

# Understanding Persuasion Cascades in Online Product Rating Systems

Hong Xie,<sup>1,2</sup> Yongkun Li,<sup>3</sup> John C.S. Lui<sup>2</sup>

<sup>1</sup>College of Computer Science, Chongqing University, China

<sup>2</sup>Department of Computer Science and Engineering, The Chinese University of Hong Kong

<sup>3</sup>School of Computer Science and Technology, University of Science and Technology of China  
hongx87@gmail.com, ykli@ustc.edu.cn, cslui@cse.cuhk.edu.hk

## Abstract

Online product rating systems have become an indispensable component for numerous web services such as Amazon, eBay, Google play store and TripAdvisor. One functionality of such systems is to uncover the product quality via product ratings (or reviews) contributed by consumers. However, a well-known psychological phenomenon called “*message-based persuasion*” lead to “*biased*” product ratings in a cascading manner (we call this the *persuasion cascade*). This paper investigates: (1) *How does the persuasion cascade influence the product quality estimation accuracy?* (2) *Given a real-world product rating dataset, how to infer the persuasion cascade and analyze it to draw practical insights?* We first develop a mathematical model to capture key factors of a persuasion cascade. We formulate a high-order Markov chain to characterize the opinion dynamics of a persuasion cascade and prove the convergence of opinions. We further bound the product quality estimation error for a class of rating aggregation rules including the averaging scoring rule, via the matrix perturbation theory and the Chernoff bound. We also design a maximum likelihood algorithm to infer parameters of the persuasion cascade. We conduct experiments on the data from Amazon and TripAdvisor, and show that persuasion cascades notably exist, but the average scoring rule has a small product quality estimation error under practical scenarios.

## Introduction

Online product rating systems have become an indispensable component for numerous web services, i.e., Amazon, TripAdvisor, Yelp, Google Play Store, ..., etc. The online product rating systems reduce the information asymmetry in the sense that product ratings reveal the product quality to consumers while exposing consumer preferences to sellers. Reducing this information asymmetry benefits both the consumers and sellers, i.e., it improves the purchasing experience (Lackermair, Kailer, and Kanmaz 2013; Li et al. 2013; Zhao et al. 2013) and it generates more revenues for sellers (Berger 2012; Luca 2016; Zaroban 2015).

In general, online product rating systems realize the idea of “*wisdom of the crowds*”. Each consumer provides ratings (or reviews) to only a subset of products, and share their experiences or opinions toward these products. These product reviews, once posted, are made public to all consumers, so

that subsequent consumers can refer to them for making purchasing decisions. One main functionality of consumer contributed ratings (or reviews) is to reveal the product quality via the collective wisdoms (Xie and Lui 2015). If consumers provide ratings to reveal their true opinions (or we call “unbiased”), a consensus will be reached when a sufficient number of ratings are collected (Xie and Lui 2015). As dictated by the statistical law of large numbers, this consensus is an accurate estimate of the product quality.

However, a well-known psychological phenomenon called the message-based persuasion (Wood 2000) can change consumers’ opinions and lead to “biased” ratings (or reviews). Evidences of this phenomenon include a number of survey studies (Agag and El-Masry 2017; Rudolph 2016; Shrestha 2016) and empirical studies (Li et al. 2013; Zhao et al. 2013). In fact, reading a product review can be interpreted as exposing to a persuasion message, which may lead to opinion changes as psychology theory has indicated (Hovland, Janis, and Kelley 1953; Wood 2000). This biased review will further serve as a persuasion message to change subsequent consumers’ opinions in a cascading manner, and we call this phenomenon *persuasion cascades*.

Let us use two simple examples to illustrate persuasion cascades. For simplicity, the overall opinion of each consumer is summarized by a binary rating set  $\{-1(\text{“bad”}), 1(\text{“good”})\}$ . Suppose that intrinsically 20% of consumers have an overall opinion of  $-1$ , while 80% have 1 toward a product. In practice, many online product rating systems use the average scoring rule to measure the intrinsic quality, i.e.,  $1 \times 0.8 - 1 \times 0.2 = 0.6$ .

**Example 1 (Unbiased ratings)** *With probability 0.8 (or 0.2) a consumer has an opinion of “good” (or “bad”) and provides a rating 1 (or  $-1$ ). Using the average scoring rule, around 100 ratings can produce an accurate estimator of the intrinsic quality (Xie and Lui 2015), i.e., 0.6. Results in (Xie and Lui 2015) also showed that we can still reveal the intrinsic quality under a small fraction of misbehaving ratings, i.e., intentionally inject 1 to promote or  $-1$  to bad-mouth a product. In this case, a small number of ratings can give a high accuracy to reveal the product quality.*

**Example 2 (Ratings under persuasion)** *For simplicity of illustration, each consumer only reads one review, i.e., the latest review, and provides the same rating as the read re-*

view (i.e., an extremely strong persuasion cascade). If the first rating is 1, then all the subsequent ratings will be 1, otherwise all the subsequent ratings will be  $-1$ . In this case, it is impossible to reveal the intrinsic quality. One misbehaving rating, i.e., intentionally inject 1 (or  $-1$ ), can lead the subsequent ratings to become 1 (or  $-1$ ).

Example 1 and 2 highlight that under these special cases of no cascades and extremely strong cascades: (1) The observed opinions (or ratings) vary from invariant of initial reviews to fully dependent on them; (2) The accuracy of product quality estimation reduces from high to low; (3) The impact of misbehavior varies from tolerable (a small fraction) to destructive (even one misbehaving rating).

This paper explores in the general cases: (1) How does the persuasion cascade influences the product quality estimation accuracy? (2) What are the impacts of initial ratings and misbehaving ratings? (3) How to infer the persuasion cascade from real-world product rating datasets, and analyze it to draw practical insights? Answering these questions are challenging due to the complicated psychological nature of persuasion cascades. One needs to make trade-offs between the model complexity and the mathematical tractability. Our contributions are:

- We develop a mathematical model to capture key factors of persuasion cascades in online product rating systems, under both the *honest rating* and *misbehavior* scenarios. Our model uses empirical findings in the digital world, and has meaningful interpretations.
- We formulate a high-order Markov chain to characterize the opinion dynamics of the persuasion cascade and prove the convergence of opinions. We apply the matrix perturbation theory and the Chernoff bound to bound the product quality estimation error for a class of rating aggregation rules including the averaging scoring rule.
- We develop a maximum likelihood algorithm to infer parameters of the persuasion cascade (i.e., the persuasion strength, etc). We conduct experiments on the Amazon and TripAdvisor datasets, and obtain a number of interesting findings. For example, persuasion cascades notably exist, but the average scoring rule has a small product quality estimation error under practical scenarios.

## Model on Product Ratings

We first model unbiased product ratings. Then, we extend to model product ratings under the message-based persuasion and model persuasion cascades. Finally, we model misbehaving product ratings.

### Unbiased Product Ratings

**Intrinsic opinions.** Without loss of generality, we focus on one product denoted by  $\mathcal{P}$ . Let  $(\mathbf{r}_i, R_i), \forall i \in \mathbb{N}_+$ , denote the  $i$ -th review for the product  $\mathcal{P}$ , where  $\mathbf{r}_i$  denotes the review texts and  $R_i \in \mathcal{M} \triangleq \{1, \dots, M\}, \forall M \in \mathbb{N}_+$ , denotes the associated numerical rating. The review text  $\mathbf{r}_i$  is a vector of words, and the rating  $R_i$  summarizes the overall opinion of  $\mathbf{r}_i$ . A higher rating means that the review is more positive. For example, a typical  $M = 5$  level cardinal rating metric

is  $\{1 = \text{“Terrible”}, 2 = \text{“Poor”}, 3 = \text{“Average”}, 4 = \text{“Good”}, 5 = \text{“Excellent”}\}$ . Let  $O_i \in \mathcal{M}$  denote the intrinsic opinion of the consumer who provides  $(\mathbf{r}_i, R_i)$ . Note that  $O_i$  is a hidden variable and  $(\mathbf{r}_i, R_i)$  is public to all consumers.

**Definition 1** A consumer who provides  $(\mathbf{r}_i, R_i)$  is “unbiased”, if  $R_i = O_i$ , otherwise she is “biased”.

Namely, a consumer is unbiased if she provides her intrinsic opinion. We consider the case that  $O_i$ 's are IID random variables with a probability mass function (pmf):

$$\mathbb{P}[O_i = m] = \alpha_m, \quad \forall m \in \mathcal{M}, i \in \mathbb{N}_+,$$

where  $\alpha_m \geq 0$  and  $\sum_{m \in \mathcal{M}} \alpha_m = 1$ . The row vector  $\boldsymbol{\alpha} \triangleq [\alpha_1, \dots, \alpha_M]$  characterizes the collective opinion of the whole consumer population toward the product  $\mathcal{P}$  and it is also a hidden vector as  $O_i$ 's are hidden. For example, consider  $M = 5$ , then  $\boldsymbol{\alpha} = [0.1, 0.1, 0.1, 0.1, 0.6]$  models that 60% consumers have the intrinsic rating of 5 (i.e., “Excellent”). Let  $\mathcal{O} \triangleq \{\boldsymbol{\alpha} \mid \boldsymbol{\alpha} \in [0, 1]^M, \sum_{m \in \mathcal{M}} \alpha_m = 1\}$  denote a space of all possible collective opinions. We next define a notation to compare collective opinions.

**Definition 2** Consider  $\boldsymbol{\alpha}, \tilde{\boldsymbol{\alpha}} \in \mathcal{O}$ , we say  $\boldsymbol{\alpha} \succeq \tilde{\boldsymbol{\alpha}}$  iff

$$\sum_{j=m}^M \alpha_j \geq \sum_{j=m}^M \tilde{\alpha}_j, \quad \forall m \in \mathcal{M}.$$

For example, consider the collective opinion  $\tilde{\boldsymbol{\alpha}} = [0.6, 0.1, 0.1, 0.1, 0.1]$ , over product  $\tilde{\mathcal{P}}$ , which models that 60% consumers have the intrinsic rating of 1 (i.e., “Terrible”). One can check that  $\boldsymbol{\alpha} \succeq \tilde{\boldsymbol{\alpha}}$ , where  $\boldsymbol{\alpha} = [0.1, 0.1, 0.1, 0.1, 0.6]$  for product  $\mathcal{P}$ . In other words, the collective opinion  $\boldsymbol{\alpha}$  is more positive than  $\tilde{\boldsymbol{\alpha}}$ , i.e., product  $\mathcal{P}$  is more favorable than  $\tilde{\mathcal{P}}$ .

**Intrinsic product quality.** Let  $A : \mathcal{O} \rightarrow [1, M]$  denote an opinion aggregation rule, which prescribes an overall product quality indicator for each collective opinion. For example, a widely practiced opinion aggregation rule is the average scoring rule, i.e.,  $A(\boldsymbol{\alpha}) = \sum_{m \in \mathcal{M}} m \alpha_m, \forall \boldsymbol{\alpha} \in \mathcal{O}$ .

**Definition 3** Given an opinion aggregation rule  $A$  and a collective opinion  $\boldsymbol{\alpha} \in \mathcal{O}$  over  $\mathcal{P}$ , we define the intrinsic quality of  $\mathcal{P}$  as  $A(\boldsymbol{\alpha})$ .

For example, given  $\boldsymbol{\alpha} = [0.1, 0.1, 0.1, 0.1, 0.6]$  and the average scoring rule  $A$ , the intrinsic of quality of  $\mathcal{P}$  is  $A(\boldsymbol{\alpha}) = 4$ .

**Assumption 1** The rating aggregation rule  $A$  satisfies that

$$\boldsymbol{\alpha} \succeq \tilde{\boldsymbol{\alpha}} \Rightarrow A(\boldsymbol{\alpha}) \geq A(\tilde{\boldsymbol{\alpha}}), \quad \forall \boldsymbol{\alpha}, \tilde{\boldsymbol{\alpha}} \in \mathcal{O}.$$

Assumption 1 states that if consumers' collective opinion toward a product is more positive than the other one, then the overall product quality should also be larger. For example, given  $\boldsymbol{\alpha} = [0.1, 0.1, 0.1, 0.1, 0.6]$ ,  $\tilde{\boldsymbol{\alpha}} = [0.6, 0.1, 0.1, 0.1, 0.1]$ , for the average scoring rule  $A$ , we have  $A(\boldsymbol{\alpha}) > A(\tilde{\boldsymbol{\alpha}})$ .

**Definition 4** Let  $\hat{\boldsymbol{\alpha}}_i \triangleq [\hat{\alpha}_{i,1}, \dots, \hat{\alpha}_{i,M}]$  denote the empirical collective opinion summarized from  $R_1, \dots, R_i$ :

$$\hat{\alpha}_{i,m} \triangleq |\{j \mid R_j = m, j \leq i\}| / i, \quad \forall m \in \mathcal{M}, i \in \mathbb{N}_+. \quad (1)$$

In real-world online product rating systems, the average scoring rule is widely applied to estimate the product quality. One of our objectives is to study the accuracy of this method under the persuasion cascade (we will model it next).

## Honest Rating under Persuasion Cascades

Here, we consider all users provide product rating (or reviews) honestly.

**Review reading behavior.** Let  $N_i \in \mathcal{N} \triangleq \{0, 1, \dots, N\}$ , where  $N \in \mathbb{N}_+$ , denote the number of reviews that the consumer reads before posting the  $i$ -th product review  $(\mathbf{r}_i, R_i)$ . Here  $N$  models the maximum number of reviews that a consumer may read and  $N_i$  is a hidden variable. We inject  $N$  null reviews to product  $\mathcal{P}$  denoted by

$$(\mathbf{r}_i, R_i) \triangleq (\mathbf{0}, 0), \quad \forall 1 - N \leq i \leq 0,$$

to deal with the problem that initially product  $\mathcal{P}$  does not have enough reviews for a consumer to read, i.e.,  $i - 1 < N_i$ . For example, if  $N = 3$ , we inject three null reviews  $(\mathbf{r}_{-2}, R_{-2}), (\mathbf{r}_{-1}, R_{-1}), (\mathbf{r}_0, R_0)$ . Suppose that  $N_1 = 3$  and  $N_2 = 2$ , then the first consumer reads three null reviews  $(\mathbf{r}_{-2}, R_{-2}), (\mathbf{r}_{-1}, R_{-1}), (\mathbf{r}_0, R_0)$ , and the second consumer reads two reviews  $(\mathbf{r}_0, R_0), (\mathbf{r}_1, R_1)$ . We consider  $N_i$ 's as IID random variables with a pmf

$$\mathbb{P}[N_i = n] = \beta_n, \quad \forall n \in \mathcal{N}, i \in \mathbb{N}_+,$$

where  $\beta_n \in [0, 1]$  and  $\sum_{n \in \mathcal{N}} \beta_n = 1$ . The row vector  $\beta \triangleq (\beta_0, \beta_1, \dots, \beta_N)$  characterizes the collective reading behavior of the whole consumer population and it is a hidden vector.

Let us focus on a typical review reading behavior, which we call *reading in reverse chronological order*. Precisely, before providing the product review  $(\mathbf{r}_i, R_i)$ , the consumer reads the set of historical reviews denoted by

$$\mathbf{H}_{i-1}^{[N_i]} \triangleq [(\mathbf{r}_{i-N_i}, R_{i-N_i}), \dots, (\mathbf{r}_{i-1}, R_{i-1})], \quad \forall i \in \mathbb{N}_+.$$

Note that the superscript  $[N_i]$  means selecting the latest  $N_i$  reviews. This reading order captures: (1) The evidence that the recency in product reviews is a critical factor to attract consumers' readings and trust (BrightLocal 2016; Shrestha 2016); (2) The feature of online product review systems that presents the reviews in a reverse chronological order. The justifications are: (1) Recent reviews are more likely to reflect the true state of a product (Gebicki 2015); (2) Recent reviews are more likely to represent the opinion of consumers with higher level of expertise (McAuley and Leskovec 2013). As we will see that this simple review reading model can already reveal a number of fundamental understandings on persuasion cascades. In general, the review reading behavior can be quite complicated due to: (1) The complex nature of psychological behavior; (2) Some systems use both the reverse chronological order and other ways, e.g., helpfulness order, etc., to present the reviews.

**Message-based persuasion.** Reading reviews can be interpreted as message-based persuasions, where reviews serve as persuasion messages. Message-based persuasions influence consumers' reviews through the persuaded opinion formation. To model it, we use  $\hat{R}_i \in \mathcal{M}$  to denote the overall opinion formed from reading reviews  $\mathbf{H}_{i-1}^{[N_i]}$ . We consider  $\hat{R}_i$ 's as independent random variables with pmf:

$$\mathbb{P}[\hat{R}_i = m | \mathbf{H}_{i-1}^{[N_i]}] = \theta_m(\mathbf{h}_{i-1}^{[N_i]}), \quad \forall m \in \mathcal{M}, i \in \mathbb{N}_+,$$

where  $\mathbf{h}_{i-1}^{[N_i]} \triangleq [R_{i-N_i}, \dots, R_{i-1}]$  denotes a vector of historical ratings associated with  $\mathbf{H}_{i-1}^{[N_i]}$ ,  $\theta_m(\mathbf{h}_{i-1}^{[N_i]}) \in [0, 1]$  and  $\sum_{m \in \mathcal{M}} \theta_m(\mathbf{h}_{i-1}^{[N_i]}) = 1$ . Namely, we model the message-based persuasion by considering the product rating only, because ratings summarize the opinions of review texts.

**Assumption 2** Let  $\theta(\mathbf{h}_{i-1}^{[N_i]}) \triangleq [\theta_1(\mathbf{h}_{i-1}^{[N_i]}), \dots, \theta_M(\mathbf{h}_{i-1}^{[N_i]})]$  denote the persuaded opinion function. It satisfies

$$\mathbf{h}_{i-1}^{[N_i]} = \mathbf{h}_{j-1}^{[N_j]} \Rightarrow \theta(\mathbf{h}_{i-1}^{[N_i]}) = \theta(\mathbf{h}_{j-1}^{[N_j]}), \quad \forall i, j \in \mathbb{N}_+.$$

Assumption 2 implies that consumers form the same persuaded opinion (in distribution) if they read two review histories with the same ratings. One simple example of  $\theta(\mathbf{h}_{i-1}^{[N_i]})$  is to follow the empirical opinions, i.e.,

$$\theta_m(\mathbf{h}_{i-1}^{[N_i]}) = \frac{\#[\text{ratings in } \mathbf{h}_{i-1}^{[N_i]} \text{ equal } m]}{\#[\text{non-zero ratings in } \mathbf{h}_{i-1}^{[N_i]}}. \quad (2)$$

Eq. (2) captures consumers' herding behavior in opinion formation from reading reviews. For simplicity of notations, we define  $\mathbf{h}_{i-1}^{[0]} \triangleq \mathbf{0}$  to capture that a consumer does not read any reviews (i.e.,  $N_i = 0$ ) or reads only null reviews, and we define the persuaded opinions  $\theta(\mathbf{0}) \triangleq \alpha$  correspondingly.

**Persuasion cascades.** After purchasing a product, a consumer's final opinion is modeled as a weighted combination of her persuaded opinion and her intrinsic opinion:

$$\begin{aligned} \mathbb{P}[R_i = m | \mathbf{H}_{i-1}^{[N_i]}] \\ \triangleq \underbrace{\gamma \mathbb{P}[\hat{R}_i = m | \mathbf{H}_{i-1}^{[N_i]}]}_{\text{persuaded opinion}} + (1 - \gamma) \underbrace{\mathbb{P}[O_i = m]}_{\text{intrinsic opinion}}, \end{aligned}$$

where  $\gamma \in [0, 1]$  models the strength of the message-based persuasion. For example, the case  $\gamma = 0$  models that consumers post intrinsic opinions, while  $\gamma = 1$  models that consumers fully follow the historical reviews. The persuasion strength increases in  $\gamma$ . This message-based persuasion effect leads to a persuasion cascade, i.e., the previous reviews influence a consumer's opinion, and this consumer's opinion will influence the opinions of the subsequent consumers. One of our objectives is to explore the impact of  $\gamma$  on the accuracy of product quality estimation using  $A(\hat{\alpha}_i)$ .

## Misbehaving Product Ratings

Misbehavior (also known as review spam (Jindal and Liu 2007)) exists in online product review systems. Some companies may hire others to provide positive reviews to promote their own products, while providing negative reviews to bad-mouth their competitors' products. We consider a  $(L, \tilde{m}, k)$ -misbehavior model, which is defined as follows.

**Definition 5**  $(L, \tilde{m}, k)$ -misbehavior is to inject a consecutive sequence of  $L$  reviews with rating  $\tilde{m}$  toward a product starting from the  $k$ -th review, where  $L \in \mathbb{N}_+$ ,  $\tilde{m} \in \mathcal{M}$ ,  $k \in \mathbb{N}_+$ .

For example, a  $(6, 5, 4)$ -misbehavior means injecting six reviews with rating 5 starting from the 4-th review, i.e.,  $R_4 =$

$5, \dots, R_9 = 5$ . The  $(L, \tilde{m}, k)$ -misbehavior is a strong review attack, because in practice, to avoid being caught or detected, the  $L$  misbehaving reviews may be divided into several sub-groups and each group will be injected after some honest reviews. If the product quality estimation is robust to this kind of  $(L, \tilde{m}, k)$ -misbehavior, then we can imagine that it will be robust to lighter review attacks as well. Thus, we aim to understand the impact of the  $(L, \tilde{m}, k)$ -misbehavior.

## Analyzing the Persuasion Cascade

We formulate a high-order Markov chain to track opinion dynamics of the persuasion cascade. We study the convergence of the collective opinions and apply the matrix perturbation theory and the Chernoff bound (Markov dependency version) to bound the error of product quality estimation.

### Characterizing Opinion Dynamics

We characterize the status of a persuasion cascade via the latest  $N$  ratings. For example, if product  $\mathcal{P}$  has  $i$  ratings, the status of the persuasion cascade is  $[R_{i-N+1}, \dots, R_i]$ . Precisely, we define a Markov chain  $(\mathcal{S}, \mathbf{P})$ . The *state space*  $\mathcal{S}$  characterizes all the possible status of a persuasion cascade, and it can be derived as  $\mathcal{S} = \bigcup_{j=0}^N \{0\}^{N-j} \times \mathcal{M}^j$ .

The persuasion cascade is at state  $\mathbf{s} \triangleq [s_1, \dots, s_N] \in \mathcal{S}$ , if the latest  $N$  historical ratings are  $s_1, \dots, s_N$ . For example, each product is initialized with no reviews, namely, the initial state is  $\mathbf{s} = [0, \dots, 0]$ . After the first consumer posts the review  $(\mathbf{r}_1, R_1)$ , the state becomes i.e.,  $\mathbf{s} = [0, \dots, 0, R_1]$ . The *transition matrix*  $\mathbf{P}$  characterizes the evolving dynamics of opinions, i.e.,  $\mathbf{P} = [P(\tilde{\mathbf{s}}|\mathbf{s}) : \mathbf{s} \in \mathcal{S}, \tilde{\mathbf{s}} \in \mathcal{S}]$ , where

$$P(\tilde{\mathbf{s}}|\mathbf{s}) \triangleq \mathbb{P}[\text{the next state is } \tilde{\mathbf{s}} | \text{current state } \mathbf{s}], \quad \forall \mathbf{s}, \tilde{\mathbf{s}} \in \mathcal{S}.$$

Each row of  $\mathbf{P}$  corresponds to all possible transitions from a state to another state, i.e.,  $P(\tilde{\mathbf{s}}|\mathbf{s}), \forall \tilde{\mathbf{s}} \in \mathcal{S}$ . To derive the closed-form of  $P(\tilde{\mathbf{s}}|\mathbf{s})$ , we define the neighboring set of  $\mathbf{s}$  as  $\mathcal{N}_{\mathbf{s}} \triangleq \{[s_2, \dots, s_N, m] | m \in \mathcal{M}\}$ . The number of neighbors for each state is  $|\mathcal{N}_{\mathbf{s}}| = M$ . Each state is only possible to transit to one of its neighbors, namely,  $P(\tilde{\mathbf{s}}|\mathbf{s}) = 0, \forall \tilde{\mathbf{s}} \notin \mathcal{N}_{\mathbf{s}}$ , and the transition happens if and only if the product  $\mathcal{P}$  receives a new rating equals. Recall the stationary assumption (i.e. Assumption 2), without loss of generality, suppose the latest  $N$  ratings are  $[R_{i-N}, \dots, R_{i-1}] = \mathbf{s}$ , then with some basic probability arguments, we have

$$P(\tilde{\mathbf{s}}|\mathbf{s}) = \sum_{n \in \mathcal{N}} \beta_n \left( \gamma \theta_{\tilde{s}_N}(\mathbf{s}^{[n]}) + (1 - \gamma) \alpha_{\tilde{s}_N} \right), \quad \forall \tilde{\mathbf{s}} \in \mathcal{N}_{\mathbf{s}},$$

where  $\mathbf{s}^{[n]} \triangleq [s_{N-n+1}, \dots, s_N], \forall n = 1, \dots, N$ , is defined as the latest  $n$  elements of  $\mathbf{s}$  and  $\mathbf{s}^{[0]} \triangleq \mathbf{0}$ .

### Honest Product Ratings

**(1) Convergence of opinions.** Suppose consumers rate honestly, i.e., there is no misbehaving reviews. Without loss of generality, we impose the following assumption on collective intrinsic opinions.

**Assumption 3** *The  $\alpha$  satisfies that  $\alpha_m > 0, \forall m \in \mathcal{M}$ .*

Assumption 3 states that there are no redundant rating levels. Technically, if there are some redundant rating levels, i.e.,  $\alpha_m = 0$  for some  $m \in \mathcal{M}$ , one can eliminate them to make Assumption 3 hold.

**Theorem 1** *Suppose Assumption 3 holds and  $\gamma < 1$ . There exists a unique constant vector  $\hat{\alpha}_{\infty} \in \mathcal{O}$  such that  $\hat{\alpha}_i \xrightarrow{a.s.} \hat{\alpha}_{\infty}$ . Furthermore,  $\hat{\alpha}_{\infty}$  can be derived as  $\hat{\alpha}_{\infty} = \pi \mathbf{W}$ , where  $\pi$  is the first row of the matrix  $\mathbf{P}^{\infty} \triangleq \lim_{i \rightarrow \infty} \mathbf{P}^i$  and  $\mathbf{W} \triangleq [W_{\mathbf{s}, m} : \mathbf{s} \in \mathcal{S}, m \in \mathcal{M}]$  denotes a  $|\mathcal{S}| \times M$  weight matrix with  $W_{\mathbf{s}, m} = \mathbf{1}_{\{s_N = m\}}$ .*

Theorem 1 states that the persuasion cascade will converge to a unique collective opinion, which is invariant of initial ratings. Furthermore, it derives a closed-form converged collective opinion, which enables us to bound product quality estimation error. Due to page limit, we present all the proofs in our technical report (Xie, Li, and Lui 2018).

**(2) Bound product quality estimation error.** In practice,  $A(\hat{\alpha}_i)$  is widely adopted to estimate the intrinsic product quality  $A(\alpha)$ . We now bound the estimation error of this method, i.e.,  $|A(\hat{\alpha}_i) - A(\alpha)|$ . Using the triangle inequality we decompose the estimation error as

$$|A(\hat{\alpha}_i) - A(\alpha)| = \epsilon_a + \epsilon_p, \quad (3)$$

where  $\epsilon_a \triangleq |A(\hat{\alpha}_i) - A(\hat{\alpha}_{\infty})|$  is defined as the *approximation error* caused by the partial information (i.e., finite number of ratings) and  $\epsilon_p \triangleq |A(\hat{\alpha}_{\infty}) - A(\alpha)|$  is defined as the *persuasion error* caused by persuasion biases.

Based on Theorem 1, we apply the matrix perturbation theory (Stewart 1990) to bound  $\epsilon_p$ . Let  $\mathbf{P}_{\mathcal{R}} \triangleq [P(\tilde{\mathbf{s}}|\mathbf{s}) : \mathbf{s} \in \mathcal{R}, \tilde{\mathbf{s}} \in \mathcal{R}]$ , where  $\mathcal{R} \triangleq \mathcal{M}^N$ . The physical meaning of  $\mathcal{R}$  is that the persuasion cascade will visit the state in  $\mathcal{R}$  infinity often, while visiting the state out of  $\mathcal{R}$  for at most once. Let  $\tilde{\mathbf{P}}_{\mathcal{R}} \triangleq [P(\tilde{\mathbf{s}}|\mathbf{s}; \gamma = 0) : \mathbf{s} \in \mathcal{R}, \tilde{\mathbf{s}} \in \mathcal{R}]$  denote the state transition matrix restricted to  $\mathcal{R}$  and  $\gamma = 0$ . Namely,  $\tilde{\mathbf{P}}_{\mathcal{R}}$  is the station transition matrix (restricted to  $\mathcal{R}$ ) when consumers provide unbiased reviews. Let  $\mathbf{E} \triangleq \mathbf{P}_{\mathcal{R}} - \tilde{\mathbf{P}}_{\mathcal{R}}$  denote the error matrix caused by the persuasion bias. The error matrix  $\mathbf{E}$  reflects the persuasion strength  $\gamma$  and it converges to a matrix with all zero entries as  $\gamma$  converges to zero.

**Theorem 2** *Consider the average scoring rule  $A$ , the persuasion error  $\epsilon_p$  can be bounded as  $\epsilon_p \leq O(\|\mathbf{E}\|)$ .*

Please refer to our technical report (Xie, Li, and Lui 2018) for details on  $O(\|\mathbf{E}\|)$ . Theorem 2 derives an upper bound for the persuasion error  $\epsilon_p$  under the average scoring rule. It implies that when the persuasion strength  $\gamma$  is small, the persuasion error  $\epsilon_p$  can be small. For other rating aggregation rules such as the median rating rule, weighted average scoring rule, etc., one can extend theorem 2 to derive similar bounds on  $\epsilon_p$ .

Based on theorem 1, we now apply the Chernoff bound (the Markovian version) (Paulin 2015) to bound the approximation error  $\epsilon_a$ . To study the convergence of a persuasion cascade, we first quantify *the distance to convergence* and *the speed of convergence*. Let the  $|\mathcal{R}|$ -dimensional row vector  $\pi_0 \triangleq [\pi_0(\mathbf{s}) : \mathbf{s} \in \mathcal{R}]$ , where  $\pi_0 \in [0, 1]^{|\mathcal{R}|}$  and  $\sum_{\mathbf{s} \in \mathcal{R}} \pi_0(\mathbf{s}) = 1$ , denote the joint distribution of the first

$N$  ratings. Namely,  $\pi_0$  characterizes the initial state of the persuasion cascade with restriction to  $(\mathcal{R}, \mathbf{P}_{\mathcal{R}})$ . Eventually, the persuasion cascade will converge to a steady-state distribution  $\pi_{\mathcal{R}}$ . To quantify how far away the persuasion cascade is from convergence, we define the distance metric as

$D(\pi_0, \pi_{\mathcal{R}}) \triangleq \sqrt{\sum_{s \in \mathcal{R}} \pi_0^2(s) / \pi_{\mathcal{R}}(s)}$ . Using the Cauchy-Schwarz inequality, one obtains that when  $\pi_0 = \pi_{\mathcal{R}}$ , the minimum distance is attained  $D(\pi_0, \pi_{\mathcal{R}}) = 1$ . Let  $\nu_{ps}$  denote the pseudo spectral gap (Paulin 2015) of the Markov chain  $(\mathcal{R}, \mathbf{P}_{\mathcal{R}})$ , where  $\nu_{ps}$  quantifies the speed of convergence of the persuasion cascade. With these two metrics, we now bound the opinion error caused by approximation.

**Theorem 3** *Suppose the number of ratings  $i$  satisfies*

$$i \geq \frac{1}{\epsilon^2} \frac{16(1 + 1/\nu_{ps}) + 40\epsilon}{\nu_{ps}} \ln \left( \frac{\sqrt{2}MD(\pi_0, \pi_{\mathcal{R}})}{\delta} \right) + N, \quad (4)$$

where  $\epsilon > 0$  and  $\delta \in [0, 1]$ . Consider  $A$  to be the average scoring rule,  $\epsilon_a \leq O(\epsilon)$  holds with probability at least  $1 - \delta$ .

Please refer to our technical report (Xie, Li, and Lui 2018) for details on  $O(\epsilon)$ . Theorem 3 states that the approximation error  $\epsilon_a$  can be arbitrarily small with a sufficiently large number of ratings under the average scoring rule. It derives “minimum number of ratings” to guarantee a small  $\epsilon_a$ . It reveals that the minimum number of ratings is critical to the estimation accuracy  $\epsilon$  and the speed of convergence (i.e., eigenvalue pseudo spectral gap  $\nu_{ps}$ ), while it is not critical to the distance to convergence (i.e.,  $\ln D(\pi_0, \pi_{\mathcal{R}})$ ), because it is proportional to the logarithmic of the distance, i.e.,  $\ln D(\pi_0, \pi_{\mathcal{R}})$ . Theorem 3 can be extended to median rule, weighted average scoring rule, etc.

**(3) Summary of honest rating.** The persuasion cascade converges to a unique collective opinion, which is invariant of the initial rating. The  $A(\hat{\alpha}_i)$  is an accurate estimator of the intrinsic quality  $A(\alpha)$  when the persuasion strength is weak and the number of ratings is sufficiently large.

## Misbehavior Attacks

We first analyze the impact of the  $(L, \tilde{m}, k)$ -misbehavior on the convergence of the persuasion cascade.

**Corollary 1** *Suppose Assumption 3 holds and  $\gamma < 1$ . The convergence of the persuasion cascade and  $\hat{\alpha}_{\infty}$  remain unchanged under any  $(L, \tilde{m}, k)$ -misbehavior with  $L < \infty$ .*

Corollary 1 states that the convergence of the persuasion cascade is invariant of any  $(L, \tilde{m}, k)$ -misbehavior with a finite length. This implies that the persuasion error  $\epsilon_p$  is invariant of the  $(L, \tilde{m}, k)$ -misbehavior. However, the misbehavior slows down the convergence of opinions. We will need to compensate this by a larger number of ratings to make the approximation error  $\epsilon_a$  small. One can apply Theorem 3 to derive the number of ratings needed to compensate.

**Summary of misbehavior.** The convergence of persuasion cascades is invariant of the misbehavior attack, except the case  $\gamma = 1$ . The misbehavior slows down the convergence.

## Inferring the Persuasion Cascade

In this section, we design maximum likelihood algorithms to infer the persuasion cascade, i.e., the parameters  $\alpha, \beta$  and  $\gamma$ , from real-world product rating datasets.

**Maximum likelihood estimation.** Without loss of generality, suppose we are given a sequence of  $I \in \mathbb{N}_+$  ratings of the product  $\mathcal{P}$  in chronological order, i.e.,  $R_1, \dots, R_I$ . We do not assume any a-priori knowledge on the persuasion strength  $\gamma$ , the review reading distribution  $\beta$ , and the intrinsic collective opinion  $\alpha$ . Our objective is to infer these parameters from  $R_1, \dots, R_I$  via maximum likelihood estimation. In the remaining parts of this paper, we consider a specific form of persuaded opinion function, i.e.,  $\theta(\mathbf{h}_i)$  satisfies (2), which capture consumers’ herding or following behavior in opinion formation from review readings. Given  $R_1, \dots, R_I$ , we derive the log-likelihood function as

$$\begin{aligned} \ell(\alpha, \beta, \gamma) &= \sum_{i=1}^I \ln \left[ \sum_{n \in \mathcal{N}} \beta_n \left( \gamma \theta_{R_i}(\mathbf{h}_{i-1}^{[n]}) + (1 - \gamma) \alpha_{R_i} \right) \right]. \end{aligned}$$

**Problem 1** *Given  $R_1, \dots, R_I$  of the product  $\mathcal{P}$ ,  $\mathcal{M}, \mathcal{N}$  and  $\theta(\mathbf{h}_i)$  which satisfies (2), select  $\alpha, \beta$  and  $\gamma$  to maximize  $\ell(\alpha, \beta, \gamma)$ . Formally,*

$$\begin{aligned} &\text{maximize}_{\alpha, \beta, \gamma} \quad \ell(\alpha, \beta, \gamma) \\ &\text{subject to} \quad \sum_{n \in \mathcal{N}} \beta_n = 1, \sum_{m \in \mathcal{M}} \alpha_m = 1, \\ &\quad \beta \in [0, 1]^{|\mathcal{M}|}, \alpha \in [0, 1]^{|\mathcal{M}|}, \gamma \in [0, 1]. \end{aligned}$$

In case we are given the ratings for multiple products, one can repeat Problem 1 for each product separately.

**Theorem 4** *Problem 1 has at least one optimal solution.*

Problem 1 has at least one optimal solution, because its feasible space is compact and its objective function  $\ell(\alpha, \beta, \gamma)$  is continuous over the feasible space.

**Inference algorithms.** Note that Problem 1 is not a concave program. In general, the optimal solution is not unique and there may exist several local optimal solutions as well. We employ gradient-based methods to search local optimal solutions. Repeating the searching algorithm for multiple times with random start, we may hit the optimal solution with high probability. For brevity, we present other gradients in our technical report (Xie, Li, and Lui 2018). With these gradients, we use the interior point algorithmic framework (Wright and Nocedal 1999) to search the optimal solution.

## Real-World Data Analytics

We apply our framework to analyze real-world datasets from Amazon and TripAdvisor. We infer the persuasion strength  $\gamma$ , the distribution of number of review readings  $\beta$  and the intrinsic opinion  $\alpha$ . We analyze the inferred  $\alpha, \beta$  and  $\gamma$  to gain practical understandings on the persuasion cascade.

## Data Analysis Setting

**Datasets.** The datasets contain historical ratings of 32,888 products in Amazon and of 11,543 hotels in TripAdvisor. We

plot the distribution of the number of ratings across items (i.e., products or hotels) in Figure 1(a). One can observe that around 80% items have less than 300 ratings. From these two datasets, we select all the items that have no less than 1000 ratings for analysis. This is to attain a balance between the data sufficiency for the inference algorithm to be accurate and the scale of the experiment dataset. In total, we select 738 products from the Amazon dataset and 507 hotels from the TripAdvisor dataset, which are depicted in Figure 1(b). One observes that most selected items have around 1000 ratings.

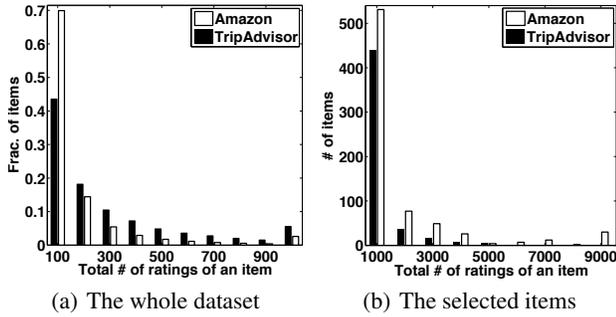


Figure 1: Distribution of number of ratings across items.

### Parameter setting and persuasion cascade inference.

We consider the following experimental settings. Amazon and TripAdvisor adopt the same five-level rating metric  $\{1, 2, 3, 4, 5\}$ . We therefore set  $\mathcal{M} = \{1, 2, 3, 4, 5\}$ . To capture the finding that around 98% consumers read less than 50 reviews of a product (Rudolph 2016), we set  $\mathcal{N} = \{0, 1, \dots, 50\}$ . We consider a simple persuaded opinion formation function  $\theta(\mathbf{h}_i)$  that satisfies (2), to capture consumers' following (or herding) behavior in persuaded opinion formation. For each selected item in Figure 1(b), we input the above parameters and the whole rating history of this item into the inference algorithm. Through this we infer persuasion strength  $\gamma$ , the review reading distribution  $\beta$  and the intrinsic opinion  $\alpha$  for each item.

### Persuasion Strengths

Note that each item may be associated with a different persuasion strength. Figure 2 shows the distribution of the inferred persuasion strength  $\gamma$  across items. Figure 2(a) depicts the pmf of  $\gamma$ , where we discretize  $[0, 1]$  into 20 sub-intervals such that the  $j$ -th sub-interval represents  $[0.05(j-1), 0.05j]$  for all  $j = 1, \dots, 20$ . From Figure 2(a), one can observe the persuasion strength has a normal-like distribution for both Amazon and TripAdvisor. The persuasion strength for most items fall into the range of 0.1 to 0.4. Only a small fraction of items are associated with a persuasion strength larger than 0.6. These observations show that the persuasion cascade widely exists. One possible reason is that consumers rely on product reviews to estimate the product quality, which makes them prone to message-based persuasions (i.e., reading reviews). Figure 2(b) depicts the cumulative distribution function (CDF) of the persuasion strength across items. One can observe that the CDF curve

for Amazon lies above that for TripAdvisor. This implies that on average, the persuasion strength in TripAdvisor is stronger than Amazon. One possible reason is that perceiving the quality of a hotel is more complicated than perceiving the quality of a product. This makes consumers more likely to follow the message-based persuasions.

**Lessons learned:** In both Amazon and TripAdvisor, the persuasion cascade widely exists. The persuasion strength has a normal-like distribution and concentrates in the range of 0.1 to 0.4. The persuasion strength in TripAdvisor is stronger than that in Amazon.

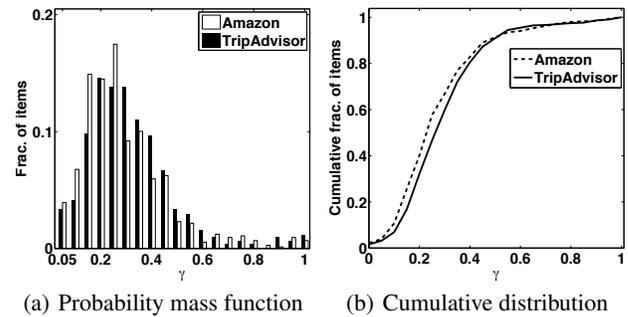


Figure 2: The distribution of  $\gamma$  across items.

### Review Reading Distributions

Based on the inferred  $\beta$  for a product, we define its average number of review readings as  $\bar{N} \triangleq \sum_{n \in \mathcal{N}} n \beta_n$ . Figure 3 shows the distribution of  $\bar{N}$  across items, where we discretize the interval  $[0, 50]$  into 10 sub-intervals such that the  $j$ -th sub-interval corresponds to  $[5(j-1), 5j]$ , for all  $j = 1, \dots, 10$ . From Figure 3(a), one can observe that the distribution of the  $\bar{N}$  across hotels in TripAdvisor has a normal-like shape. Most hotels attract consumers to read around 35 reviews (on average). However, for Amazon, there are a large fraction of products, which attract consumers to read only around 5 reviews (on average). One reason is that these products may have a lot of public information and thus, consumers only read a small number of reviews to estimate their quality. For both Amazon and TripAdvisor, only around 2% items attract consumers to read around 50 reviews. This observation is consistent with the survey study (Rudolph 2016). From Figure 3(b) one can observe that the CDF curve of  $\bar{N}$  for Amazon lies above that for TripAdvisor. This means that consumers tend to read more reviews in TripAdvisor than Amazon. One reason is that the public information of hotels are less than products in Amazon, and thus consumers rely more on the product reviews to estimate the quality of hotels.

**Lessons learned:** The average number of review readings concentrates around 35. Consumers tend to read more reviews in TripAdvisor than Amazon.

### Opinion and Product Quality Estimation

From Figure 1 we observe that around 90% of the items (i.e., products or hotels) have less than 500 ratings. We therefore

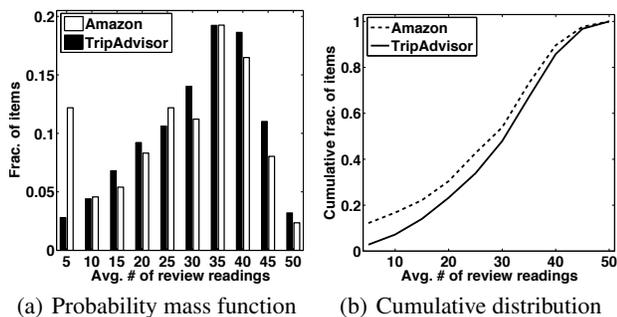


Figure 3: The distribution of  $\tilde{N}$  across items.

explore: (1) Given a small number of ratings (i.e.,  $i \leq 500$ ), how accurate is  $\hat{\alpha}_i$  in reflecting  $\alpha$ ? (2) How accurate is  $A(\hat{\alpha}_i)$  in estimating the intrinsic quality  $A(\alpha)$ ?

We use the inferred  $\alpha$  as the intrinsic opinion. For each selected item, we compute the associated empirical collective opinion  $\hat{\alpha}_i$  for each given number of ratings  $i = 1, \dots, 500$ . Given,  $\alpha$  and  $\hat{\alpha}_i$ , we define the relative opinion error as  $E_o \triangleq \|\hat{\alpha}_i - \alpha\|_1 / \|\alpha\|_1$ , where  $\|\cdot\|_1$  denote the taxicab norm. The opinion error  $E_o \in [0, 2]$ , and the smaller the  $E_o$ , the more accurate  $\hat{\alpha}_i$  is. We define the relative product quality estimation error as  $E_q \triangleq |A(\hat{\alpha}_i) - A(\alpha)| / |A(\alpha)|$ . The product quality error  $E_q \in [0, M - 1]$ , and the smaller the  $E_q$ , the more accurate  $A(\hat{\alpha}_i)$  is. We set  $A$  to be the average scoring rule. Figure 4 shows  $E_o$  and  $E_q$  (average across items) when we vary the number of ratings  $i$  from 1 to 500. From Figure 4, one can observe that when the number of ratings is small, both the opinion error  $E_o$  and the quality error  $E_q$  are large. As we increase the number of ratings,  $E_o$  and  $E_q$  decrease. Namely, the larger the number of ratings, the higher the accuracy. When the number of ratings is less than 100, increasing the number of ratings decreases the error (both  $E_o$  and  $E_q$ ) significantly. Further increasing the number of ratings from 100 to 500 only decreases the error (both  $E_o$  and  $E_q$ ) slightly. When the number of ratings is around 200, the quality error  $E_q$  is around 5%, while the opinion error is around 20%. In other words, 200 ratings can achieve an accurate estimation of the product quality, while the estimation of the opinion requires more ratings.

**Lessons learned:** Consider the average scoring rule. When the number of ratings is small, the product quality estimation error is much smaller than the opinion estimation error. Around two hundred ratings can produce an accurate estimator on the product quality. The average scoring rule is effective at reducing the persuasion bias.

## Related Work

A variety of works studied the benefits of online product review systems. Chevalier *et al.* (Chevalier and Mayzlin 2006) showed that product reviews can increase the sales of sellers. Mudambi *et al.* (Mudambi and Schuff 2010) showed that product reviews serve as an important information source for consumers to make purchase decisions. This argument was further justified by Lackermeir *et al.* (Lacker-

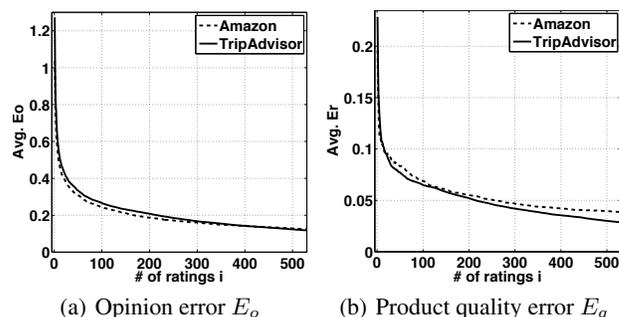


Figure 4: The opinion error and average rating error.

mair, Kailer, and Kanmaz 2013) and Li (Li *et al.* 2013). A number of survey studies (BrightLocal 2016; Rudolph 2016; Shrestha 2016) were conducted to investigate how consumers use the product reviews.

Consumers' reviews are subject to biases due to a number of reasons. One the product side, Guo *et al.* revealed the category-wise biases (Guo and Dunson 2015). On the product review system side, the system interfaces (Cosley *et al.* 2003) or the recommendation algorithms (Shafto and Nasraoui 2016) can lead to opinion biases. On the consumer side, the opinion biases can be caused by the evolving dynamics of their preferences (Koren 2009), or the change in their experience (McAuley and Leskovec 2013). A variety of effective methods, e.g., (Cosley *et al.* 2003; Guo and Dunson 2015; Koren 2009; McAuley and Leskovec 2013; Shafto and Nasraoui 2016), have been proposed to mitigate these biases.

A variety of works explore the review (or opinion) formation from psychological perspectives. Wu *et al.* (Wu and Huberman 2008) found that users tend to post reviews with opinions different from historical reviews. Zhang *et al.* (Zhang, Zhao, and Lui 2017) identified the assimilate and contrast phenomenon in providing product reviews and they applied the "assimilate-contrast" theory to mitigate this bias. A number of experiments (Muchnik, Aral, and Taylor 2013) revealed that consumer's reviews are subject to social influence biases, users tend to provide higher ratings when they are exposed to higher historical ratings, which disagrees with the work of Jurca *et al.* (Jurca *et al.* 2010). Krishnan *et al.* (Krishnan *et al.* 2014) and Wang *et al.* (Wang, Wang, and Wang 2014) developed models to quantify the effect of social influence biases and designed algorithms to mitigate the social influence bias. However, a latest work by Adomavicius *et al.* (Adomavicius *et al.* 2016) showed that the social influence bias can be dominated by the personalized recommendation bias. Their observations motivated us to investigate the dominate biases source. The survey studies (BrightLocal 2016; Rudolph 2016; Shrestha 2016) revealed that consumers read reviews in decision making. This uncovers a typical psychological phenomenon called the message-based persuasion (Wood 2000). As psychology theory (Hovland, Janis, and Kelley 1953; Wood 2000) have indicated, message-based persuasion is powerful in changing a consumer's opinions. We explore it in online ratings.

## Conclusion

This paper develops a data analytic framework to analyze persuasion cascades in online rating systems. We develop a mathematical model to capture key factors of persuasion cascades and we prove the convergence of opinions. We apply the matrix perturbation theory and the Chernoff bound to bound the product quality estimation error for a class of rating aggregation rules including the average scoring rule. We develop a maximum likelihood algorithm to infer the persuasion cascade. We conduct experiments on the data from Amazon and TripAdvisor, and show that persuasion cascades notably exist, but the average scoring rule has a small product quality estimation error under practical scenarios.

## Acknowledgments

The work of John C.S. Lui was supported in part by the GRF Funding 14200117. We would like to thank anonymous reviewers for their comments to improve the paper.

## References

- Adomavicius, G.; Bockstedt, J.; Curley, S. P.; and Zhang, J. 2016. Understanding effects of personalized vs. aggregate ratings on user preferences. In *Proc. of INTRS WORKSHOP*.
- Agag, G. M., and El-Masry, A. A. 2017. Why do consumers trust online travel websites? drivers and outcomes of consumer trust toward online travel websites. *Journal of Travel Research* 56(3):347–369.
- Berger, J. 2012. *Bad Reviews Can Boost Sales. Here's Why*. Harvard Business Review.
- BrightLocal. 2016. *Local Consumer Review Survey*. BrightLocal.
- Chevalier, J. A., and Mayzlin, D. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of marketing research* 43(3):345–354.
- Cosley, D.; Lam, S. K.; Albert, I.; Konstan, J. A.; and Riedl, J. 2003. Is seeing believing?: How recommender system interfaces affect users' opinions. In *Proc. ACM CHI*.
- Gebicki, M. 2015. *Should we trust TripAdvisor reviews?* Traveller.
- Guo, F., and Dunson, D. B. 2015. Uncovering systematic bias in ratings across categories: A bayesian approach. In *Proc. of ACM RecSys*.
- Hovland, C. I.; Janis, I. L.; and Kelley, H. H. 1953. Communication and persuasion; psychological studies of opinion change.
- Jindal, N., and Liu, B. 2007. Review spam detection. In *Proc. of WWW*.
- Jurca, R.; Garcin, F.; Talwar, A.; and Faltings, B. 2010. Reporting incentives and biases in online review forums. *ACM Transactions on the Web (TWEB)* 4(2):5.
- Koren, Y. 2009. Collaborative filtering with temporal dynamics. In *Proc. of ACM KDD*.
- Krishnan, S.; Patel, J.; Franklin, M. J.; and Goldberg, K. 2014. A methodology for learning, analyzing, and mitigating social influence bias in recommender systems. In *Proc. of ACM RecSys*.
- Lackermair, G.; Kailer, D.; and Kanmaz, K. 2013. Importance of online product reviews from a consumer's perspective. *Advances in Economics and Business* 1(1):1–5.
- Li, M.; Huang, L.; Tan, C.-H.; and Wei, K.-K. 2013. Helpfulness of online product reviews as seen by consumers: Source and content features. *International Journal of Electronic Commerce* 17(4):101–136.
- Luca, M. 2016. Reviews, reputation, and revenue: The case of yelp.com.
- McAuley, J. J., and Leskovec, J. 2013. From amateurs to connoisseurs: Modeling the evolution of user expertise through online reviews. In *Proc. of WWW*.
- Muchnik, L.; Aral, S.; and Taylor, S. J. 2013. Social influence bias: A randomized experiment. *Science* 341(6146):647–651.
- Mudambi, S. M., and Schuff, D. 2010. What makes a helpful online review? a study of customer reviews on amazon.com. *MIS Quarterly* 34(1):185–200.
- Paulin, D. 2015. Concentration inequalities for markov chains by marton couplings and spectral methods. *Electronic Journal of Probability* 20.
- Rudolph, S. 2016. *50 Stats You Need to Know About Online Reviews [Infographic]*. Business 2 Community.
- Shafto, P., and Nasraoui, O. 2016. Human-recommender systems: From benchmark data to benchmark cognitive models. In *Proc. of ACM RecSys*.
- Shrestha, K. 2016. *50 Stats You Need to Know About Online Reviews [Infographic]*. Vendasta.
- Stewart, G. W. 1990. Matrix perturbation theory.
- Wang, T.; Wang, D.; and Wang, F. 2014. Quantifying herding effects in crowd wisdom. In *Proc. of ACM KDD*.
- Wood, W. 2000. Attitude change: Persuasion and social influence. *Annual review of psychology* 51(1):539–570.
- Wright, S. J., and Nocedal, J. 1999. Numerical optimization. *Springer Science* 35(67-68):7.
- Wu, F., and Huberman, B. A. 2008. How public opinion forms. In *International Workshop on Internet and Network Economics*, 334–341. Springer.
- Xie, H., and Lui, J. C. S. 2015. Mathematical modeling and analysis of product rating with partial information. *ACM Trans. Knowl. Discov. Data* 9(4):26:1–26:33.
- Xie, H.; Li, Y.; and Lui, J. C. 2018. *Understanding Persuasion Cascades in Online Product Rating Systems*. <https://1drv.ms/b/s!AkqQNKuLPUBeE3OSVZqPwngsqg>.
- Zaroban, S. 2015. *Product reviews boost revenue per online visit 62%*. Digital Commerce.
- Zhang, X.; Zhao, J.; and Lui, J. C. S. 2017. Modeling the assimilation-contrast effects in online product rating systems: Debiasing and recommendations. In *Proc. of ACM RecSys*.
- Zhao, Y.; Yang, S.; Narayan, V.; and Zhao, Y. 2013. Modeling consumer learning from online product reviews. *Marketing Science* 32(1):153–169.