

PROMOTIONAL REVIEWS: AN EMPIRICAL INVESTIGATION
OF ONLINE REVIEW MANIPULATION

Dina Mayzlin¹

USC Marshall School of Business

Yaniv Dover

Tuck School of Business, Dartmouth College

Judith Chevalier

Yale School of Management and NBER

September 11, 2013

¹The authors contributed equally, and their names are listed in reverse alphabetical order. We thank the Wharton Interactive Media Initiative, the Yale Whitebox Center, and the National Science Foundation (Chevalier) for providing financial support for this project. We thank Steve Hood, Sr. Vice President of Research at STR for helping us with data collection. We also thank David Godes and Avi Goldfarb for detailed comments on the paper. We also thank numerous seminar participants for helpful comments. All errors remain our own.

Abstract

Firms' incentive to manufacture biased user reviews impede review usefulness. We examine the differences in reviews for a given hotel between two sites: Expedia.com (only a customer can post a review) and TripAdvisor.com (anyone can post). We argue that the net gains from promotional reviewing are highest for independent hotels with single-unit owners and lowest for branded chain hotels with multi-unit owners. We demonstrate that the hotel neighbors of hotels with a high incentive to fake have more negative reviews on TripAdvisor relative to Expedia; hotels with a high incentive to fake have more positive reviews on TripAdvisor relative to Expedia.

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 ability of consumers to evaluate unobservable product quality. However, one important impediment to the usefulness of reviews in revealing product quality is the possible existence of fake or “promotional” online reviews. Specifically, reviewers with a material interest in consumers’ purchase decisions 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 persuasion and advertising (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 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 any material connection between himself and the manufacturer.² Relatedly, in February 2012, the UK Advertising Standards Authority ruled that travel review website TripAdvisor must cease claiming 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 TripAdvisor site was genuine, and when we understood that might not be the case, we concluded that the claims were misleading.”³

¹Much 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), and Chintagunta et al. (2010).

²The 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>)

³www.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; as shown in previous literature (see Li and Hitt (2008) and Dellarocas and Wood (2007)), individuals who had an extremely positive or negative experience with a product may be particularly inclined to post reviews. In this paper, we do not attempt to classify whether any particular review is fake, and instead 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 a website where faking is difficult and a website where faking is relatively 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 she 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. Purchasing a hotel night through Expedia requires the reviewer to undertake a credit card transaction on Expedia.com. Thus, the reviewer is not anonymous to the website host, potentially raising the probability of detection of any fakery.⁴ We also explore the robustness of our results using data from Orbitz.com, where reviews can be either “verified” or “unverified.”

We present a simple analytical model in the Appendix that examines the equilibrium levels of

⁴As 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 higher potential for fraud in TripAdvisor’s business model relative to Expedia.

manipulation of two horizontally-differentiated competitors who are trying to convince a consumer to purchase their product. The model demonstrates that the cost of review manipulation (which we relate to reputational risk) determines the 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 model as well as on the previous literature we examine the following hypotheses: 1) hotels with a neighbor are more likely to receive negative fake reviews than more isolated hotels, 2) small owners are more likely to engage in review manipulation than hotels owned by companies that own many hotel units, 3) independent hotels are more likely to engage in review manipulation (post more fake positive reviews for themselves and more fake negative reviews for their competitors) than branded chain hotels, and 4) 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 one- and two-star (the lowest) reviews to total reviews. We ask whether the difference in this ratio for TripAdvisor vs. Expedia is higher for hotels with a neighbor within a half kilometer vs. hotels without a neighbor. Either difference alone would be problematic. TripAdvisor and Expedia reviews could differ due to differing populations at the site. Possibly, hotels with and without neighbors could have different distributions of true quality. 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- and two-star reviews to total reviews for TripAdvisor vs. Expedia for hotels that are close geographic neighbors of hotels with small owners vs. large owners, close neighbors of independent hotels vs. chain-affiliated hotels, and neighbors of hotels with large management companies versus small management companies. That is, we measure whether the neighbor of hotels with small owners fare worse on TripAdvisor than on Expedia, for example, than the neighbors of hotels owned by large multi-unit entities. We also measure the ratio of five-star (the highest) reviews to total reviews for TripAdvisor vs. Expedia for independent vs. chain hotels, hotels with small owners vs. large owners, and hotels with large management companies versus small management companies. Thus, our empirical exercise is a joint test of the hypotheses that promotional reviewing take place on TripAdvisor and that the incentive to post false reviews is a function of organizational form. Our identifying assumption

is that TripAdvisor and Expedia users do not differentially value hotel ownership and affiliation characteristics and the ownership and affiliation characteristics of neighbors. In our specifications, we control for a large number of hotel observable characteristics that could be perceived differently by TripAdvisor and Expedia consumers. We discuss robustness to selection on unobservables that may be correlated with ownership and affiliation characteristics.

The results are largely consistent with our hypotheses. That is, we find that the presence of a neighbor, neighbor characteristics (such as ownership, affiliation and management structure), and own hotel characteristics affect the measures of review manipulation. The mean hotel in our sample has a total of 120 reviews on TripAdvisor, of which 37 are 5-star. We estimate that an independent hotel owned by a small owner will generate an incremental 7 more fake positive Tripadvisor reviews than a chain hotel with a large owner. The mean hotel in our sample has thirty 1- and 2-star reviews on TripAdvisor. Our estimates suggest that a hotel that is located next to an independent hotel owned by a small owner will have 6 more fake negative Tripadvisor reviews compared to an isolated hotel.

The paper proceeds as follows. In Section 2 we discuss the prior literature. In Section 3 we describe the data and present summary statistics. In Section 4 we discuss the theoretical relationship between ownership structure and the incentive to manipulate reviews. 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 Prior Literature

Broadly speaking, our paper is informed by the literature on the 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, an information sender can influence the receiver's decision by optimally choosing the information structure (Crawford and Sobel (1982), Chakraborty and Harbaugh (2010), and Dziuda (2011) 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 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 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 by honest consumers.

The potential for biased reviews to affect consumer responses to user reviews has been recognized in the popular press. 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)).⁵ 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

⁵Similarly, 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.

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 more 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 strong substitutes; for example, in this paper, we hypothesize that two hotels in the same location are generally substitutes.⁶

A burgeoning computer science literature has attempted to empirically examine the issue of fakery by creating textual analysis algorithms to detect fakery. For example, Ott et al. (2011) create an algorithm to identify 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. They found a number of reliable differences in the language patterns of the fake reviews. One concern with this approach is that 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, it is an ongoing challenge for textual analysis methodologies to provide durable mechanisms for detecting fake reviews.⁷ 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 review as “helpful.” That is, Kornish (2009) uses

⁶In 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.

⁷One can think of the issue here as being similar to the familiar “arms race” between spammers and spam filters.

a correlation between review sentiment and usefulness votes as an indicator of manipulation. This approach isolates one possible type of review manipulation and is vulnerable to the critique that there may be other (innocent) reasons for a correlation between review sentiment and usefulness votes: if most people who visit a product’s page are positively inclined towards the product, the positive reviews may be on average considered to be more useful.

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).

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 affect consumer purchase behavior (see, for example, Chevalier and Mayzlin (2006) and Luca (2012)). 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 et al. (2010) use data from a major online travel agency in China to demonstrate a correlation between traveler reviews and online sales.

3 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. While TripAdvisor is primarily a review site, transactions-based sites such as Expedia and Orbitz also contain reviews.

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.⁸ In October of 2011, we “scraped” data on all hotels in these cities from TripAdvisor and Expedia. TripAdvisor and Expedia were co-owned at the time of our data collection activities but maintained separate databases of customer reviews at the two sites. As of December 2011, TripAdvisor derived 35 percent of its revenues from click-through advertising sold to Expedia.⁹ Thus, 35 percent of TripAdvisor’s revenue derived from customers who visited Expedia’s site immediately following their visit to the TripAdvisor site.

Some hotels are not listed on both sites, and some hotels do not have reviews on one of the sites (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 obtained data from Smith Travel Research, 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 2931 hotels matched between TripAdvisor, Expedia, and STR with reviews on both sites. Our biggest hotel city is Atlanta with 160 properties, and our smallest is Toledo, with 10 properties.

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 3.52 versus 3.95 for Expedia. Based on these summary statistics, it appears that hotel reviewers are more critical than reviewers in other previously studied 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 for a sample of 2387 books.

Review characteristics are similar if we use reviews, rather than hotels as the unit of observation.

⁸We dropped Las Vegas, as these hotels tend to have an extremely large number of reviews at both sites relative to hotels in other cities; these reviews are often focused on the characteristics of the casino rather than the hotel. Many reviewers may legitimately, then, have views about a characteristic of the hotel without ever having stayed at the hotel.

⁹Based on information in S-4 form filed by Tripadvisor and Expedia with SEC on July 27, 2011 (see <http://ir.tripadvisor.com/secfiling.cfm?filingID=1193125-11-199029&CIK=1526520>)

Table 1: User Reviews at TripAdvisor and Expedia

	Mean	Standard deviation	Minimum	Maximum
Number of TripAdvisor reviews	119.58	172.37	1	1675
Number of Expedia reviews	42.16	63.24	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			
Share of TripAdvisor 2-star reviews	0.11			
Share of Expedia 1-star reviews	0.07			
Share of Expedia 2-star reviews	0.08			
Share of TripAdvisor 5-star reviews	0.31			
Share of Expedia 5-star reviews	0.44			
Total number of hotels	2931			

Notes: The table reports summary statistics for user reviews for 2931 hotels with reviews at both TripAdvisor and Expedia collected in October of 2011.

Our data set consists of 350,485 TripAdvisor reviews and 123,569 Expedia reviews. Of all reviews, 8.0% of TripAdvisor reviews are 1s, 8.4% are 2s, and 38.1% are 5s. For Expedia, 4.7% of all review are 1s, 6.4% are 2s, and 48.5% of all reviews are 5s. Note that these numbers differ from the numbers in Table 1 because hotels with more reviews tend to have better reviews. Thus, the share of all reviews that are 1s or 2s is lower than the mean share of 1-star reviews or 2-star reviews for hotels. Since the modal review on TripAdvisor is a 4-star review, in most of our analyses we consider “negative” reviews to be 1- or 2-star reviews.

We 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. These locations are used to determine whether or not a hotel has a neighbor.

Importantly, we use STR data to construct the various measures of organizational form that we use for each hotel in the data set. We consider the ownership, affiliation, and management of a hotel. A hotel’s affiliation is the most observable attribute of a hotel to a consumer. Specifically, a hotel can have no affiliation (“an independent”) or it can be a unit of a branded chain. In our data, 17% of hotels do not have an affiliation. The top 5 parent companies of branded chain hotels in our sample are: Marriott, Hilton, Choice Hotels, Intercontinental, and Best Western. However, an important feature of hotels is that affiliation is very distinct from ownership. A chain hotel unit can be a franchised unit or a company-owned unit. In general, franchising is the primary organizational form for the largest hotel chains in the US. For example, International Hotel Group (Holiday Inn) and Choice Hotels are made up of more than 99% franchised units. Within the broad category of franchised units, there is 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, 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. Furthermore, while independent hotels do not have a parent brand, they are in some cases operated by large multi-unit owners. In our sample, 16% of independent hotels and 34% of branded chain hotels are owned by a multi-unit owners. Thus affiliation and ownership are distinct attributes of a hotel.

Owners often, though not always, subcontract day to day management of the hotel to a man-

agement company. Typically, the management company charges 3 to 5 percent of revenue for this service, although agreements which involve some sharing of gross operating profits have become more common in recent years.¹⁰ In some cases, the parent brand operates a management company. For example, Marriott provides management services for approximately half of the hotels not owned by Marriott but operated 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 brand), a Doubletree, and a Hampton Inn (both Hilton brands). While a consumer can clearly observe a hotel's affiliation, 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 a hotel and characteristics of the hotel's neighbors. The first nine rows in Table 2 provides summary measures of the hotel's own characteristics. We construct dummies for whether a hotel's affiliation is independent (vs. part of a branded chain). We also construct a dummy for whether the hotel has a multi-unit owner. For example, chain-affiliated 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. In our data, the modal hotel is a chain member, but operated by a small owner. For some specifications, we will also include a dummy variable that takes the value of one if the hotel is operated by a large multi-unit management company. This is the case for 35% of independent hotels and for 55% of branded chain hotels in our data.

We then characterize the neighbors of the hotels in our data. The summary statistics for these measures are given in the bottom four rows in Table ???. 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 all 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. In our data 49% 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

¹⁰See O'Fallon and Rutherford (2010).

Table 2: Own and Neighbor 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	0.29
Hotel has a neighbor	0.76	0.72	0.77
Hotel has an independent neighbor	0.31	0.50	0.27
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

Total Hotels in Sample = 2931

Notes: Table shows summary information about brand affiliation, ownership, and management characteristics for 2931 hotels sampled with reviews at TripAdvisor and Expedia and their neighbors within 0.5km

operated by a multi-unit management entity, which is the case for 59% of hotels in our sample.

In our specifications, we will be measuring the difference between a hotel’s reviews on TripAdvisor and Expedia. The explanatory variables of interest are the neighbor characteristics, the ownership and affiliation status, and the ownership and affiliation status of the neighbors. However, it is important that our specifications also include a rich set of observable hotel characteristics to control for the possibility that TripAdvisor and Expedia users value hotels with different characteristics differently. We obtain a number of characteristics. First, we include the “official” hotel rating for the hotel. At the time of our study, these official ratings were reported in common by TripAdvisor and Expedia and are based on the amenities of the hotel. From STR, we obtain a different hotel classification system; hotels are categorized as “Economy Class”, “Luxury Class”, “Midscale Class”, “Upper Midscale Class”, “Upper Upscale Class” and “Upscale Class.” We use dummy variables to represent these categories in our specifications. We also obtain the “year built” from STR and use it to construct a hotel age variable (censored at 100 years old). Using STR categorizations, we also construct dummy variables for “all suites” hotels, “convention” hotels, and a dummy that takes the value of one if the hotel contains at least one restaurant. Even within the same city, hotels have different location types. In all of our specifications, we include dummies for airport locations, resort locations, and interstate/suburban locations, leaving urban locations as the excluded type.

4 Theoretical Relationship between Ownership Structure and Review Manipulation

Previous literature on promotional reviewing (see Mayzlin (2006), Dellarocas (2006)) models review generation as a mixture of unbiased reviews and reviews surreptitiously generated by competing firms. The consumer, upon seeing a review, must discount the information taking into account the equilibrium level of review manipulation.

In the Appendix we present a simple model that is closely related to the previous models of promotional reviews but also allows the cost of review manipulation to differ across firms, a new key element in the current context. In the model firms engage in an optimal level of review manipulation (which includes both fake positive reviews for self and fake negative reviews for competitors). The cost of review manipulation is related to the probability of getting caught, which in turn increases in each fake review that is posted. This model yields the following intuitive result: an increase in

the firm's cost of review manipulation decreases the amount of manipulation in equilibrium. Note that this also implies that if the firm's competitor has lower cost of review manipulation, the firm will have more negative manufactured reviews.

The model reflects the fact that in practice the primary cost of promotional reviews from the firm's perspective is the risk that the activity will be publicly exposed. The penalties that an exposed firm faces range from government fines, possibility of lawsuits, and penalties imposed by the review-hosting platform. We use the literature on reputational incentives and organizational form to argue that this cost is also affected by the size of the entity. In this regard, our analysis is related to Blair and Lafontaine (2005) and Jin and Leslie (2009) who examine the incentive effects of reputational spillovers among co-branded entities. Our analysis is also related to Pierce and Snyder (2008), Bennett et al. (forthcoming), and Ji and Weil (2009). Bennett et al. (forthcoming) show that competition leads vehicle inspectors to cheat and pass vehicles that ought to fail emissions testing. Pierce and Snyder (2008) show that larger chains appear to curb cheating behavior from their inspectors; inspectors at a large chain are less likely to pass a given vehicle than are inspectors who work for independent shops. Similarly, Ji and Weil (2009) show that company-owned units of chains are more likely to adhere to labor standards laws than are franchisee-owned units. While our analysis is related to this prior literature, we exploit the rich differences in organizational form (chain vs. independent, large owner vs. small owner, and large management company vs. small management company) particular to the hotel industry.

Before we formulate our hypotheses on the effect of entity size on review manipulation, we note a few important details on the design of travel review sites. In particular, note that reviews on these sites are hotel-specific, rather than chain or owner specific. That is, a Hampton Inn in Cambridge, MA has unique reviews, distinct from the reviews of a Hampton Inn in Atlanta, GA. If one wants to enhance the reputation of both hotels positively, one must post positive reviews of both hotels separately on the site. If one wants to improve the attractiveness of these hotels relative to their neighbors, one must post negative reviews for the individual neighbors of each hotel separately on the site. These design features make it unlikely that reviews would generate positive reputational spillovers across hotels - that a fake review by one unit of a multi-unit entity is more productive because it creates positive reputational spillovers for other units in the entity. Note also that while the presence of positive spillovers is conceivable in the case of chain-affiliated hotel posting positive fake reviews about itself (an improved customer review at one Hampton Inn, for example, could

possibly benefit another Hampton Inn), it seems very unlikely in the case of the ownership variable since co-ownership is not visible to the customers. Thus, it seems inconceivable that a positive review for, say, Archon Hospitality's Memphis Fairfield Inn would improve the reputation of its Memphis Hampton Inn. Positive spillovers are also less likely to arise in the case of negative competitor reviews. Posting a negative review of one hotel will likely only benefit that hotel's neighbors, not other hotels throughout the chain.

In contrast to the discussion above, there are sizable negative spillovers associated with promotional reviews. Each incremental promotional review posted increases the probability of getting caught. A larger entity suffers a greater penalty from being caught undertaking fraudulent activities due to negative spillovers across various units of the organization. Specifically, if an employee of a multi-unit entity gets caught posting or soliciting fake reviews, any resulting government action, lawsuit, or retribution by the review site will implicate the entire organization. Because of this spillover, many larger entities have "social media policies," constraining the social media practices of employees or franchisees.¹¹

To make this concrete: suppose that the owner of Archon Hospitality, which owns 41 hotels in our sample under various nameplates, were contemplating posting a fake positive review about an Archon Hotel. As discussed above, the benefit of the fake review would likely only accrue to the one hotel about which the fake review was posted. To benefit another hotel, another fake review would have to be posted. However, the probability of getting caught increases in each fake review that is posted. If the owner of Archon were caught posting a fake review about one hotel, the publicity and potential TripAdvisor sanctions would spill over to all Archon hotels. Hence the cost of posting a fake review increases in the number of hotels in the ownership entity, but the benefit of doing so does not.

This mechanism is also demonstrated in a recent case. The Irish hotel Clare Inn Hotel and Suites, part of the Lynch Hotel Group, was given the "red badge" by TripAdvisor warning customers that the hotel manipulated reviews after it was uncovered that a hotel executive solicited positive reviews. TripAdvisor also removed reviews from other Lynch Hotel Group hotels, and the treatment of Lynch Hotel Group was covered by news media in Ireland. Although the Lynch Hotel Group hotels are not co-branded under a common nameplate, TripAdvisor took action against the whole hotel group

¹¹For example, Hyatt's social media policy instructs Hyatt employees to "Avoid commenting on Hyatt...only certain authorized individuals may use social media for Hyatt as business purposes...your conduct may reflect upon the Hyatt brand." (http://www.constangy.net/nr_images/hyatt-hotels-corporation.pdf, accessed April 10 2013).

given the common ownership and management of the hotels.¹² Thus, the key assumption underlying our ownership/affiliation specifications is that the reputational benefit of posting a fake review only accrues to one hotel, while the cost of posting the fake review (getting caught) multiplies in the number of hotels in the ownership or affiliation entity. Hence smaller entities have a bigger incentive to post fake reviews. In terms of our model, the larger entity bears a higher δ and γ , and hence will fake fewer reviews in equilibrium based on Proposition 1.

There is an additional incentive issue that applies specifically to ownership and works in the same direction as the mechanism that we highlight. Drawing on the literature on the separation of ownership and control, we hypothesize that owner-operated hotels have a greater incentive to engage in review manipulation (either positively for themselves or negatively for their neighbors). Owner-operators are residual claimants of hotel profitability and employee-operators are not. Thus, owner-operators would have more incentive to post fake reviews because owner-operators have sharper incentives to generate hotel profitability. An employee of a large ownership entity would have little to gain in terms of direct profit realization from posting fake reviews but would risk possible sanctions from the entity for undertaking fake reviewing activity.

In our paper, we consider the differential incentives of multi-unit entities using three measures of entity type. First, we consider ownership entities that are large multi-unit owners versus small owners. For example, this measure captures the distinction between an owner-operator Hampton Inn versus a Hampton Inn owned by a large entity such as Archon Hospitality. Our ownership hypotheses suggest that an owner-operator will have more incentive to post promotional reviews than will an employee of a large entity. Second, we consider independent hotels versus hotels operating under a common nameplate. As discussed above, affiliation is a distinct characteristic from ownership; independent hotels can be owner-operated but can also be owned by a large ownership entity. We hypothesize that units of branded hotels will have less incentive to post promotional reviews than will independents. As discussed above, brand organizations actively discourage promotional reviewing by affiliates (with the threat of sanctions) because of the chain-wide reputational implications of being caught. Third, we consider management by a large management company versus management by a smaller entity. Again in this case, a review posted by the entity will benefit only one unit in the entity while the cost of being caught can conceivably spill over to the entire entity. Unlike owners, hotel management companies are not residual claimants and unlike franchise operations,

¹²<http://www.independent.ie/national-news/hotel-told-staff-to-fake-reviews-on-TripAdvisor-2400564.html>

do not always engage in profit-sharing. Thus, while we examine hotel management companies in our analysis, it is less clear that they have a strong enough stake in the hotel to influence reviewing behavior.

In summary, we argue that the ownership and affiliation structure of the hotel affects the costs of the promotional reviewing activity, which in turn affects the equilibrium level of manufactured reviews. Specifically, based on our simple model and the discussion above, we make the following three theoretical claims:

1. A firm that is located close to a competitor will have more fake negative reviews than a firm with no close neighbors.
2. A firm that is part of a smaller entity will have more positive fake reviews.
3. A firm that is located close to a smaller entity competitor will have more fake negative reviews.

5 Methodology and Results

As Section 3 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.¹³ 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 neighbor and ownership/affiliation conditions.

¹³Before 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.

We use the claims of Section 4 to argue that the incentives to post fake reviews will differ across different neighbor and ownership/affiliation 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:

$$\frac{NStarReviews_{ij}^{TA}}{TotalReviews_{ij}^{TA}} - \frac{NStarReviews_{ij}^{Exp}}{TotalReviews_{ij}^{Exp}} = X_{ij}B_1 + OwnAf_{ij}B_2 + Nei_{ij}B_3 +$$

$$NeiOwnAf_{ij}B_4 + \sum \gamma_j + \varepsilon_{ij} \quad (1)$$

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/2-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, as discussed above, 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.), hotel age, location type dummies (airport, suburban, etc), and dummies for convention hotels, the presence of a hotel restaurant, and all suites hotels. Nei_{ij} is an indicator variable indicating the presence of a neighbor within 0.5km. $OwnAf_{ij}$ contains the own-hotel ownership and affiliation characteristics. In our primary specifications, these include the indicator variable for independent and the indicator variable for membership in a large ownership entity. $NeiOwnAf_{ij}$ contains the variables measuring the ownership and affiliation characteristics of other hotels within 0.5km. Specifically, we include 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.

Our cleanest specifications examine the effect of Nei_{ij} and $NeiOwnAf_{ij}$ variables on review manipulation. Following Claim 1 in Section 4, we hypothesize that a hotel with at least one neighbor will have more fake negative reviews (have a higher share of 1-star/2-star reviews on TripAdvisor than on Expedia) than a hotel with no neighbor. In addition, using Claim 3 from Section 4, we

hypothesize that the neighbor effect will be exacerbated when the firm has an independent neighbor, and that the neighbor effect will be mitigated when the firm has a multi-unit owner or multi-unit management company neighbor.

We then turn to the effects of own-hotel organizational and ownership characteristics ($OwnAf_{ij}$) on the incentive to manipulate reviews. Following the discussion in Section 4, we hypothesize that an entity that is associated with more properties has more to lose from being caught manipulating reviews: the negative reputational spillovers are higher. Hence, 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.

Our interpretation of these results relies on our maintained assumption that TripAdvisor and Expedia users value hotels with different ownership and affiliation characteristics similarly. An important alternative explanation for our results is that there are important differences in tastes of TripAdvisor and Expedia users for unobserved characteristics that are correlated with our ownership and neighbor variables. For example, one explanation for a finding that independent hotels have a higher share of positive reviews on TripAdvisor is that the TripAdvisor population likes independent hotels more than the Expedia population. We discuss this alternative hypothesis at length in the robustness section below. Here we note that this alternative explanation is much more plausible a priori for some of our results than for others. In particular, we find the alternative hypothesis less plausible for the specifications for which the neighbor variables are the variables of interest. For the neighbor specifications, the alternative hypothesis suggests that, for example, some consumers will systematically dislike a Fairfield Inn whose neighbor is an owner-operated Days Inn relative to a Fairfield Inn whose neighbor is a Days Inn owned by a large entity like Archon, and that this difference in preferences is measurably different for TripAdvisor and Expedia users.

Note that our empirical methodology is similar to the approach undertaken in 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) infer arms embargo violations if weapon-making companies' stocks react to changes in conflict intensity. 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.

5.1 Main Results

In this Section we present the estimation results of the basic differences in differences approach to identify review manipulation. Table 3 presents the results of the estimation of Equation (1). Heteroskedasticity robust standard errors are used throughout.

We first consider to the specification where the dependent variable is the difference in the share of 1- and 2-star reviews. Our dependent variable is thus $\frac{1+2StarReviews_{ij}^{T.A}}{TotalReviews_{ij}^{T.A}} - \frac{1+2StarReviews_{ij}^{Exp}}{TotalReviews_{ij}^{Exp}}$. This is our measure of negative review manipulation. We begin with the simplest specification: we examine the difference between negative reviews on TripAdvisor and Expedia for hotels that do and do not have neighbors within 0.5km. This specification includes all of the controls for hotel characteristics (X_{ij} in Equation 1), but does not include the $OwnAf_{ij}$ and $NeiOwnAf_{ij}$ characteristics. The results are in Column 1 of Table 3. The results show a strong and statistically significant effect of the presence of a neighbor on the difference in negative reviews on TripAdvisor vs. Expedia. The coefficient estimate suggests that hotels with a neighbor have an increase of 1.9 percentage points in the share of 1-star and 2-star reviews across the two sites. This is a large effect given that the average share of 1- and 2-star reviews is 25% for a hotel on TripAdvisor.

We continue with our analysis of negative reviews by examining ownership and affiliation characteristics. We include in the specification all of the own hotel ownership characteristics and the neighbor owner characteristics ($OwnAf_{ij}$ and $NeiOwnAf_{ij}$). For these negative review manipulation results, we do not expect to see any effects of the hotel's own organizational structure on

Table 3: Estimation Results of Equation 1

		Difference in share of 1- and 2-star reviews	Difference in share of 1- and 2-star reviews	Difference in share of 5- star reviews
X_{ij}	Site rating	-0.0067 (0.0099)	-0.0052 (0.0099)	-0.0205** (0.0089)
	Hotel age	0.0004*** (0.0002)	0.0003* (0.0002)	0.0002 (0.0002)
	All Suites	0.0146 (0.0092)	0.0162* (0.0092)	0.0111 (0.0111)
	Convention Center	0.0125 (0.0086)	0.0159* (0.0091)	-0.0385*** (0.0113)
	Restaurant	0.0126 (0.0093)	0.0114 (0.0092)	0.0318*** (0.0099)
	Hotel tier controls?	Yes	Yes	Yes
	Hotel location controls?	Yes	Yes	Yes
$OwnAf_{ij}$	Hotel is Independent		0.0139 (0.0110)	0.0240** (0.0103)
	Multi-unit owner		-0.0011 (0.0063)	-0.0312*** (0.0083)
Nei_{ij}	Has a neighbor	0.0192** (0.0096)	0.0296** (0.0118)	-0.0124 (0.0119)
$NeiOwnAf_{ij}$	Has independent neighbor		0.0173* (0.0094)	-0.0051 (0.0100)
	Has multi-unit owner neighbor		-0.0252*** (0.0087)	-0.0040 (0.0097)
γ_j	City-level fixed effects?	YES	YES	YES
	Num. of observations	2931	2931	2931
	R-squared	0.05	0.06	0.12

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

Regression estimates of Equation (1).

The dependent variable in all specifications is the share of reviews that are N star for a given hotel at TripAdvisor minus the share of reviews for that hotel that are N star at Expedia.

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

its share of 1- and 2-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 3. As before, our coefficient estimates suggest that the presence of any neighbor within 0.5km significantly increases the difference in the 1- and 2-star share across the two sites. 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 the presence of a multi-unit owner hotel within 0.5km results in 2.5 percentage point decrease in the difference in the share of 1- and 2-star reviews across the two sites, relative to having only single-unit owner neighbors. This negative effect is statistically different from zero at the 1 percent confidence level. As expected, the hotel’s own ownership and affiliation characteristics do not have a statistically significant relationship to the presence of 1-star and 2-star reviews. The presence of an independent hotel within 0.5km results in an additional increase of 1.7 percentage point in the difference in the share of 1-star and 2-star reviews across the two sites. Our point estimates imply that having an independent neighbor versus having no neighbor results in a 4.7 percentage point increase in and 1- and 2 star reviews (3.0 percentage points for having any neighbor plus 1.7 for the neighbor being independent). These estimated effects are large given that the average share of 1- and 2-star reviews is 25% for a hotel on TripAdvisor.

Of course, the neighbor characteristics are the characteristics of interest in the 1- and 2-star review specifications. However, our specifications include the hotel’s own ownership characteristics as control variables. The estimated coefficients for the hotel’s own ownership characteristics are small in magnitude and statistically insignificant. This is consistent with our manipulation hypotheses but seem inconsistent with the alternative hypothesis of differences in preferences for ownership characteristics across TripAdvisor and Expedia users.

We next turn to the specification where the dependent variable is the difference in the share of 5-star reviews. That is, the dependent variable is $\frac{5StarReviews_{ij}^{TA}}{TotalReviews_{ij}^{TA}} - \frac{5StarReviews_{ij}^{Exp}}{TotalReviews_{ij}^{Exp}}$. 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 2.4 percentage points higher difference in the share of 5-star reviews across the two sites than branded chain hotels. This effect is statistically different from zero at the five percent confidence level. Since hotels on TripAdvisor have on average a 31% share of 5-star reviews, the magnitude of the effect is reasonably large. However, 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.

We also measure the disparity across sites in preferences for hotels with multi-unit owners. 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 a 3.1 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 120 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 large or statistically significant. In interpreting these results, it is important to remember that the ownership characteristic is virtually unobservable to the consumer; it measures the difference between, for example, an Archon Hospitality Fairfield Inn and an owner-operator Fairfield Inn. Nonetheless, it is plausible that TripAdvisor and Expedia users differentially value hotel characteristics that are somehow correlated with the presence of an owner-operator (and not included in our regression specifications). We return to this issue below.

For the 5-star specifications, the hotel's own ownership characteristics are the variables of interest, rather than the neighbor variables. Here, we find the estimated coefficients of the neighbor characteristics to be small and statistically insignificant. This finding is consistent with our manipulation hypothesis but seems inconsistent with the alternative hypothesis that TripAdvisor and Expedia users have systematically different preferences for hotels with different kinds of neighbors.

What do our results suggest about the extent of review manipulation on an open platform such as TripAdvisor overall? Note that we cannot identify the baseline level of manipulation on TripAdvisor that is uncorrelated with our characteristics. Thus, we can only provide estimates for the difference between hotels of different characteristics. However, 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). Recall that the average hotel in our sample has 120 reviews, of

which 37 on average are 5-star. Our estimates suggest that we would expect about 7 more positive TripAdvisor 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 a completely isolated hotel and case d) being located near an independent hotel that is owned by a small owner. For the average hotel with 120 reviews, thirty 1-star and 2-star reviews would be expected as a baseline. Our estimates suggest that there would be a total of 6 more fake negative reviews on TripAdvisor in case d versus case c.

Our main results focus on the presence of neighbors and the ownership and affiliations of hotels and their neighbors. However, hotels differ structurally not only in their ownership but also in their management. As explained above, some hotel units have single unit owners, but these owners outsource day to day management of the hotels to a management company. In our sample of 2931 hotels, of the 2029 that do not have multi-unit owners, 767 do outsource management to multi-unit managers. As we explain in Section 4, 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. As in the case of multi-unit owners, posting of fake reviews by an employee of a management company could, if detected, have negative implications for the management company as a whole. 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). This implies that hotel neighbors of hotels with multi-unit managers should have fewer 1- and 2-star reviews on TripAdvisor while hotels with multi-unit managers should have fewer 5-star reviews on TripAdvisor, once again if we assume that the share of Expedia reviews stays the same.

In the first column in Table 4, we use the share difference in 1- and 2-star reviews as the dependent variable. Here, as before, we have no predictions for the own hotel characteristics (and none are statistically different from zero). We do have predictions for neighbor characteristics. As before, we find that having any neighbor is associated with having more 1- and 2-star reviews, a 3.8 percentage point increase. As before, an independent hotel neighbor is associated with more negative reviews on TripAdvisor relative to Expedia and having a large owner chain neighbor is associated with fewer negative reviews on TripAdvisor. The presence of a large management company neighbor is associated with fewer negative reviews on TripAdvisor, although the effect is not statistically significant at standard confidence level. The presence of a large owner neighbor and the presence of a large management company neighbor are quite positively correlated. A test of the joint significance

shows that the two variables are jointly significant in our specification at the 1 percent level.

In the second column of Table 4, we examine 5-star reviews. Here, as before, the neighbor characteristics are uninformative. As before, independent hotels have more 5-star reviews on TripAdvisor relative to Expedia and that hotels with a large owner company have fewer 5-star reviews. In addition, the results show that a hotel that is managed by a multi-unit management company has a statistically significant 2.1 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 continue to have significantly more 5-star reviews on TripAdvisor relative to Expedia and hotels with multi-unit owners have fewer 5-star reviews. This result is important because, like a multi-unit owner company, management by a multi-unit management company is invisible to the consumer. Thus, altogether, there is suggestive evidence that, like larger owner companies, larger management companies are associated with less review manipulation.

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 (2006) 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 context 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 or 2 star review should be treated by customers as more credible than a 5-star review.

5.2 Results for One-Time Reviewers

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 markers of fake reviews. Nonetheless, for robustness, 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

Table 4: Management Company Specifications

		Difference in share of 1- and 2- star reviews	Difference in share of 5-star reviews
X_{ij}	Site rating	-0.0047 (0.0100)	-0.0183** (0.0090)
	Hotel age	0.0003** (0.0002)	0.0002 (0.0002)
	All Suites	0.0169* (0.0091)	0.0144 (0.0112)
	Convention Center	0.0163* (0.0090)	-0.0363*** (0.0113)
	Restaurant	0.0110 (0.0092)	0.0323*** (0.0099)
	Hotel tier controls?	YES	YES
	Hotel location controls?	YES	YES
$OwnAf_{ij}$	Hotel is Independent	0.0141 (0.0111)	0.213** (0.0104)
	Multi-unit owner	-0.0014 (0.0064)	-0.0252*** (.0086)
	Multi-unit management company	0.0022 (0.0077)	-0.0211 ** (0.0092)
Nei_{ij}	Has a neighbor	0.0379*** (0.0142)	-0.0098 (0.0140)
$NeiOwnAf_{ij}$	Has independent neighbor	0.0173* (0.0094)	-0.006 (0.0100)
	Has multi-unit owner neighbor	-0.0169* (0.0097)	0.0004 (0.0114)
	Has multi-unit management company neighbor	-0.0183 (0.0125)	-0.0059 (0.0136)
γ_j	City-level fixed effects?	YES	YES
	Num. of observations	2931	2931
	R-squared	0.06	0.12

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

Regression estimates of Equation (1). The dependent variable in all specifications is the share of reviews that are N star for a given hotel at TripAdvisor minus the share of reviews for that hotel that are N star at Expedia. Heteroskedasticity robust standard errors in parentheses. All neighbor effects calculated for 0.5km radius.

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; yet 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, 23.0% of all TripAdvisor reviews are posted by one-time reviewers. These reviews are more likely to be extreme compared to the entire TripAdvisor sample: 47.6% of one-time reviewers are 5-star versus 38.1% in the entire TripAdvisor sample. There are more negative outliers as well: 24.3% of one-time reviews are 1-star and 2-star versus 16.4% 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 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 $\frac{\text{one-time Reviews}_{ij}^{TA}}{\text{Total Reviews}_{ij}^{TA}}$. This captures the incidence of these suspicious reviews and includes potential positive as well as negative manipulation. The most striking result is that one-time reviews are 8.8 percentage points more common for independent hotels. This is consistent with our earlier results, but also could be attributable to legitimate customer reviewing preferences. Also consistent with our earlier results, we find a negative impact of multi-unit owner on one-time reviewing activity, and a negative impact of multi-unit owner neighbors. 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 (although this effect is insignificant); 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 $\frac{\text{one-time } N\text{Star Reviews}_{ij}^{TA}}{\text{one-time Reviews}_{ij}^{TA}} - \frac{N\text{Star Reviews}_{ij}^{Exp}}{\text{Total Reviews}_{ij}^{Exp}}$. That is, we look at the difference between the share of N -star reviews among “suspicious” reviews on TripAdvisor and the overall share of N -Star reviews on Expedia. Ideally, we might want to compare one-time reviews on TripAdvisor to one-time reviews on Expedia. Unfortunately, Expedia’s reviewer identification features render identifying one-time reviewers impossible. Column 2 shows the case where $N = 1$ or 2, our specification that focuses on the characteristics of neighbor hotels. The presence of a neighbor is associated with a 2.9 percentage point increase in the share of one-time reviews that are 1 or 2 stars. This effect is statistically significant at the ten percent level. The effect of indepen-

Table 5: Results for TripAdvisor one-time contributor reviewers

		Share of one-time contributor user reviews	Difference in share of 1 and 2 star reviews	Difference in share of 5 star reviews
X_{ij}	Site rating	-0.0176*** (0.0061)	-0.0175 (0.0113)	-0.0083 (0.0102)
	Hotel age	0.0003** (0.0001)	0.00005 (0.0002)	0.0002 (0.0002)
	All Suites	0.0086 (0.0065)	-0.0147 (0.0137)	0.0035 (0.0150)
	Convention Center	-0.0177** (0.0082)	0.0532*** (0.0147)	-0.0716*** (0.0170)
	Restaurant	0.0376*** (0.0068)	0.0079 (0.0126)	0.0329*** (0.0128)
	Hotel tier controls?	YES	YES	YES
	Hotel location controls?	YES	YES	YES
Own_{ij}	Hotel is Independent	0.0881*** (0.0079)	-0.0035 (0.0135)	0.0082 (0.0123)
	Multi-unit owner	-.0135** (0.0052)	0.0109 (0.0102)	-0.0239** (0.0117)
Nei_{ij}	Has a neighbor	-0.0091 (0.0080)	0.0285* (0.0156)	-0.0093 (0.0159)
$NeiOwnAf_{ij}$	Has independent neighbor	0.0002 (0.0066)	0.0203 (0.0133)	0.0027 (0.0130)
	Has multi-unit owner neighbor	-0.0144** (0.0062)	-0.0150 (0.0125)	-0.0038 (0.0132)
γ_j	City-level fixed effects?	YES	YES	YES
	Num. of observations	2874	2874	2874
	R-squared	0.35	0.05	0.07

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

Estimation of Equation 1 with the sample restricted to hotels that have at least one review by a one-time contributor (the reviewer has only submitted one review on TripAdvisor).

The dependent variable in the first column is the share of reviews by one-time contributors among all TripAdvisor reviews for a given hotel. The dependent variable in the other two columns is the share of reviews by one-time contributors that are N star for a given hotel at TripAdvisor minus the share of reviews for that hotel that are N star at Expedia.

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

dent neighbors and multi-unit neighbors are positive and negative, respectively, in accordance with our model and previous results. However, these effects are not significant at standard confidence levels. It is possible that these results are weak in part because one-time reviews are “suspicious.” TripAdvisor has a policy whereby hotels can contest suspicious reviews and TripAdvisor may, at its discretion, remove contested “suspicious” reviews from the site. Negative reviews by one-time reviewers may be more likely to be expunged from the site. In Column 3, we examine 5-star reviews, the specifications in which we focus on own-hotel characteristics. 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.

Overall, these results confirm our prior results that manipulation of reviews takes place in a way that is consistent with predicted hotel incentives. However, our results for “suspicious” reviews are not as compelling as our results for all “reviews.” Of course, with this analysis we are forced to construct the left hand side variable using a smaller subset of reviews, which may be noisy. Further, if fakers are sophisticated in their attempt to avoid detection, they may be avoiding these suspicious reviewing activities.

5.3 Robustness Checks

Perhaps the main concern with our results is the potential for selection on unobservables. That is, TripAdvisor and Expedia users may differ in their taste for hotel characteristics. We have included many such possible characteristics in our specifications (hotel age, hotel tier, hotel location type, etc.). Thus, differences in tastes for these included characteristics are not a problem for our analysis; we have controlled for these in our specification. However, it is possible that consumer tastes differ across the two websites for unobservable characteristics. This is a concern if the unobservable characteristics are correlated with the ownership, affiliation, and neighbor variables of interest. This could in principle lead to significant measured impacts of our ownership, affiliation, and neighbor variables even if ownership, affiliation, and neighbor characteristics are not associated with review manipulation. A priori, we find this alternative hypothesis less plausible for any specifications in which the variables of interest are neighbor variables. It seems unlikely, for example, that TripAdvisor users systematically dislike (relative to Expedia users) hotels whose hotel neighbors are franchisees that operate a single hotel. A priori, we find selection on unobservables to be a more

plausible concern for specifications in which the variables of interest are own-hotel ownership and affiliation.

To investigate selection on unobservables, we undertake the following exercise. Recall that our base specifications include a rich set of control variables. We re-estimate the base specifications in Table 3, maintaining the neighbor, ownership, and affiliation variables but removing all of the control variables. We compare the result of this no-controls specification to our basic results including all of the control variables. We examine how much (if at all) inclusion of the control variables attenuates the coefficients for the variables of interest. If unobservable characteristics are positively correlated with observable characteristics, one might expect that the inclusion of additional controls, if they were available, would further attenuate the coefficients on the variables of interest.¹⁴ The no-control specifications are shown in Columns 1 through 3 of Table 6. Comparing Table 3 to Table 6, for the neighbor specification shown in Column 1, reestimation excluding all control variables actually produces a smaller point estimate of the neighbor effect. Thus, inclusion of a set of control variables does not attenuate the results at all. This finding has been interpreted in the literature as assuaging concerns about selection on unobservables. Similarly, the full neighbor ownership-affiliation specification in Column 2 of Table 3 can be compared with the specifications with no control variables in Table 6. The independent neighbor variable has a stronger measured impact on review differences in the regression with the controls versus the regression without controls. Again, inclusion of controls does not attenuate the independence effect. The owner-neighbor effect does attenuate from -0.031 to -0.025 with the inclusion of the control variables. However, our specifications contain a very rich set of control variables. If we could hypothetically perform a regression that contained all of the unobservables, and if these unobservables were as powerful as the observable control variables in attenuating the ownership effect, the ownership effect would still remain substantial in magnitude. Thus, for the neighbor specifications in Columns 1 and 2 of Table 3, we conclude that selection on unobservables is unlikely to be a major explanation for our results.

A priori, selection on unobservables is more plausible for the five star specifications examining a hotel's own characteristics. Own hotel ownership and affiliation are plausibly correlated with characteristics that TripAdvisor and Expedia customers could value differently. Again, we examine this issue by comparing the no controls specifications in Column 3 of Table 6 to the full controls specifications in Column 3 of Table 3. Here, the alternative hypothesis of selection on unobservables

¹⁴See Altonji et al. (2005) for a more formal discussion of this test.

Table 6: Robustness specifications: Specifications with no controls

		Difference in Share of 1- and 2 -star reviews	Difference in Share of 1- and 2- star reviews	Difference in Share of 5- star reviews
X_{ij}	Site rating	–	–	–
	Hotel age	–	–	–
	All Suites	–	–	–
	Convention Center	–	–	–
	Restaurant	–	–	–
	Hotel tier controls	NO	NO	NO
	Hotel location controls	NO	NO	NO
Own_{ij}	Hotel is Independent	–	0.0093 (0.0092)	0.0429*** (0.0103)
	Multi-unit owner	–	-0.0181** (0.0075)	-0.0642*** (0.0084)
Nei	Has a neighbor	0.0118 (0.0079)	0.0324*** (0.0098)	-0.0177 (0.0109)
$NeiOwnAf_{ij}$	Has independent neighbor	–	0.0022 (0.0082)	-0.0069 (0.0091)
	Has multi-unit owner neighbor	-	-0.0310*** (0.0083)	-0.0211** (0.0093)
γ_j	City level fixed effects?	NO	NO	NO
	Num of observations	2931	2931	2931
	R-squared	0.001	0.01	0.04

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

Estimation of Equation 1 excluding control variables.

The dependent variable in all specifications is the share of reviews that are N star for a given hotel at TripAdvisor minus the share of reviews that are N star at Expedia.

Heteroskedasticity robust standard errors in parentheses.

All neighbor effects calculated for 0.5 km radius.

Results in this Table can be compared to the “base” specifications in Table 3 to measure whether and if the inclusion of control variables in Table 3 leads to substantial attenuation of the coefficients of interest.

is more difficult to reject. Both the multi-unit owner dummy and the independent hotel dummy are attenuated by approximately 50% when controls are added to the regression. Thus, our interpretation of these coefficients as evidence for review manipulation relies on the included hotel characteristics being more powerful than omitted hotel characteristics in explaining the difference in reviewer behavior on TripAdvisor and Expedia.

Our analysis of Table 6 is one strategy to examine the importance of omitted hotel characteristics. In Appendix Table 8, we take another approach to examining omitted hotel characteristics. Here, we reexamine the base specifications of Table 3, including 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 3. Here, the neighbor variables of interest are all of roughly the same magnitude and significance as in our base specifications. 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; the ownership variable remains of the expected sign and statistically significantly different from zero.

Given the importance of our negative review specifications, we next turn to a few robustness checks that examine the robustness of our neighbor ownership and affiliation results. Throughout, we have used 1 and 2 star reviews as our marker of “negative” reviews. We chose this specification in part due to the summary statistics outlined in Table 1. While 31% of reviews on TripAdvisor are 5s, 1s and 2s together only account for 25% of reviews. Hence, a firm attempting to denigrate its competitor will often be able to do so effectively with either 1- or 2-star promotional reviews. Furthermore, a scan of web blogs, etc. suggests that hoteliers complain to TripAdvisor about fake 2-star negative reviews from competitors and that TripAdvisor has sometimes deemed such reviews as fake and removed them.¹⁵ Nonetheless, we provide robustness results where we examine the basic specification in Equation 1 above, but consider only determinants of 1-star reviews. This is shown in the first column of Table 7. The results are similar to the base 1- and 2-star results in Column 2 of Table 3; the own-hotel ownership and affiliation characteristics have little explanatory power and are insignificant. The independent neighbor and large company owner neighbor coefficients are

¹⁵See, for example, “Fake Review Number Two” in <http://TripAdvisorwatch.wordpress.com/trip-advisor-fake-reviews/>.

similar in magnitude and significance to the main specification. The “having any neighbor within 0.5km” indicator variable has a smaller coefficient (although still of the hypothesized sign) but is not statistically significant at standard confidence levels.

We also examine the robustness of our results by altering the radius that we use to define neighbors. In our base specifications in Table 3, we define a neighbor as a hotel that is very close to the hotel of interest— within 0.5km kilometer of the hotel of interest. Under this definition, 76% of the hotels in our sample have a neighbor. In Columns 2 through 4 of Table 7 we re-estimate the specification of Column 2 of Table 4, but using different radii to define neighbors— 0.3km, 0.7km, and 0.9km. Under the narrower radius definition of 0.3km, 65% of hotels have a neighbor. Under the wider radius definitions of 0.7km and 0.9km, 82% and 85% of the hotels in our sample have a neighbor, respectively. These varying radii specifications are similar to our base specification. As the radius widens, the dummy for having any neighbor appears to diminish in magnitude and significance (as nearly every hotel has a neighbor), while the neighbor characteristics maintain or even increase explanatory power. Thus, we conclude that our results change with the radius size in a sensible way.

Finally, we examine the robustness of our results to our particular choices of review site. Specifically, we examine the relationship between our results and the results that would obtain by replacing the data from Expedia analyzed above with data from another site, Orbitz.com. Orbitz.com, like Expedia, is primarily a travel booking site that hosts user reviews. Orbitz is a less popular site than Expedia; Orbitz had approximately 60 percent fewer page views than Expedia in 2012.¹⁶ In addition, whereas we expect there to be a large overlap between TripAdvisor and Expedia audiences due to the companies’ co-marketing efforts at the time of data collection (see our discussion above), we do not have the same expectations for TripAdvisor and Orbitz. We provide details of our analysis in Online Appendix. In summary, while we find that that our Orbitz results are qualitatively similar to the Expedia results presented in the paper, the magnitude of some of the Orbitz results is smaller. Overall, we take these results as suggestive that our findings are robust when examining alternative sites.

¹⁶Data from ComScore, found at <http://www.newmediatrendwatch.com/markets-by-country/17-usa/126-online-travel-market>, accessed April 11, 2013.

Table 7: Specifications with negative reviews as dependent variable

		Difference in share of 1-star reviews	Difference in share of 1- and 2-star reviews	Difference in share of 1- and 2-star reviews	Difference in share of 1- and 2-star reviews
X_{ij}	Site rating	-0.0177** (0.0076)	-0.0055 (0.0099)	-0.0050 (0.0099)	-0.0048 (0.0099)
	Hotel age	0.0005*** (0.0001)	0.0003 ** (0.0002)	0.0003** (0.0002)	0.0003** (0.0002)
	All Suites	0.0091 (0.0076)	0.0159* (0.0092)	0.0156* (0.0091)	0.0158* (0.0091)
	Convention Center	0.0104 (0.0073)	0.0166* (0.0091)	0.0170* (0.0091)	0.0163* (0.0091)
	Restaurant	0.0039 (0.0076)	0.0110 (0.0092)	0.0091 (0.0092)	0.0091 (0.0092)
	Hotel tier controls?	YES	YES	YES	YES
	Hotel location controls?	YES	YES	YES	YES
Own_{ij}	Hotel is Independent	0.0117 (0.0100)	0.0126 (0.0110)	0.0113 (0.0109)	0.0114 (0.0109)
	Multi-unit owner	-0.0025 (0.0047)	-0.0012 (0.0063)	-0.0015 (0.0063)	-0.0020 (0.0063)
Nei_{ij}	Has a neighbor	0.0095 (0.0106)	0.0258** (0.0102)	0.0125 (0.0132)	0.0137 (0.0147)
$NeiOwnAf_{ij}$	Has independent neighbor	0.0192** (0.0081)	0.0109 (0.0099)	0.0176* (0.0095)	0.0207** (0.0093)
	Has multi-unit owner neighbor	-0.0204*** (0.0075)	-0.0262*** (0.0084)	-0.0252*** (0.0093)	-0.0262*** (0.0096)
γ_j	City-level fixed effects?	YES	YES	YES	YES
	Num. of observations	2931	2931	2931	2931
	Neighbor radius?	0.5km	0.3km	0.7km	0.9km
	R-squared	0.09	0.05	0.05	0.05

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

Estimation of Equation (1). The dependent variable in all specifications is the share of reviews that are N star for a given hotel at TripAdvisor minus the share of reviews that are N star at Expedia. Heteroskedasticity robust standard errors in parentheses. The radius for which neighbors are calculated for a given hotel is given in the table.

6 Conclusion and directions for future work

We propose a novel methodology for empirically detecting review manipulation. In particular, we examine the difference in review distributions 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 hotels with next-door neighbors have more negative reviews on Tripadvisor, and this effect is exacerbated if the neighbor is an independent hotel with a small owner. That is, we find evidence for negative review manipulation. We also observe review patterns that are consistent with positive manipulation: we find that independent hotels with small owners and small management companies have more positive reviews on Tripadvisor. While we find evidence for both negative and positive manipulation, throughout the paper we emphasize the fact that our results on negative manipulation are more robust to selection issues than our positive manipulation results. We conclude from our results that promotional reviewing is sufficiently economically important that actors that are differentially situated economically will indulge in promotional reviewing to a measurably different extent.

Our paper also contributes to the literature on incentives and organizational form. Our unusually rich data set 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; we hypothesize that the large franchisee company is less incentivized to engage in this type of fraudulent activity than a small franchisee. In our paper, we advance the literature on ownership by utilizing data on these complex ownership structures. We show that larger organizations appear to be measurably better at curbing cheating.

While it is not our primary goal, our paper also contributes to the literature on fake review detection. Previous methodologies in the computer science literature infer that reviews are more likely to be fake if they contain certain textual markers of fakery (such as not using spatial language). We have noted that a concern with these methodologies is that manipulating the textual markers in response to detection algorithms is relatively inexpensive. In contrast, the organization form of a hotel and its neighbors are very difficult to alter. Our results suggest that a detection algorithm could incorporate these factors in assessing the probability that a given review is fake.

Our paper also has implications for user review system design. Our results suggest that pro-

motional reviews are less common on Expedia than on TripAdvisor. Thus, the policy of verifying reviews does limit promotional reviews. However, this limitation comes at a cost: there are far fewer reviews on Expedia than on TripAdvisor. While the policy used by Orbitz (and now Amazon) of marking verified and unverified reviews is an interesting compromise, it may discourage unverified reviews and does not fully solve the review site's problem of whether to fully incorporate unverified reviews into summary data.

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 "traveler" could claim to have witnessed a bed bug infestation. This is an interesting issue for future work.

In this work, 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 suggest important avenues for future work.

7 Appendix

7.1 A Simple Model

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.¹⁷ 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 ratings are the following:

$$s_A = q_A + e_{A,A} - e_{B,A} \quad (2)$$

$$s_B = q_B + e_{B,B} - e_{A,B} \quad (3)$$

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 ($e_{A,A}$) and the negative effort by its competitor ($e_{B,A}$). Note that while we model the benefit of effort (the signal generated from user

¹⁷The case where only firms, but not the consumers, observe each other's true quality yields similar results, but is considerably more complicated.

Table 8: Chain Fixed Effects Robustness Check

		Difference in share of 1- and 2-star reviews	Difference in share of 1- and 2-star reviews	Difference in share of 5- star reviews
X_{ij}	Site rating	-0.0007 (0.0100)	-0.0067 (0.0101)	-0.0193** (0.0089)
	Hotel age	0.0003 (0.0002)	0.0002 (0.0002)	-0.00006 (0.0002)
	All Suites	0.0107 (0.0090)	0.0112 (0.0091)	0.0097 (0.0123)
	Convention Center	0.0185** (0.0090)	0.0193** (0.0092)	-0.0263** (0.0113)
	Restaurant	0.0085 (0.0095)	0.0077 (0.0095)	0.0271*** (0.0100)
	Hotel tier controls?	Yes	Yes	Yes
	Hotel location controls?	Yes	Yes	Yes
	Chain-level fixed effects?	Yes	Yes	Yes
$OwnAf_{ij}$	Hotel is Independent	–	0.0053 (0.0135)	0.0079 (0.0119)
	Multi-unit owner	–	0.0053 (0.0067)	-0.0194** (0.0086)
Nei_{ij}	Has a neighbor	0.0205** (0.0096)	0.0304*** (0.0118)	-0.0121 (0.0119)
$NeiOwnAf_{ij}$	Has independent neighbor	–	0.0162* (0.0094)	-0.0071 (0.0099)
	Has multi-unit owner neighbor	–	-0.0253*** (0.0088)	-0.0018 (0.0097)
γ_j	City-level fixed effects?	YES	YES	YES
	Num. of observations	2931	2931	2931
	R-squared	0.06	0.06	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.

reviews) as linear in the self-promotion effort for the sake of simplicity, in reality the benefit is more likely to be concave in effort. That is, since a rating can't be higher than 5-stars, an increase in the number of manufactured positive reviews is likely to have diminishing marginal returns. Similarly, since a competitor's rating can't be lower than 1 star, an increase in the number of manufactured negative reviews is likely to have diminishing marginal returns.

4. We model the manipulation effort as costly to the firm. We can think of this cost as the cost of effort associated with writing fake reviews as well as 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, which may differ for different types of firms. We assume that the cost of writing reviews is a convex function of the effort. That is, compare the cost of writing the first manufactured review to the cost of writing the 30th review. While the first review can reflect the owner's own authentic writing style, the 30th review must be dissimilar from the reviews that preceded it in order to avoid detection. 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: $C(e_{i,i}, e_{i,j}) = \frac{\delta_i}{2}(e_{i,i})^2 + \frac{\gamma_i}{2}(e_{i,j})^2$. Here δ_i signifies the damage caused to the firm i if it caught doing self-promotion, and γ_i the damage if it is posting negative reviews for its competitor.
5. **Stage IV**: Finally, the consumer chooses the product that maximizes 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

$$E[q_A|s_A] - tx - p_A \geq E[q_B|s_B] - t(1-x) - p_B \quad (4)$$

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. The equilibrium concept here is Perfect Bayesian Nash Equilibrium.

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

firms' qualities are:

$$E[q_A|s_A] = s_A - \hat{e}_{A,A}^* + \hat{e}_{B,A}^* \quad (5)$$

$$E[q_B|s_B] = s_B - \hat{e}_{B,B}^* + \hat{e}_{A,B}^* \quad (6)$$

where $\hat{e}_{A,A}^*$ 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

$$\hat{x} = \frac{1}{2} + \frac{E[q_A|s_A] - E[q_B|s_B] + p_B - p_A}{2t} \quad (7)$$

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 **III**:

$$\Pi_{A,Stage\ 3}^* = \max_{e_{A,A}, e_{A,B}} \left(p_A E_{s_A, s_B, \varepsilon_A, \varepsilon_B} \left[\frac{1}{2} + \frac{E[q_A|s_A] - E[q_B|s_B] + p_B - p_A}{2t} \right] - \delta_A \frac{e_{A,A}^2}{2} - \gamma_A \frac{e_{A,B}^2}{2} \right) \quad (8)$$

$$\Pi_{B,Stage\ 3}^* = \max_{e_{B,B}, e_{B,A}} \left(p_B E_{s_A, s_B, \varepsilon_A, \varepsilon_B} \left[\frac{1}{2} + \frac{E[q_B|s_B] - E[q_A|s_A] + p_A - p_B}{2t} \right] - \delta_B \frac{e_{B,B}^2}{2} - \gamma_B \frac{e_{B,A}^2}{2} \right) \quad (9)$$

Substituting (5) and (6) into (8) and (9), and taking the expectation, we can re-write the firms' maximization problem as the following:

$$\Pi_{A,Stage\ 3}^* = \max_{e_{A,A}, e_{A,B}} \left(p_A \left[\frac{1}{2} + \frac{e_{A,A} + e_{A,B} - \hat{e}_{A,A}^* - \hat{e}_{A,B}^* + c_A + p_B - p_A}{2t} \right] - \delta_A \frac{e_{A,A}^2}{2} - \gamma_A \frac{e_{A,B}^2}{2} \right) \quad (10)$$

$$\Pi_{B,Stage\ 3}^* = \max_{e_{B,B}, e_{B,A}} \left(p_B \left[\frac{1}{2} - \frac{e_{B,B} + e_{B,A} - \hat{e}_{B,B}^* - \hat{e}_{B,A}^* + c_B + p_A - p_B}{2t} \right] - \delta_B \frac{e_{B,B}^2}{2} - \gamma_B \frac{e_{B,A}^2}{2} \right) \quad (11)$$

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 presents the optimal manipulation levels for the firms:

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

$$e_{A,A}^* = \frac{p_A}{2\delta_{At}}; e_{A,B}^* = \frac{p_A}{2\gamma_{At}} \quad (12)$$

$$e_{B,B}^* = \frac{p_B}{2\delta_{Bt}}; e_{B,A}^* = \frac{p_B}{2\gamma_{Bt}} \quad (13)$$

Proof. To solve for the optimal promotional levels, we 1) derive the first order conditions of firm A's profit function by differentiating Equation (10) with respect to $e_{A,A}$ and $e_{A,B}$ and by differentiating Equation (11) with respect to $e_{B,B}$ and $e_{B,A}$, and 2) simultaneously solve the system of the four resulting equations. This yields a unique solution since $\frac{\partial^2 \Pi_{A,Stage 3}^*}{\partial^2 e_{A,A}} < 0$, $\frac{\partial^2 \Pi_{A,Stage 3}^*}{\partial^2 e_{A,B}} < 0$, $\frac{\partial^2 \Pi_{B,Stage 3}^*}{\partial^2 e_{B,B}} < 0$ and $\frac{\partial^2 \Pi_{B,Stage 3}^*}{\partial^2 e_{B,A}} < 0$. \square

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) *An increase in the reputational costs of manipulation decreases 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.*

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.

In the Online Appendix we derive the comparative statics under endogenous prices.

References

- J. Altonji, T. Elder, and C. Taber. Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools. *Journal of Political Economy*, 113(1):151–184, 2005.
- 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.
- V. Bennett, L. Pierce, J. Snyder, and M. Toffel. Customer-Driven Misconduct: How Competition Corrupts Business Practices. *Management Science*, forthcoming.
- 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.
- C. Dellarocas and C. Wood. The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. *Management Science*, 54(3):460–476, 2007.
- 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.
- W. Dziuda. Strategic Argumentation. *Journal of Economic Theory*, 146(4):1362–1397, 2011.
- A. Harmon. Amazon glitch unmask 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 Journal of Economics*, 118(3):843–877, 2003.
- M. Ji and D. Weil. Does Ownership Structure Influence Regulatory Behavior? The Impact of Franchising on Labor Standards Compliance? Boston U. School of Management Research Paper No. 2010-21, 2009.
- G. Jin and P. Leslie. Reputational Incentives for Restaurant Hygiene. *American Economic Journal: Microeconomics*, 1(1):236–267, 2009.
- 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.
- X. Li and L. Hitt. Self-Selection and Information Role of Online Product Reviews. *Information Systems Research*, 19(4):456–474, 2008.
- M. Luca. Reviews, reputation, and revenue: The case of yelp.com. Harvard Business School Working Paper, 2012.
- 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.
- P. Nelson. Advertising as information. *Journal of Political Economy*, 82(4):729–754, 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.
- L. Pierce and J. Snyder. Ethical Spillovers in Firms: Evidence from Vehicle Emissions Testing. *Management Science*, 54(11):1891–1903, 2008.
- P. Resnick and R. Zeckhauser. Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay’s Reputation System. *The Economics of the Internet and E-Commerce*, 11, 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.
- Q. Ye, R. Law, B. Gu, and W. Chen. The influence of user-generated content on traveler behavior: An empirical investigation on the effects of e-word-of-mouth to hotel online bookings. *Computers in Human Behavior*, 27(2):634–639, 2010.