



9-1-2014

Understanding the Yelp Review Filter: An Exploratory Study

David Kamerer

Loyola University Chicago, dkamerer@luc.edu

Follow this and additional works at: https://ecommons.luc.edu/communication_facpubs



Part of the [Other Communication Commons](#)

Recommended Citation

Kamerer, D. (2014). Understanding the Yelp review filter: an exploratory study. *First Monday*, Vol. 19, No. 9. <http://dx.doi.org/10.5210/fm.v19i9.5436>.

This Article is brought to you for free and open access by the Faculty Publications at Loyola eCommons. It has been accepted for inclusion in School of Communication: Faculty Publications and Other Works by an authorized administrator of Loyola eCommons. For more information, please contact ecommons@luc.edu.



This work is licensed under a [Creative Commons Attribution-Noncommercial-No Derivative Works 3.0 License](#).

© First Monday, 2014.

© David Kamerer, 2014.



Understanding the Yelp review filter: An exploratory study by David Kamerer

Abstract

Reviews on Yelp.com can be an important factor in driving customers to a business. However, many business owners have expressed concern with Yelp's review filtering system, which was created to flag low-quality or fake reviews. This study performs a content analysis of a subset of Yelp restaurant and religious organization reviews, visible and filtered, exploring signals from the reviews or the reviewers that might explain the filtering process. The study finds that factors intrinsic to the review itself are not related to filtering, but factors related to the reviewer are strong predictors. The Yelp system is much more likely to filter reviews from occasional, isolated reviewers than from prolific, socially connected reviewers.

Contents

[Introduction](#)
[Economic theory applied to reviews](#)
[About Yelp](#)
[Complaints about Yelp](#)
[Yelp's position](#)
[Research questions](#)
[Method](#)
[Results](#)
[Discussion and conclusion](#)

Introduction

"Please help the business people from the Internet Mafia," said a business owner, complaining about Yelp.com to the U.S. Federal Trade Commission (FTC). "Yelp is the thug of the internet," wrote another. Yet another business owner complained that 42 out of 47 reviews of its business abruptly vanished (Kang, 2013). These were among almost 700 complaints against Yelp filed at the FTC.

Business owners know: online reviews are powerful. Research supports this view as well. Positive reviews on Yelp have been empirically shown to drive business success. One study found that a one-star increase in a restaurant's Yelp rating brought about a five-to-nine percent increase in revenue (Luca, 2011). Another study found a one-half star increase in a restaurant's Yelp rating was linked to selling out 49 percent more frequently (Anderson and Magruder, 2012).

It's not just bad reviews that upset business owners, however. Almost all reviews posted to Yelp are viewable online, but some — those that Yelp's software views as "not recommended" — are hidden from most viewers and are not counted towards the star rating. When a business finds a positive review that has been filtered, the owner is certain to take exception.

Yelp has steadfastly maintained that the review filter is an algorithm designed to keep fake reviews from hurting the integrity of its database. The company further says that no individual has the power to move reviews into or out of the filter. But Yelp's transparency — the company allows anyone to see the filtered reviews — has undermined trust. Business owners cry foul when positive reviews are filtered, and also when negative reviews are not. But there's also a more serious accusation: that Yelp is a pay-for-play scam, and that the company manipulates the review filter to reward people who buy advertising packages and to punish those who don't.

This view is widely held in the small business community, and has grown with class-action lawsuits against Yelp, media coverage from aggrieved business owners, and complaints against Yelp filed at the Federal Trade Commission.

This study examines the Yelp review filter to better understand how it works. The inquiry is prompted by the chasm between Yelp's stated position and the many vocal complaints from businesses about Yelp.

Economic theory applied to reviews

Online reviews are best understood through economic models. They are part of the Web 2.0, which is characterized by sites where consumers can publish content on socially-connected platforms. While many aspects of social publishing may be appealing to consumers, one key characteristic of this activity relevant to publishers is that it creates free content. This model, where status updates, tweets and blog posts create the content that builds the inventory of viewers, which is then monetized by ads, has fueled the growth of Facebook, Twitter, *Huffington Post* and many other sites. In this case, Yelp economically benefits from crowd-sourced reviews because the company does not have to pay its own staff to write reviews. Instead, it has created a community with social incentives, such as "friending" and review endorsements (these include "useful," "funny" and "cool"). Yelp's most active reviewers are granted "elite" status on the site. Members of the Yelp Elite Squad are invited to private events and receive rewards from Yelp to reinforce their behavior.

Consumers turn to reviews in part to reduce information asymmetry. This concept has been developed by economists to challenge the assumption of a rational market (Akerlof, 1970). Typically, sellers are information-rich while buyers lack information. Thus, "buyer behavior is crucially dependent upon the information that is available before and after purchase. In order to make choices, buyers need to at least know the price and quality of the various alternatives that they are considering" (Nayyar, 1990).

Lacking information about the best choice of business to patronize, consumers use many signals to reduce information asymmetry. Traditionally, this has included consulting owned media (such as a Web site or brochure) and paid media (advertising). For example, a restaurant may use words like "organic," "free-range" or "local" to communicate the type and quality of food offered to diners.

Digital ads, however, have limited effectiveness, particularly in getting people to respond to a call to action. While they may work as a branding tactic, industry-wide, banners receive a less than one-tenth of one percent click rate (Chapman, 2011). In a study conducted by Edelman for Adobe, 53 percent of respondents reported that they think "most marketing is a bunch of BS," 68 percent said that online advertising was annoying, and 54 percent said they didn't think that online banners worked (Adobe, 2012).

As social media have become mass media channels, peer-based or crowd-sourced messages such as online reviews have become an important force in shaping opinions. Edelman's (2012) annual Trust Barometer research has consistently shown that we're most likely to trust "a person like me." Crowd-sourced reviews have surfaced across product categories, including consumer products (Amazon), movies (Rotten Tomatoes) and hotels (Trip Advisor). One study found in a quasi-experimental comparison of identical books on Amazon and Barnes and Noble Web sites that more positive reviews for a book on one site were related to higher sales of that book on that site (Chevalier and Mayzlin, 2006). Reviews reduce information asymmetry between buyer and seller. As such, in cases where there is relatively little asymmetry, consumers value reviews less. For example, in one study, positive reviews were helpful for independent restaurants, but not for chain restaurants. This study also found that chain restaurants have lost market share as the online review ecosystem has grown in importance (Luca, 2011). Both of these findings speak to the value of reviews in bridging a competitive gap between chains, which may be familiar to consumers, and independent restaurants, which are likely not.

Filtered reviews on Yelp were found to be more extreme in rating (more likely to be one–star or five–star) than visible reviews (Luca and Zervas, 2013). The study also found that restaurants are more likely to commit review fraud when their reputations are weak (*i.e.*, they have few reviews or poor reviews), and that restaurants are more likely to leave fake reviews for competition when they are in a highly competitive marketplace.

Another economic concept related to online reviews is signal cost. Sometimes the cost is low, such as when anyone can write a review (this is true on Amazon and Yelp, for example). However, a review posted on Open Table, where you must have made reservations in order to be eligible to write a review, has a higher signal cost. Similarly, Expedia Travel has a “paid and stayed” requirement for reviewers, imposing a higher signal cost. On Angie’s list, where you must pay for a subscription, the signal cost for all reviews is rolled into an annual membership fee. Sites with lower signal costs have been found to be more likely to host corrupt or fake reviews (Ott, *et al.*, 2012). Sites with higher signal costs reduce the overall number of reviews written, trading the lower volume for higher levels of trust.

Another economic concept applied to reviews is the exposure benefit. All things equal, a site with more viewers has a higher exposure benefit. Thus, sites with low signal costs and high exposure benefits are most likely to have corrupt reviews. In a study of six different travel review sites, the authors observed that a higher signal cost did reduce the number of corrupt reviews (Ott, *et al.*, 2012).

Yelp reviews feature low signal cost (anyone can create a free account and write a review) with high exposure benefit (Yelp has the largest audience of all review sites). Based on these two factors, there is likely to be a high proportion of corrupt reviews on Yelp, relative to other sites.

About Yelp

Yelp is the largest standalone review site online, home of more than 40 million reviews and garnering more than 100 million unique visitors per month (Stoppelman, 2013).

A business may be completely passive in managing its presence on Yelp, or may take a more active role. A third option is to also advertise on Yelp.

At the simplest level, the business can learn from the reviews to improve its performance.

The business owner can also claim its business. This allows the owner to annotate the page by adding photos, information on location, hours of operation and other relevant details. The business owner also gains access to analytics to track views on the site. Additionally, the owner can respond to reviews or create Yelp Deals or a mobile check–in offer to encourage customer visits.

The business can also advertise on Yelp. Fundamentally, Yelp is an advertising company and it monetizes its operations by selling its inventory of reviewers and readers to advertisers. Advertising packages start at US\$300/month, and are based upon a guaranteed number of impressions. Advertisers can upload slide shows and a video to their Yelp page and can add a “call to action” button (such as “reserve now”). Ads are served on relevant searches by location and category, appearing in competitive searches. Yelp also removes competitor ads removed from the pages of advertisers.

The ad programs deliver poor value, according to online marketer Rakesh Agrawal (2012). Yelp’s advertising charges for impressions rather than clicks. Across the Internet, impressions vary in cost but average about US\$2.52 per 1000 impressions, or CPM (Lipsman, 2010). Because of the abundance of online ad outlets, CPMs are often well below US\$1. According to Agrawal (2012), “At a time when much online advertising is being sold for 60 cents per thousand impressions (CPMs), Yelp is charging some local advertisers \$600 per 1,000 impressions ... At the high end, it’s a \$600 CPM. At the low end, that’s a still eye–popping \$367 CPM — more than 10–times the rate of a Super Bowl ad.” Yelp ad packages also require a three, six or 12–month contract, forcing the business to continue the program even if it is not effective.

Complaints about Yelp

Many business owners resent Yelp as an advertising platform. They may not like to pay for ads when they can’t control how people rate their business or what ads may appear on their page. And they may not like dealing with pushy ad salespeople.

But the strongest emotions are reserved for Yelp’s review filter. The filter controls which ads are easily viewed and counted towards the overall star rating. Unlike other review sites, Yelp makes it possible to see “suspect” reviews (to view filtered reviews, scroll to the bottom of a business’ page. Click on the light gray link that says “... reviews that are not currently recommended”). When a business owner sees a positive review that has been filtered, he or she wants to have it made visible. And the business owner may also wish to hide negative reviews in the filter.

Many small business owners claim that Yelp uses the review filter to reward advertisers and to punish everyone else. Many of these complaints allege extortion. They claim that a Yelp salesperson will make promises of moving positive reviews out of the filter when a business commits to a Yelp advertising plan. Alternately, say the complainers, the salesperson may threaten to move negative reviews out of the filter for a business that declines the plans.

These are dominant themes in complaints filed against Yelp at the Federal Trade Commission. These complaints were obtained in a Freedom of Information request by MuckRock.com (2012) (<https://www.muckrock.com/>), an investigative Web site. Forbes (2013) summarizes some of these complaints [1]:

“Consumer states that [Yelp] told him to pay 300.00 for a ‘membership’ that would give him more positive reviews. Consumer noticed that a couple of positive reviews popped up immediately after paying the money.”

“Although the site cites that disclosing those [filtering] criteria will allow circumventing their built–in checks and balances, in practice it gives them a complete power to manipulate the ratings of any business.”

“The site constantly and relentlessly filters out the good and excellent reviews displaying less positive reviews causing the rating to downgrade.”

Over four years, almost 700 businesses complained about Yelp to the FTC. In addition to the extortion charges, there are complaints about the cost of advertising packages, high–pressure sales people, listing of home addresses and phone numbers in the database, out–of–date listings, poor customer service and refusal to remove reviews, even when the facts were proven to be false.

These complaints have also been the subject of journalistic investigations. In an article entitled “Yelp and the business of extortion 2.0,” The East Bay (California) *Express* interviewed six business owners who claimed that their “Yelp sales representatives promised to move or remove negative reviews if their business would advertise. In another six instances, positive reviews disappeared — or negative ones appeared — after owners declined to advertise.” The article also noted “And in at least one documented instance, a business owner who refused to advertise subsequently received a negative review from a Yelp employee” (Richards, 2010).

The claim of extortion has also been taken up by other journalists and has even been featured as a storyline on the television show “The People’s Court” (Jones, 2010; Keeling, 2013).

The database of complaints against Yelp at the FTC also includes complaints of insistent ad sales calls, apparent violations of the Federal Do–Not–Call laws.

Yelp has also become the subject of at least three class action lawsuits. Says the complaint in one, “Yelp, however, regularly manipulates the content on Yelp.com listing pages, despite Yelp’s mantra of ‘Real people. Real reviews.’ As a result, business listings on Yelp.com are in fact biased in favor of businesses that buy Yelp advertising ... As part of Yelp’s regular practices, the company asks business owners to pay for ‘protection’ from bad reviews (in the form of advertising dollars) while Yelp controls whether bad reviews are posted in the first place — the classic scheme of offering ‘protection’ from a problem that the ‘protector’ himself creates” (*Cats and Dogs Animal Hospital, et al. v. Yelp! Inc.*, 2010).

Further evidence of anti–Yelp sentiment from businesses is manifest in Web sites, such as Yelp–sucks.com, the “Yelp survival guide,” at BOSSHi.com, and Facebook groups such as “I hate Yelp” (Yelp–Sucks, 2014; BOSSHi, 2014; I Hate Yelp, 2014).

While it’s easy to see a trust deficit surrounding Yelp’s practices, thus far Yelp has not lost a major legal challenge (Chang, 2011). This is in part due to the expansive treatment online businesses receive under Section 230 of the Communications Decency Act, which generally absolves businesses for liability from content posted by third parties (Title 47, Section 230). But Yelp’s critics have also failed to prove their cases.

Yelp's position

When a site contains 40 million reviews, a few fake reviews are bound to get through. These might include a merchant writing a positive review of his or her own business, a merchant writing a negative review of a competitor's business, or a disgruntled employee/ex-spouse/troll writing a negative review of almost anything.

Yelp is also threatened by a cottage industry of reviewers-for-hire, doing business online at places like Craigslist, Fiverr.com and through the Amazon Mechanical Turk marketplace.

Yelp has a duty to protect the integrity of its database, as do other review sites. But Yelp has drawn a disproportionate amount of criticism for its approach to filtering reviews. By comparison, Amazon, another large review database, reserves the right to remove reviews in four different categories: objectionable material, promotional content, inappropriate content and off-topic information (Amazon, 2014). You could argue that Yelp's practice is more transparent because if you want, you can read the suspect reviews and decide for yourself.

Yelp relies on algorithm-derived software to sort the reviews into live or filtered. The company has repeatedly stated that it does not manually filter reviews. On a Reddit AMA ("ask me anything") CEO Jeremy Stoppelman said, "There has never been any amount of money you could pay us to manipulate reviews. We do have an algorithm that highlights the most useful and reliable reviews on our site which is about 75 percent of contributed content" (Wasserman, 2013). Wrote Yelp spokesperson Vince Sollitto: "Some news outlets have recently run stories or columns re-hashing the sensational allegation that Yelp manipulates reviews and ratings to reward advertisers or punish non-advertisers. Let me be clear: This claim is not — and has never been — true" (Fell, 2013).

Yelp's explanation of filtering is shown to consumers every time they request a review that's not recommended (filtered). To see this video, scroll to the bottom of the first page for any business that's on Yelp. When you click on the grey "reviews that are not currently recommended" the video is on the top of the next page served. In part, the voiceover says:

"Every Yelp review is automatically evaluated by Yelp's recommendation software, based up quality, reliability and user activity on Yelp. More often than not, those useful reviews come from active members of the Yelp community. Our recommendation software is looking for reviews written by people whose opinions, positive or negative, we feel are helpful and reliable to consumers. Currently, about 75 percent of all reviews are recommended. We try not to highlight reviews written by users we don't know much about, or those that might be fakes, or unhelpful rants or raves" (YouTube, 2013).

Research questions

This study investigates how the Yelp review filter functions in a small subset of reviews on the site. The study analyzes how the filter functions at the reviewer level. The larger question is: does Yelp filter the business or the reviewer? Most complaints against Yelp assume that Yelp filters the business. For example, a critic may believe that a business that declined an advertising package might see negative reviews moved from the filter to the live database of reviews. The first question, does Yelp filter the business, is difficult to answer because it requires information about whether the business is a Yelp advertiser, and then time-based analysis based upon when the business began advertising or declined an advertising offer. But it's much simpler to measure whether Yelp filters the reviewer. Essentially, this is Yelp's position. Of the two questions, it's more easily tested. Many of the signals about the reviewer are posted on the Yelp Web site. Some, such as the reviewer's IP address, are not. But there is plenty of information about reviewer activity and whether reviews are filtered, so we will begin there.

We will investigate two research questions:

1. Are the review-based signals related to the review being filtered?
2. Are the reviewer-based signals related to having the review being filtered?

Method

Data for this study were collected via content analysis of reviews written in two categories: restaurants located on Chicago's north side and religious organizations across Chicago. The grouping of restaurants was chosen for having a low concentration of chain restaurants; previous research has shown that reviews are less important for chain restaurants because consumers often know what to expect before dining at one. The religious organizations category was chosen for contrast; additionally, the review volume in this category is much lower than it is for restaurants.

The restaurants sample was selected in January 2014, and was created in two stages. First, researchers established a sampling frame of 809 restaurants in the Yelp database that met the criteria of being restaurants on Chicago's north side. The sample was created using a skip interval of five and a random start point created by the Web site random.org (<http://www.random.org/>), yielding 162 eligible restaurants. Once the restaurants were selected, the total number of reviews for each restaurant was tabulated. This number was created by adding all visible reviews to all filtered reviews. Then one review was selected from this restaurant using a number generated by the random.org site.

This scheme was adopted as a relatively random way to gain access to the database of reviews. It should be noted that, since one review was pulled from each restaurant in the sample, that heavily reviewed or popular restaurants are under-represented in the study. Because we're not interested in reviews at the restaurant level (and instead interested in the reviewer), this is only a minor sampling limitation.

The religious organizations sample was drawn in August 2014. Because there are fewer reviews in this category, the list was sorted from most- to least-reviewed. The sample was built from the top down, using random.org to choose one review from each organization. For the religious organizations sample, $n=128$.

The variables related to the review in question were coded as follows:

- Word count of review (each review was pasted into Microsoft Word, which calculated the number of words)
- Number of stars of review (1–5 stars)
- Review filtered? (yes/no)

The variables related to the reviewer were coded as follows:

- Yelper name
- Number of Yelp friends
- Number of Yelp reviews by reviewer
- Avatar picture provided (yes/no)
- Number of fans of Yelper
- Number of reviews written by Yelper at each star level
- Number of "useful," "funny," and "cool" votes for Yelper

After the sample was created, preliminary analysis showed a low percentage of filtered reviews (nine percent). In order to increase the incidence of filtered reviews, researchers revisited the sample and added cases likely to be filtered. This was done by calculating the 93rd percentile review for each organization in the sample and coding that review. The resulting filtered reviews were then added to the database, yielding a total sample of 240 reviews. The resulting database, then, is weighted to include a disproportionate percentage of filtered reviews. While Yelp reports about 25 percent of all reviews to be filtered, in this study the percentage is 35 percent for restaurants and 38 percent for religious organizations.

The principal investigator contributed to the database and trained two additional coders. Because each coder only reported manifest content (such as writing down the number of stars of a review or the number of reviews written by an individual, no intercoder reliability check was needed. In essence, no judgments were recorded; no interpretations were made. The coders simply transferred numbers and text from the Yelp database to the study database.

These data were keypunched into Microsoft Excel and then imported to IBM SPSS 20 for analysis.

Table 5 shows the content variables in the restaurant review sample. Not one variable is significantly related to filtering. This analysis is repeated with reviews of religious organizations in Table 6 with the same result; there is no relationship between the content variables and whether reviews were filtered.

Table 5: Content variables and filtering of restaurant reviews.					
variable	visible	filtered	t	df	sig
Negative emotions	1.34	0.91	0.876	221	0.382
Exclusion	2.80	2.89	-0.251	221	0.382
Motion	1.76	1.76	-0.039	221	0.969
See	0.73	0.50	1.642	221	0.102
Hear	0.43	0.49	-0.418	221	0.676
Feel	0.60	0.39	1.603	221	0.110

Table 6: Content variables and filtering of religious organization reviews.					
variable	visible	filtered	t	df	sig
Negative emotions	0.71	0.65	0.347	125	0.729
Exclusion	2.35	1.96	-0.251	221	0.382
Motion	1.76	1.76	1.158	125	0.249
See	1.06	1.06	-0.01	125	0.992
Hear	0.66	0.8	-0.765	125	0.446
Feel	0.49	0.72	-1.308	125	0.193

To further explore the differences between the words used in filtered versus visible reviews, the review database was sorted from low to high. This approach might allow greater insight into dominant words used in extremely negative or extremely positive reviews, both places where a fake review is more likely to be found. Figure 5 presents all one-star reviews in the restaurant database that were filtered, while Figure 6 presents all one-star restaurant reviews that were not filtered.




Figure 5: Word cloud for all one-star restaurant reviews that were filtered.



Figure 6: Word cloud for all one–star restaurant reviews that were visible.

While the two word clouds are similar, words about food items, such as “noodles,” “oil,” “gluten,” and “pasta,” dominate the filtered reviews. In the visible one–star reviews, service–related words, such as “manager,” “served,” “order,” and “minutes” are more prominent. Again, trigger words, such as “loved” or “hated,” are not prominent.

Figure 7 presents a word cloud of all filtered five–star reviews in the religious organizations database, while Figure 8 presents all visible five–star reviews for religious organizations. Again, the word cloud analyses fail to discriminate the differences between filtered and unfiltered reviews. If there are any strong rhetorical differences between the two groups of words, they are not prominent and are buried inside the word clouds.



Figure 7: Word cloud for all religious organization filtered five–star reviews.



Figure 8: Word cloud for all religious organization visible five–star reviews.

Summarizing qualitative and quantitative analyses of review content, we have found no explanations to explain review filtering, other than that visible reviews are generally longer than filtered reviews.

Next, turning to reviewer–based variables, Table 7 shows the signals related to the ethos of the people who wrote the reviews in the sample. This analysis shows a clear demarcation between reviewers whose writing is filtered versus those whose isn’t. People who write reviews that Yelp filters have fewer network connections than those whose reviews are visible. In the restaurant review database, the average number of friends for those who write visible reviews is 112, versus seven for those whose reviews are filtered; in the religious organizations database, people who wrote visible reviews have an average of 249 friends while people who wrote filtered reviews have an average of one friend.

Table 7: Reviewer signals on Yelp.			
	Restaurant	Reviewers	Religious organizations
	Filtered	Visible	Filtered
			Reviewers
			Visible

Friends	7	112	1	249
Reviews	10	210	11	247
Avatar	35%	93%	37%	87%
Fans	0	18	0	32
Useful	6	644	7	1,106
Funny	2	402	3	809
Cool	2	458	3	823
N	85	147	49	79

People whose reviews are visible are more prolific on Yelp; in the restaurant database the average number of reviews written for this group is 210, while the group whose reviews are filtered is only 10. In the religious organizations database, people who authored visible reviews averaged 249 reviews, while people who wrote filtered reviews averaged 11.

Even something as simple as posting an avatar picture makes a difference in how your content is treated by Yelp: for those whose work is visible, 93 percent of restaurant reviewers have posted an avatar photo; for those whose work is filtered, only 35 percent have posted an avatar. In the religious organizations database, 87 percent of the visible reviews were accompanied by an avatar picture, while only 37 percent of the filtered reviews had an avatar. Other signals from the Yelp community also appear to be related to having your work visible: these include having fans and having your work being voted "useful," "funny," or "cool." Visible reviews enjoy many orders of magnitude more social feedback than those that are filtered.

To summarize, people whose reviews are live in the database are prolific reviewers and have built many social signals into their Yelp personae. Generally, they are known and trusted by the community. People whose reviews are filtered have written few reviews and are seen by the Yelp community as isolated and unknown. These results are consistent across two very different categories of reviews on Yelp.

Discussion and conclusion

In this study we have attempted, in a limited manner, to test the question: does Yelp's review filter function as Yelp says it does?

When looking at evidence related to the actual reviews, there appears to be only minor differences between reviews that are filtered and those that are presented in the live database. While we would expect that filtered reviews are likely to be short, in this study they were only a little bit shorter than visible reviews (by an average of 20 words). And, while previous research found that filtered reviews are more extreme, more likely to be one or five stars, in this study, filtered reviews average 3.5 stars and non-filtered reviews averaged 3.6 stars in the restaurant database and 4.5 stars and 4.7 stars respectively in the religious organizations database — fairly similar star ratings (Luca and Zervas, 2013). These moderate-to-high averages share similar standard deviations so the measures of central tendency are not masking some unseen extreme behavior in the data.

Further, the word clouds generated by filtered reviews lack any "trigger" words that might betray these reviews as fake. This is seen in all comparisons, whether we look at all reviews, or only one- or five-star reviews. Simply put, the word clouds generated by filtered reviews look very similar to word clouds made by visible reviews.

Quantitative analysis using the LIWC software also failed to discriminate between filtered and visible reviews. If there are differences between filtered and visible reviews, we have yet to find the variables that explain them.

This evidence suggests that the actual content of a review is a minor factor in whether a review is filtered by Yelp.

Moving to social signals that connect the reviewer to the Yelp community, we find strong support for filtering based on the online reputation of the reviewer. People who are strongly linked to the Yelp community through review creation, social interaction and endorsements are unlikely to have their reviews filtered. Individuals who don't write much, aren't endorsed and who don't fill out their user profiles are much more likely to have their reviews filtered.


With these results in mind, it becomes easier to understand why some reviews are filtered: the owner of a new restaurant asks her friends to write some positive reviews to boost business. A disgruntled ex-employee wants to hurt his former employer so he writes a negative review. Or someone writes a negative review of a business based upon the business owner's political orientation. In all of these cases, the writer is not likely to be a trusted member of the Yelp community, so these reviews are largely filtered.

Does Yelp make the wrong call on the filtering of some reviews? Undoubtedly. The size of the review database makes this a certainty. Humans, after all, are still smarter than algorithms. But the wrong calls that are certain to exist must be considered in a proper context. To Yelp, a small fraction of reviews wrongly filtered is an acceptable error. But to a business owner, each review matters very much.

Does Yelp selectively filter reviews based upon the business they're written for? That question is beyond the scope of this study. The data here support the notion that prolific, socially connected reviewers have their reviews published, while people who don't build their network on Yelp and don't write many reviews are much more likely to have them filtered. A look at the review database from the level of an individual business, over time, may yield different results.

Sentiment against Yelp is strong among the business community, and some of that could be mitigated through better customer service. In the complaints to the FTC, there are many frivolous claims. But some appear to be valid. Yelp's inbound communication is poor, and the company appears to be unresponsive about changing out-of-date locations, hours, closed businesses and other time-based issues, such as when a restaurant keeps its name but changes management. In this case, should old reviews be removed? To a business, these are important issues. Ultimately, Yelp must address this trust deficit if the company is to keep its dominant position in review space.

Businesses can lessen their dependence on Yelp reviews by practicing good public relations tactics: by becoming a high-performance company, so the reviews will trend positively, and being responsive to complaints. The Yelp system does allow businesses to reach out to unhappy reviewers, but this should be done with care. Messaging strategies for businesses include building a strong Web site with good search engine optimization. This increases the likelihood that owned media would be the first thing the customer sees. Businesses should also claim their space on other relevant review sites, such as those offered by Google and Bing, and use those to drive traffic to the owned media.

The review space is fast moving and Yelp could change its algorithm in a moment, without notice. For this reason alone, it would be worthwhile to replicate this study. Future studies should work with larger samples, across locations and business categories. Because Yelp is an advertising company, it would be useful to take a direct look at how advertisers and non-advertisers are treated by the company. Indeed, the company should welcome such a study. 

About the author

David Kamerer serves as assistant professor in the School of Communication at Loyola University Chicago.
E-mail: david [at] davidkamerer [dot] com

Acknowledgements

The author would like to acknowledge the assistance of Rebecca Browne and Anna Sherman in developing the database used for this study.

Notes

1. Complaints are shown with original spelling, grammar and punctuation.

2. <http://www.liwc.net/>.

References

- Adobe, 2012. "U.S. study reveals online marketing is failing with consumers" (24 October), at <http://www.adobe.com/aboutadobe/pressroom/pressreleases/201210/102412AdobeAdvertisingResearch.html>, accessed 7 July 2014.
- Rakesh Agrawal, 2012. "Yelp advertising is a rip-off for small advertisers" (6 February), at <http://venturebeat.com/2012/02/06/yelp-advertising-is-a-rip-off-for-small-advertisers/>, accessed 7 July 2014.
- George A. Akerlof, 1970. "The market for 'lemons': Quality uncertainty and the market mechanism," *Quarterly Journal of Economics*, volume 84, number 3, pp. 488–500.
doi: <http://dx.doi.org/10.2307/1879431>, accessed 22 August 2014.
- Amazon, 2014. "General review creation guidelines," at <http://www.amazon.com/gp/community-help/customer-reviews-guidelines>, accessed 7 July 2014.
- Michael Anderson and Jeremy Magruder, 2012. "Learning from the crowd: Regression discontinuity estimates of the effects of an online review database," *Economic Journal*, volume 122, number 563, pp. 957–989.
doi: <http://dx.doi.org/10.1111/j.1468-0297.2012.02512.x>, accessed 22 August 2014.
- BOSSHI, 2014. "Why yelp sucks — Yelp survival guide," at <http://www.bosshi.com/why-yelp-sucks/2014>, accessed 7 July 2014.
- Cats and Dogs Animal Hospital, et al. v. Yelp! Inc.*, 2010. "U.S. District Court, Central District of California, Case number CV10–1340 VBF (SSX)" (22 February), at <http://www.dmlp.org/sites/citmedialaw.org/files/2010-02-23-Cats%20and%20Dogs%20Animal%20Hospital%20Complaint.pdf>, accessed 7 July 2014.
- Andrea Chang, 2011. "Yelp wins dismissal of class-action lawsuits," *Los Angeles Times* (26 October), at <http://latimesblogs.latimes.com/technology/2011/10/class-action-lawsuits-against-yelp-dismissed.html>, accessed 7 July 2014.
- Mike Chapman, 2011. "What clicks worldwide? Rich media can help lift click-through — but only a little," *Adweek* (31 May), at <http://www.adweek.com/news/advertising-branding/what-clicks-worldwide-132085>, accessed 7 July 2014.
- Judith A. Chevalier and Dina Mayzlin, 2006. "The effect of word of mouth on sales: Online book reviews," *Journal of Marketing Research*, volume 43, number 3, pp. 345–354.
doi: <http://dx.doi.org/10.1509/jmkr.43.3.345>, accessed 22 August 2014.
- Edelman, 2012. "Credibility of government officials, CEOs plummets," at <http://www.edelman.com/insights/intellectual-property/2012-edelman-trust-barometer/the-state-of-trust/credibility-of-government-officials-ceos-plummets/>, accessed 7 July 2014.
- Jason Fell, 2013. "Yelp continues to defend against claims of review manipulation," *Entrepreneur* (24 May), at <http://www.entrepreneur.com/article/226832>, accessed 7 July 2014.
- Paula Forbes, 2013. "FTC complaints about Yelp allege extortion, libel, more," *Eater* (23 January), at <http://eater.com/archives/2013/01/23/ftc-complaints-about-yelp-allege-extortion-libel-more.php>, accessed 7 July 2014.
- I Hate Yelp, 2014. "I Hate Yelp community," at <https://www.facebook.com/IHateYelpcom>, accessed 7 July 2014.
- Ashby Jones, 2010. "Real people. Real reviews. Real extortion scheme?" *Wall Street Journal, Law Blog*, at <http://blogs.wsj.com/law/2010/02/26/real-people-real-reviews-real-extortion-scheme/>, accessed 7 July 2014.
- Inkoo Kang, 2013. "Businesses: 'Yelp is the thug of the Internet'," *MuckRock* (23 January), at <https://www.muckrock.com/news/archives/2013/jan/23/businesses-yelp-thug-of-the-internet/>, accessed 7 July 2014.
- Brock Keeling, 2013. "Video: Defendant alleges Yelp extortion on 'The People's Court'" (22 July), at http://sfist.com/2013/07/22/video_angry_defendant_bashes_yelps.php, accessed 7 July 2014.
- Linguistic Inquiry and Word Count, at <http://www.liwc.net>, accessed 11 August 2014.
- Andrew Lipsman, 2010. "The *New York Times* Ranks as top online newspaper according to May 2010 U.S. comScore media matrix data," *comScore* (16 June), at http://www.comscore.com/Insights/Press_Releases/2010/6/The_New_York_Times_Ranks_as_Top_Online_Newspaper_According_to_May_2010_U.S._comScore_Media_Matrix_Data, accessed 7 July 2014.
- Michael Luca, 2011. "Reviews, reputation, and revenue: The case of Yelp.com," *Harvard Business School, Working Paper*, number 12–016 (4 October), at <http://hbswk.hbs.edu/item/6833.html>, accessed 7 July 2014.
- Michael Luca and Georgios Zervas, 2013. "Fake it till you make it: Reputation, competition, and Yelp review fraud," *Harvard Business School, Working Paper*, number 14–006 (17 September), at <http://people.hbs.edu/mluca/fakeitillyoumakeit.pdf>, accessed 7 July 2014.
- MuckRock, 2012. "FOI request: FTC complaints for www.yelp.com," at <https://www.muckrock.com/foi/united-states-of-america-10/ftc-complaints-for-wwwyelpcom-1645/>, accessed 7 July 2014.
- Praveen R. Nayyar, 1990. "Information asymmetries: A source of competitive advantage for diversified service firms," *Strategic Management Journal*, volume 11, number 7, pp. 513–519.
doi: <http://dx.doi.org/10.1002/smj.4250110703>, accessed 22 August 2014.
- Myle Ott, Claire Cardie, and Jeff Hancock, 2012. "Estimating the prevalence of deception in online review communities," *WWW '12: Proceedings of the 21st International Conference on World Wide Web*, pp. 201–210.
doi: <http://dx.doi.org/10.1145/2187836.2187864>, accessed 22 August 2014.
- Katherine Richards, 2010. "Yelp and the business of extortion 2.0," *East Bay Express* (18 February), at <http://www.eastbayexpress.com/oakland/yelp-and-the-business-of-extortion-20/Content?oid=1176635>, accessed on July 7 2014.
- Jeremy Stoppelman, 2013. "Yelp.com welcomes 100 million unique visitors in January 2013" (6 February), at <http://officialblog.yelp.com/2013/02/yelpcom-welcomes-100-million-unique-visitors-in-january-2013.html>, accessed 7 July 2014.
- Yla R. Tausczik and James W. Pennebaker, 2010. "The psychological meaning of words: LIWC and computerized text analysis methods," *Journal of Language and Social Psychology*, volume 29, number 1, pp. 24–54.
doi: <http://dx.doi.org/10.1177/0261927X09351676>, accessed 22 August 2014.
- U.S. Code, "Title 47 (Telecommunications), Section 230," at <http://www.gpo.gov/fdsys/pkg/USCODE-2011-title47/pdf/USCODE-2011-title47-chap5-subchap11-part1-sec230.pdf>, accessed 7 July 2014.
- Todd Wasserman, 2013. "Yelp CEO Jeremy Stoppelman takes on critics in freewheeling Reddit AMA" (8 November), at <http://mashable.com/2013/11/08/yelp-jeremy-stoppelman-reddit-ama/>, accessed on 28 May 2014.
- Yelp-Sucks.com, 2014. At Yelp-Sucks.com, accessed 7 July 2014.
- YouTube, 2013. "Why does Yelp recommend reviews?" (13 November), at <https://www.youtube.com/watch?v=PniMEnM89iY>, accessed 7 July 2014.

Editorial history

Received 7 July 2014; revised 16 August 2014; accepted 20 August 2014.

Understanding the Yelp review filter: An exploratory study
by David Kamerer.
First Monday, Volume 19, Number 9 - 1 September 2014
<http://firstmonday.org/ojs/index.php/fm/rt/prINTERfriendly/5436/4111>
doi: <http://dx.doi.org/10.5210/fm.v19i9.5436>