Understanding Persuasion Cascades in Online Product Rating Systems: Modeling, Analysis, and Inference

HONG XIE and MINGZE ZHONG, Chongqing University
YONGKUN LI, University of Science and Technology of China
JOHN C. S. LUI, The Chinese University of Hong Kong

36

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 article 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 both synthetic data and real-world data from Amazon and TripAdvisor. Experiment results show that our inference algorithm has a high accuracy. Furthermore, persuasion cascades notably exist, but the average scoring rule has a small product quality estimation error under practical scenarios.

CCS Concepts: • **Information systems → Web mining**; **Web applications**;

Additional Key Words and Phrases: Online rating systems, persuasion cascades, high order Markov chain, product quality estimation

ACM Reference format:
Hong Xie, Mingze Zhong, Yongkun Li, and John C. S. Lui. 2021. Understanding Persuasion Cascades in Online Product Rating Systems: Modeling, Analysis, and Inference. *ACM Trans. Knowl. Discov. Data* 15, 3, Article 36 (April 2021), 29 pages.
https://doi.org/10.1145/3440887

Hong Xie was supported in part by National Nature Science Foundation of China (61902042), Chongqing Natural Science Foundation (cstc2020jcyj-msxmX0652), and the Fundamental Research Funds for the Central Universities (2020CDJ-LHZZ-057). John C.S. Lui was supported in part by the GRF 14200420.

Authors' addresses: H. Xie and M. Zhong, Chongqing Key Laboratory of Software Theory and Technology, Chongqing University, No. 174 Shazhengjie, Shapingba, Chongqing, China, 400044; emails: xiehong2018@cqu.edu.cn, zmz502071957@gmail.com; Y. Li, School of Computer Science and Technology, University of Science and Technology of China, No. 96, JinZhai Road Baohe District, Hefei, Anhui, China, 230027; email: ykli@ustc.edu.cn; J. C. S. Lui, Department of Computer Science and Engineering, The Chinese University of Hong Kong, Shatin, NT, Hong Kong SAR, China, 000000; email: cslui@cse.cuhk.edu.hk.

1556-4681/2021/04-ART36 $15.00
https://doi.org/10.1145/3440887

1 INTRODUCTION

Online product rating systems have become an indispensable component for numerous web services, e.g., Amazon, TripAdvisor, Yelp, and Google Play Store. 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 [13, 14, 38] and it generates more revenues for sellers [3, 15, 35]. Last but not least, many experts believe that online product rating systems will play an increasingly important role in the future online commerce [24, 34].

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 [33]. 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 [33]. 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 [28] can change consumers' opinions and lead to "biased" ratings (or reviews). Evidences of this phenomenon include a number of survey studies [2, 4, 20, 23, 26] and empirical studies [14, 38]. 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 [9, 28]. 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 [33], i.e., 0.6. Results in [33] 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 Cascades). For simplicity, each consumer only reads one review, i.e., the latest review, and provides the same rating as the read review (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).

Examples 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) and the impact of misbehavior varies from tolerable (a small fraction) to destructive (even one misbehaving rating).

This article 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 tradeoffs 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). Experiments on synthetic data show that our inference algorithm has a high accuracy.
— 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.

This article organizes as follows. Section 2 presents the model to characterize the message-based persuasion. Section 3 analyzes the persuasion cascade theoretically. Section 4 develops a maximum likelihood algorithm to infer the persuasion cascade. Section 5 presents the verification and evaluation studies on synthetic datasets. Section 6 infers and analyzes the persuasion in Amazon and TripAdvisor. Section 7 discusses the related work. Section 8 presents the proof to lemmas and theorems. Section 9 concludes the article.

2 MODEL ON PRODUCT REVIEWS

We first model unbiased product reviews. Then, we extend to model product reviews under the message-based persuasion and model persuasion cascades. Finally, we model misbehaving reviews. Table 1 summarized the main notations of this article.

2.1 Unbiased Product Reviews

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 of product $\mathcal{P}$, where $\mathbf{r}_i$ denotes the review texts and $R_i \in \mathcal{M} \triangleq \{1, \ldots, M\}, \forall M \in \mathbb{N}_+$, 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 = "Terrible," 2 = "Poor," 3 = "Average," 4 = "Good," and 5 = "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."

Table 1. Main Notations

$\mathcal{P}$	notation for a product
$\mathbf{r}_i$	the i-th review of product $\mathcal{P}$
R_i	the rating associated with $\mathbf{r}_i$
$\mathcal{M}$ (M)	the rating metric (the number of rating points)
O_i	the intrinsic opinion of the consumer who provides $(\mathbf{r}_i, R_i)$
$\boldsymbol{\alpha}$	the collective opinion of the whole consumer population toward product $\mathcal{P}$
$\mathcal{O}$	a space of all possible collective opinions
A	an opinion aggregation rule
$\widehat{\boldsymbol{\alpha}}_i$	an empirical collective opinion
N_i	number of reviews that the consumer reads before posting $(\mathbf{r}_i, R_i)$.
$\mathcal{N}$	a set of all possible number of review readings
$\boldsymbol{\beta}$	the review reading distribution
$\mathbf{H}_{i-1}^{[N_i]}$	the set of N_i latest historical reviews assigned before $(\mathbf{r}_i, R_i)$,
$\widehat{R}_i$	the overall opinion formed from reading reviews $\mathbf{H}_{i-1}^{[N_i]}$.
$\mathbf{h}_{i-1}^{[N_i]}$	a vector of historical ratings associated with $\mathbf{H}_{i-1}^{[N_i]}$,
$\theta_m(\mathbf{h}_{i-1}^{[N_i]})$	probability mass function of $\widehat{R}_i$
$\gamma \in [0, 1]$	the strength of the message-based persuasion.
$\widehat{\boldsymbol{\alpha}}_\infty$	a collective opinion vector
$(\mathcal{S}, \mathbf{P})$	(state space, transition probability) of a Markov chain
$\mathbf{s}, \tilde{\mathbf{s}}$	state of the persuasion cascade
$P(\tilde{\mathbf{s}}\|\mathbf{s})$	transition probability from $\mathbf{s}$ to $\tilde{\mathbf{s}}$
$\mathcal{N}_\mathbf{s}$	neighbor set of $\mathbf{s}$
$\boldsymbol{\pi}$	a probability (row) vector
$\mathbf{W}$	a $\|\mathcal{S}\| \times M$ weight matrix
$\mathcal{R}$	recurrent state set
$\mathcal{T}$	transient state set
ϵ_a	the *approximation error*caused by the partial information
ϵ_p	the *persuasion error* caused by persuasion biases
$\tilde{\mathbf{P}}_\mathcal{R}$	the state transition matrix restricted to $\mathcal{R}$
$\tilde{\boldsymbol{\pi}}$	the steady-state distribution under $\gamma = 0$.
$\boldsymbol{\pi}_0$,	the joint distribution of the first N ratings
$D(\boldsymbol{\pi}_0, \boldsymbol{\pi}_\mathcal{R})$	the distance metric
ν_{ps}	the pseudo spectral gap
ϵ, δ	the error, confidence level
ϵ_{sd}	the steady-state distribution error
I	the number of ratings for infer the persuasion cascade
$\mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \gamma)$	the likelihood function
$\ell(\boldsymbol{\alpha}, \boldsymbol{\beta}, \gamma)$	the log-likelihood function
$\widehat{\gamma}$	the inferred persuasion strength.
$\bar{N}, \bar{\gamma}$	average number of review readings, persuasion strength
E_o, E_q	the relative opinion, product quality estimation error

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 $\boldsymbol{\alpha} \triangleq [\alpha_1, \ldots, \alpha_M]$ 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 $\mathcal{O}$ denote a space of all possible collective opinions

$$\mathcal{O} \triangleq \left\{ \boldsymbol{\alpha} \,\middle|\, \boldsymbol{\alpha} \in [0,1]^M, \sum_{m\in\mathcal{M}} \alpha_m = 1 \right\}.$$

We next define a notation to compare the collective opinions over different products.

Definition 2. Consider $\boldsymbol{\alpha}, \tilde{\boldsymbol{\alpha}} \in \mathcal{O}$, we say $\boldsymbol{\alpha} \succeq \tilde{\boldsymbol{\alpha}}$ if and only if

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

Definition 2 defines a way to compare two collective opinion vectors. In fact, it treats each collective opinion vector as a probability distribution over $\mathcal{M}$, and compares two collective opinion vectors via the stochastic ordering of the associated probability distributions. A larger collective opinion vector corresponds to a more positive collective opinion. 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 $A : \mathcal{O} \to [1, M]$ 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 $A(\boldsymbol{\alpha}) > A(\tilde{\boldsymbol{\alpha}})$. Note that Assumption 1 is not restricted. In fact, many rating aggregation rules such as majority rule and median rule satisfy Assumption 1.

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

$$\widehat{\alpha}_{i,m} \triangleq \frac{|\{j | R_j = m, j \leq i\}|}{i}, \quad \forall m \in \mathcal{M}, i \in \mathbb{N}_+. \tag{1}$$

In real-world online product review systems, a simple method, i.e., average scoring rule $A(\widehat{\boldsymbol{\alpha}}_i)$, 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).

2.2 Product Reviews Under Persuasion Cascades

Review reading behavior. Let $N_i \in \mathcal{N} \triangleq \{0, 1, \ldots, N\}$, 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

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

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 $N_1=3$ 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

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

where $\beta_n \geq 0$ and $\sum_{n \in \mathcal{N}} \beta_n = 1$. The row vector $\boldsymbol{\beta} \triangleq (\beta_0, \beta_1, \ldots, \beta_N)$ 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

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

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 [4, 23]; and (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 [7]; and (2) recent reviews are more likely to represent the opinion of consumers with higher level of expertise [16]. 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; and (2) some systems use both the reverse chronological order and other ways, e.g., helpfulness order, to present the reviews. We leave it as a future work to capture these factors.

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

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

where $\mathbf{h}_{i-1}^{[N_i]} \triangleq [R_{i-N_i}, \ldots, R_{i-1}]$ 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 product ratings summarize the opinions of review texts.

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

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

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 $\boldsymbol{\theta}(\mathbf{h}_{i-1}^{[N_i]})$ is to follow the empirical opinions, i.e.,

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

Equation (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 $\boldsymbol{\theta}(\mathbf{0}) \triangleq \boldsymbol{\alpha}$ 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:

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

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 γ. 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 γ on the accuracy of a product quality estimation method $A(\widehat{\boldsymbol{\alpha}}_i)$.

2.3 Misbehaving Product Reviews

Misbehavior (also known as review spam [10]) 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 $(L, \mathcal{F})$-misbehavior model, which is defined as follows.

Definition 5. $(L, \mathcal{F})$-misbehavior is to inject L reviews with ratings characterized by set $\mathcal{F}$, where $L \in \mathbb{N}_+$ and $\mathcal{F}$ is defined as

$$\mathcal{F} \triangleq \left\{\left(i_\tau, \widetilde{R}_{i_\tau}\right) : \tau = 1, \ldots, L\right\}$$

with $\widetilde{R}_{i_\tau} \in \mathcal{M}$ denoting the τ-th misbehaving rating and $i_\tau \in \mathbb{N}_+$ denoting the corresponding index.

Note that this is a deterministic misbehavior model. Note that the analytical method developed in Section 3 can be easily extended probabilistic models of misbehaving ratings, i.e., with certain probability a rating is a misbehaving rating. This probabilistic model captures a fraction of misbehaving ratings. For brevity, we omit it. A special case of the $(L, \mathcal{F})$-misbehavior is the $(L, \tilde{m}, k)$-misbehavior model, which injects a consecutive sequence of L reviews with rating $\tilde{m}$ toward a product starting from the k-th review, where $\tilde{m} \in \mathcal{M}, k \in \mathbb{N}_+$. For example, a $(6, 5, 4)$-misbehavior means injecting six reviews with rating 5 starting from the 4-th review, i.e., $R_4 = 5, \ldots, R_9 = 5$. 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.

2.4 Motivating Questions

Formally, this article aims to investigate the following interesting questions to understand a persuasion cascade:

(1) *Will the collective opinion reaches a consensus in the honest and misbehavior scenarios? Precisely, does there exist a unique vector* $\widehat{\boldsymbol{\alpha}}_\infty \in \mathcal{O}$ *such that* $\mathbb{P}\left[\lim_{i\to\infty} \widehat{\boldsymbol{\alpha}}_i = \widehat{\boldsymbol{\alpha}}_\infty\right] = 1$ *(written as* $\widehat{\boldsymbol{\alpha}}_i \xrightarrow{a.s.} \widehat{\boldsymbol{\alpha}}_\infty$*)?*

(2) *How accurate is the average scoring rule (widely practiced) in estimating product quality under persuasion cascades? Precisely, will* $|A(\widehat{\boldsymbol{\alpha}}_i) - A(\boldsymbol{\alpha})|$ *be small?*

(3) *How to infer and understand the persuasion cascade in real-world online product review systems? Precisely, how to design algorithms to infer* $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, *and* γ, *from product review datasets, and draw insights from them?*

3 ANALYZING THE PERSUASION CASCADE

We first formulate a high-order Markov chain to track opinion dynamics of the persuasion cascade. We then 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 under both the honest and the misbehaving settings.

3.1 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 $[R_{i-N+1}, \ldots, R_i]$. Precisely, we define a Markov chain $(\mathcal{S}, \mathbf{P})$: (1) The *state space* $\mathcal{S}$ characterizes all the possible status of a persuasion cascade and (2) The *state transition probabilities* $\mathbf{P}$ characterizes the evolving dynamics of opinions. The $\mathcal{S}$ can be derived as

$$\mathcal{S} = \bigcup_{j=0}^{N} \underbrace{\{0\} \times \cdots \times \{0\}}_{N-j} \times \underbrace{\mathcal{M} \times \cdots \times \mathcal{M}}_{j}.$$

The persuasion cascade is at state

$$\mathbf{s} \triangleq [s_1, \ldots, s_N] \in \mathcal{S},$$

if the latest N historical ratings are $s_1, \ldots, s_N$. For example, each product is initialized with no reviews, namely, the initial state is $\mathbf{s} = [0, \ldots, 0]$. After the first consumer posts the review $(\mathbf{r}_1, R_1)$, the state becomes i.e., $\mathbf{s} = [0, \ldots, 0, R_1]$. In general, when the product $\mathcal{P}$ has i ratings, the state of the persuasion cascade is $\mathbf{s} = [R_{i-N+1}, \ldots, R_i]$. The total number of states is

$$|\mathcal{S}| = \sum_{j=0}^{N} M^j = \frac{M^{N+1} - 1}{M - 1}.$$

The matrix $\mathbf{P}$ characterizes all possible state transitions of a persuasion cascade, i.e., $\mathbf{P} = [P(\tilde{\mathbf{s}}|\mathbf{s}) : \mathbf{s} \in \mathcal{S}, \tilde{\mathbf{s}} \in \mathcal{S}]$, where

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

Namely, $\mathbf{P}$ characterizes the evolving dynamics of the persuasion cascade. 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-from of $P(\tilde{\mathbf{s}}|\mathbf{s})$, we define neighbors of a state.

Definition 6. We define the neighboring set of $\mathbf{s}$ as

$$\mathcal{N}_{\mathbf{s}} \triangleq \left\{[s_2, \dots, s_N, m] \middle| m \in \mathcal{M}\right\}, \qquad \mathbf{s} \in \mathcal{S}.$$

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 derive the state transition probabilities as

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

where $\mathbf{s}^{[n]}$ is defined as the latest n elements of $\mathbf{s}$, i.e.,

$$\mathbf{s}^{[n]} \triangleq [s_{N-n+1}, \dots, s_N], \qquad \forall n = 1, \dots, N,$$

and $\mathbf{s}^{[0]} \triangleq \mathbf{0}$. We can then summarize the state transition probabilities as

$$P(\tilde{\mathbf{s}}|\mathbf{s}) = \begin{cases} \sum_{n \in \mathcal{N}} \beta_n \left(\gamma \theta_{\tilde{s}_N}(\mathbf{s}^{[n]}) + (1-\gamma)\alpha_{\tilde{s}_N}\right), & \text{if } \tilde{\mathbf{s}} \in \mathcal{N}_{\mathbf{s}}, \\ 0, & \text{if } \tilde{\mathbf{s}} \notin \mathcal{N}_{\mathbf{s}}. \end{cases}$$

Each row of the state transition matrix $\mathbf{P}$ has at most $|\mathcal{N}_{\mathbf{s}}| = M$ non-zero entries. It is a sparse matrix because $M << |\mathcal{S}|$. We will refer to this Markov chain as $(\mathcal{S}, \mathbf{P})$.

3.2 Convergence of Opinions

Suppose consumers rate honestly, i.e., these is no misbehaving reviews. *We now study the convergence of opinions, i.e., the existence and uniqueness of a constant vector $\widehat{\boldsymbol{\alpha}}_\infty \in \mathcal{O}$ such that $\widehat{\boldsymbol{\alpha}}_i \xrightarrow{a.s.} \widehat{\boldsymbol{\alpha}}_\infty$.* Let us begin with the notion of steady-state distribution of $(\mathcal{S}, \mathbf{P})$, which characterize the opinion dynamics of the persuasion cascade asymptotically.

Definition 7. A probability (row) vector $\boldsymbol{\pi} \triangleq [\pi(\mathbf{s}) : \mathbf{s} \in \mathcal{S}]$, where $\boldsymbol{\pi} \in [0,1]^{|\mathcal{S}|}$ and $\sum_{\mathbf{s} \in \mathcal{S}} \pi(\mathbf{s}) = 1$, is a steady-state distribution for the Markov chain $(\mathcal{S}, \mathbf{P})$, if $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$.

Remark: The convergence of $(\mathcal{S}, \mathbf{P})$ to a unique steady-state distribution $\boldsymbol{\pi}$ means that the persuasion cascade will be at each state $\mathbf{s} \in \mathcal{S}$ with probability $\pi(\mathbf{s})$. In other words, the steady-state characterizes the stationary behavior of the persuasion cascade, which implies the following theorem.

THEOREM 1. *Suppose $(\mathcal{S}, \mathbf{P})$ converges to a unique steady-state distribution $\boldsymbol{\pi}$. There exists a unique constant vector $\widehat{\boldsymbol{\alpha}}_\infty \in \mathcal{O}$ such that $\widehat{\boldsymbol{\alpha}}_i \xrightarrow{a.s.} \widehat{\boldsymbol{\alpha}}_\infty$. Furthermore,*

$$\widehat{\boldsymbol{\alpha}}_\infty = \boldsymbol{\pi}\mathbf{W}, \tag{3}$$

where $\mathbf{W}$ denotes a $|\mathcal{S}| \times M$ weight matrix

$$\mathbf{W} \triangleq [W_{\mathbf{s},m} : \mathbf{s} \in \mathcal{S}, m \in \mathcal{M}], \quad W_{\mathbf{s},m} = \mathbf{1}_{\{s_N = m\}}.$$

Remark: All proofs of lemmas and theorems are in Section 8. Theorem 1 states that the stationary behavior of the persuasion cascade (i.e., steady-state distribution $\boldsymbol{\pi}$ of $(\mathcal{S}, \mathbf{P})$), not only implies the convergence of the empirical collective opinion $\widehat{\boldsymbol{\alpha}}_i$ to a unique $\widehat{\boldsymbol{\alpha}}_\infty$, but also implies a closed-form $\widehat{\boldsymbol{\alpha}}_\infty$ in terms of $\boldsymbol{\pi}$.

Now we establish conditions to guarantee the the stationary behavior of the persuasion cascade (i.e., steady-state distribution $\boldsymbol{\pi}$ of $(\mathcal{S}, \mathbf{P})$). Without loss of generality, we impose the following assumptions on collective intrinsic opinions.

Assumption 3. *The $\boldsymbol{\alpha}$ 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 then to make Assumption 3 hold, through which our analysis applies.

Theorem 2. *Suppose Assumption 3 holds and $\gamma < 1$. The Markov chain $(\mathcal{S}, \mathbf{P})$ has a unique steady-state distribution $\boldsymbol{\pi}$, which has the following structure:*

$$\begin{cases} \pi(\mathbf{s}) = 0, & \text{if } \mathbf{s} \in \mathcal{T}, \\ \pi(\mathbf{s}) > 0, & \text{if } \mathbf{s} \in \mathcal{R}, \end{cases}$$

where $\mathcal{R} \triangleq \mathcal{M}^N$ and $\mathcal{T} \triangleq \mathcal{S} \setminus \mathcal{R}$. Furthermore, $\boldsymbol{\pi}_{\mathcal{R}} \triangleq [\pi(\mathbf{s}) : \mathbf{s} \in \mathcal{R}]$ is a unique solution of

$$\boldsymbol{\pi}_{\mathcal{R}} = \boldsymbol{\pi}_{\mathcal{R}} \mathbf{P}_{\mathcal{R}}, \qquad \sum_{\mathbf{s} \in \mathcal{R}} \pi(\mathbf{s}) = 1, \tag{4}$$

where $\mathbf{P}_{\mathcal{R}} \triangleq [P[\tilde{\mathbf{s}}|\mathbf{s}] : \mathbf{s} \in \mathcal{R}, \tilde{\mathbf{s}} \in \mathcal{R}]$.

Remark: Theorem 2 states that asymptotically the persuasion cascade will reach the same steady-state distribution $\boldsymbol{\pi}$ no matter what the initial ratings or reviews are. This implies that the persuasion cascade will converge to the same collective opinion, which is invariant of initial ratings. Furthermore, the steady-state distribution is characterized by a linear system (4), which enables us to bound product quality estimation error.

3.3 Bounding the Product Quality Estimation Error

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

$$\begin{aligned} |A(\widehat{\boldsymbol{\alpha}}_i) - A(\boldsymbol{\alpha})| &\leq |A(\widehat{\boldsymbol{\alpha}}_i) - A(\widehat{\boldsymbol{\alpha}}_\infty)| + |A(\widehat{\boldsymbol{\alpha}}_\infty) - A(\boldsymbol{\alpha})| \\ &= \epsilon_a + \epsilon_p, \end{aligned} \tag{5}$$

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

(1) Bounding the persuasion error. Based on Theorem 2, we apply the matrix perturbation theory [25] to bound ϵ_p. 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 γ and it converges to a matrix with all entries being 0 as γ converges to zero. Let the row vector $\tilde{\boldsymbol{\pi}} \triangleq [\tilde{\pi}(\mathbf{s}) : \mathbf{s} \in \mathcal{S}]$ denote the steady-state distribution under $\gamma = 0$. Namely, $\tilde{\boldsymbol{\pi}}$ corresponds to the intrinsic collective opinion

$$\boldsymbol{\alpha} = \tilde{\boldsymbol{\pi}} \mathbf{W}. \tag{6}$$

Comparing with Equation (3), one can observe that ϵ_p arises from the "error" in the stationary behavior of the persuasion cascade (i.e., steady-state distribution). Thus, we start from understanding the steady-state distribution error caused by the persuasion biases. First, $\tilde{\boldsymbol{\pi}}$ has the following closed-form.

Corollary 1. *The $\tilde{\boldsymbol{\pi}}$ can be derived as*

$$\tilde{\pi}(\mathbf{s}) = \begin{cases} 0, & if\, \mathbf{s} \in \mathcal{T}, \\ \prod_{n=1}^{N} \alpha_{s_n}, & if\, \mathbf{s} \in \mathcal{R}. \end{cases}$$

Now, we apply the matrix perturbation theory to bound the steady-state distribution error.

Theorem 3. *Given the steady-state distribution $\boldsymbol{\pi}$ under $\gamma > 0$ and $\tilde{\boldsymbol{\pi}}$ under $\gamma = 0$. We bound the error between them as*

$$\frac{||\boldsymbol{\pi} - \tilde{\boldsymbol{\pi}}||}{||\tilde{\boldsymbol{\pi}}||} \leq \epsilon_{sd},$$

where ϵ_{sd} denotes the steady-state distribution error

$$\begin{aligned} \epsilon_{sd} \triangleq\ & ||[(\mathbf{P}_{\mathcal{R}} - \mathbf{I})(\mathbf{P}_{\mathcal{R}}^T - \mathbf{I}) + \mathbf{e}^T\mathbf{e}]^{-1}|| \times \\ & \Big(||\mathbf{E}|| \times ||\tilde{\mathbf{P}}_{\mathcal{R}} - \mathbf{I}|| + ||\mathbf{E}^T|| \times ||\tilde{\mathbf{P}}_{\mathcal{R}}^T - \mathbf{I}|| + ||\mathbf{E}||^2 + ||\mathbf{E}^T||^2 + ||\mathbf{E}|| \times ||\mathbf{E}^T||\Big), \end{aligned}$$

the notation $||\cdot||$ denotes a general matrix norm, $\mathbf{e} \triangleq [1, \ldots, 1]$ denotes a $|\mathcal{R}|$ dimensional row vector with all 1 elements and $\mathbf{I}$ denotes a $|\mathcal{R}| \times |\mathcal{R}|$ identity matrix.

Remark. Theorem 3 derives a closed-form upper bound for the steady-state distribution error. It implies that the steady-state distribution error $||\boldsymbol{\pi} - \tilde{\boldsymbol{\pi}}||/||\tilde{\boldsymbol{\pi}}||$ can be arbitrarily small when the persuasion strength γ is sufficiently small. In other words, when the persuasion strength is small, the stationary behavior of the persuasion cascade will slightly different from the one with no persuasion biases (i.e., providing intrinsic reviews). As a consequence, Theorem 3 implies the following upper bound on the collective opinion error.

Corollary 2. *The collective opinion error can be bounded as*

$$||\boldsymbol{\alpha} - \widehat{\boldsymbol{\alpha}}_\infty|| \leq \epsilon_{sd}||\tilde{\boldsymbol{\pi}}|| \times ||\mathbf{W}||.$$

Remark: Corollary 2 derives a closed-form upper bound for the opinion error. It implies that when the persuasion strength is small, the empirical opinion of the persuasion cascade converges to an opinion, which is slightly different from the intrinsic opinion. In other words, the empirical opinion $\widehat{\boldsymbol{\alpha}}_i$ would be an accurate estimator for the intrinsic opinion when the number of ratings i is sufficiently large. Corollary 2 implies the following upper bound on the persuasion error ϵ_p.

Corollary 3. *Consider the average scoring rule A, the persuasion error ϵ_p can be bounded as*

$$\epsilon_p \leq \epsilon_{sd}||\tilde{\boldsymbol{\pi}}|| \times ||\mathbf{W}\mathbf{m}^T||,$$

where $\mathbf{m} \triangleq [1, \ldots, M]$ denotes an M-dimensional row vector.

Remark: Corollary 3 derives a closed-form upper bound for the persuasion error ϵ_p under the average scoring rule. For other rating aggregation rules such as the median rating rule and weighted average scoring rule, one can apply Corollary 2 to derive similar bounds on ϵ_p. It implies that when the persuasion strength γ is small, the persuasion error ϵ_p can be small (we will study the error numerically in Section 5). In other words, the simple product estimation method, i.e., $A(\widehat{\boldsymbol{\alpha}}_i)$, would eventually be accurate, when the number of ratings i is sufficiently large. We next derive the *minimum number of ratings* to guarantee the convergence of opinions, which implies the approximation error ϵ_a as well.

(2) Bounding the approximation error. Based on Theorem 2, we now apply the Chernoff bound (the Markovian version) [19] to bound the approximation error ϵ_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

$$\boldsymbol{\pi}_0 \triangleq [\pi_0(\mathbf{s}) : \mathbf{s} \in \mathcal{R}],$$

where $\boldsymbol{\pi}_0 \in [0,1]^{|\mathcal{R}|}$ and $\sum_{s\in\mathcal{R}} \pi_0(\mathbf{s}) = 1$, denote the joint distribution of the first N ratings. Namely, $\boldsymbol{\pi}_0$ characterizes the initial state of the persuasion cascade with restriction to $(\mathcal{R}, \mathbf{P}_{\mathcal{R}})$. Eventually, the persuasion cascade will converges to a steady-state distribution $\boldsymbol{\pi}_{\mathcal{R}}$. To quantify how far away the persuasion cascade is from convergence, we define the distance metric as

$$D(\boldsymbol{\pi}_0, \boldsymbol{\pi}_{\mathcal{R}}) \triangleq \sqrt{\sum_{\mathbf{s}\in\mathcal{R}} \frac{\pi_0^2(\mathbf{s})}{\pi_{\mathcal{R}}(\mathbf{s})}}.$$

Using the Cauchy–Schwarz inequality, one obtains that when $\boldsymbol{\pi}_0 = \boldsymbol{\pi}_{\mathcal{R}}$, the minimum distance is attained $D(\boldsymbol{\pi}_0, \boldsymbol{\pi}_{\mathcal{R}}) = 1$. Let ν_{ps} denote the pseudo spectral gap [19] of the Markov chain $(\mathcal{R}, \mathbf{P}_{\mathcal{R}})$, where ν_{ps} quantifies the speed of convergence of the persuasion cascade. With these two metrics, we now bound the opinion error caused by approximation, i.e., $||\widehat{\boldsymbol{\alpha}}_i - \widehat{\boldsymbol{\alpha}}_\infty||$.

Theorem 4. *Suppose the number of ratings i satisfies*

$$i \geq \frac{1}{\epsilon^2} \frac{16(1 + 1/\nu_{ps}) + 40\epsilon}{\nu_{ps}} \ln\left(\frac{\sqrt{2}MD(\boldsymbol{\pi}_0, \boldsymbol{\pi}_{\mathcal{R}})}{\delta}\right) + N, \tag{7}$$

where $\epsilon > 0$ *and* $\delta \in [0,1]$. *It holds with probability at least* $1-\delta$ *that*

$$||\widehat{\boldsymbol{\alpha}}_i - \widehat{\boldsymbol{\alpha}}_\infty|| \leq \epsilon ||\mathbf{c}||,$$

where $\mathbf{c} \triangleq [1,\ldots,1]$ *denotes a M-dimensional row vector.*

Remark: Theorem 4 states the "minimum number of ratings" to guarantee that the empirical opinion $\widehat{\boldsymbol{\alpha}}_i$ is an accurate estimator of the converged opinion $\widehat{\boldsymbol{\alpha}}_\infty$. It reveals that the minimum number of ratings is critical to the estimation accuracy ϵ and the speed of convergence (i.e., pseudo spectral gap ν_{ps}), while it is not critical to the distance to convergence (i.e., $\ln D(\boldsymbol{\pi}_0, \boldsymbol{\pi}_{\mathcal{R}})$), because it is proportional to the logarithmic of the distance, i.e., $\ln D(\boldsymbol{\pi}_0, \boldsymbol{\pi}_{\mathcal{R}})$. We will study the minimum number of ratings quantitatively in Section 5. Theorem 4 implies the following approximation error ϵ_a.

Corollary 4. *Consider A to be the average scoring rule. Suppose the number of ratings satisfies (7). With probability at least* $1-\delta$, *it holds that*

$$\epsilon_a \leq \epsilon ||\mathbf{c}|| \times ||\mathbf{m}||.$$

Remark: Corollary 4 states that the approximation error ϵ_a can be arbitrarily small with a sufficiently large number of ratings under the average scoring rule. For other rating aggregation rules such as the median rating rule, one can apply Theorem 4 to derive similar bounds for ϵ_a. Combining Equation (5), Corollary 3, and Corollary 4, one can obtain the error bound for the product estimation error $|A(\widehat{\boldsymbol{\alpha}}_i) - A(\boldsymbol{\alpha})|$.

(3) Summary of honest case. The persuasion cascade converges to a unique collective opinion, which is invariant of the initial rating, except the extremal case $\gamma = 1$. The $A(\widehat{\boldsymbol{\alpha}}_i)$ is an accurate estimator of the intrinsic quality $A(\boldsymbol{\alpha})$ when the persuasion strength is weak and the number of ratings is large.

3.4 Impact of Misbehavior Attacks

We first analyze the impact of the $(L, \mathcal{F})$-misbehavior on the convergence of the persuasion cascade. As a corollary of Theorem 2, we have the following result.

Corollary 5. *Suppose Assumption 3 holds and $\gamma < 1$. The steady-state distribution of the Markov chain $(\mathcal{S}, \mathbf{P})$ remains unchanged under any $(L, \mathcal{F})$-misbehavior with $L < \infty$.*

Remark: Corollary 5 states that the convergence of the persuasion cascade is invariant of any $(L, \mathcal{F})$-misbehavior with a finite length. This implies that the persuasion error ϵ_p is invariant of the $(L, \mathcal{F})$-misbehavior.

Corollary 6. *Suppose Assumption 3 holds and $\gamma < 1$. The persuasion error ϵ_p remains unchanged under any $(L, \mathcal{F})$-misbehavior with $L < \infty$.*

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 ϵ_a small (will study it numerically in Section 5). One can apply Theorem 4 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.

4 INFERRING THE PERSUASION CASCADE

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

4.1 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., $R_1, \ldots, R_I$. We do not assume any a-priori knowledge on the persuasion strength γ, the review reading distribution $\boldsymbol{\beta}$, and the intrinsic collective opinion $\boldsymbol{\alpha}$. Our objective is to infer these parameters from $R_1, \ldots, R_I$ via maximum likelihood estimation. In the remaining parts of this article, we consider a specific form of persuaded opinion function, i.e., $\boldsymbol{\theta}(\mathbf{h}_i)$ satisfies Equation (2), which capture consumers' herding or following behavior in opinion formation from review readings. Given $R_1, \ldots, R_I$, we derive the likelihood function of $\boldsymbol{\alpha}, \boldsymbol{\beta}, \gamma$ as:

$$\begin{aligned}\mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \gamma) &\triangleq \mathbb{P}[R_1, \ldots, R_I; \boldsymbol{\alpha}, \boldsymbol{\beta}, \gamma] \\ &= \prod_{i=1}^{I} \left(\sum_{n \in \mathcal{N}} \beta_n \left[\gamma \theta_{R_i} \left(\mathbf{h}_{i-1}^{[n]} \right) + (1-\gamma) \alpha_{R_i} \right] \right).\end{aligned}$$

Then the log-likelihood function can be derived as

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

Formally, we state our inference problem as follows.

PROBLEM 1. *Given $R_1, \ldots, R_I$ of the product $\mathcal{P}$, $\mathcal{M}, \mathcal{N}$ and $\boldsymbol{\theta}(\mathbf{h}_i)$ which satisfies (2), select $\boldsymbol{\alpha}, \boldsymbol{\beta}$ and γ to maximize the log-likelihood function $\ell(\boldsymbol{\alpha}, \boldsymbol{\beta}, \gamma)$. Formally,*

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

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

THEOREM 5. *Problem 1 has at least one optimal solution.*

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

We next identify some properties of the objective function, which may enable us to design efficient algorithms to locate the optimal solution.

LEMMA 1. *The log-likelihood function $\ell(\boldsymbol{\alpha}, \boldsymbol{\beta}, \gamma)$ is concave with respect to $\boldsymbol{\alpha}$, concave with respect to $\boldsymbol{\beta}$ and concave with respect to γ, respectively.*

Lemma 1 states that given any two of $\boldsymbol{\alpha}, \boldsymbol{\beta}$, and γ, the log-likelihood function is concave with the remaining one. Note that there are a variety of efficient algorithms to solve concave optimization problems. Hence, Lemma 1 inspires one possible efficient inference algorithm: update $\boldsymbol{\alpha}, \boldsymbol{\beta}$, and γ alternately, and each time apply concave optimization algorithms to update one of them.

4.2 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. By repeating the searching algorithm for multiple times with random start, we may hit the optimal solution with high probability. Let us derive the gradient of $\ell(\boldsymbol{\alpha}, \boldsymbol{\beta}, \gamma)$ as:

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

With these gradients, we use the interior point algorithmic framework [29] to search the optimal solution. For brevity, we omit the details, because it only involves some minor modifications of the standard interior point algorithm [29].

5 Synthetic Data Analytics and Implications

We first conduct synthetic data analytics to systematically verify our theoretical findings and examine the impact of various factors on the convergence of the persuasion cascade and on the accuracy of product quality estimation. Then, we evaluate the accuracy of our inference algorithm.

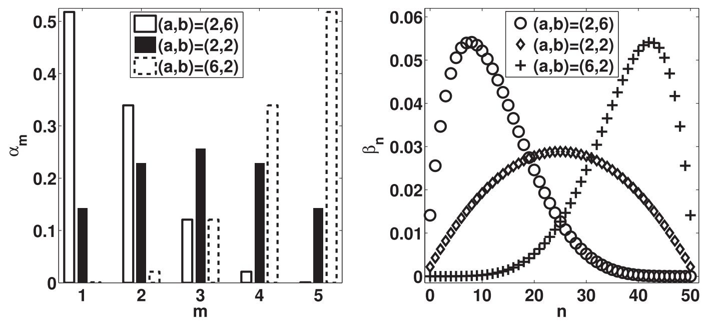

(a) $\boldsymbol{\alpha}$ with $M = 5$ (b) $\boldsymbol{\beta}$ with $N = 50$

Fig. 1. Examples of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ ($M = 5, N = 50$).

5.1 Synthetic Datasets

We set $M = 5$, which is consistent with real-world online product review systems. We set $N = 50$ to capture the finding that around 88% consumers read less than 10 reviews of a product [20]. Thus, the number of states of the underlying Markov chain $(\mathcal{S}, \mathbf{P})$ is

$$|\mathcal{S}| \approx 5^{50} \approx 8 \times 10^{34}.$$

Namely, it is computationally expensive to calculate the exact steady-state distribution $\boldsymbol{\pi}$. We consider a persuaded opinion formation function $\boldsymbol{\theta}(\mathbf{h}_i)$ that satisfies Equation (2), to capture consumers' herding or following behavior in persuaded opinion formation. Without loss of generality, we focus on generating ratings for one product, i.e., product $\mathcal{P}$. In particular, we instantiate the associated intrinsic collective opinion $\boldsymbol{\alpha}$ and review reading distribution $\boldsymbol{\beta}$ via discretizing the beta distribution $Beta(a, b)$ with parameters $a, b \in \mathbb{R}_+$:

$$\alpha_m \propto x^{a-1}(1-x)^{b-1}|_{x=\frac{m}{M+1}}, \qquad m \in \mathcal{M},$$
$$\beta_n \propto x^{a-1}(1-x)^{b-1}|_{x=\frac{n+1}{N+2}}, \qquad n \in \mathcal{N}.$$

One can select different pairs of (a, b) to model different intrinsic collective opinions $\boldsymbol{\alpha}$ and model different review reading distributions $\boldsymbol{\beta}$. In particular, we select three pairs of (a, b) illustrated in Figure 1. From Figure 1(a), one can observe that the parameter $(a, b) = (2, 6)$ (or $(a, b) = (2, 2)$ or $(a, b) = (6, 2)$) models that most consumers have low (or medium or high) intrinsic ratings. In other words, the product is of low (or medium or high) overall quality. Similarly, from Figure 1(b), one can observe that the parameter $(a, b) = (2, 6)$ (or $(a, b) = (2, 2)$ or $(a, b) = (6, 2)$) models that most consumers read a small (or medium or large) number of reviews.

With the above parameters, we simulate our model (details in Algorithm 1) to synthesize product ratings.

5.2 Opinion Dynamics and Biases

Impact of initial ratings: We consider one typical instance of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ generated by discretizing $Beta(2, 2)$ (depicted in Figure 1). We study a strong persuasion strength $\gamma = 0.9$, because if the persuasion cascade converges under strong persuasion strengths, it also implies the convergence under weak persuasion strengths. We vary the initial rating R_1 from 1 to 5. We input these parameters into Algorithm 1 to generate $I = 35{,}000$ product ratings. We evaluate the empirical opinion $\widehat{\boldsymbol{\alpha}}_i$ via Equation (1). Figure 2 shows the value of $\hat{\alpha}_{i,1}, \hat{\alpha}_{i,2}, \hat{\alpha}_{i,3}$, and $\hat{\alpha}_{i,4}$ as we vary the number of ratings i from 1 to 35,000, where the empirical $\alpha_{i,m}$ corresponds to $\widehat{\alpha}_{i,m}$. Note that we omit $\hat{\alpha}_{i,5}$ because its convergence can be implied by that of $\hat{\alpha}_{i,1}, \hat{\alpha}_{i,2}, \hat{\alpha}_{i,3}$, and $\hat{\alpha}_{i,4}$. From Figure 2, we

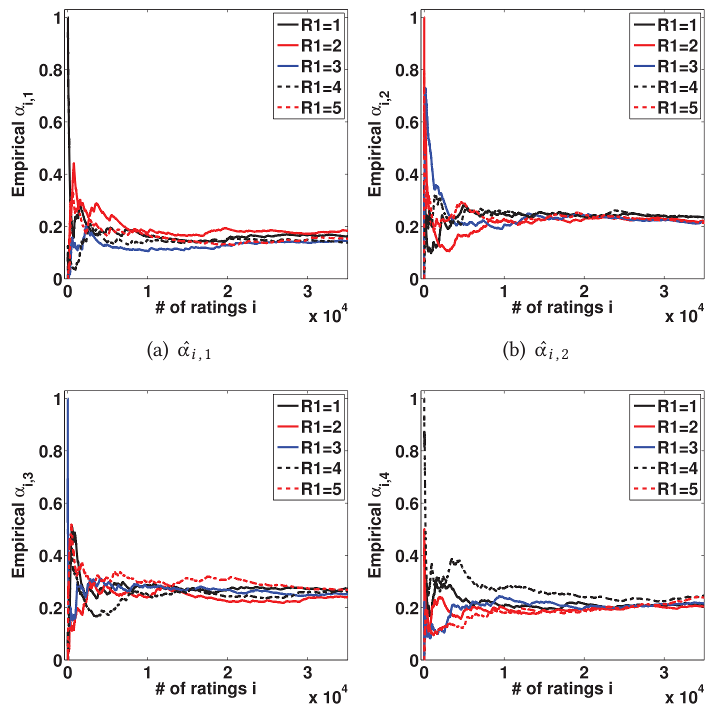

(a) $\hat{\alpha}_{i,1}$ (b) $\hat{\alpha}_{i,2}$

(c) $\hat{\alpha}_{i,3}$ (d) $\hat{\alpha}_{i,4}$

Fig. 2. Impact of the initial rating R_1 on convergence of $\widehat{\boldsymbol{\alpha}}_i$ ($\boldsymbol{\alpha}, \boldsymbol{\beta}$ are generated by $Beta(2, 2)$ and $\gamma = 0.9$).

ALGORITHM 1: : Synthesize Product Ratings

Require: $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, γ ,M, N, $\boldsymbol{\theta}(\mathbf{h}_i)$ satisfies (2), R_1, I.
Ensure: I ratings for product $\mathcal{P}$, i.e., $R_1, \ldots, R_I$.

$R_{-i} \leftarrow 0$ for all $i = 0, 1, \ldots, N-1$
for $i = 2$ to I **do**
  Generate the number of review readings n from $\boldsymbol{\beta}$.
  $\mathbf{h}_{i-1}^{[n]} \leftarrow [R_{i-n}, \ldots, R_{i-1}]$.
  Plug $\mathbf{h}_{i-1}^{[n]}$ into (2) to obtain $\boldsymbol{\theta}(\mathbf{h}_{i-1}^{[n]})$.
  Generate R_i according to the pmf $\gamma\boldsymbol{\theta}(\mathbf{h}_{i-1}^{[n]}) + (1-\gamma)\boldsymbol{\alpha}$.
end for

observe that the empirical collective opinion $\widehat{\boldsymbol{\alpha}}_i$ converges to a unique value no matter what the initial rating R_1 is. This verifies our theoretical findings. It is interesting to observe that around 10000 ratings are enough to guarantee the convergence, even though the number of the states of the underlying Markov chain is $|\mathcal{S}| \approx 8 \times 10^{34}$.

Impact of persuasion strengths: We extend the last experiment to vary the persuasion strength from 0 to 0.9. For brevity, we select $R_1 = 1, R_1 = 5$, and $\widehat{\alpha}_{i,1}$ to study. Figure 3 shows the value of $\widehat{\alpha}_{i,1}$ across the number of ratings i. One can observe that when there is no persuasion bias, i.e., $\gamma = 0$, we only need hundreds of ratings to guarantee the convergence of opinions. As γ increases, we need more ratings to guarantee the convergence. In particular, when the persuasion strength is strong, i.e., $\gamma = 0.9$, we need around 10,000 ratings. It is interesting to observe that eventually the

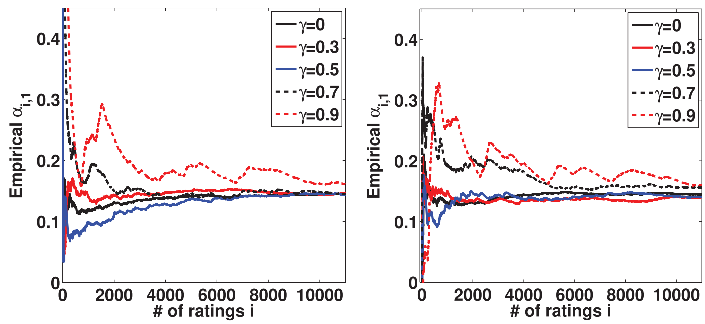

(a) $R_1 = 1$ (b) $R_1 = 5$

Fig. 3. Impact of the persuasion strength γ on the convergence of $\widehat{\boldsymbol{\alpha}}_i$ ($\boldsymbol{\alpha}, \boldsymbol{\beta}$ are generated by $Beta(2, 2)$ and $R_1 = 1$ or 5).

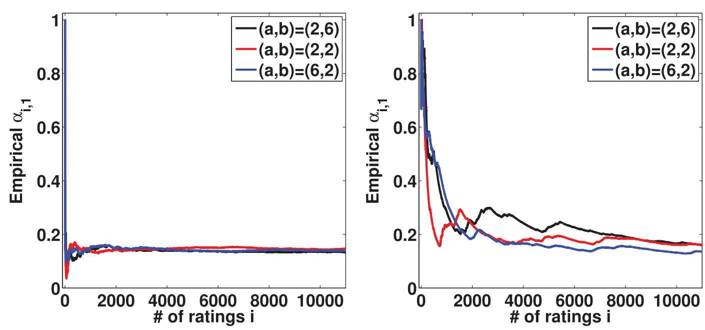

(a) $\gamma = 0.3$ (b) $\gamma = 0.9$

Fig. 4. Impact of the review reading distribution $\boldsymbol{\beta}$ on the convergence of $\widehat{\boldsymbol{\alpha}}_i$ ($\boldsymbol{\alpha}$ is generated by $Beta(2, 2)$, $R_1 = 1$ and $\gamma = 0.3$ or 0.9).

value of $\widehat{\alpha}_{i,1}$ converges to similar values under different values of γ. In other words, in the long run the persuasion bias will be averaged out, but it slows down the speed of convergence significantly. This observation also suggests that real-world online product review system operators should focus on speeding up the convergence, because most products only have hundreds of ratings.

Impact of review reading distributions: We extend the last experiment settings to consider three instances of $\boldsymbol{\beta}$ depicted in Figure 1(b). For brevity, we select $\gamma = 0.3$, $\gamma = 0.9$, $\widehat{\alpha}_{i,1}$, and $R_1 = 1$ to study. From Figure 4, one can observe that under different review reading distributions $\boldsymbol{\beta}$, the value of $\widehat{\alpha}_{i,1}$ converges to similar values. It further verifies our finding that the persuasion biases will eventually be averaged out.

Lessons learned: The collective opinion of a persuasion cascade converges and this convergence is invariant of initial ratings. The persuasion bias will eventually be averaged out; however, it slows down the speed of convergence.

5.3 Product Quality Estimation Accuracy

We set A as the average scoring rule and we aim to study the convergence of the product quality, i.e., $A(\widehat{\boldsymbol{\alpha}}_i)$. We use the same experiment setting as that in Section 5.2. Figure 5(a) shows the convergence of $A(\widehat{\boldsymbol{\alpha}}_i)$, where we vary the persuasion strength from 0 to 0.9. One can observe

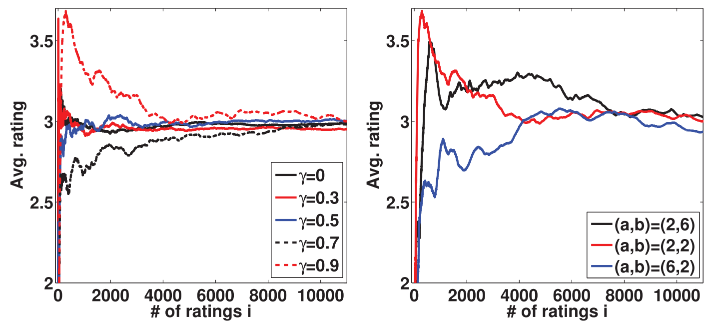

(a) $\boldsymbol{\beta}$ generated by $Beta(2,2)$ (b) $\gamma = 0.9$

Fig. 5. Impact of γ and $\boldsymbol{\beta}$ on the convergence of the average rating ($\boldsymbol{\alpha}$ is generated by $Beta(2, 2)$, $R_1 = 1$).

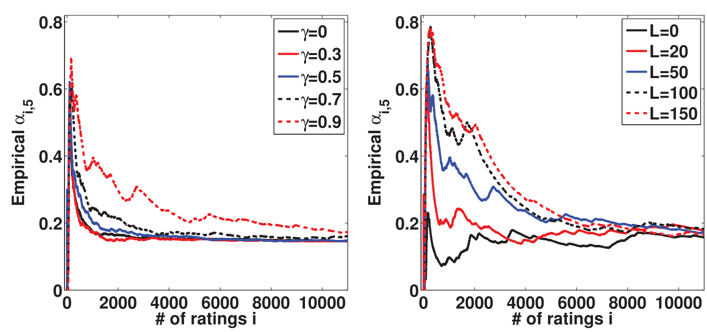

(a) $L = 50$ (b) $\gamma = 0.9$

Fig. 6. Impact of $(L, \tilde{m}, k)$-misbehavior on the convergence of $\widehat{\boldsymbol{\alpha}}_i$ ($\boldsymbol{\alpha}, \boldsymbol{\beta}$ generated by $Beta(2, 2)$, $R_1 = 1, \tilde{m} = 5, k = 50$).

that under different persuasion strengths, the average rating $A(\widehat{\boldsymbol{\alpha}}_i)$ converges to similar values. In other words, asymptotically $A(\widehat{\boldsymbol{\alpha}}_i)$ is an accurate estimator of the product quality. The persuasion bias, however, slows down the convergence significantly. It suggests that online product review system operators should focus on speeding up the convergence. Figure 5(b) shows the impact of the review reading distribution on the convergence of $A(\widehat{\boldsymbol{\alpha}}_i)$. One can observe that under different $\boldsymbol{\beta}$ (i.e., different review reading distributions), the average rating $A(\widehat{\boldsymbol{\alpha}}_i)$ converges to similar values with similar speeds of convergence.

Lessons learned: Using the average scoring rule to estimate product quality, the estimator $A(\widehat{\boldsymbol{\alpha}}_i)$ is asymptotically accurate; however, the persuasion bias slows down the speed of convergence.

5.4 Impact of Misbehavior Attacks

To understand the robustness against misbehavior, we consider a company launches a strong misbehavior attack, i.e., $(L, \tilde{m}, k)$-misbehavior attack, where $\tilde{m} = 5$ and $k = 50$, to promote its products. We use the same experiment setting as that in Section 5.2. For brevity, we only focus on $\widehat{\alpha}_{i,5}$. Figure 6 shows the impact of the $(L, 5, 50)$-misbehavior when we vary the persuasion strength γ and the length L. One can observe that in the short run, the misbehavior can increase the positive opinion $\widehat{\alpha}_{i,5}$ significantly, especially when the persuasion strength γ is strong or the number of misbehaving ratings L is large. However, in the long run, the bias caused by the misbehavior attack will eventually be averaged out, and $\widehat{\alpha}_{i,5}$ converges to similar values.

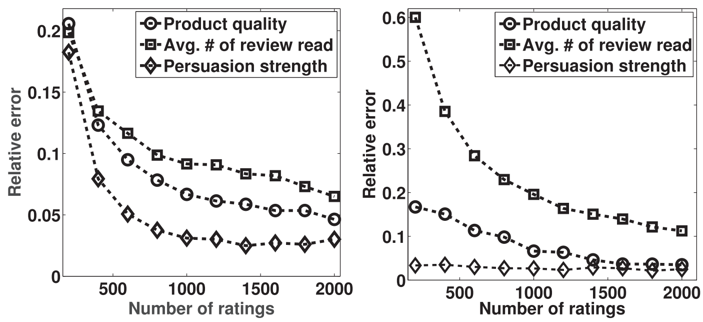

(a) Honest rating (b) Misbehaving rating

Fig. 7. The accuracy of our inference algorithm under both honest rating and misbehaving rating settings ($\boldsymbol{\alpha}, \boldsymbol{\beta}$ generated by $Beta(2, 2)$, $\gamma = 0.9$, $R_1 = 1$, $\tilde{m} = 5$, $k = 50$, $L = 50$).

Lessons learned: The opinion bias caused by misbehavior attacks will eventually be averaged out, however, the misbehavior attack slows down the convergence.

5.5 Evaluating the Inference Algorithm

Now we evaluate the accuracy of our inference algorithm (stated in Section 4). We run our inference algorithm (stated in Section 4) on synthetic data to infer $\boldsymbol{\alpha}, \boldsymbol{\beta}$, and γ. Note that in the synthetic dataset, the ground truth of $\boldsymbol{\alpha}, \boldsymbol{\beta}$, and γ are given. Observe that $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are vectors, and we measure the inference accuracy of them in terms of the intrinsic quality $A(\boldsymbol{\alpha})$ and the average number of review readings $\sum_{n \in \mathcal{N}} n\beta_n$. We consider that A is the average scoring rule. Figure 7 shows the relative error in intrinsic quality $A(\boldsymbol{\alpha})$, average number of review readings $\sum_{n \in \mathcal{N}} n\beta_n$ and the persuasion strength γ, under both the honest rating and misbehaving rating settings. One example of the relative error is $|\gamma - \widehat{\gamma}|/\gamma$, where $\widehat{\gamma}$ denotes the inferred persuasion strength. From Figure 7, one can observe that the relative error decreases as we increase the number of ratings. This means that our inference algorithm becomes more accurate as more ratings can be used to conduct the inference. This also gives evidences on the convergence of our inference algorithm. Figure 7(a) shows that under the honest rating scenario, the relative error of inferring the persuasion strength is less than 5% when the number of ratings is around 800, implying a high inference accuracy. Furthermore, the relative errors of inferring the product quality and the average number of review readings are around 5% and 7%, respectively, when the number of ratings is 2,000, implying a high inference accuracy. Figure 7(b) shows that under the misbehaving scenario, the relative error in intrinsic quality and persuasion strength can also be around 5% when the number of ratings is around 2,000. And the relative error in average number of review readings is around 10%. This implies that our inference algorithm is highly robust against misbehaving ratings.

Lessons learned. Our inference algorithm is accurate in uncovering the intrinsic quality, average number of review readings and the persuasion strength under both the honest rating and misbehaving rating settings.

6 REAL-WORLD DATA ANALYTICS AND GUIDANCES

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

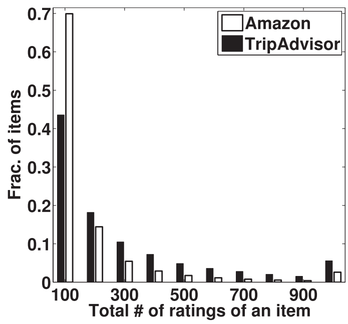

(a) The whole dataset

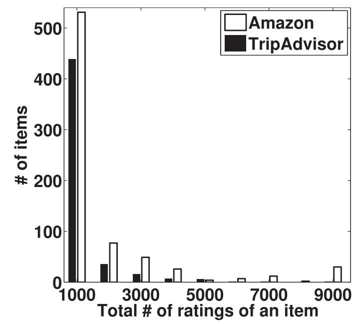

(b) The selected items

Fig. 8. The distribution of the number of ratings across items.

6.1 Data Analysis Setting

Datasets. The datasets contain historical ratings (with time stamps) of 32,888 products in Amazon and of 11,543 hotels in TripAdvisor. The dataset was crawled by Xie et al. [33] in 2013 and it was used in several previous works [30, 32]. We plot the distribution of the number of ratings across items (i.e., products or hotels) in Figure 8(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 1,000 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 8(b). One can observe that most selected items have around 1,000 ratings.

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 $\mathcal{M} = \{1, 2, 3, 4, 5\}$. A survey study [20] reported that around 84% consumers read less than 10 reviews of a product. Based on this survey study, we set $\mathcal{N} = \{0, 1, \ldots, 15\}$ to avoid over parameterization of the model. We consider a simple persuaded opinion formation function $\boldsymbol{\theta}(\mathbf{h}_i)$ that satisfies Equation (2), to capture consumers' following (or herding) behavior in persuaded opinion formation. For each selected item in Figure 8(b), we input the above parameters and the whole rating history of this item into the inference algorithm (presented in Section 4). Through this we infer persuasion strength γ (will be shown in Section 6.2) ,the review reading distribution $\boldsymbol{\beta}$ (will be shown in Section 6.3) and the intrinsic opinion $\boldsymbol{\alpha}$ (will be shown in Section 6.4) for each item.

6.2 Persuasion Strengths

Note that each item may be associated with a different persuasion strength. Figure 9 shows the distribution of the inferred persuasion strength γ across items. Figure 9(a) depicts the pmf of γ, 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 $j = 1, \ldots, 20$. From Figure 9(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 0.1–0.4. Only a very small fraction of items are associated with a persuasion strength larger than 0.6. Around 10% items has a very weak persuasion strength, i.e., less than 0.05. These observations show that the persuasion effect exists for most users/items. Furthermore, these observations match well with the survey study [20], which reported that 88% users trust rating or reviews and they incorporate them into their decisions. Figure 9(b) depicts the cumulative distribution function (CDF) of the persuasion strength across items. One can observe that the CDF

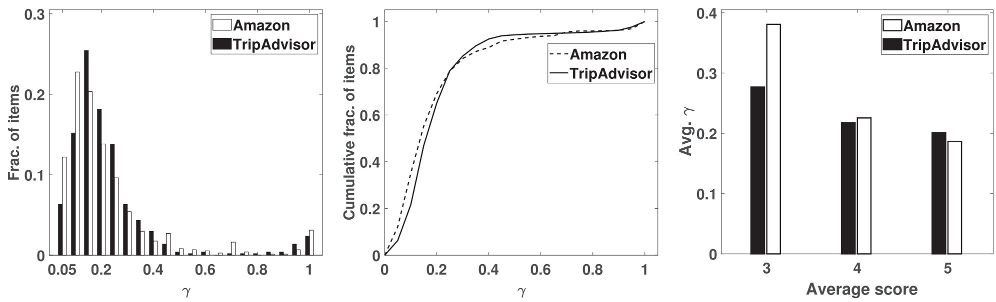

(a) Probability mass function (b) Cumulative distribution function (c) $\bar{\gamma}$ vs. avg. score

Fig. 9. The the variation of the persuasion strength across items and average score of products.

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. Figure 9(c) shows the variation of persuasion strength across the average score of products. Note that there is no product with average score being smaller than 2. The bar with index 3/4/5 depicts the average persuasion strength of all products with average score in the range $(2, 3]/(3, 4]/(4, 5]$. From Figure 9(c), one can observe that as the average score increases, the persuasion strength decreases. This observation matches well with the survey study [20], which reported that 95% users suspect reviews when there are no bad ratings or reviews. For a product with average score around three, it contains both good ratings and bad ratings. Hence, as indicated by the survey study [20], users trust the ratings or reviews more resulted in a stronger persuasion strength. However, for a product with average score around five, it nearly has no bad ratings. Hence, users trust the ratings or reviews less resulted in a weaker persuasion strength.

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 0.1–0.4. The persuasion strength in TripAdvisor is stronger than that in Amazon.

As the average score of a product increases, the persuasion strength decreases.

6.3 Review Reading Distributions

Number of Review Readings: Based on the inferred $\boldsymbol{\beta}$ for a product, we define its average number of review readings

$$\bar{N} \triangleq \sum_{n \in \mathcal{N}} n\beta_n.$$

Figure 10 shows the distribution of $\bar{N}$ across items, where we discreterize the interval $[0, 15]$ into 10 sub-intervals such that the j-th sub-interval corresponds to $[2(j-1), 2j-1]$, for all $j = 1, \ldots, 8$.

From Figure 10(a), one can observe that the distributions of the $\bar{N}$ across hotels in TripAdvisor and products in Amazon are in normal-like shapes. Around 82% users in Amazon only read up to eleven reviews, and around 90% users in TripAdvisor read up to thirteen reviews. This observation roughly agrees with the survey study [20], which reported that around 84% users reads up to ten reviews. From Figure 10(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.

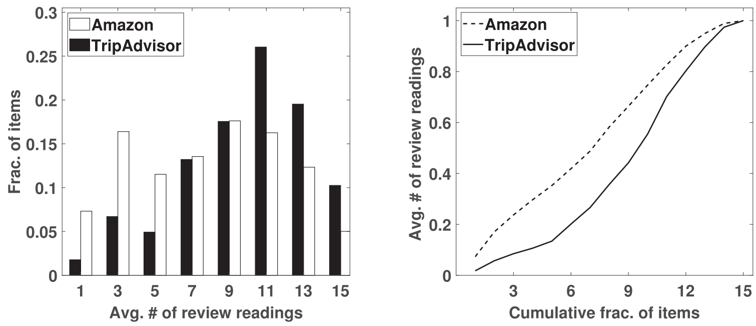

(a) Probability mass function (b) Cumulative distribution function

Fig. 10. The distribution of the number of review readings $\bar{N}$ across items.

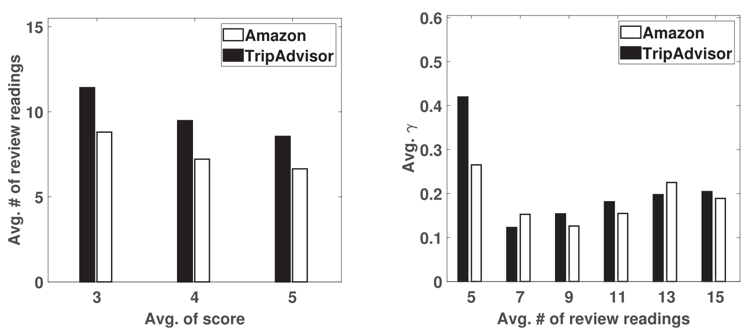

(a) $\bar{N}$ vs. average score (b) $\bar{\gamma}$ vs. $\bar{N}$

Fig. 11. Variation of $\bar{N}$ across average score of products and average number of review readings.

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.

Variation of $\bar{N}$: Given the average number of review readings $\bar{N}$, we define the corresponding average persuasion strength as

$$\bar{\gamma} \triangleq \frac{\text{summation of } \gamma \text{ over items having } \bar{N}}{\#[\text{items have } \bar{N}]}.$$

Figure 11(a) shows the variation of number of review readings across the average score of products. Note that there is no product with average score being smaller than 2. The bar with index 3/4/5 depicts the average of $\bar{N}$ over all products with average score in the range $(2,3]/(3,4]/(4,5]$. From Figure 11(a), one can observe that as the average score increases, the number of review readings decreases. This observation is supported by a survey study [20], which reported that star rating is the number one factor used by consumers to judge a product. For a product with average score around three, it contains both good ratings and bad ratings. Users need to read more reviews to weight the pros and cons of a product. However, for a product with average score around five, it nearly has no bad ratings. Hence, users read less reviews to form an opinion about a product.

Figure 11(b) shows $\bar{\gamma}$ across $\bar{N}$, where we use the same discretization of $[0, 15]$ as that in Figure 10(a). To illustrate, the bar with index 5/7 depicts the average of persuasion strength over all products with $\bar{N}$ in the range $(0,5]/(6,7]$. From Figure 10(a), one can observe that the persuasion strength under $\bar{N} \leq 5$ is larger than that under $\bar{N} \geq 6$. Furthermore, the persuasion strength varies

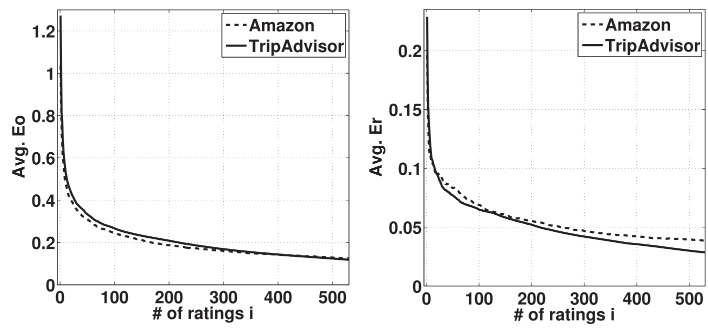

(a) Opinion error E_o (b) Product quality error E_q

Fig. 12. The opinion error and average rating error.

slightly when $\bar{N} \geq 6$. This observation is supported by the survey study [20], 73% of consumers form an opinion by reading up to six reviews. Forming an opinion with only a small number of reviews means that these reviews are really persuasive resulting in a larger persuasion strength.

Lessons learned: The average number of review readings concentrates around 11 for both Amazon and TripAdvisor. Consumers tend to read more reviews in TripAdvisor than Amazon. As the average score of a product increases, the number of review readings decreases.

6.4 The Opinion and Product Quality Estimation Error

From Figure 8, we observe that around 90% of the items (i.e., products or hotels) have less than 500 ratings. We therefore explore: *(1) Given a small number of ratings (i.e, $i \leq 500$), how accurate is $\widehat{\boldsymbol{\alpha}}_i$ in reflecting $\boldsymbol{\alpha}$? (2) How accurate is $A(\widehat{\boldsymbol{\alpha}}_i)$ in estimating the intrinsic quality $A(\boldsymbol{\alpha})$?*

Note that the intrinsic opinion $\boldsymbol{\alpha}$ for each item has been inferred in Section 6.1. For each selected item, we compute the associated empirical collective opinion $\widehat{\boldsymbol{\alpha}}_i$ for each given number of ratings $i = 1, \ldots, 500$. Given, $\boldsymbol{\alpha}$ and $\widehat{\boldsymbol{\alpha}}_i$, we define the relative opinion error as

$$E_o \triangleq \frac{||\widehat{\boldsymbol{\alpha}}_i - \boldsymbol{\alpha}||_1}{||\boldsymbol{\alpha}||_1},$$

where $||\cdot||_1$ denote the taxicab norm. The opinion error $E_o \in [0, 2]$, and the smaller the E_o, the more accurate $\widehat{\boldsymbol{\alpha}}_i$ is. We define the relative product quality estimation error as

$$E_q \triangleq \frac{|A(\widehat{\boldsymbol{\alpha}}_i) - A(\boldsymbol{\alpha})|}{|A(\boldsymbol{\alpha})|}.$$

The product quality error $E_q \in [0, M-1]$, and the smaller the E_q, the more accurate $A(\widehat{\boldsymbol{\alpha}}_i)$ is. We set A to be the average scoring rule. Figure 12 shows E_o and E_q (average across items) when we vary the number of ratings i from 1 to 500. From Figure 12, 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.

7 RELATED WORK

A variety of works studied the benefits of online product review systems. Chevalier et al. [5] showed that product reviews can increase the sales of sellers. Mudambi et al. [18] showed that product reviews serve as an important information source for consumers to make purchase decisions. This argument was further justified by Lackermair et al. [13] and Li [14]. A number of survey studies [4, 20, 23] were conducted to investigate how consumers use the product reviews.

They revealed that around 90% consumers reading product reviews before making decisions. Furthermore, consumers only read a small subset of the reviews and most of them only trust the latest reviews. These observations give us important insights to model the persuasion cascades.

Consumers' reviews are subject to biases due to a number of reasons. One the product side, Guo et al. revealed the category-wise biases [8]. On the product review system side, the system interfaces [6] or the recommendation algorithms [22] can lead to opinion biases. On the consumer side, the opinion biases can be caused by the evolving dynamics of their preferences [11], or the change in their experience [16]. A variety of effective methods, e.g., [6, 8, 11, 16, 22], have been proposed to mitigate these biases.

To gain deeper understanding in the personal biases in product reviews, a variety of works explore the review (or opinion) formation from psychological perspectives. Zhang et al. [37] 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 [17, 21] revealed that consumer's reviews are subject to social influence biases. Krishnan et al. [12] and Wang et al. [27] 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. [1] 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 [4, 20, 23] revealed that consumers read reviews in decision making. This uncovers a typical psychological phenomenon called the message-based persuasion [28]. As psychology theory [9, 28] have indicated, message-based persuasion is powerful in changing a consumer's opinion or attitudes. Our work explores this message-based persuasion in online rating systems.

Lastly, we discuss several related works from a modeling perspective. Xie et al. [31, 33] proposed a rating model, where rating are statistically independent. Their study revealed insights on the minimum number of reviews needed such that a consensus is reached. However, their model does not capture important factors that historical ratings (or reviews) influence users' opinions. Wang et al. [27] developed an additive generative model to quantify the impact of historical ratings on subsequent users' opinions. The model assumes that users' opinions is influenced by all previous ratings. However, a number of survey studies [4, 20, 23] revealed that most users only trust recent ratings (or reviews). Our high-order Markov chain model captures this rating recency factor. Furthermore, our high order Markov chain model enables us to study convergence of opinions theoretically, while their model can only allow one to study such questions via simulating the model. Studying the convergence of opinions theoretically is important, because it reveals important insights on product quality estimation accuracy, defending misbehavior ratings, and the like. Zhang et al. [36, 37] proposed a latent factor based model to quantify the impact of historical ratings on subsequent users' opinions. Their model also does not capture the rating recency factor [4, 20, 23], while our high order Markov chain model captures it. Also, their can only enable one to study such questions via simulating the model, which our model enables us to study such questions theoretically.

8 PROOF FOR LEMMAS AND THEOREMS

We first state a lemma to characterize the states of the Markov chain $(\mathcal{S}, \mathbf{P})$, which will be used a lot in our proof.

Lemma 2. *Suppose Assumption 3 holds and $\gamma < 1$. The Markov chain $(\mathcal{S}, \mathbf{P})$ is an ergodic unichain with recurrent states denoted by $\mathcal{R}$ and transient sates denoted by $\mathcal{T}$:*

$$\mathcal{R} = \mathcal{M}^N, \quad \mathcal{T} = \mathcal{S} \setminus \mathcal{R}.$$

Proof. We need to show that $\mathcal{R}$ contains a class of recurrent states and $\mathcal{T}$ contains all the transient states (might be of multiple classes). Let us first show that $\mathcal{R}$ form a class of states, i.e., any state $\mathbf{s} \in \mathcal{R}$ communicates with any other state $\tilde{\mathbf{s}} \in \mathcal{R}$ and does not communicates with any other state in $\mathcal{T}$. Without loss of generality, suppose $R_{i-N} = s_1, \ldots, R_{i-1} = s_N$. Note that Assumption 3 holds and $\gamma < 1$, which imply that R_i will be of level $m, \forall m \in \mathcal{M}$, rating with a positive probability. Then with a positive probability, we can have $R_i = \tilde{s}_1$. Similarly, with a positive probability we can have $R_{i+n} = \tilde{s}_{1+n}$, for all $n = 1, \ldots, N-1$. These N ratings $R_i, \ldots, R_{i+N-1}$ form a "path," through which the persuasion cascade transits from state $\mathbf{s}$ to state $\tilde{\mathbf{s}}$. We can construct a similar path, through which the the persuasion cascade transits from state $\tilde{\mathbf{s}}$ to state $\mathbf{s}$. It is easy to observe that it is only possible to construct a path such that the persuasion cascade transits from a state in $\mathcal{T}$ to a state in $\mathcal{R}$. Thus we conclude that $\mathcal{R}$ form a class of states. Furthermore, given a state $\mathbf{s} \in \mathcal{R}$, any state $\tilde{\mathbf{s}} \in \mathcal{R}$ is accessible from $\mathbf{s}$ and $\mathbf{s}$ is accessible from $\tilde{\mathbf{s}}$. Thus $\mathcal{R}$ form a class of recurrent states. Note that all the states in $\mathcal{R}$ have the same period. Consider a special class of states $\mathbf{s} = [m, \ldots, m]$, $\forall m \in \mathcal{M}$. Then there exists a self-loop transition for such states, i.e., receive a rating m and the next state is still $\tilde{\mathbf{s}} = [m, \ldots, m]$. The period for such states is 1. Thus the period of $\mathcal{R}$ is 1. We therefore conclude that $\mathcal{R}$ is an esodic class of states. Note that all states in $\mathcal{T}$ are transit, because the persuasion cascade never transits back to such states. □

Proof of Theorem 1: Note that the Markov chain $(\mathcal{S}, \mathbf{P})$ converges to a steady state distribution $\boldsymbol{\pi}$. Then the persuasion cascade will at each state $\mathbf{s}$ with probability $\pi(\mathbf{s})$. Note that persuasion cascade at each state $\mathbf{s}$ corresponds to that a rating s_N is earned. Hence, under the steady-state distribution $\boldsymbol{\pi}$, each rating will be of level m with probability $\sum_{\mathbf{s} \in \mathcal{S}} \pi(\mathbf{s}) \mathbf{1}_{\{s_N = m\}}$. This implies that

$$\widehat{\alpha}_{i,m} \xrightarrow{a.s.} \sum_{\mathbf{s} \in \mathcal{S}} \pi(\mathbf{s}) \mathbf{1}_{\{s_N = m\}}.$$

Write it in matrix form, we complete this proof. □

Proof of Theorem 2: Lemma 2 implies that the Markovian chain has a unique steady-state distribution $\boldsymbol{\pi}$. Furthermore, $\mathcal{T}$ contains all the transient states, thus $\pi(\mathbf{s}) = 0$ for all $\mathbf{s} \in \mathcal{T}$. Under the steady-state distribution, the persuasion cascade will at each recurrent state with a positive probability, thus $\pi(\mathbf{s}) > 0$ for all $\mathbf{s} \in \mathcal{R}$. Then it follows that $\boldsymbol{\pi}_{\mathcal{R}}$ satisfies Equation (4). □

Proof of Corollary 1: Note that for the consumers provide unbias reviews. Namely, each rating will be of level m with probability α_m. Then, in the steady state, the persuasion cascade will be at each state $\mathbf{s}\mathcal{R}$ with probability

$$\tilde{\pi}(\mathbf{s}) = \prod_{n=1}^{N} \alpha_{s_n}.$$

On the other hand, one can easily check that $\tilde{\boldsymbol{\pi}}$ (with restrict to $\mathcal{R}$) satisfies Equation (4). □

Proof of Theorem 3: Let $\mathbf{e} \triangleq [1, \ldots, 1]$ denote a $|\mathcal{R}|$-dimensional row vector. First, we can write the linear system (4) in a matrix form as

$$\begin{bmatrix} \mathbf{P}_{\mathcal{R}}^T - \mathbf{I} \\ \mathbf{e} \end{bmatrix} \boldsymbol{\pi}_{\mathcal{R}}^T = \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}$$

Note that the coefficient matrix is a $(|\mathcal{R}| + 1) \times |\mathcal{R}|$ matrix. We thus multiple each side of the above equation with a $|\mathcal{R}| \times (|\mathcal{R}| + 1)$ matrix $[\mathbf{P}_{\mathcal{R}} - \mathbf{I} \;\; \mathbf{e}^T]$ to obtain that

$$\begin{bmatrix} \mathbf{P}_{\mathcal{R}} - \mathbf{I} & \mathbf{e}^T \end{bmatrix} \begin{bmatrix} \mathbf{P}_{\mathcal{R}}^T - \mathbf{I} \\ \mathbf{e} \end{bmatrix} \boldsymbol{\pi}_{\mathcal{R}}^T = \begin{bmatrix} \mathbf{P}_{\mathcal{R}} - \mathbf{I} & \mathbf{e}^T \end{bmatrix} \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}.$$

Then it follows that $[(\mathbf{P}_{\mathcal{R}} - \mathbf{I})(\mathbf{P}_{\mathcal{R}}^T - \mathbf{I}) + \mathbf{e}^T\mathbf{e}]\boldsymbol{\pi}_{\mathcal{R}}^T = \mathbf{e}^T$ Similarly, we have that

$$\left[(\tilde{\mathbf{P}}_{\mathcal{R}} - \mathbf{I})(\tilde{\mathbf{P}}_{\mathcal{R}}^T - \mathbf{I}) + \mathbf{e}^T\mathbf{e}\right] \tilde{\boldsymbol{\pi}}_{\mathcal{R}}^T = \mathbf{e}^T. \tag{8}$$

Note the linear system (4) has a unique solution when $\gamma = 0$. Thus, the inverse of the coefficient matrix exists, i.e., we have

$$\left[(\tilde{\mathbf{P}}_{\mathcal{R}} - \mathbf{I})(\tilde{\mathbf{P}}_{\mathcal{R}}^T - \mathbf{I}) + \mathbf{e}^T\mathbf{e}\right]^{-1}. \tag{9}$$

Note that $\mathbf{P}_{\mathcal{R}} = \tilde{\mathbf{P}}_{\mathcal{R}} + \mathbf{E}$. Then we can further have

$$\begin{aligned} & \left[(\tilde{\mathbf{P}}_{\mathcal{R}} + \mathbf{E} - \mathbf{I})(\tilde{\mathbf{P}}_{\mathcal{R}}^T + \mathbf{E}^T - \mathbf{I}) + \mathbf{e}^T\mathbf{e}\right] \boldsymbol{\pi}_{\mathcal{R}}^T = \mathbf{e}^T \\ \Leftrightarrow \quad & \left[(\tilde{\mathbf{P}}_{\mathcal{R}} - \mathbf{I})(\tilde{\mathbf{P}}_{\mathcal{R}}^T - \mathbf{I}) + \mathbf{e}^T\mathbf{e}\right] \boldsymbol{\pi}_{\mathcal{R}}^T + \left[\mathbf{E}\left(\tilde{\mathbf{P}}_{\mathcal{R}}^T - \mathbf{I} + \frac{\mathbf{E}^T}{2}\right) + \left(\tilde{\mathbf{P}}_{\mathcal{R}} - \mathbf{I} + \frac{\mathbf{E}}{2}\right)\mathbf{E}^T\right] \boldsymbol{\pi}_{\mathcal{R}}^T = \mathbf{e}^T. \end{aligned} \tag{10}$$

Combine Equations (8), (9), and (10) and apply the matrix perturbation theory [25], i.e., perturbing the linear system, we have

$$\begin{aligned} \frac{||\boldsymbol{\pi}_{\mathcal{R}} - \tilde{\boldsymbol{\pi}}_{\mathcal{R}}||}{||\boldsymbol{\pi}_{\mathcal{R}}||} &\leq ||[(\tilde{\mathbf{P}}_{\mathcal{R}} - \mathbf{I})(\tilde{\mathbf{P}}_{\mathcal{R}}^T - \mathbf{I}) + \mathbf{e}^T\mathbf{e}]^{-1}||(||\mathbf{E}|| \times ||\tilde{\mathbf{P}}_{\mathcal{R}} - \mathbf{I} + \mathbf{E}^T|| + ||\mathbf{E}^T|| \times ||\tilde{\mathbf{P}}_{\mathcal{R}}^T - \mathbf{I}||) \\ &\leq ||[(\mathbf{P}_{\mathcal{R}} - \mathbf{I})(\mathbf{P}_{\mathcal{R}}^T - \mathbf{I}) + \mathbf{e}^T\mathbf{e}]^{-1}|| \\ &\quad \times \left(||\mathbf{E}|| \times ||\tilde{\mathbf{P}}_{\mathcal{R}} - \mathbf{I}|| + ||\mathbf{E}^T|| \times ||\tilde{\mathbf{P}}_{\mathcal{R}}^T - \mathbf{I}|| + ||\mathbf{E}||^2 + ||\mathbf{E}^T||^2 + ||\mathbf{E}|| \times ||\mathbf{E}^T||\right). \end{aligned}$$

We can then complete this proof by noting that $||\boldsymbol{\pi} - \tilde{\boldsymbol{\pi}}|| = ||\boldsymbol{\pi}_{\mathcal{R}} - \tilde{\boldsymbol{\pi}}_{\mathcal{R}}||$ and $||\tilde{\boldsymbol{\pi}}|| = ||\tilde{\boldsymbol{\pi}}_{\mathcal{R}}||$. □

Proof of Corollary 2: Note that $\boldsymbol{\alpha} = \tilde{\boldsymbol{\pi}}\mathbf{W}$ and $\widehat{\boldsymbol{\alpha}}_{\infty} = \boldsymbol{\pi}\mathbf{W}$. Then it follows that

$$||\boldsymbol{\alpha} - \widehat{\boldsymbol{\alpha}}_{\infty}|| = ||(\tilde{\boldsymbol{\pi}} - \boldsymbol{\pi})\mathbf{W}|| \leq ||\tilde{\boldsymbol{\pi}} - \boldsymbol{\pi}|| \times ||\mathbf{W}|| \leq \epsilon_{sd}||\tilde{\boldsymbol{\pi}}|| \times ||\mathbf{W}||.$$

This proof is then complete. □

Proof of Corollary 3: Note that $A(\boldsymbol{\alpha}) = \boldsymbol{\alpha}\mathbf{m}^T = \tilde{\boldsymbol{\pi}}\mathbf{W}\mathbf{m}^T$. Also note that $A(\widehat{\boldsymbol{\alpha}}_{\infty}) = \widehat{\boldsymbol{\alpha}}_{\infty}\mathbf{m}^T = \boldsymbol{\pi}\mathbf{W}\mathbf{m}^T$. Then it follows that

$$\begin{aligned} |A(\boldsymbol{\alpha}) - A(\widehat{\boldsymbol{\alpha}}_{\infty})| = ||\tilde{\boldsymbol{\pi}}\mathbf{W}\mathbf{m}^T - \boldsymbol{\pi}\mathbf{W}\mathbf{m}^T|| &\leq ||\tilde{\boldsymbol{\pi}} - \boldsymbol{\pi}|| \times ||\mathbf{W}\mathbf{m}^T|| \\ &\leq \epsilon_{sd}||\tilde{\boldsymbol{\pi}}|| \times ||\mathbf{W}\mathbf{m}^T||. \end{aligned}$$

This proof is then complete. □

Proof of Theorem 4: Note that

$$\widehat{\alpha}_{i,m} = \sum_{j=1}^{i} \mathbf{1}_{\{R_j=m\}}/i = \sum_{j=1}^{i} \mathbf{1}_{\{\mathrm{S}_j^{[1]}=m\}}/i,$$

where $\mathrm{S}_j \triangleq [S_{j,1},\dots,S_{j,N}] \in \mathcal{S}$ denotes the state of the Markov chain $(\mathcal{S},\mathbf{P})$ corresponding to the j-th rating R_j, i.e., $\mathrm{S}_j^{[1]} \triangleq S_{j,N} = R_j$. Also note that $\widehat{\alpha}_{\infty,m} = \lim_{i\to\infty}\sum_{j=1}^{i}\mathbf{1}_{\{\mathrm{S}_j^{[1]}=m\}}/i$. Then it follows that

$$|\widehat{\alpha}_{i,m} - \widehat{\alpha}_{\infty,m}| > \epsilon \Leftrightarrow \left|\sum_{j=1}^{i} \mathbf{1}_{\{\mathrm{S}_j^{[N]}=m\}} - i\lim_{k\to\infty}\frac{\sum_{j=1}^{k}\mathbf{1}_{\{\mathrm{S}_j^{[N]}=m\}}}{k}\right| > i\epsilon.$$

Apply Lemma 2, we obtain that the Markov chain $(\mathcal{S},\mathbf{P})$ is an ergodic unichain with recurrent states $\mathcal{R}$. Also note that initially, each product is at a transient state and once it receives N ratings, it hit one recurrent state with distribution $\boldsymbol{\pi}_0$. Then, apply the Chernoff bound (Markovian version) [19] with restrict to the recurrent states, i.e., the Markov chain $(\mathcal{R},\mathbf{P}_{\mathcal{R}})$, we have

$$\begin{aligned}\mathbb{P}[|\widehat{\alpha}_{i,m} - \widehat{\alpha}_{\infty,m}| > \epsilon] &\le \sqrt{2}D(\boldsymbol{\pi}_0,\boldsymbol{\pi}_{\mathcal{R}})\exp\left(-\frac{1}{2}\frac{(i-N)^2\epsilon^2\nu_{ps}}{8(i-N+1/\nu_{ps})+20(i-N)\epsilon}\right)\\ &\le \sqrt{2}D(\boldsymbol{\pi}_0,\boldsymbol{\pi}_{\mathcal{R}})\exp\left(-\frac{1}{2}\frac{(i-N)\epsilon^2\nu_{ps}}{8(1+1/\nu_{ps})+20\epsilon}\right)\end{aligned}$$

Then to make

$$\sqrt{2}D(\boldsymbol{\pi}_0,\boldsymbol{\pi}_{\mathcal{R}})\exp\left(-\frac{1}{2}\frac{(i-N)\epsilon^2\nu_{ps}}{8(1+1/\nu_{ps})+20\epsilon}\right) \le \frac{\delta}{M},$$

we only need

$$i \ge \frac{1}{\epsilon^2}\frac{16(1+1/\nu_{ps})+40\epsilon}{\nu_{ps}}\ln\left(\frac{\sqrt{2}MD(\boldsymbol{\pi}_0,\boldsymbol{\pi}_{\mathcal{R}})}{\delta}\right)+N.$$

By the union bound we have

$$\mathbb{P}\left[\max_{m\in M}|\widehat{\alpha}_{i,m} - \widehat{\alpha}_{\infty,m}| \le \epsilon\right] \ge 1 - \sum_{m\in\mathcal{M}}\mathbb{P}[|\widehat{\alpha}_{i,m} - \widehat{\alpha}_{\infty,m}| > \epsilon] = 1-\delta.$$

Note that

$$||\widehat{\boldsymbol{\alpha}}_i - \widehat{\boldsymbol{\alpha}}_\infty|| \le \max_{m\in M}|\widehat{\alpha}_{i,m} - \widehat{\alpha}_{\infty,m}| \times ||\mathbf{c}|| = \epsilon||\mathbf{c}||.$$

This proof is then complete. □

Proof of Corollary 4: Note that $A(\widehat{\boldsymbol{\alpha}}_i) = \widehat{\boldsymbol{\alpha}}_i\mathbf{m}^T$. Also note that $A(\widehat{\boldsymbol{\alpha}}_\infty) = \widehat{\boldsymbol{\alpha}}_\infty\mathbf{m}^T$. Then it follows that

$$|A(\widehat{\boldsymbol{\alpha}}_i) - A(\widehat{\boldsymbol{\alpha}}_\infty)| = ||\widehat{\boldsymbol{\alpha}}_i\mathbf{m}^T - \widehat{\boldsymbol{\alpha}}_\infty\mathbf{m}^T|| \le \epsilon||\mathbf{c}||\times||\mathbf{m}||.$$

This proof is then complete. □

Proof of Corollary 5: Note that the $(L,\mathcal{F})$-misbehavior does not affect the transition matrix of the Markov chain $(\mathcal{S},\mathbf{P})$. Hence Lemma 2 still holds. Lemma 2 implies this corollary. □

Proof of Corollary 6: Note that Lemma 2 still holds. Lemma 2 implies the steady-state distribution of the Markov chain $(\mathcal{S},\mathbf{P})$ is invariant of the misbehavior. □

Proof of Theorem 5: Note that the domains of $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$ and γ are compact. Also note that the constraints are linear functions. That the feasible space is also compact. Note that the objective function is continuous in the interior of the feasible space. On the boundary of the feasible space, the objective function is either semi-continuous or goes to $-\infty$ (can not be optimal on these boundaries). This proof is then complete. □

Proof of Lemma 1: Without loss of generality, let us consider $\boldsymbol{\alpha}$. Note that the objective function is a summation of log functions. Each log function is concave with respect to $\boldsymbol{\alpha}$, because $\sum_{n \in \mathcal{N}} \beta_n[\gamma\theta_{R_i}(\mathbf{h}_{i-1}^{[n]}) + (1-\gamma)\alpha_{R_i}]$ is affine with respect to $\boldsymbol{\alpha}$. Similarly, we conclude this lemma for $\boldsymbol{\beta}$ and γ. □

9 CONCLUSION

This article develops a data analytic framework to analyze persuasion cascades in online rating systems. We develop a mathematical model to captures 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.

REFERENCES

[1] Gediminas Adomavicius, Jesse Bockstedt, Shawn P. Curley, and Jingjing Zhang. 2016. Understanding effects of personalized vs. aggregate ratings on user preferences. In *Proceedings of the 2016 INTRS WORKSHOP.*

[2] Gomaa M. Agag and Ahmed A. El-Masry. 2017. Why do consumers trust online travel websites? Drivers and outcomes of consumer trust toward online travel websites. *Journal of Travel Research* 56, 3 (2017), 347–369.

[3] Jonah Berger. 2012. *Bad Reviews Can Boost Sales. Here's Why.* Harvard Business Review.

[4] BrightLocal. 2016. *Local Consumer Review Survey.* BrightLocal.

[5] Judith A. Chevalier and Dina Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research* 43, 3 (2006), 345–354.

[6] Dan Cosley, Shyong K. Lam, Istvan Albert, Joseph A. Konstan, and John Riedl. 2003. Is seeing believing?: How recommender system interfaces affect users' opinions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* 585–592.

[7] Michael Gebicki. 2015. *Should We Trust TripAdvisor Reviews?* Traveller.

[8] Fangjian Guo and David B. Dunson. 2015. Uncovering systematic bias in ratings across categories: A Bayesian approach. In *Proceedings of the 9th ACM Conference on Recommender Systems.* 317–320.

[9] Carl I. Hovland, Irving L. Janis, and Harold H. Kelley. 1953. *Communication and Persuasion: Psychological Studies of Opinion Change.* Yale University Press.

[10] Nitin Jindal and Bing Liu. 2007. Review spam detection. In *Proceedings of the International Conference on World Wide Web.*

[11] Yehuda Koren. 2009. Collaborative filtering with temporal dynamics. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* 447–456.

[12] Sanjay Krishnan, Jay Patel, Michael J. Franklin, and Ken Goldberg. 2014. A methodology for learning, analyzing, and mitigating social influence bias in recommender systems. In *Proceedings of the 8th ACM Conference on Recommender Systems.* 137–144.

[13] Georg Lackermair, Daniel Kailer, and Kenan Kanmaz. 2013. Importance of online product reviews from a consumer's perspective. *Advances in Economics and Business* 1, 1 (2013), 1–5.

[14] Mengxiang Li, Liqiang Huang, Chuan-Hoo Tan, and Kwok-Kee Wei. 2013. Helpfulness of online product reviews as seen by consumers: Source and content features. *International Journal of Electronic Commerce* 17, 4 (2013), 101–136.

[15] Michael Luca. 2016. Reviews, reputation, and revenue: The case of Yelp. com. Harvard Business School Working Papers 12-016, Harvard Business School.

[16] Julian John McAuley and Jure Leskovec. 2013. From amateurs to connoisseurs: Modeling the evolution of user expertise through online reviews. In *Proceedings of the 22nd International Conference on World Wide Web.*

[17] Lev Muchnik, Sinan Aral, and Sean J. Taylor. 2013. Social influence bias: A randomized experiment. *Science* 341, 6146 (2013), 647–651.

[18] Susan M. Mudambi and David Schuff. 2010. What makes a helpful online review? A study of customer reviews on Amazon.Com. *MIS Quarterly* 34, 1 (2010), 185–200.

[19] Daniel Paulin. 2015. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electronic Journal of Probability* 20, 79 (2015), 1–32.

[20] Stacey Rudolph. 2016. *50 Stats You Need to Know About Online Reviews [Infographic].* Business 2 Community.

[21] Matthew J. Salganik, Peter Sheridan Dodds, and Duncan J. Watts. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311, 5762 (2006), 854–856.

[22] Patrick Shafto and Olfa Nasraoui. 2016. Human-recommender systems: From benchmark data to benchmark cognitive models. In *Proceedings of the 10th ACM Conference on Recommender Systems.*

[23] Khusbu Shrestha. 2016. *50 Stats You Need to Know About Online Reviews [Infographic].* Vendasta.

[24] Signpost. 2017. *The Future of Commerce: The Impact of Location-Based Reviews on Revenue.* Signpost.

[25] Gilbert W. Stewart. 1990. Matrix Perturbation Theory. Academic Press.

[26] ThomsonReuters. 2016. *Consumers Overwhelmingly Trust Online Reviews, Says New FindLaw and Super Lawyers Survey.* ThomsonReuters.

[27] Ting Wang, Dashun Wang, and Fei Wang. 2014. Quantifying herding effects in crowd wisdom. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* 1087–1096.

[28] Wendy Wood. 2000. Attitude change: Persuasion and social influence. *Annual Review of Psychology* 51, 1 (2000), 539–570.

[29] Stephen J. Wright and Jorge Nocedal. 1999. Numerical optimization. *Springer Science* 35, 67–68 (1999), 7.

[30] Hong Xie, Yongkun Li, and John C.S. Lui. 2019. Understanding persuasion cascades in online product rating systems. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence.* Vol. 33. 5490–5497.

[31] Hong Xie and John C.S. Lui. 2014. Mathematical modeling of group product recommendation with partial information: How many ratings do we need?*Performance Evaluation* 77 (2014), 72–95.

[32] H. Xie and J. C. S. Lui. 2015. A data driven approach to uncover deficiencies in online reputation systems. In *Proceedings of the 2015 IEEE International Conference on Data Mining.*

[33] Hong Xie and John C. S. Lui. 2015. Mathematical modeling and analysis of product rating with partial information. *ACM Transactions on Knowledge Discovery from Data* 9, 4 (2015), 26:1–26:33.

[34] Meital Yachin. 2017. *Online Reviews – Trends, Innovations & Strategy 2017.* Carmelon Digital.

[35] Stefany Zaroban. 2015. *Product Reviews Boost Revenue Per Online Visit 62%.* Digital Commerce.

[36] Xiaoying Zhang, Hong Xie, Junzhou Zhao, and John C. S. Lui. 2019. Understanding assimilation-contrast effects in online rating systems: Modelling, debiasing, and applications. *ACM Transactions on Information Systems* 38, 1 (Oct. 2019), 25 pages. DOI: https://doi.org/10.1145/3362651

[37] Xiaoying Zhang, Junzhou Zhao, and John C. S. Lui. 2017. Modeling the assimilation-contrast effects in online product rating systems: Debiasing and recommendations. In *Proceedings of the 11th ACM Conference on Recommender Systems.* 98–106.

[38] Yi Zhao, Sha Yang, Vishal Narayan, and Ying Zhao. 2013. Modeling consumer learning from online product reviews. *Marketing Science* 32, 1 (2013), 153–169.

Received November 2019; revised September 2020; accepted December 2020