated for 0.5km radius.

Table 7: Specifications with 1-star reviews as dependent variable

		Difference in share of 1 stars reviews	Difference in share of 1 stars reviews	Difference in share of 1 stars reviews
X_{ij}	Site rating	-0.0087 **	-0.0092 **	-0.0100 **
		(0.0054)	(0.0054)	(0.0054)
	Hotel quality-tier controls?	YES	YES	YES
	Hotel age	0.00049 ***	0.0005 ***	0.0005 ***
		(0.00013)	(0.0001)	(0.0001)
Own_{ij}	Hotel is Independent dummy	0.011	0.012	0.0031
		(0.0096)	(0.0096)	(0.0123)
	Hotel is part of a large-owner chain dummy	-0.00078	-0.0023	0.0002
		(0.0048)	(0.0125)	(0.0049)
	Hotel is managed by a large management company dummy	-0.0061	–	–
		(0.0067)		
$NeighOwn_{ij}$	Has a hotel neighbor dummy	0.0152	0.0125	0.0101
		(0.0124)	(0.011)	(0.0104)
	Has independent hotel neighbor dummy	0.0191 **	0.0204 ***	0.0182 **
		(0.0078)	(0.0079)	(0.0079)
	Has large-owner chain hotel neighbor dummy	-0.0132	-0.0169 **	-0.0206 ***
		(0.0083)	(0.0079)	(0.0072)
	Has a hotel neighbor managed by a large management company dummy	-0.0140	–	–
		(0.0108)		
	Has a hotel neighbor with no-affiliation dummy	–	-0.0089	–
			(0.0087)	
γ_j	City-level fixed effects?	YES	YES	YES
	Num. of observations	2931	2931	2931
	R-squared	0.09	0.09	0.09

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

Heteroskedasticity robust standard errors in parentheses. All neighbor effects calculated for 0.5km radius.

Table 8: Summary Statistics for Orbitz Reviews

	No. of Hotels	Mean	Std. Dev	Minimum	Maximum
Total Orbitz Verified Reviews	2569	32.59	47.75	1	628
Share of Verified Reviews 1 star	2569	0.093	0.16	0	1
Share of Verified Reviews 5 star	2569	0.241	0.20	0	1

difference between verified and unverified reviews are, nonetheless, interesting. For all Orbits verified reviews, 32% are 1s or 5s. For Orbitz unverified reviews, 58% are either 1s or 5s. Thus, unverified reviews are much more extreme than verified reviews. This, of course, could be due to unverified promotional reviews. However, it could also be the case that travelers who did not book through Orbitz nonetheless post reviews on Orbitz if they have had extreme experiences.

Table 9 repeats the regression specifications of Table 4, replacing Orbitz verified reviews with Expedia reviews. Regressions results are very similar to the results found in Table 4. As in Table 4, we find that Independent hotels have more 5 star reviews, and hotels from large ownership entities have fewer. In the Orbitz specification, the magnitude of the independence effect is somewhat larger than in our Expedia specifications, while the magnitude and significance of the multi-unit owner effect is smaller. Turning to 1 star reviews, we find, as in Table 4, that the presence of a neighbor has a positive but insignificant effect on 1 star reviews. Having an independent neighbor is associated with more 1 star reviews. As compared to Table 4, this effect is similar in magnitude, but is only statistically significant at the 15 percent confidence level. We also find, as in Table 4, that neighbors belonging to a large ownership entity are associated with fewer 1 star reviews. This effect is statistically significant at the ten percent level.

.

6 Conclusion and Limitations

We propose a novel methodology for empirically detecting review manipulation. In particular, we examine the difference in review distribution across Expedia and Tripadvisor, sites with different reviewer identity verification policies, and across different competitive/ownership conditions. Consistent with our theoretical claims, we find that an increase in hotel incentives to manipulate reviews results in an increase in our measures of manipulation. Substantively, we find that independent ho-

Table 9: TripAdvisor versus Orbitz Results

		Difference in Share of 5 star reviews	Difference in Share of 1 star reviews
X_{ij}	Site rating	-0.0076	-0.0029
		(0.0069)	(0.0048)
	Hotel age	-0.00095***	0.00044***
		(0.0002)	(0.00014)
	Hotel tier controls?	YES	YES
Own_{ij}	Hotel is Independent	0.048***	0.007
		(0.011)	(0.0106)
	Multi-unit owner	-0.013	-0.0030
		(0.009)	(0.0050)
	Has a neighbor	-0.0085	0.0109
		(0.012)	(0.0107)
$NeighOwn_{ij}$	Has independent neighbor	0.0076	0.0122
		(0.0106)	(0.0090)
	Has multi-unit owner neighbor	0.0188*	-0.0131*
		(0.0103)	(0.0084)
γ_j	City level fixed effects?	YES	YES
	Num of observations	2569	2569
	R-squared	0.04	0.05

***$p<0.1$, **$p<0.05$, *$p<0.10$

Heteroskedasticity robust standard errors in parentheses. All neighbor effects calculated for 0.5k radius.

tels engage in more review manipulation (both positive and negative), while hotels with multi-unit owners as well as hotels that are managed by a multi-unit management companies engage in less review manipulation (in the former case we find the effect for positive and negative manipulation, while in the latter we find the effect only in the case of positive manipulation). One important strength of our proposed methodology compared to earlier attempts is that our method does not require us to identify any particular review as fake or real, an inherently noisy and difficult task. Instead, we confine ourselves to examining differences between distributions.

Finally, we find that while the amount of review manipulation is economically significant, it is still small relatively to the total amount of reviewing activity. Why don't hotels engage in more intense review manipulation, given the fact that the mechanical costs of faking a review are low? Aside from any ethical concerns that the hotels have in engaging in this activity, we hypothesize that engaging in this activity exposes firms to reputational risks. The fact that the over-all level of manipulation activity seems to be relatively low is consistent with the notion that risks are perceived as relatively high. This perhaps explains how an open platform like Tripadvisor, that does not verify reviewer identity, can survive in the market. The obvious advantage of an open platform is that it allows the site to draw customers form all other sites (as well as from offline) as opposed to only restricting the reviews to its own customers. The downside is the degrading effect of review manipulation on the informational value of the site. Our empirical results show that the hotels are essentially able to self-police so that while they engage in some manipulation, the amount is not big enough to overwhelm the informational value of the site.

There are a number of limitations of this work. Perhaps the biggest limitation is that we do not observe manipulation directly but must infer it. This issue is of course inherent in doing research in this area. In the paper we deal with this limitation by building a strong case that the effects that we examine are due to review manipulation and not due to other unobserved factors. The second important limitation is that our measure of review manipulation does not include any content analysis. That is, one could imagine that one way in which a hotel could increase the impact of a fake review is by making particularly strong claims in the text of the review. For example, to hurt a competitor, a competitor could claim to be a traveler who witnessed a bed bug infestation. This is an interesting issue for future work.

Another limitation of this work is that we are unable to measure the impact that this manipulation has on consumer purchase behavior. Do consumers somehow detect and discount fake

reviews? Do they discount all reviews to some extent? Do they make poor choices on the basis of fake reviews? These questions are also left for future work.

7 Appendix

7.1 Proofs

Proof of Proposition 1:

The formulas in the Proposition are derived by taking the F.O.C.s of Equation (9) with respect to $e_{A,A}$ and $e_{A,B}$, and taking the F.O.C.s of (10) with respect to $e_{B,B}$ and $e_{B,A}$.

Endogenizing the Prices. As we argue in the main body of the paper, the firm does not expect manipulation to change its market share in expectation, given the optimal discounting by the consumer. Hence, the maximization problem in the second stage is the following:

$$\Pi^*_{A,Stage\,2} = \underset{p_A}{max} p_A \left[\frac{1}{2} + \frac{p_B - p_A}{2t}\right] - \delta_A \frac{\left(e^*_{A,A}\right)^2}{2} - \gamma_A \frac{\left(e^*_{A,B}\right)^2}{2} \tag{14}$$

$$\Pi^*_{B,Stage\,2} = \underset{p_B}{max} p_B \left[\frac{1}{2} + \frac{p_A - p_B}{2t}\right] - \delta_A \frac{\left(e^*_{B,B}\right)^2}{2} - \gamma_A \frac{\left(e^*_{B,A}\right)^2}{2} \tag{15}$$

After the appropriate substitutions (Proposition 1 provides $e^*_{A,A}$, etc.), taking the first order conditions, and some algebra, we have the following expressions for the equilibrium prices:

$$p_A = \frac{12t^3\delta_A\gamma_A\delta_B\gamma_B + 2t^2\delta_A\gamma_A(\delta_B+\gamma_B)\mu_s^2}{12t^2\delta_A\gamma_A\delta_B\gamma_B + 4\mu_s^2 t\left[(\gamma_A+\delta_A)\delta_B\gamma_B + (\gamma_B+\delta_B)\delta_A\gamma_A\right] + \mu_s^4\left[(\gamma_A+\delta_A)(\delta_B+\gamma_B)\right]} \tag{16}$$

$$p_B = \frac{12t^3\delta_A\gamma_A\delta_B\gamma_B + 2t^2\delta_B\gamma_B(\delta_A+\gamma_A)\mu_s^2}{12t^2\delta_A\gamma_A\delta_B\gamma_B + 4\mu_s^2 t\left[(\gamma_A+\delta_A)\delta_B\gamma_B + (\gamma_B+\delta_B)\delta_A\gamma_A\right] + \mu_s^4\left[(\gamma_A+\delta_A)(\delta_B+\gamma_B)\right]} \tag{17}$$

For simplicity, let's assume that $\delta_A = \gamma_A = \rho$ and $\delta_B = \gamma_B = 1$. We want to show that an increase in ρ (an increase in the reputational costs) results in less promotion on the part of firm A. Once

again, from Proposition 1 we know that in stage 3:

$$e^*_{A,A} = \frac{p_A \mu_s}{2\delta_A t}; e^*_{A,B} = \frac{p_A \mu_s}{2\gamma_A t} \tag{18}$$

$$e^*_{B,B} = \frac{p_B \mu_s}{2\delta_B t}; e^*_{B,A} = \frac{p_B \mu_s}{2\gamma_B t} \tag{19}$$

We take a derivative of these expressions, taking into account the fact that the prices are endogenous. That is, we can show that $\frac{\partial e^*_{A,A}}{\partial \rho} = \frac{\partial e^*_{A,B}}{\partial \rho} = \frac{\mu_s}{2t} \frac{\left[\frac{\partial p_A}{\partial \rho}\rho - p_A\right]}{\rho^2} < 0$.

Proof of Proposition 2:

Consider the firms' maximization problem:

$$\Pi^*_{A,Stage\,3} = \max_{e_{A,A}, e_{A,B}} \left(p_A \left[\frac{1}{2} + \frac{\mu_s(e_{A,A} + e_{A,B} - \hat{e}^*_{A,A} - \hat{e}^*_{A,B} + c_A) + p_B - p_A}{2t} \right] - \delta_A \frac{e^2_{A,A}}{2} - \gamma_A \frac{e^2_{A,B}}{2} \right) \tag{20}$$

$$\Pi^*_{B,Stage\,3} = \max_{e_{B,B}, e_{B,A}} \left(p_B \left[\frac{1}{2} - \frac{\mu_s(e_{B,B} + e_{B,A} - \hat{e}^*_{B,B} - \hat{e}^*_{B,A} + c_B) + p_A - p_B}{2t} \right] - \delta_B \frac{e^2_{B,B}}{2} - \gamma_B \frac{e^2_{B,A}}{2} \right) \tag{21}$$

The only difference here is that the consumer's inference ($\hat{e}^*_{A,A}$, $\hat{e}^*_{A,B}$, etc) will be different since the consumers can not observe the firm's cost function. Taking the derivative with respect to the promotion levels, it is clear that the optimal promotion level does not depend on the consumer's inference. That is, as before,

$$e^*_{A,A} = \frac{p_A \mu_s}{2\delta_A t}; e^*_{A,B} = \frac{p_A \mu_s}{2\gamma_A t} \tag{22}$$

$$e^*_{B,B} = \frac{p_B \mu_s}{2\delta_B t}; e^*_{B,A} = \frac{p_B \mu_s}{2\gamma_B t} \tag{23}$$

The consumer's inference will be a weighted average of the two types' optimal promotion levels, as is given in the Proposition. The result on share distortions follows directly from the fact that the consumer over-discounts the reviews for high-cost firm and under-discounts the review for low-cost

firm.

References

B. Anand and R. Shachar. (Noisy) Communication. *Quantitative Marketing and Economics*, 5(3): 211–237, 2009.

M. Anderson and J. Magruder. Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database. *The Economic Journal*, doi: 10.1111/j.1468-0297.2012.02512.x, 2012.

S. Anderson and R. Renault. Advertising Content. *American Economic Review*, 96(1):93–113, 2006.

K. Bagwell and G. Ramey. Advertising and Corrdination. *Review of Economic Studies*, 61(1): 153–171, 1994.

J. Berger, A. Sorensen, and S. Rasmussen. Positive Effects of Negative Publicity: When Negative Reviews Increase Sales. *Marketing Science*, 29(5):815–827, 2010.

R. Blair and F. Lafontaine. *The Economics of Franchising*. Cambridge University Press, 2005.

A. Chakraborty and R. Harbaugh. Persuasion by Cheap Talk. *American Economic Review*, 100(5): 2361–2382, 2010.

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

P. Chintagunta, S. Gopinath, and S. Venkataraman. 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, 2010.

V. Crawford and J. Sobel. Strategic Information Transmission. *Econometrica*, 50(6):1431–1451, 1982.

C. Dellarocas. Strategic manipulation of internet opinion forums: Implications for consumers and firms. *Management Science*, 52(20):1577–1593, 2006.

S. Dellavigna and E. La Ferrara. Detecting Illegal Arms Trade. *American Economic Journal*, 2(4): 26–57, 2010.

M. Duggan and S. Levitt. Winning Isn't Everything: Corruption in Sumo Wrestling. *American Economic Review*, 92(5):1594–1605, 2002.

Amy Harmon. Amazon glitch unmasks war of reviewers. *New York Times*, February 14 2004.

I. Horstmann and S. Moorthy. Advertising Spending and Quality for Services: The Role of Capacity. *Quantitative Marketing and Economics*, 1(3):337–365, 2003.

N. Hu, I. Bose, N. Koh, and L. Liu. Manipulation of online reviews: An analysis of ratings, readability, and sentiments. *Decision Support Systems*, 52:674–684, 2012.

B. Jacob and S. Levitt. Rotten Apples: An Investigation of the Prevalence and Predictors of Teacher Cheating. *Quarterly Jounral of Economics*, 118(3):843–877, 2003.

N. Jindal and B. Liu. Review Spam Detection. In *WWW '07 Proceedings of the 16th international conference on World Wide Web*, pages 1189 – 1190, 2007.

E. Kamenica and M. Gentzkow. Bayesian Persuasion. *American Economic Review*, 101(October): 2590–2615, 2011.

R. Kihlstrom and M. Riordan. Advertising as a Signal. *Journal of Political Economy*, 92(3):427–450, 1984.

L. Kornish. Are User Reviews Systematically Manipulated? Evidence from the Helpfulness Ratings. Leeds School of Business Working Paper, 2009.

D. Mayzlin. Promotional chat on the internet. *Marketing Science*, 25(2):155–163, 2006.

D. Mayzlin and J. Shin. Uninformative Advertising as an Invitation to Search. *Marketing Science*, 30(4):666–685, 2011.

M. Melnik and J. Alm. Does a seller's ecommerce reputation matter? Evidence from ebay auctions. *Journal of Industrial Economics*, 50(3):337–349, 2002.

P. Milgrom and J. Roberts. Price and Advertising Signals of Product Quality. *Journal of Political Economy*, 94(August):796–821, 1986.

B. Liu Mukherjee, A. and N. Glance. Spotting fake review groups in consumer reviews. *International World Wide Web Conference Committee*, April 16-20, 2012., 2012.

P. Nelson. Advertising as Information. *Journal of Political Economy*, 78(3):311–329, 1974.

M. O'Fallon and D. Rutherford. *Hotel Management and Operations*. John Wiley and Sons, 2010.

M. Ott, Y. Choi, C. Cardie, and J. Hancock. Finding Deceptive Opinion Spam by Any Stretch of the Imagination. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 309–319, 2011.

P. Resnick and R. Zeckhauser. *Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System*, volume 11, pages 127–157. 2002.

P. Resnick, R. Zeckhauser, J. Swanson, and K. Lockwood. The value of reputation on eBay: A controlled experiment. *Experimental Economics*, 9(2):79–101, 2006.

Scott and Orlikowski. Reconfiguring relations of accountability: Materialization of social media in the travel sector. *Accounting, Organizations, and Society*, 37:26–40, 2012.

I. Vermeulen and D. Seegers. "tried and tested: The impact of online hotel reviews on consumer consideration". *Tourism Management*, 30(1):123–127., 2009.

R. Law Ye, Q. and B. Gu. The impact of online user reviews on hotel room sales. *International Journal of Hospitality Management*, 28(1):180–182, 2009.

N. Zmuda. Ann taylor investigation shows ftc keeping close eye on blogging. *Advertising Age*, April 2010.