An Evaluation of Optimal Anonymity in Online Review Platforms

Ji Hong Park
Faculty Advisor: Stanislav Rabinovich
April 23, 2017

Submitted to the Department of Economics at Amherst College in partial fulfillment of the requirements for the degree of Bachelor of Arts with honors.

Abstract

Biases in online review systems make buyers discount information from reviews, reducing the informativeness of reviews. My thesis analyzes the implications of these biases for the optimal choice of reviewer anonymity, and shows how the optimal choice of anonymity depends on market and product characteristics. Using a Bayesian learning model, I find that markets with more company sellers minimize aggregate bias by adopting more anonymous systems and markets with more individual sellers do so by adopting more non-anonymous systems. My predictions are consistent with the choice of anonymity in existing review platforms.

Keywords: Anonymity, Bayesian Learning, Information Economics, Online Review Systems, Reviews.

Acknowledgements

I would like to thank Professor Rabinovich for being a wonderful advisor for the entire thesis-writing process. His encouragement helped me gain confidence in my research, and I enjoyed the project more than anything else. I would like to thank Professor Ishii for introducing me to the economics of Industrial Organization and advising me to find a topic that I was genuinely interested in. I would also like to thank Professor Kingston and Professor Reyes for being awesome mentors and teachers during my time at Amherst. I am glad I studied economics, and I am excited to continue research on law and economics in law school.

To my parents, thank you for your love and support. To Tomo and Muling, I cherish the moments we spent studying together and laughing at economics jokes. To Alena, Chico, Cody, Jamie, and Megan, thank you all for being the best group of thesis writers!

Table of Contents

1. Introduction **1**

2. Background **4**

2.1 Literature Review 4
2.2 Comparison of Review Systems in Online Review Platforms 8

3. Model **13**

3.1 Model Setup 13
3.2 Updating and the Likelihood function 17
3.3 Baseline Model with No Bias 19
3.4 Model of Review System with Fake Reviews and No Review Inflation 20
3.5 Model of Review System with Review Inflation and No Fake Reviews 24
3.6 Model of Review System with both Fake Reviews and Review Inflation 26

4. Main Results **27**

4.1 The Optimal Level of Anonymity and Comparative Statics 27
4.1.1 Special case: $\mu_G = 0, \pi_G = 1, \pi_B = 0$ 29
4.1.2 Special case: $\mu = \mu_G = \mu_B$ 31
4.2 Welfare Implications 33

5. Discussion **34**

5.1 Real World Characteristics that correspond to differences in $\boldsymbol{\phi}$ 35
5.2 Real World Characteristics that correspond to differences in $\boldsymbol{\mu}$ 40

6. Conclusion **42**

7. Appendix **44**

7.1 Derivation of Trade-Offs between $\boldsymbol{\gamma}$ and $\boldsymbol{\sigma}$ 44
7.2 Derivation of Comparative Statics 45

8. Bibliography **47**

1. Introduction

In 2014, Airbnb changed policies regarding its review system, allowing both hosts and guests to see reviews of one another only after both parties completed reviews.[1] The change in Airbnb's policy reflects its concern about the bias that results when hosts or guests are allowed to see the other party's review before leaving theirs. The bias - the review inflation resulting from a reviewer's fear of retaliation for a negative review - is caused by a lack of anonymity in Airbnb's review system. Airbnb attempted to mitigate the costs of a more non-anonymous review system, the inflation of reviews, which lowers the informational value of reviews. This story illustrates a concern with minimizing bias in reviews caused by differing degree of reviewer anonymity. In this thesis, I focus on this aspect of review system design: the degree of anonymity granted to reviewers. I provide a theoretical framework for evaluating the costs and benefits of anonymity, and link the resulting optimal degree of anonymity to market and product characteristics.

User reviews in online platforms play an important role for both buyers and sellers. Reviews update a buyer's prior belief on the type of the seller and the quality of goods or services he or she will experience through a transaction with the seller. As reviews accumulate, they form the seller's reputation and distinguish bad sellers from good sellers. For instance, travelers (buyers) can identify Airbnb hosts (sellers) that are slow at communicating with visitors or provide sub-standard accommodations from reading reviews written by previous guests. An effective review system should enable buyers to begin or to continue transactions with sellers who produce good quality goods and also sort out the bad sellers, increasing market efficiency. This is done through facilitating the information

[1] Airbnb is a two-way review platform, in which both buyers and sellers leave reviews of the other party. Before 2014, buyers or sellers, whoever wrote reviews after the other party have left the reviews, could see those reviews before writing their own.

sharing between buyers, which allow buyers to overcome information asymmetry with sellers through reviews.[2] Reviews play an especially important role in peer-to-peer markets, such as markets for ride sharing or house sharing, where "large numbers of fragmented buyers and sellers" exchange goods or services (Einav 2015). Whereas sellers in traditional markets have reputation built on brand image and advertisement, sellers in peer-to-peer markets are relatively less established, making it difficult for buyers to infer the reputation of sellers. Sellers depend on reviews to build their reputation, since traditional marketing methods such as word-of-mouth marketing or mass media advertisement are not as effective when the goods are produced in smaller scale locally (for instance, Airbnb hosts who have multiple listings usually have them in the same city).[3]

Existing review platforms vary widely in the degree of anonymity granted to reviewers, as I will describe in more detail in section 2.2. Airbnb's review system is the least anonymous system, since the identity of the reviewer is visible to both the sellers and other buyers who read reviews. Angie's List's review system is more anonymous than that of Airbnb, as it protects the reviewer's identity from other buyers, but not from sellers. Finally, Opentable adopts the most anonymous review system, publishing all reviews under the name "Opentable Diner."

Consider a review system designer who seeks to minimize the magnitude of aggregate bias present in reviews, and thereby to maximize the informativeness of reviews. Previous literature has shown that the review system is potentially subject to two types of bias. *Fake reviews* occur when reviews are not written by actual buyers. The fake reviews

[2] Informational asymmetry refers to a situation where one party of a transaction has more information about the quality of the good than the other party.
[3] When products or services are produced on a small scale, it is difficult to verify the quality of products or services unless buyers leave some kind of feedback regarding their experience.

can be either positive (written by an owner of a business) or negative (written by a competitor). *Review inflation* occurs when buyers are compelled to leave positive reviews of sellers despite having negative experiences. The two problems occur to a different degree in anonymous and non-anonymous review systems, as I will describe in the following paragraph. Therefore, a review system designer needs to trade off these two sources of bias when choosing the degree of anonymity.

Under a more anonymous review system, sellers are more likely to write fake reviews due to the difficulty of identifying and punishing reviewers who write fraudulent reviews. However, buyers are more likely to leave honest reviews, since they have less incentive to distort their experience when they do not have to disclose their identity when leaving reviews. On the other hand, under a more non-anonymous review system, sellers are less likely to fake reviews but buyers are more likely to inflate their experience out of a fear of retaliation, resulting in review inflation (Zervas et al. 2015). Non-anonymous review systems are more vulnerable to seller coercion because sellers who bribe the buyer to write good reviews can later verify if the buyer wrote positive reviews of their product or service. Thus, a decrease in anonymity reduces fake reviews but increases review inflation.

I set up a Bayesian learning model incorporating this trade-off. Reviews left by previous buyers provide information to future buyers about the seller's quality. A likelihood function determines the informativeness of these reviews, and therefore the speed with which buyers update their beliefs. I show that this likelihood negatively depends on both of the two biases mentioned above – fake reviews and review inflation. Furthermore, I posit that a decrease in anonymity lowers the proportion of fake reviews but raises the degree of review inflation, and therefore there are both costs and benefits to anonymity. I show that the optimal degree of anonymity depends on the elasticity of review inflation with respect to fake reviews. This elasticity can be interpreted as the price of eliminating fake reviews relative to eliminating review inflation.

Finally, I interpret this elasticity in terms of real-world market and product characteristics. I argue that markets with more company sellers rather than individual sellers have a higher elasticity of review inflation with respect to fake reviews and therefore the optimal review platform features a high degree of anonymity. The opposite holds true for markets with a high proportion of individual sellers. These predictions of the model are consistent with observed patterns on actual review systems.

This paper proceeds as follows. In section 2, I briefly overview existing literature on online review platforms and the degree of anonymity in existing review platforms. In section 3, I describe the model and show that both types of bias reduce informativeness of reviews. In section 4, I solve for the optimal degree of anonymity subject to the constraint that the two biases are negatively related. I also conduct comparative statics analysis. In section 5, I use my model to evaluate the observed anonymity choices by various platforms and suggest the optimal level of anonymity in platforms based on each platform's market and product characteristics.

2. Background

2.1 Literature Review

In this section, I will discuss the previous research on online review platforms to show the conditions under which fake reviews and review inflation occur in review systems and why the two types of biases occur to different degrees depending on the degree of anonymity of review systems. Moreover, I will examine the effects of the presence of review biases on informativeness of reviews when buyers are aware of the presence of review biases, which is a crucial aspect of my model that will be discussed in later sections.

Fake Reviews

I will first discuss previous literature on fake reviews, one of the two main types of problems surrounding the review system. The occurrence of fake reviews is associated with 1) the initial reputation of the business, 2) the number of competitors surrounding a business, and 3) the extent to which the business is established. Luca and Zervas (2016) examined the conditions under which fake reviews are likely to be generated through data from Yelp's filtering system. First, they show that businesses with lower ratings are more likely to subsequently receive positive fake reviews. Second, they found that as restaurants face more competition, they are more likely to receive negative fake reviews. Third, they show that chain restaurants, which usually have more established reputation than individually owned restaurants, and restaurants with high ratings are less likely to generate positive fake reviews.

Previous research also supports that the occurrence of fake review may depend on the anonymity of review systems. Mayzlin et al. (2014) show that the proportion of fake reviews is different in anonymous and non-anonymous review system since the reputational risks sellers face vary according to the review system design. The paper argues that under an anonymous review system, sellers face less risk of getting caught with fake reviews, and therefore have more incentive to fake reviews. To show this, they compare hotels cross-listed in TripAdvisor, a more anonymous review system, and Expedia.com, a review system that is less anonymous. The proportion of one or two-star reviews, which are likely to be negative fake reviews, is higher in TripAdvisor than in Expedia.com, since sellers are incentivized to leave bad reviews of their competitors in a more anonymous review system.

Moreover, hotels with neighbors that have more incentive to fake reviews, such as hotels located next to small, independent hotels that have lower reputational risks associated with fake reviews, have higher proportion of one or two-star reviews in TripAdvisor than in Expedia.com. Similarly, for positive fake reviews, hotels that have higher incentive to engage in fake reviews have higher proportion of five-star reviews in TripAdvisor than in Expedia.com. My analysis in section 5 incorporates this finding to argue that, in markets with more individual sellers, fake reviews are more responsive to anonymity.

Review Inflation

On the other hand, review inflation is likely to occur in review system that is less anonymous. Zervas and Byers (2015) showed that, among the properties cross-listed in Airbnb (non-anonymous) and Tripadvisor's (anonymous) vacation rental section, properties have 14% more high ratings (4.5 or higher) in Airbnb. By comparing rental housings cross-listed in the two review systems, Zervas and Byers controlled for selection bias on the side of sellers that could have resulted in difference in average ratings across the two platforms. They attribute the cause of review inflation in Airbnb to the bilateral review system that leads reviewers to write positive reviews fearing retaliation or wanting the other user to reciprocate positive reviews.[4]

Why are buyers dishonest, and inflate their experiences under a non-anonymous review system? Through a field experiment on Airbnb, Fradkin et al. (2015) showed that a strategic interaction between buyers and sellers and a socially-induced reciprocity both contribute to the upward bias in reviews under a non-anonymous system. To examine the extent to which strategic reciprocity increases review inflation, the experimenters treated a

[4] The difference in average ratings could also be the result of high proportion of negative fake reviews in Tripadvisor, as explained in the previous paragraph.

"simultaneous reveal" policy that promised reviewers not to reveal their review to the other party until both parties completed reviews. They find that this policy increased the fraction of buyers leaving reviews and decreased the proportion of five-star ratings. This shows that under a non-anonymous two-way review system with no simultaneous reveal, the first reviewer (buyer or seller) writes positive reviews of the other party wanting the second reviewer to reciprocate positive reviews or fearing that the other party would retaliate once they leave negative reviews. This is supported by their analysis of reviews not written under a simultaneous reveal policy, which showed that both buyers and sellers retaliate against negative reviews.

Review biases and Review Informativeness

Next, I will discuss how the two types of biases lead buyers to discount information in reviews. Knowing that there are a greater proportion of positive fake reviews in a more anonymous system, buyers might not consider the information from positive reviews seriously. Specifically, Mayzlin (2006) evaluates the effectiveness of promotional reviews as a signal of seller's quality when the receivers of the signal are unable to identify if actual buyers or sellers disguised as buyers wrote the reviews. Assuming that there are two states of the world, where firm A is superior to firm B and vice versa, the paper uses a Bayesian updating process to describe how reviews update an uninformed buyer's prior belief about the reputation of the firms as the buyer reads positive reviews. The paper finds that there could be two equilibria: First, when the fixed cost of fake reviews is sufficiently low, an inferior firm produces more promotional chats but the reviews are still persuasive since there are more positive reviews of the superior firm in the equilibrium. In this equilibrium, since positive reviews are informative to buyers despite the fact that there are fake reviews, buyers are more likely to buy products with positive reviews online and inferior firms benefit from promotional fake reviews. The second equilibrium is when the fixed cost of fake reviews is too high. In this equilibrium, both superior and inferior firms' optimal strategy is not to write fake reviews at all.

The conclusion of Mayzlin (2006) connects to my paper's proposition that the optimal choice of anonymity depends on the costs associated with eliminating review biases. However, the paper only talked about positive fake reviews written by sellers and failed to mention the negative fake reviews written by competitors as well as dishonest reviews written by buyers - other types of biases that my paper addresses. Moreover, although the paper mentions consumer's welfare loss caused by the noise in the review system, it failed to discuss how such biases could be reduced through altering the design of the review system. Motivated by this literature, I focus on the optimal anonymity choice.

2.2 Comparison of Review Systems in Online Review Platforms

In this section, I examine the choice of anonymity in existing review platforms. I document that markets with more company sellers have chosen more anonymous review system, and markets with more individual sellers have chosen more non-anonymous review system.

I categorized platforms based on the type of seller in the market in which the platform operates. The sellers can be largely categorized into two types: individual sellers and company sellers. Individual sellers operate businesses in a small scale and do not have established brands as much as company sellers. An example of markets with high proportion of individual sellers than company sellers is peer-to-peer market. On the other hand, company sellers, such as hotel chains and restaurant franchise that are reviewed in popular review sites have reputation built on transaction histories and advertisements of their products. An example of markets with more company sellers is the hotel market.

I found that peer-to-peer market review platforms, such as Airbnb, Angie's List and Uber, where majority of sellers are individuals, tend to be less anonymous. Airbnb's review system is closest to a fully non-anonymous system, since it requires both guests and hosts to verify their identity through email address, phone number and often government ID.[5] Angie's List and Uber keep reviewer information anonymous to other buyers, but the reviewer information is visible to sellers. Moreover, Uber and Airbnb allow only buyers who have had transaction with the seller to leave reviews of the seller, eliminating a chance for sellers to leave reviews of themselves or their competitors. This is another way that review platforms in markets with more individual sellers verify information about reviewers, making the system more non-anonymous. I speculate that Airbnb and Uber chose a review system closer to a non-anonymous system because they operate in a market where the cost of eliminating fake review is low relative to the cost of eliminating review inflation. These platforms chose a system that arguably yields less fake reviews but are more vulnerable to review inflation. This will be formalized in my model in section 4, where I show that this choice would maximize the elimination of aggregate bias resulting from both review inflation and fake reviews given the constraints that platforms face in markets with more individual sellers.

The vulnerability of these systems to review inflation is evidenced by the strict punishment policies that peer-to-peer market platforms adopted to prevent sellers from coercing buyers to write dishonest reviews. For instance, Airbnb states that hosts who

[5] This information is not visible to other users of the platform, but users are aware that the reviewer identity is verified by the platform.

promise a monetary exchange for a positive review and guests who use reviews to threaten the hosts are to be reported. Moreover, Airbnb changed its review policy in 2014 so that both the hosts and the guests cannot see the reviews before both parties completed reviews. These policies suggest that indeed a non-anonymous system is vulnerable to review inflation. To what extent these policies actually prevent strategic interaction between buyers and sellers is an open question.

On the other hand, review sites where majority of sellers are branded companies or established service providers, such as Amazon, Tripadvisor and Opentable, opted for a more anonymous system. These platforms allow reviewers to use nicknames instead of real names. Opentable guarantees greatest anonymity to reviewers, allowing reviews to be written under the name "Opentable Diner" or "VIP." Sellers registered in Opentable are upscale, fine-dining restaurants in contrast to Yelp, a less anonymous system, where registered restaurants range from small, local cafes to chain restaurants. I speculate that platforms that operate in markets with more company sellers face a higher cost of eliminating fake review relative to eliminating review inflation and thus opted for a more anonymous review system. A more anonymous review system arguably yields less review inflation since buyers are honest when evaluating the quality of goods and services provided by sellers but is more vulnerable to fake reviews. In section 4, I will formalize this in my model and I show that this choice would maximize the elimination of aggregate bias resulting from both review inflation and fake reviews given the constraints that platforms face in markets with more company sellers.

The vulnerability of these systems to fake reviews is evidenced by the built in verification system that Expedia.com and Opentable adopted to prevent sellers from leaving fake reviews. Expedia.com only allows buyers who had booked a trip through its website within six months at the time of writing reviews to leave reviews. Moreover, Opentable only allows buyers who made reservation through its site to leave reviews, so that restaurant owners cannot leave promotional reviews of their business or negative reviews of their competitors.

Overall, empirical evidence show that when faced with a trade-off between eliminating fake reviews and review inflation, review platforms choose the degree of anonymity of review system based on the relative costs of eliminating the two types of bias. This intuition will be formalized in my model in section 3.6 and 4.

In the diagram below, existing platforms are categorized based on buyers' need for information sharing (which affects the level of review inflation independent of the anonymity of the platform) and the extent of reputational risk associated with writing fake reviews (which affects the level of fake reviews independent of the anonymity of the platform). Platforms that have more sellers that have incentive to fake reviews chose a more non-anonymous system, while platforms that have fewer sellers that have incentive to fake reviews and more buyers who are likely to be dishonest chose a more anonymous system.

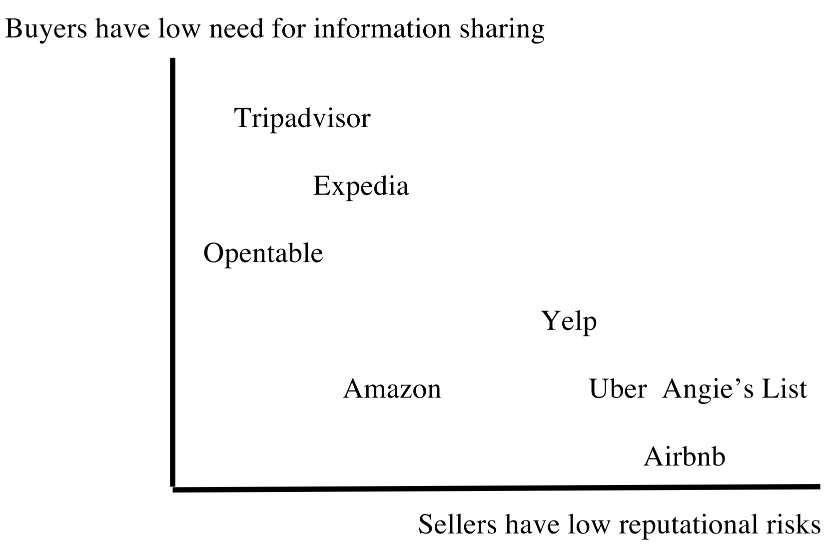

Diagram 2.1 Analysis of informational concerns in review platforms

Platform Name	Type of product provided through platform	Anonymity	Type of Seller	Type of Review System	Review Quality Control
Amazon	Retail products	Anonymous	Mix of company and individual sellers	One-way	Amazon verified purchase, Helpful votes (Buyers can vote on helpful reviews)
Angie's List	Home improvement	Anonymous But reviewer information visible to sellers	Individual service providers	One-way	Runs verification process to eliminate fake reviews
Airbnb	House - sharing	Non-Anonymous	Individual service providers	Two-way	Simultaneous revealing of guest and host reviews, Helpful Votes
Tripadvisor	Hotel reservation, Vacation rentals	Anonymous	Company sellers (chain hotels)	One-way	Reviewer badge, Helpful votes, Punishment policy for fake reviews
Expedia	Hotel reservation	Anonymous	Company sellers	One-way	Only buyers who recently booked through Expedia can leave reviews, Helpful votes
Yelp	Restaurant information	Non-Anonymous	Mix of company and individual sellers	One-way	Nominate 'Elite' reviewers for review volume and quality, No system of transaction verification
Opentable	Restaurant reservation	Anonymous	Mix of company and individual sellers	One-way	Helpful votes, Indicates reviewer's membership period
Uber	Ride sharing	Anonymous	Individual sellers	Two-way	None Requires riders to leave ratings/reviews before they book the next ride, eliminating selection bias

Table 2.2 Comparison of review platforms

3. Model

3.1 Model Setup

Model Environment

I set up a three-stage model to understand how reviews update a group of buyers' belief about a seller's quality. There are two types of agents: buyers and sellers. Buyers share a common goal of identifying the seller's type through their interactions with each seller.

Two buyers, buyer 1 and 2, interact with one seller. In period 0, nature determines the type of seller (T= G or B), who then produces in periods 1 and 2. Buyers have a common belief about the probability that the seller is a good type (α_0), and $0 < \alpha_0 < 1$. The seller type is fixed once it is determined. However, the product quality, which can be high or low, that each seller produces, can differ in each period.

Each period (t=1, 2), a good seller produces high quality goods with a probability of π_G and low quality goods with a probability of $1 - \pi_G$. A bad seller produces high quality goods with a probability of π_B and low quality goods with a probability of $1 - \pi_B$. A good seller has a higher probability of producing high quality goods than a bad seller: $\pi_G > \pi_B$.

The prior (α_0) is a belief that buyers have about the seller's quality (or type) before any buyer buys a good from the seller. We assume that the prior is high enough that buyers start out buying from the seller in period 1. In other words, the prior is greater than the threshold value (which I will characterize later) on which buyers base their purchase decisions: $\alpha_0 \geq \bar{\alpha}$. Think of $\bar{\alpha}$ as a minimum reputation that a seller requires for buyers to continue transactions with the seller.

Buyer's choices

In period 1, the first buyer buys, observes the quality of the good and leaves a review stating either that the good is high quality (positive review of the product) or that the good is low quality (negative review of the product). At this stage, the seller or its competitor can also write a fake review, pretending to be a buyer. Reviews can be only of two kinds: Positive or Negative. Reviews update buyers' prior belief on the probability of the seller being a good type. We assume that the first buyer's prior is high enough to make them buy the goods from the seller.

In period 2, the second buyer reads a review and decides whether to buy the product from the seller based on the updated belief on seller's reputation, or posterior α_1. In the model, a fake review is a review that the second buyer reads that was not written by the first buyer (such as a review written by the seller). A dishonest review is a review written by the first buyer that does not reflect the true quality of their experience with the seller.

Buyers make decision to buy or not based on the expected utility of buying a product from a particular seller. When the buyer's belief about the seller's reputation is above the threshold, buyer's expected utility buying from the seller is greater than zero. To illustrate how buyers calculate expected utility, assume that buyer's utility of consuming high quality good is $u_H > 0$ and buyer's utility of consuming low quality good is $u_L < 0$. The utility of not buying is zero $u_{Not\ Buying}= 0$.

Why would buyers want to buy from a good seller instead of a bad seller?

I assume that buyer's utility from making transaction with a good seller is positive and buyer's utility from making transaction with a bad seller is negative.[6]

$$E(u_G) = \pi_G \times u_H + (1 - \pi_G) \times u_L > 0$$

[6] The buyer's utility, u_H and u_L, are utilities net of the price.

$$E(u_B) = \pi_B \times u_H + (1 - \pi_B) \times u_L < 0$$

Buyers' optimal choice is to buy when their expected utility is greater than zero. Since buyers are not aware of the seller's type, buyer calculates their expected utility based on their belief about the seller's reputation.

At any given moment, the probability of the good being high quality is:

$$\Pr(H) = \pi(\alpha) = \alpha \times \pi_G + (1 - \alpha) \times \pi_B$$

Where α is a current belief that the seller is a good type. To simplify the notation, I define $\pi(\alpha)$ as a probability of the good being high quality.

Similarly, the probability of buying low quality goods is:

$$\Pr(L) = 1 - \pi(\alpha) = \alpha\,(1 - \pi_G) + (1 - \alpha)(1 - \pi_B)$$

At any given moment, buyer's expected utility is:

$$E(u) = \pi(\alpha) \times u_H + [1 - \pi(\alpha)] \times u_L$$

Therefore, in period 2, the second buyer buys from the seller if their expected utility based on the posterior is greater than zero.

$$\text{Buy if } \pi(\alpha_1) \times u_H + [1 - \pi(\alpha_1)] \times u_L \geq 0$$

$$\text{Do not buy if } \pi(\alpha_1) \times u_H + (1 - \pi(\alpha_1)) \times u_L < 0$$

At threshold value of α, the expected utility from buying is zero. This determines $\bar{\alpha}$.

$$\pi(\bar{\alpha}) \times u_H + (1 - \pi(\bar{\alpha})) \times u_L = 0$$

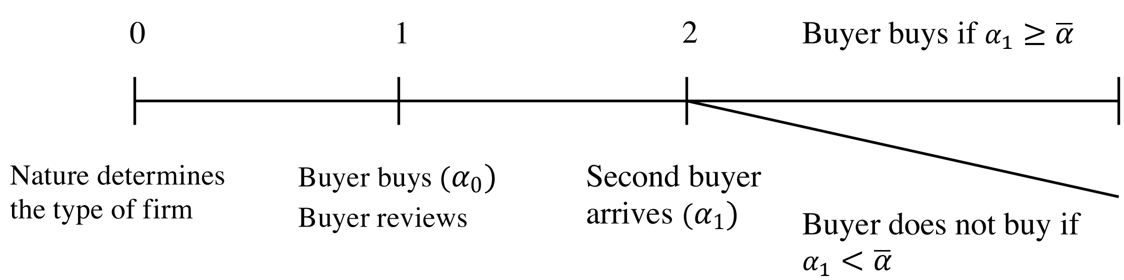

Discussion of the model

In this section, I discuss several key model assumptions. In our model, there are two buyers interacting with one seller. We can think of the first buyer as representing a first group of buyers who transact with a seller without reading reviews. In the real world, after the buyers' make transactions, they leave reviews that can be either honest or dishonest. Some buyers choose not to leave reviews, but the model assumes that the first buyer always leaves a review.[7] Moreover, sellers or their competitors can leave fake reviews without making transactions.

Similarly, we can think of the second buyer as representing a group of buyers who read reviews left by the first generation of buyers or fake reviews left by the seller. In the real world, the review that the second buyer bases his purchase decision in stage 2 would be an average of the reviews rather than each review individually. Review inflation results when dishonest reviews written by the first generation of buyers are upward biased.

Moreover, in our model, sellers are not making any decisions. This means that the seller cannot change the quality of his product in response to reviews. Therefore, buyers can learn about the type of seller through reviews.

Lastly, I assume that the price is exogenously given and the price for high and low quality goods is the same. The reason why the price is identical regardless of the product quality is that buyers cannot determine the quality of the goods before making transactions. Since buyers cannot buy different quality goods at different prices, transactions involving low-quality goods are not utility maximizing. Therefore, to maximize their utility from buying, buyers seek to identify a good seller who produces high quality goods more often

[7] In Section 6, I explain why I did not choose to model the buyer's choice to write reviews.

than a bad seller. As I will discuss in section 4.2, we evaluate the market efficiency based on buyer's welfare that in turn depends on the informativeness of reviews.

3.2 Updating and the Likelihood function

Buyers update their prior (α_0) based on the review that the first buyer leaves. The more bias-free the review system is, the more efficient it is at informing the second buyer of the seller's type based on the first buyer's experience.

The posterior, or the second buyer's belief on the probability of seller being a good type is proportional to prior multiplied by likelihood, the probability that the good seller will get a particular review. The likelihood shows how a particular review updates buyer's belief about the seller. For instance, if the likelihood that the review is positive for a good seller Pr (R=H | T=G) is high relative to Pr (R=H | T=B), it indicates that you are more likely to see a high review from a good seller than a bad seller. So, a high review is more likely a signal of a good seller than a bad seller. Therefore, the difference in how each review system updates prior to posterior is reflected in the different likelihood function:

$$\text{Posterior} \propto \text{Prior} \times \text{Likelihood}$$

By Bayes' rule, the posterior equals prior multiplied by the likelihood divided by the probability of getting a particular review. The probability of getting a particular review Pr (R=H or L) is a weighted sum of Pr (R=H | G) and Pr (R=H | B). For convenience, I will refer to Pr (R=H or L) as the data used to update prior.

Pr (Seller is good | Review= H or L)

$$= \Pr(\text{Seller is good}) \times \frac{\Pr(\text{Review=H or L} \mid \text{Seller is Good})}{\Pr(\text{Review=H or L})}$$

There are two types of bias in the review systems that are reflected in the likelihood function: fake reviews and reviews written as a result of coercion. The first type of bias, fake reviews, results when agents other than buyers write reviews. The key variable of interest is γ, which denotes the fraction of reviews that are fake $(0 \leq \gamma \leq 1)$. Later, this will be the key choice variable, since the more anonymous the system is, the greater the fraction of fake reviews. In addition, the fraction of positive reviews among fake reviews is μ $(0 \leq \mu \leq 1)$. This may depend on seller type. Denote μ_G as the fraction of positive review among fake review of good sellers and μ_B as a fraction of positive review among fake review of bad sellers.[8]

The second type of bias results when buyers leave dishonest reviews. Due to reasons that were explained in the introduction, buyers are compelled to write positive reviews despite having negative experiences. We define σ= Pr (Positive review | Buyer's experience was negative).

Both types of review bias introduce noise in the review system. They make the second buyer discount the information they receive from reviews written by the first buyer. The extent to which the second buyer discounts reviews will depend on the magnitude of bias, γ and σ.

The general version of likelihood function that reflects both types of review bias is:

$$
\begin{aligned}
&\Pr(\text{Review} = \text{H or L} \mid \text{Seller is Good}) \\
&\qquad = \Pr(\text{Reviews are not fake}|\text{Seller is good}) \\
&\qquad \times[\Pr\,(\text{Review} = \text{H or L}|\text{Seller is good and reviews are not fake, buyers are honest})\times \Pr\,(\text{Honest}) \\
&\qquad + \Pr\,(\text{Review} = \text{H or L}|\text{Seller is good and reviews are not fake, buyers are dishonest})] \\
&\qquad \times \Pr\,(\text{Dishonest}) \\
&\qquad + \Pr\,(\text{Reviews are fake}|\text{Seller is good})\times \Pr\,(\text{Review} = \text{H or L}|\text{Seller is good and reviews are fake})
\end{aligned}
$$

[8] Since I want to focus on the choice of γ, I hold other parameters, such as μ, as exogenous. In Mayzlin (2006)'s model, the fraction of positive fake reviews of good sellers (μ_G) is different from that of bad sellers (μ_B) because in the equilibrium, bad sellers benefit more from promotional reviews than good sellers.

$$\Pr(\mathrm{H|G}) = (1-\gamma)\times[\pi_G + (1-\pi_G)\sigma] + \gamma\times\mu_G$$

$$\Pr(\mathrm{L|G}) = (1-\gamma)\times[(1-\pi_G)(1-\sigma)] + \gamma\times(1-\mu_G)$$

For instance, when the second buyer reads a positive review, the buyer's prior is updated as follows.

$$\alpha_1(G|H)$$

$$= \alpha_0\times\frac{\overbrace{(1-\gamma)\times[\pi_G+(1-\pi_G)\sigma]+\gamma\times\mu_G}^{\Pr(H|G)}}{\alpha_0\times[\underbrace{(1-\gamma)\times[\pi_G+(1-\pi_G)\sigma]+\gamma\times\mu_G}_{\Pr(H|G)}]+(1-\alpha_0)\times[\underbrace{(1-\gamma)\times[\pi_B+(1-\pi_B)\sigma+\gamma\times\mu_B}_{\Pr(H|B)}]}$$

3.3 Baseline model with no bias

In a perfect world with no fake reviews $\gamma = 0$ and where all buyers are honest $\sigma = 0$, there is no noise in the review system. Therefore, the probability of getting a positive review for a good seller given that the seller is good equals the probability of getting high quality goods from a good seller.

$$\Pr(\text{Review} = \text{H} \mid \text{Seller is Good}) = \pi_G$$

Since $\gamma = 0$ and $\sigma = 0$, when the second buyer reads a positive review, the buyer's prior is updated as follows.

$$\alpha_1 = \alpha_0\times\frac{\pi_G}{\alpha_0\times\pi_G + (1-\alpha_0)\times\pi_B}$$

Since $\pi_G > \pi_B$,

$$\frac{\pi_G}{\alpha_0\times\pi_G + (1-\alpha_0)\times\pi_B} > 1$$

Therefore, the second buyer's belief about the seller's reputation increases upon reading a positive review.

To illustrate the case when the second buyer reads a negative review, the probability of getting a negative review for a good seller given that the seller is good equals the probability of getting low quality goods from a good seller.

$$\text{Pr(Review} = \text{L | Seller is Good)} = 1 - \pi_G$$

When the second buyer reads a negative review, the buyer's prior is updated as follows.

$$\alpha_1 = \alpha_0 \times \frac{(1 - \pi_G)}{\alpha_0 \times (1 - \pi_G) + (1 - \alpha_0) \times (1 - \pi_B)}$$

Since $\pi_G > \pi_B$,

$$\frac{(1 - \pi_G)}{\alpha_0 \times (1 - \pi_G) + (1 - \alpha_0) \times (1 - \pi_B)} < 1$$

Therefore, upon reading a negative review, the second buyer's belief about the reputation of the seller decreases.

The baseline model illustrates a situation where the updating process is most accurate. Since there is no noise in the review system, the information that the first buyer delivers through reviews is read by the second buyer without discounting. The most the buyers can update the prior from good or bad reviews are when the reviews are unbiased. In the next sub section, I will discuss how the degree of updating would decrease as biases are introduced in the review system.

3.4 Model of Review System with Fake Reviews and no Review Inflation

Now, I would like to illustrate a case where only one type of review bias is present: fake reviews (γ). This section illustrates a case in which fake reviews are possible, but review inflation is shut down: σ is assumed to be identically zero. I aim to show why a platform designer would like to reduce fake reviews if they are the only type of review bias.

Since there is only one type of review bias, γ, there is no trade-off between different types of biases. In other words, an increase or decrease in γ does not change σ. In this case, the goal of the platform is to minimize the proportion of fake reviews to increase the informativeness of reviews.

The likelihood function under review system with only Fake Reviews is:

$$\text{P (Review= H | Seller is Good)} = (1-\gamma)\times\pi_G \times +\gamma\times\mu_G$$

The following result shows the effect of positive fake reviews on the degree of updating.

Proposition 1.

Under a review system with only fake reviews, as the proportion of fake reviews (γ) increases, buyers discount the information from reviews.

Proof:

For simplicity, we assume that $\mu_G = \mu_B = \mu$. The case where $\mu_G \leq \mu_B$ is analogous.

Upon observing a positive review, the second buyer's posterior would be:

$$\alpha_1(G|H) = \frac{\alpha_0\times[(1-\gamma)\times\pi_G + \gamma\times\mu]}{\alpha_0\times[(1-\gamma)\times\pi_G + \gamma\times\mu] + (1-\alpha_0)\times[(1-\gamma)\times\pi_B + \gamma\times\mu]}$$

$$= \alpha_0\times\frac{(1-\gamma)\times\pi_G+\gamma\times\mu}{\alpha_0\times[(1-\gamma)\times\pi_G+\gamma\times\mu]+(1-\alpha_0)\times[(1-\gamma)\times\pi_B+\gamma\times\mu]}$$

$$= \alpha_0\times\frac{(1-\gamma)\times\pi_G+\gamma\times\mu}{(1-\gamma)[\alpha_0\times\pi_G+(1-\alpha_0)\times\pi_B]+\gamma\times\mu}$$

$$= \alpha_0\times\frac{(1-\gamma)\times\pi_G+\gamma\times\mu}{(1-\gamma)\times\pi\,(\alpha_0)+\gamma\times\mu}$$

As we discussed above, when the first buyer observes a high quality good and leaves a positive review, the posterior increases because the likelihood divided by P (H) is greater than one.

$$\frac{\alpha_1(G|H)}{\alpha_0} = \frac{(1-\gamma)\times\pi_G + \gamma\times\mu}{(1-\gamma)\times\pi\ (\alpha_0) + \gamma\times\mu} > 1$$

When there are fake reviews, positive review increases posterior to a lesser extent than it does under a bias-free review system, because the likelihood divided by P (H) is a decreasing function of γ. As proportion of fake review increases, the value of information gained from positive review falls.

Taking the derivative of α_1/α_0 with respect to γ, we obtain:

$$\frac{d}{d\gamma}\frac{\alpha_1(G|H)}{\alpha_0} = \frac{(\mu-\pi_G)\times[(1-\gamma)\pi(\alpha_0)+\gamma\times\mu] - \left(\mu-\pi(\alpha_0)\right)\times[(1-\gamma)\pi_G+\gamma\times\mu]}{[(1-\gamma)\times\pi\ (\alpha_0)+\gamma\times\mu]^2}$$

$$= \frac{\mu\times(1-\gamma)\times[\pi(\alpha_0)-\pi_G]+\mu\times\gamma\times[\pi(\alpha_0)-\pi_G]}{[(1-\gamma)\times\pi\ (\alpha_0)+\gamma\times\mu]^2}$$

$$= \frac{\mu\times[\pi(\alpha_0)-\pi_G]}{[(1-\gamma)\times\pi\ (\alpha_0)+\gamma\times\mu]^2} \begin{cases} < 0 \text{ when } \mu > 0 \\ = 0 \text{ when } \mu = 0 \end{cases}$$

■

Since Bayesian consumers are aware that there are fake reviews when reading reviews, they discount the information based on their belief about the probability that the given review is a fake review.

In our model, if prior (α_0) was above the threshold value ($\bar{\alpha}$), the second buyer continues to buy from the seller unless the posterior (α_1) is below the threshold. The posterior can fall below threshold if, for instance the prior was close to the threshold and the value of the negative review is sufficiently high to decrease posterior significantly.

The following result shows the effect of negative fake reviews on the degree of updating.

Proposition 2. If the proportion of fake review (γ) is sufficiently large, the second buyer may continue to make transaction with a seller that he would not have made under a bias-free review system.

Proof:

Upon observing low quality reviews in a review system with only fake reviews, the posterior is:

$$\alpha_1(G|L) = \frac{\alpha_0 \times [(1-\gamma)\times(1-\pi_G) + \gamma\times(1-\mu)]}{\alpha_0 \times [(1-\gamma)\times(1-\pi_G) + \gamma\times(1-\mu)] + (1-\alpha_0)\times[(1-\gamma)\times(1-\pi_B) + \gamma\times(1-\mu)]}$$

$$= \alpha_0 \times \frac{[(1-\gamma)\times(1-\pi_G)+\gamma\times(1-\mu)]}{\alpha_0\times[(1-\gamma)\times(1-\pi_G)+\gamma\times(1-\mu)]+(1-\alpha_0)\times[(1-\gamma)\times(1-\pi_B)+\gamma\times(1-\mu)]}$$

$$= \alpha_0 \times \frac{[(1-\gamma)\times(1-\pi_G)+\gamma\times(1-\mu)]}{(1-\gamma)\times(1-\pi\,(\alpha_0)\,)+\gamma\times(1-\mu)}$$

When the first buyer observes a low quality good, the posterior decreases because the likelihood divided by P (L) is less than one.

$$\frac{\alpha_1(G|L)}{\alpha_0} = \frac{(1-\gamma)\times(1-\pi_G) + \gamma\times(1-\mu)}{(1-\gamma)\times(1-\pi\,(\alpha_0)) + \gamma\times(1-\mu)} < 1$$

When there are fake reviews, negative review decreases posterior to a lesser extent than it does under a bias-free review system, because the likelihood divided by P (L) is an increasing function of γ. In other words, the value of information gained by negative review falls.

Taking the derivative of α_1/α_0 with respect to γ, we obtain:

$$\frac{d}{d\gamma}\frac{\alpha_1(G|L)}{\alpha_0} = \frac{(\pi_G-\mu)\times[(1-\gamma)(1-\pi(\alpha_0)+\gamma\times(1-\mu)]-(\pi(\alpha_0)-\mu)\times[(1-\gamma)(1-\pi_G)+\gamma\times(1-\mu)]}{[(1-\gamma)\times(1-\pi\,(\alpha_0)+\gamma\times(1-\mu)]^2}$$

$$= \frac{(1-\mu)\times[\pi_G-\pi\,(\alpha_0)]}{[(1-\gamma)\times(1-\pi\,(\alpha_0)+\gamma\times(1-\mu)]^2} > 0 \text{ when } (1-\mu) > 0$$

$$= 0 \text{ when } \mu = 1 \qquad \blacksquare$$

Under a review system with fake reviews, the second buyer discounts the negative review by the first buyer as their belief about the proportion of fake review increases. The weakening of influence reviews have on posterior is problematic in the case of negative reviews, because if γ is big enough, the second buyer can decide to buy the product regardless of a negative review, thinking that the negative review is more likely to be a fake review written by a competitor than an honest review written by a buyer. This de-valuation of information attained through reviews can potentially decrease buyer's welfare by encouraging buyers to continue transaction with a bad seller. Overall, fake reviews slow the process of identifying bad seller through negative reviews.

Therefore, in a world with no review inflation, the review platform designer would want to choose the lowest level of fake review as possible to maximize his utility, which depends on the informativeness of reviews. This means that the platform designer who wants to maximize the updating speed, would choose $\gamma = 0$.

3.5 Model of Review System with Review Inflation and no Fake Reviews

In the previous section, I discussed a world in which there is no review inflation and the platform designer's optimal choice of γ under such condition. In this section, I will discuss a case in which there are no fake reviews but there is review inflation to see how review inflation, the second type of bias, affects the informativeness of reviews. γ is identically set to zero, but some buyers inflate the quality of goods fearing retaliation by sellers, resulting in $\sigma > 0$. Review inflation, denoted by σ, represents the degree to which the review that the second buyer reads is different from the actual experience of the first buyer.

The likelihood function of a review system with review inflation but no fake review is:

$$\text{P (Review= H | Seller is Good)} = \pi_G + (1 - \pi_G)\sigma$$

The following result shows the effect of review inflation on the degree of updating.

Proposition 3. Under a review system with only review inflation, the second buyer discounts the information from reviews as buyers inflate their experience ($\boldsymbol{\sigma > 0}$).

Proof:

$$\alpha_1(\mathrm{G|H}) = \alpha_0 \times \frac{\pi_G + (1 - \pi_G)\sigma}{\alpha_0 \times [\pi_G + (1 - \pi_G)\sigma] + (1 - \alpha_0) \times [(\pi_B + (1 - \pi_B)\sigma)]}$$

When the first buyer observes a high quality good and leaves a positive review, the posterior increases because the likelihood divided by P (H) is greater than one.

$$\frac{\pi_G + (1 - \pi_G)\sigma}{\alpha_0 \times [\pi_G + (1 - \pi_G)\sigma] + (1 - \alpha_0) \times [\pi_B + (1 - \pi_B)\sigma]} > 1$$

When there is review inflation, positive review increases posterior to a lesser extent than it does under a bias-free review system, because the likelihood divided by P (H) is a decreasing function of σ.

Taking the derivative of α_1/α_0 with respect to σ, we obtain:

$$\frac{d}{d\sigma}\frac{\alpha_1(G|H)}{\alpha_0} = \frac{(1 - \alpha_0)(\pi_B - \pi_G)}{[\alpha_0 \times \{\pi_G + (1 - \pi_G)\sigma\} + (1 - \alpha_0) \times \{\pi_B + (1 - \pi_B)\sigma\}]^2} < 0$$

■

In other words, the value of information gained by positive review falls. This also applies to the case of negative reviews, because the likelihood function P (L | G) divided by

P (L) is an increasing function of σ. Therefore, a review system designer who is unrestricted in choice of σ would like $\sigma = 0$. From this result and the result discussed in section 3.4, we conclude that both fake reviews and review inflation, respectively, reduce the informativeness of reviews and contribute to the aggregate bias in the review system.

3.6 Model of Review System with both Fake Reviews and Review Inflation

To increase the informativeness of reviews, the platform designer would want to set both γ and σ to zero. From now on, however, I am going to assume a decreasing relationship between γ and σ. If the review system is close to non-anonymous, or when the degree of fake review (γ) is low, the degree of review inflation (σ) is high because a non-anonymous review system creates an environment conducive to strategic interaction between buyers and sellers. On the other hand, if the review system is close to anonymous, or when the degree of fake review (γ) is high, the degree of review inflation is low (σ). Since $\frac{d\sigma}{d\gamma} < 0$, the review system designer cannot set both the proportion of fake reviews (γ) and review inflation (σ) to zero. The platform designer faces a trade-off between reducing fake reviews and review inflation.

The likelihood function of review system with both review inflation and fake reviews is:

$$\text{P (Review= H | Seller is Good)} = (1-\gamma)\times[\pi_G + (1-\pi_G)\sigma] + \gamma\times\mu_G$$

Therefore, in a world with both review inflation and fake reviews, the review platform designer would have to consider both the level of fake reviews and review inflation in their decision to choose the degree of anonymity in review systems to maximize their utility. The platform designer's optimal choice of γ, which determines the degree of anonymity of the system, would influence the level of review inflation (σ).

4. Main Results

In the previous section, I showed that platform designer faces a trade-off between reducing fake reviews and review inflation. The trade-off is due to the assumption that $\frac{d\sigma}{d\gamma} < 0$. In this section, I will formalize the trade-off using a function representing the relationship between γ and σ and solve an optimization problem.

The platform designer seeks to maximize the informativeness of good reviews, or the degree of updating through the likelihood function:

$$\max_{\sigma,\gamma} \frac{\alpha_1(G|H)}{\alpha_0}$$

Subject to the constraint that $\sigma = \sigma(\gamma)$.

The key variable of choice is γ, which represents the level of anonymity.

4.1 The Optimal Level of Anonymity and Comparative Statics

In this section, I will illustrate how the platform designer chooses the optimal level of anonymity (γ) when they face a trade-off between reducing the degree of fake reviews and review inflation.

An example of a functional form that satisfies these assumptions and which I will use for my comparative statics is the following: $\sigma(\gamma) = \gamma^{-\phi}$. ϕ parameterizes how sensitive σ is to γ, or the price of reducing fake reviews relative to reducing review inflation.

To see this, note that

$$\frac{\gamma}{\sigma}\frac{d\sigma}{d\gamma} = -\phi$$

We seek to maximize the likelihood function with respect to γ, subject to the constraint that σ is the aforementioned function of γ. The optimal choice of γ therefore satisfies the first-order condition:

$$\frac{d}{d\gamma}\frac{\alpha_1(G|H)}{\alpha_0} =$$

$$\frac{\frac{d\sigma}{d\gamma}[(1-\sigma)^2\times(\pi_B-\pi_G)+\gamma\times(1-\gamma)\{\mu_B\times(1-\pi_G)-\mu_G\times(1-\pi_B)\}-\mu_B\times[\pi_G+(1-\pi_G)\sigma]-\mu_G\times[\pi_B+(1-\pi_B)\sigma]}{\alpha_0\times[(1-\gamma)\times[\pi_G+(1-\pi_G)\sigma]+\gamma\times\mu_G]+(1-\alpha_0)\times[(1-\gamma)\times[\pi_B+(1-\pi_B)\sigma+\gamma\times\mu_B]^2}$$

I will use this first-order condition in my comparative statics analysis, which determines how the optimal γ depends on key parameters. The first, and most important, parameter of interest is ϕ, which measures the sensitivity of review inflation to the elimination of fake reviews. In section 5.1, I will address how ϕ, and therefore the optimal level of anonymity, is expected to vary across markets, depending on market characteristics. The second parameter of interest is μ, which measures what fraction of fake reviews is positive rather than negative. This number has implications for the optimal level of anonymity, when the objective is, e.g. maximizing the informativeness of positive reviews. As I explain below, this number captures whether fake reviews are posted mostly by sellers (as promotional reviews) or by their competitors, and thus also varies depending on market and product characteristics.

For tractability, I will focus on two special cases:

(1) A case where we assume that $\pi_G = 1, \pi_B = 0, \mu_G = 0$ and $\mu_B > 0$. This case would result when bad sellers have more incentive to leave promotional reviews than good sellers.

(2) A case where we assume that the fraction of positive fake review of good and bad sellers is the same ($\mu = \mu_G = \mu_B$). This case would result if good sellers and bad sellers were equally incentivized to leave promotional reviews. An example of such case would be a nascent industry in which both types of sellers are not well established and therefore have low reputational risk of fake reviews. The main results are very similar in case (1) and (2).

Both cases assume that $\mu_B \geq \mu_G$. This is motivated by the analysis in Mayzlin (2006), who argued that inferior firms (bad sellers) benefit more from leaving promotional reviews and therefore have higher proportion of fake positive reviews than good sellers. If $\mu_B < \mu_G$, the results could be reversed since fake reviews are informative about the seller type.

4.1.1 Special case (1): $\mu_G = 0, \pi_G = 1, \pi_B = 0$

We will first consider a case where the proportion of positive fake review of good seller is sufficiently low compared to the proportion of positive fake review of bad sellers. So, $\mu_G = 0$ and $\mu_B > 0$.

The likelihood function divided by the data used for updating looks like:

$$\alpha_1(\mathrm{G|H}) = \alpha_0 \times \frac{(1-\gamma)}{[\alpha_0 \times (1-\gamma)] + (1-\alpha_0) \times [\,\sigma \times (1-\gamma) + \mu_B \times \gamma]}$$

$$= \alpha_0 \times \frac{1}{\alpha_0 + (1-\alpha_0) \times [\,\sigma + \mu_B \times \frac{\gamma}{1-\gamma}]}$$

In order to find the value of γ that maximizes the updating of prior, or minimize the noise in the review system, differentiate the likelihood function divided by the data used for updating with respect to γ.

$$\frac{d}{d\gamma}\left(\frac{1}{\alpha_0 + (1-\alpha_0) \times [\,\sigma + \mu_B \times \frac{\gamma}{1-\gamma}]}\right) = 0$$

$$\frac{\partial \sigma}{\partial \gamma} + \frac{1}{(1-\gamma)^2} \times \mu_B = 0$$

$$-\frac{\partial \sigma}{\partial \gamma} \times (1-\gamma)^2 = \mu_B$$

This first-order condition for γ has a very intuitive interpretation.

The left-hand side of the equation represents the marginal benefit of anonymity. As review system moves closer to anonymous, review inflation decreases proportional to $(1-\gamma)^2$. In other words, as fake review increases, the probability that the reviews written by buyers (represented as proportion of non-fake reviews, $1-\gamma$) are dishonest decreases. For the reviews written by buyers, the review bias is smaller when the system gets closer to an anonymous review system.

The right-hand side of the equation represents the marginal cost of anonymity. As the review system moves closer to anonymous, the fake positive reviews of bad sellers increase proportional to μ_B. Positive fake reviews on bad sellers can mislead buyers to continue transaction with a bad seller. Therefore, the increase in fake review reduces the information delivered by positive reviews.

The optimal level of anonymity (γ) would therefore depend on the marginal benefit and marginal cost of anonymity dependent on market characteristics, such as the relative price of eliminating fake reviews to eliminating review inflation ϕ, and the proportion of positive fake reviews, μ.

(1) Comparative statics with respect to $\boldsymbol{\phi}$

As mentioned previously, we will assume the functional form

$$\sigma(\gamma) = \gamma^{-\phi}$$

which then implies

$$\frac{\partial \sigma}{\partial \gamma} = -\phi\gamma^{-(1+\phi)}$$

and so, the first-order condition for γ becomes

$$\mu_B = \phi\gamma^{-(1+\phi)} \times (1-\gamma)^2$$

Differentiating this equation with respect to ϕ, we obtain:

Proposition 4. $\frac{d\gamma}{d\phi} > 0$**. In other words, the more sensitive is review inflation to the proportion of fake reviews, the higher the optimal degree of anonymity in the system.**

***Proof*: See Appendix.**

Increase in ϕ increases the marginal benefit of anonymity, since increase in anonymity would lead to a greater reduction in review inflation.

(2) Comparative statics with respect to μ

As shown earlier, the first order condition for γ is

$$\mu_B = -\frac{\partial \sigma}{\partial \gamma} \times (1-\gamma)^2$$

Which, given our functional form assumption, reduces to

$$\mu_B = \phi\gamma^{-(1+\phi)} \times (1-\gamma)^2$$

Differentiating this equation with respect to μ_B we obtain:

Proposition 5. $\frac{d\gamma}{d\mu_B} < 0$**. The greater the proportion of bad seller's positive fake reviews, the lower the optimal degree of anonymity in the system.**

***Proof*: See Appendix.**

This indicates that in markets where the fraction of bad seller's positive fake review is high, such as in peer-to-peer markets where sellers have more incentive to write promotional reviews, the platform designer would better off by choosing a review system that is close to non-anonymous, if the goal of the platform designer is to maximize the informativeness of positive reviews (the reverse is true for maximizing the informativeness of negative reviews).

4.1.2 Special case (2): $\mu = \mu_G = \mu_B$

Let us now consider a case where the good and bad sellers have same proportion of positive fake reviews. I will show that the main results from above are robust to this modification. The first-order condition for γ now becomes

$$\frac{d\sigma}{d\gamma}[(1-\sigma)^2\times(\pi_B-\pi_G)+\gamma\times(1-\gamma)\{\mu\times(\pi_B-\pi_G)\}]=\mu\times(\pi_G-\pi_B)\times(1-\sigma)$$

$$-\left[\frac{d\sigma}{d\gamma}\times(1-\gamma)\times\{1-\gamma(1-\mu)\}\right]=\mu\times(1-\sigma)$$

The left-hand side of the equation represents the marginal benefit of anonymity. As review system moves closer to anonymous, review inflation decreases proportional to $[(1-\gamma)\times\{1-\gamma(1-\mu)\}]$. In other words, as fake review increases, the probability that the reviews written by buyers (represented as proportion of non-fake reviews, $1-\gamma$) are dishonest decreases. For the reviews written by buyers, the review bias is smaller when the system gets closer to an anonymous review system.

The right-hand side of the equation represents the marginal cost of anonymity. As the review system becomes more anonymous, fake positive reviews increase proportional to μ. Since σ is a decreasing function of γ, $(1-\sigma)$ increases as γ increases, which amplify the cost of increase in positive fake reviews. Positive fake reviews on bad sellers can mislead buyers to continue transaction with a bad seller. Therefore, the increase in fake review reduces the information delivered by positive reviews.

(1) Comparative statics with respect to ϕ

As mentioned previously, we will assume the functional form

$$\sigma(\gamma)=\gamma^{-\phi}$$

The first order condition for γ becomes

$$\mu\times(1-\gamma^{-\phi})=\left[\phi\gamma^{-(1+\phi)}\times(1-\gamma)\times(1-\gamma+\gamma\mu)\right]$$

$$= \phi\gamma^{-(1+\phi)}\times(1-\gamma)^2 + \phi(1-\gamma)\,\gamma^{-\phi}\mu$$

Differentiating this equation with respect to ϕ, we obtain the same result as *Proposition 4*.

$\frac{d\gamma}{d\phi} > 0$.

***Proof*: See Appendix.**

(2) Comparative statics with respect to μ

As shown earlier, the first order condition for γ is

$$\mu\times\left(1-\gamma^{-\phi}\right) = \phi\gamma^{-(1+\phi)}\times(1-\gamma)^2 + \phi(1-\gamma)\,\gamma^{-\phi}\mu$$

Differentiating this equation with respect to μ, we obtain similar result as *Proposition 5*.

$\frac{d\gamma}{d\mu} < 0$. The only difference is that in special case 2, we assumed that $\mu = \mu_G = \mu_B$.

Therefore, the greater the proportion of positive fake reviews, the lower the optimal degree of anonymity in the review system.

***Proof*: See Appendix.**

4.2 Welfare Implications

So far, we have discussed ways to maximize the informativeness of reviews. In this section, I will explain why maximizing the likelihood function is related to increasing buyer's welfare (surplus) and thus, the total surplus in the market. In proposition 2, we showed that if the proportion of fake review (γ) is sufficiently large, the second buyer might continue to make transaction with a seller that he would not have under a bias-free review system. This holds the same for review inflation (σ), the second type of bias that decreases the amount of updating from prior to posterior.

Considering that the review results in the maximum amount of updating in a bias-free review system, any system with review bias can result in increased number of unwanted transactions, namely transactions with a bad seller and decreased number of desirable transactions, or transactions with a good seller. In case of positive reviews, decrease in the degree of updating due to review bias can discourage buyers who would have otherwise bought goods from a good seller not to buy the goods. In case of negative reviews, decrease in the degree of updating can make buyers continue buying from a bad seller. In both cases, buyers' utility from transaction is not maximized, creating inefficiency in the market. This is why the speed of updating is a reduced-form measure of buyer's utility.

In my model, I assume that the first buyer's prior is high enough to start buying from the seller, but this assumption can be relaxed in the real world where different buyers have different belief threshold from which they start buying. In other words, the buyer's prior can be less than the threshold required to start buying, and the buyer will start buying once his posterior is above the threshold upon reading positive reviews. In this case, the decrease in the degree of updating due to review bias can discourage potential buyers from start buying goods from a good seller.

I primarily focused on maximizing the informativeness of positive reviews in section 4. One could similarly look at maximizing the informativeness of negative reviews, and the results of the comparative statics with respect to ϕ and μ are robust whether or not I want to maximize informativeness of positive or negative fake reviews.

5. Discussion

As explained earlier, a review system designer faces a tradeoff between fake review and review inflation in choosing the degree of anonymity, which is represented as the choice of the degree of fake review in our model. Since review systems with higher number of fake reviews (review system that is closer to anonymous) have lower review inflation and review systems with lower number of fake reviews (review system that is closer to non-anonymous) have higher review inflation, the platform designer has to choose the optimal level of anonymity that would result in the least amount of aggregate bias. Here, aggregate bias is the sum of bias generated by review inflation and fake reviews in the review system. Both biases contribute to reducing the informativeness of reviews, and the platform designer's objective is to achieve the lowest possible aggregate bias, or maximize the informativeness of reviews.

A review system designer would choose the degree of anonymity that would yield the least amount of aggregate bias given the relative prices of fake reviews and review inflation they face, which in turn depends on the types of goods and sellers in the market.

5.1 Real World Characteristics that correspond to differences in ϕ

In section 4, we discussed that markets with high price of eliminating fake reviews relative to eliminating review inflation would choose a more anonymous system (proposition 4). In this section, I will focus on what market characteristic is related to ϕ, which can be interpreted as a measure of relative price of eliminating fake reviews to eliminating review inflation. As mentioned in section 2, there are markets with more individual sellers than company sellers, and vice-versa. I argue that the value of ϕ depends on the relative proportion of individual and company sellers in the market in which the review platform operates.

Markets with low ϕ

I claim that in markets with more individual sellers than company sellers, the price of eliminating fake reviews relative to eliminating review inflation is low. First, I argue that in markets with more individual sellers, only a small decrease in anonymity results in a large reduction in fake reviews; second, the extent to which review inflation increases in response to decrease in anonymity is smaller.

First, I argue that in markets with more individual sellers than company sellers, such as peer-to-peer markets, eliminating fake reviews requires only a small decrease in anonymity. In peer-to-peer markets, majority of sellers are individuals who provide their labor, service or products through the network of buyers and sellers facilitated by the platform and the sellers' reputation is not well established unlike company sellers. If there is a low probability of being caught, individual sellers have high incentive to fake reviews to create positive reputation. Therefore, fake reviews are very elastic with respect to anonymity. Even a slight decrease in anonymity would greatly reduce the number of fake reviews in the review system.

Next, I argue that in markets with more individual sellers than company sellers, a decrease in anonymity increases review inflation to a lesser extent than in markets with more company sellers. Consider the buyers' incentive to write honest reviews in this type of market. For buyers, the value of information gained through reviews is higher when the seller quality is unknown or when there are a greater number of individual sellers. Reviews are an important way to overcome the informational asymmetry between buyers and sellers, and the existence of informational asymmetry in markets with more individual sellers compel buyers to leave honest reviews to effectively identify bad sellers (such as bad Uber drivers). Individual buyers are aware that the marginal effect of their positive or negative review of the seller would be bigger in markets with more individual sellers, since the number of buyers that interact with a given seller is smaller and therefore a seller has lower chance of getting reviewed. Although I did not explicitly discuss buyer's objective function, I speculate that buyers have incentive to maintain the quality of reviews in review platforms, and the incentive is greater for markets with greater proportion of individual sellers to company sellers.

Therefore, buyers are more likely to leave honest reviews to make review system accountable in markets with more individual sellers, regardless of the anonymity of the review platform. Taken together, these imply that a relatively large reduction in fake reviews, which only requires a small decrease in anonymity, would result in a small increase in review inflation. In markets with more less-established individual sellers, the relative price of eliminating fake reviews in units of eliminating review inflation is low.

Markets with high ϕ

I claim that in markets with more company sellers than individual sellers, the price of eliminating fake reviews relative to eliminating review inflation is high. First, I show that in markets with more company sellers, a small decrease in anonymity results in a small reduction in fake reviews; second, the extent to which review inflation increases in response to decrease in anonymity is greater.

First, I show that in markets with more company sellers than individual sellers, small decrease in anonymity results in only small reduction in the number of fake reviews. Since company sellers have less incentive to fake reviews than individual sellers, the number of fake review in review platform in this type of market is low regardless of anonymity of the review system. Company sellers are hesitant to leave fake reviews because the marginal benefit of fake review, the benefit from promotional reviews or negative review of competitor, does not exceed its marginal cost, the reputational risk of review manipulation. (Imagine Sheraton being caught in a review scandal!) Thus, in markets with more company sellers, a further reduction in fake reviews requires the platform to be anonymous to a greater extent than in markets with more individual sellers.

Next, I show that in markets with more company sellers, a decrease in anonymity results in a greater increase in review inflation than in markets with more individual sellers. This is because in markets with more company sellers, buyers have less incentive to write honest reviews as the need for information sharing among buyers is not as high as in markets with more individual sellers. On buyer's side, the benefit of leaving honest reviews is collective information sharing to overcome informational asymmetry problem, and the cost is revelation of reviewer identity associated with positive or negative reviews. Both markets face the cost of information sharing, revelation of reviewer identity; but the benefit is higher in markets with more individual sellers. So, buyers are more likely to leave honest reviews in markets with more individual sellers. Therefore, to reduce fake reviews in markets with more company sellers, the review system requires a greater decrease in anonymity than in markets with more individual sellers, which in turn causes high increase in review inflation.

All in all, in markets with more company sellers, the relative price of eliminating fake reviews in units of eliminating review inflation is high.

Platform's Choice of anonymity depending on the value of ϕ

Now that I discussed market characteristics associated with different values of ϕ, I would like to discuss the choice of anonymity in review platforms operating in different markets based on the value of ϕ. Proposition 4 states that market with high ϕ would choose a more anonymous system and market with low ϕ would choose a less anonymous system.

Combined with the above discussion, proposition 4 implies that review platforms in markets with greater proportion of company sellers would choose more anonymous system. The review platforms in markets with greater proportion of individual sellers would choose non-anonymous system.

The below diagram illustrates that different constraints (cost of eliminating fake reviews in units of eliminating review inflation) that platforms face according to market characteristics lead to a different optimal level of anonymity. The points along the indifference curve represent a constant aggregate level of bias caused by the combination of fake reviews and review inflation. We can think of elimination of fake review and review inflation both as goods in the market, and the platform designer as choosing combination of goods that can reduce the greatest amount of aggregate bias given their budget. The platform designer's budget is the set of all combinations of fake review and review inflation feasible to the platform designer. In our model, the budget is set of all combinations of σ and γ for which $\sigma = \gamma^{-\phi}$, and ϕ determines the slope of the budget constraint. Since platforms all desire to maximize the likelihood, the point of tangency between the budget constraint and the indifference curve determines the choice of anonymity.

Elimination of Review Inflation: lower σ

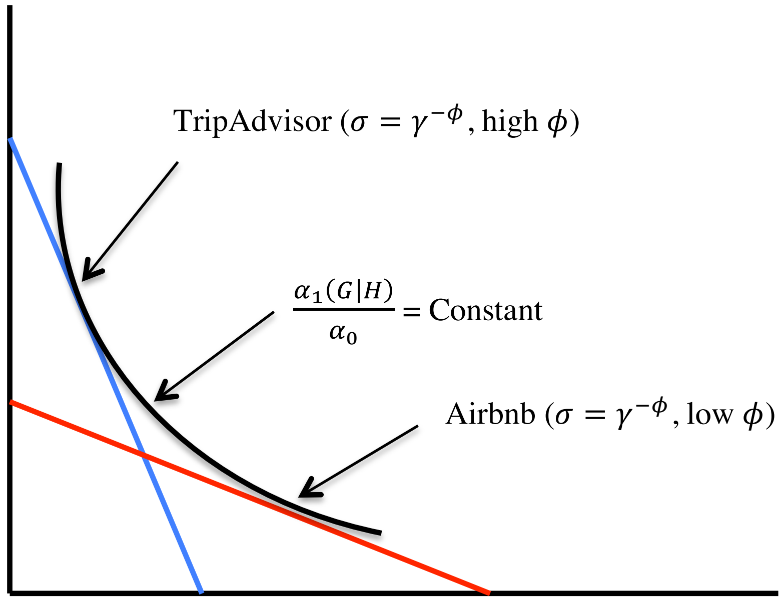

Elimination of Fake Reviews: lower γ

Diagram 5.1 Optimal choice of review inflation and fake reviews

As you can see from the diagram, my argument is consistent with the existing review platform's choice of anonymity described in section 2.2. Platforms such as Tripadvisor, which operate in markets with more company sellers (high ϕ), adopted an anonymous review system, whereas platforms such as Airbnb, which operate in markets with more individual sellers (low ϕ), adopted a non-anonymous review system.

5.2 Real World Characteristics that correspond to differences in μ

Another market characteristic to consider apart from the proportion of individual vs. company sellers is the product substitutability of markets in which the platform operates. Depending on the degree of product substitutability, sellers have high or low incentive to write negative fake reviews of the competitors' products. I speculate that markets with less room for product differentiation, or markets in which goods are easily substitutable, are likely to have higher proportion of negative fake reviews and thus relatively lower proportion of positive fake reviews (low μ).[9] Sellers are likely to leave negative reviews of their competitors' products to increase their market share. Examples of markets with high product substitutability are markets for consumer non-durables, such as toilet papers, cereal or plastic cups.[10]

On the other hand, markets in which goods are not easily substitutable are likely to have lower proportion of negative fake reviews and thus relatively higher proportion of positive fake reviews (high μ). An example of markets with low product substitutability is market for books. Chevalier and Mayzlin (2006) argue that books on similar topics are likely to be complements, not substitutes, so there is less fake negative reviews written by competing sellers than in other markets.

Proposition 5 states that the greater the proportion of bad seller's positive fake reviews, the lower the optimal degree of anonymity in the system. From this, we can conclude that platforms operating in markets with lower product substitutability would maximize the informativeness of positive reviews by choosing a more non-anonymous system and platforms operating in markets with higher product substitutability would maximize the informativeness of positive reviews by choosing a more anonymous review system. Due to the difficulty in defining the degree of product substitutability in existing platforms, I have not done substantial analysis of the real-world examples of platforms with different values of μ. The analysis of the consistency of the results (proposition 5) with the

[9] Here, I discuss the relative proportion of positive and negative fake reviews, because sellers have incentive to promote their product (write positive fake reviews) regardless of product substitutability, but sellers differ in their incentive to write negative fake reviews depending on the substitutability of their products.

[10] Product substitutability depends on how narrowly we define markets. For instance, we can say that market for ready-to-eat cereals has low product substitutability but market for organic, whole-grain ready-to-eat cereals has high product substitutability. In my analysis, I consider broader definition of products (ex) market for cereals instead of market for organic, whole-grain cereals).

choice of anonymity of existing review platforms would be a meaningful avenue of research.

6. Conclusion

Review platform is an important space for buyers to share information about the type of seller and the quality of products they expect to buy from particular sellers. The informativeness of reviews depends on the amount of aggregate bias in the review system, which hinges on the review system design. In my paper, I find that the optimal choice of anonymity, an aspect of review system design that determines the degree of bias in the review system, differs based on the market and product characteristic in which the review platform operates. Specifically, I look at the difference in the relative price of reducing fake reviews and review inflation in markets with high proportion of individual sellers or company sellers.

My model successfully captures an important trade-off between reducing fake reviews and review inflation, the two most prevalent types of review bias that lead buyers to discount information delivered through reviews. A negative relationship between the degree of review inflation and fake reviews leads the platform designer to choose a review system that would minimize the aggregate bias. One source of bias I did not discuss in my model is the selection bias that can influence the amount of positive or negative reviews depending on the characteristic of reviewers that decide to write reviews. Zervas et al. (2015) note that reviewers who volunteer to write reviews tend to write extremely positive or negative reviews, but the direction of net bias is positive. The selection bias contributes to review inflation in non-anonymous review system, where reviewers who had negative experiences tend to omit reviews. Therefore, my treatment of review inflation already implicitly incorporates selection bias.

The method I used to characterize market has advantages as well as limitations. One advantage of distinguishing markets based on the proportion of individual and company sellers is that it allows me to discuss the optimal review system design of peer-to-peer markets, most of which have high proportion of individual sellers.

However, I did not take into consideration the difference in market size, level of competition and revenue model of peer-to-peer review platforms that can affect the choice of anonymity of the platform designer. For instance, review platforms that depend on advertisements would prefer a more anonymous review system to increase the number of reviews, regardless of the quality of reviews. I deliberately put aside the profit incentive to focus on one objective of review platform designer: maximizing the informativeness of reviews. I assumed that when designing a review system, the platform designer's incentives are aligned with that of buyers, and an effective review system maximizes buyers' welfare through maximizing the informativeness of reviews. Incorporating alternative objective functions for the platform designer would be a fruitful avenue for future research.

My analysis shows that the choice of anonymity in existing platforms is consistent with the results of my model. This could be the result of review platform designers' responsiveness to the informational constraints they face in the type of market they operate in. To match the growth of size and relevance of peer-to-peer market in our economy, review system designers should continue their efforts to maximize the informativeness of reviews. An innovation in review system design would enable the platform to reduce one type of review bias without increasing the degree to which the other type of bias occurs.

7. Appendix

7.1 Derivation of Trade-Offs between γ and σ

$$\alpha_1(G|H)$$

$$=\alpha_0\times\frac{(1-\gamma)\times[\pi_G+(1-\pi_G)\sigma]+\gamma\times\mu_G}{\alpha_0\times[(1-\gamma)\times[\pi_G+(1-\pi_G)\sigma]+\gamma\times\mu_G]+(1-\alpha_0)\times[(1-\gamma)\times[\pi_B+(1-\pi_B)\sigma+\gamma\times\mu_B]}$$

Our task is to derive the likelihood function divided by the data used for updating by γ.

For simplicity, let us denote numerator and denominator as:

N_G =P (H | G)=$(1-\gamma)\times[\pi_G+(1-\pi_G)\sigma]+\gamma\times\mu_G$
N_B = P (H | B) = $(1-\gamma)\times[\pi_B+(1-\pi_B)\sigma]+\gamma\times\mu_B$

D = $\alpha_0\times[(1-\gamma)\times[\pi_G+(1-\pi_G)\sigma]+\gamma\times\mu_G]+(1-\alpha_0)\times[(1-\gamma)\times[\pi_B+(1-\pi_B)\sigma+\gamma\times\mu_B]$
 = $\alpha_0\times N_G+(1-\alpha_0)\times N_B$

$$\frac{dN_G}{d\gamma}=-[\pi_G+(1-\pi_G)]+(1-\gamma)\left[(1-\pi_G)\times\frac{d\sigma}{d\gamma}\right]+\mu_G$$

$$\frac{dD}{d\gamma}=\alpha_0\times\frac{dN_G}{d\gamma}+(1-\alpha_0)\times\frac{dN_B}{d\gamma}$$

Using the quotient rule, differentiate numerator and denominator by γ,

$$\frac{d\frac{\alpha_1(G|H)}{\alpha_0}}{d\gamma}=\frac{\frac{dN_G}{d\gamma}\times D-\frac{dD}{d\gamma}\times N_G}{D^2}$$

$$=\frac{\frac{dN_G}{d\gamma}\times[\alpha_0\times N_G+(1-\alpha_0)\times N_B]-[\alpha_0\times\frac{dN_G}{d\gamma}+(1-\alpha_0)\times\frac{dN_B}{d\gamma}]\times N_G}{[\alpha_0\times N_G+(1-\alpha_0)\times N_B]^2}=0$$

Multiplying $[\alpha_0\times N_G+(1-\alpha_0)\times N_B]^2$ both sides, we get,

$$\alpha_0\times N_G\times\frac{dN_G}{d\gamma}+(1-\alpha_0)\times N_B\times\frac{dN_G}{d\gamma}-\left[\alpha_0\times N_G\times\frac{dN_G}{d\gamma}-\ (1-\alpha_0)\times N_G\times\frac{dN_B}{d\gamma}\right]=0$$

$$N_B\times\frac{dN_G}{d\gamma}=N_G\times\frac{dN_B}{d\gamma}$$

Replacing N_G and N_B:

$$\{(1-\gamma)\times[\pi_B+(1-\pi_B)\sigma]+\gamma\times\mu_B\}\times\{-[\pi_G+(1-\pi_G)]+(1-\gamma)\left[(1-\pi_G)\times\frac{d\sigma}{d\gamma}\right]+\mu_G\}$$
$$=\{(1-\gamma)\times[\pi_G+(1-\pi_G)\sigma]+\gamma\times\mu_G\}\times\{-[\pi_B+(1-\pi_B)]+(1-\gamma)\left[(1-\pi_B)\times\frac{d\sigma}{d\gamma}\right]+\mu_B\}$$

Solving this, we get:

$$\frac{d\sigma}{d\gamma}[(1-\sigma)^2\times(\pi_B-\pi_G)+\gamma\times(1-\gamma)\{\mu_B\times(1-\pi_G)-\mu_G\times(1-\pi_B)\}$$
$$=\mu_B\times[\pi_G+(1-\pi_G)\sigma]-\mu_G\times[\pi_B+(1-\pi_B)\sigma]$$

7.2 Derivation of Comparative Statics (Proofs of Proposition 4 and 5)

Special Case (1)

Derivation of first-order condition with respect to ϕ

$$\mu_B=\phi\gamma^{-(1+\phi)}\times(1-\gamma)^2$$

Taking a natural log of the equation on both sides,

$$\ln\mu=\ln\phi-(1+\phi)\ln\gamma+2\ln(1-\gamma)$$

Differentiating with respect to ϕ:

$$0=\frac{1}{\phi}-\ln\gamma-(1+\phi)\frac{1}{\gamma}\frac{d\gamma}{d\phi}-\frac{2}{1-\gamma}\frac{d\gamma}{d\phi}$$

$$\left[\frac{1+\phi}{\gamma}+\frac{2}{1-\gamma}\right]\frac{d\gamma}{d\phi}=\frac{1}{\phi}-\ln\gamma$$

Since the left hand side of the equation is positive: $\frac{1}{\phi}-\ln\gamma>0$

And $\frac{1+\phi}{\gamma}+\frac{2}{1-\gamma}>0$ we conclude that $\frac{d\gamma}{d\phi}>0$.

Derivation of first-order condition with respect to μ_B

$$\mu_B=\phi\gamma^{-(1+\phi)}\times(1-\gamma)^2$$

Differentiating with respect to μ_B:

$$1=[-\phi\times(1+\phi)\gamma^{-(2+\phi)}(1-\gamma)^2-2\phi\gamma^{-(1+\phi)}(1-\gamma)]\frac{d\gamma}{d\mu_B}$$

Since $0\le\gamma\le1$ and $\phi>0$ we can conclude that $\frac{d\gamma}{d\mu_B}<0$.

Special Case (2)

Derivation of first-order condition with respect to ϕ

$$\mu\times\left(1-\gamma^{-\phi}\right)=\phi\gamma^{-(1+\phi)}\times(1-\gamma)\times\{1-\gamma(1-\mu)\}]$$

Taking the natural log both sides:

$$\ln\mu+\ln(1-\gamma^{-\phi})=\ln\phi-(1+\phi)\ln\gamma+\ln(1-\gamma)+\ln\{1-\gamma(1-\mu)\}$$

Taking the derivative with respect to ϕ:

$$\frac{1}{1-\gamma^{-\phi}}\times\phi\gamma^{-(1+\phi)}\times\frac{d\gamma}{d\phi}=\frac{1}{\phi}-\ln\gamma-[(1+\phi)\frac{1}{\gamma}+\frac{1}{1-\gamma}+\frac{1-\mu}{1-\gamma+\gamma\mu}]\times\frac{d\gamma}{d\phi}$$

$$\frac{1}{\phi}-\ln\gamma=[\frac{1}{1-\gamma^{-\phi}}\times\phi\gamma^{-(1+\phi)}+(1+\phi)\frac{1}{\gamma}+\frac{1}{1-\gamma}+\frac{1-\mu}{1-\gamma+\gamma\mu}]\times\frac{d\gamma}{d\phi}$$

From this, we conclude that $\frac{d\gamma}{d\phi}>0$.

Derivation of first-order condition with respect to μ

$$\mu\times\left(1-\gamma^{-\phi}\right)=\phi\gamma^{-(1+\phi)}\times(1-\gamma)^2+\phi(1-\gamma)\,\gamma^{-\phi}\mu$$

Rearranging the equation with respect to μ:

$$\mu\times\left(1-\gamma^{-\phi}-\phi(1-\gamma)\,\gamma^{-\phi}\right)=\phi\gamma^{-(1+\phi)}\times(1-\gamma)^2$$

$$\mu=\frac{\phi\gamma^{-(1+\phi)}\times(1-\gamma)^2}{1-\gamma^{-\phi}-\phi(1-\gamma)\,\gamma^{-\phi}}$$

From the equation above, we can conclude that the numerator is a decreasing function of γ and the denominator is an increasing function of γ.

Therefore, $\frac{\phi\gamma^{-(1+\phi)}\times(1-\gamma)^2}{1-\gamma^{-\phi}-\phi(1-\gamma)\,\gamma^{-\phi}}$ is a decreasing function of γ.

If we differentiate the above equation with respect to μ, we get:

$$1=\frac{d}{d\gamma}\left[\frac{\phi\gamma^{-(1+\phi)}\times(1-\gamma)^2}{1-\gamma^{-\phi}-\phi(1-\gamma)\,\gamma^{-\phi}}\right]\frac{d\gamma}{d\mu}$$

With the reasons discussed above, we conclude that $\frac{d\gamma}{d\mu}<0$.

8. Bibliography

Chevalier, Judith A, and Dina Mayzlin. 2006. “The Effect of Word of Mouth on Sales: Online Book Reviews.” *Journal of Marketing Research* 43 (3): 345–54.

Einav, Liran, Chiara Farronato, and Jonathan Levin. 2015. “Peer-to-Peer Markets.” Working Paper 21496. National Bureau of Economic Research. http://www.nber.org/papers/w21496.

Fradkin, Andrey, Elena Grewal, Dave Holtz, and Matthew Pearson. 2015. “Bias and Reciprocity in Online Reviews: Evidence from Field Experiments on Airbnb.” In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, 641–641. ACM. http://dl.acm.org/citation.cfm?id=2764528.

Luca, Michael, and Georgios Zervas. 2016. “Fake It till You Make It: Reputation, Competition, and Yelp Review Fraud.” *Management Science*. http://pubsonline.informs.org/doi/abs/10.1287/mnsc.2015.2304.

Mayzlin, Dina. 2006. “Promotional Chat on the Internet.” *Marketing Science* 25 (2): 155–163.

Mayzlin, Dina, Yaniv Dover, and Judith Chevalier. 2014. “Promotional Reviews: An Empirical Investigation of Online Review Manipulation.” *The American Economic Review* 104 (8): 2421–2455.

Zervas, Georgios, Davide Proserpio, and John Byers. 2015. “A First Look at Online Reputation on Airbnb, Where Every Stay Is Above Average.” SSRN Scholarly Paper ID 2554500. Rochester, NY: Social Science Research Network. https://papers.ssrn.com/abstract=2554500.