

Input Online Review Data and Related Bias in Recommender Systems

Selwyn Piramuthu^{1,3,*}, Gaurav Kapoor¹, Wei Zhou^{2,3}, Sjouke Mauw⁴

¹Information Systems and Operations Management, University of Florida
Gainesville, Florida 32611-7169, USA

²ESCP Europe, Paris, France

³RFID European Lab, Paris, France

⁴Université du Luxembourg, Faculté des Sciences, de la Technologie et de la Communication (FSTC),
6, rue Richard Coudenhove-Kalergi, L-1359, Luxembourg

**corresponding author email: selwyn@ufl.edu; phone: 352-392-8882*

Abstract

A majority of extant literature on recommender systems assume the input data as a given to generate recommendations. Both implicit and/or explicit data are used as input in these systems. The existence of various challenges in using such input data including those associated with strategic source manipulations, sparse matrix, state data, among others, are sometimes acknowledged. While such input data are also known to be rife with various forms of bias, to our knowledge no explicit attempt is made to correct or compensate for them in recommender systems. We consider a specific type of bias that is introduced in online product reviews due to the sequence in which these reviews are written. We model several scenarios in this context and study their properties.

Key words: sequential bias, online reviews

1 Introduction

Potential customers sometimes have the option of using recommender systems (e.g., amazon.com, buzzillions.com, consumersearch.com, digg.com, Google AdSense, Netflix challenge, prorevs.com, slash-dot.com) as convenient (although not completely reliable) automated sources of information in situations where there is a lack of other alternatives. These systems are for the most part used to supplement rather than to supplant the real thing which is recommendation from a known and completely reliable expert source.

Source of input for recommender systems include (implicit) past behavior (e.g., consumer transaction data, bookmark, page view time, from and to link for a Web page, social network) and

(explicit) customer reviews. Both implicit and explicit data complement each other in terms of information content since the former records the behavior (i.e., customer A bought widget X) while the latter records details of this customer’s (dis-)satisfaction with this purchase. Recommendations are generated based on (dis-)similarity between the characteristics of the user being recommended to and others in the database as well as (dis-)similarity between item of interest and related items. Several methods are used in the process including collaborative filtering (e.g., amazon.com) and content filtering (e.g., Music Genome Project used in pandora.com). Collaborative filtering uses the (dis-)similarity information across users and items (e.g., [17]). Content filtering, on the other hand, is based on the characteristics of users and items. Adomavicius and Tuzhilin ([1]) provide an excellent overview of this general area.

Given the popularity of recommender systems, several facets of such systems have been extensively studied including mining user-generated review data for implicit as well as explicit patterns, attacks, interface design, among others (e.g., [6], [7], [9], [18]). Other than attacks, which explicitly manipulate input data to achieve an intended recommendation (e.g., manipulate reviews so an item of interest enters or leaves the set of highly recommended items), other aspects of input data (e.g., bias) have not received their fair share of attention from researchers in this area.

Bias in user-generated reviews can take several forms including personal (based on past experience, interest, attitude), extreme reviews (overly positive or negative), context (e.g., review of a camera’s resolution characteristics can be positively or negatively biased based on its use - pictures for high-resolution printing vs. posting low-resolution pictures online), temporal (early vs. late adopters of a product may have different perspectives on the same product), awareness effect ([8]), herd behavior ([4]), and confirmation bias ([3]).

Sequential bias is a variant of first-impression bias (e.g., primacy-recency effect) and is also influenced by pre-existing (positive, negative) bias. Thus, the role played by first impression bias cannot be overestimated ([5], [16]). Therefore, the review that is first seen by a prospective customer of the product of interest plays a significant role in purchase decisions that follow. These reviews are quite influential since prospective purchasers of reviewed products rely heavily on these reviews in making their purchase decisions (e.g., [21]). The sequence in which reviews are written play an appreciable role in how the reviews that follow later in the sequence are written. For example, if a reviewer is favorable to the product reviewed, she might be biased to write stronger reviews to somehow compensate for the effects of existing negative reviews and *vice versa*. The reviews thus written are biased, in part, due to their position in the sequence of reviews. In turn, the recommender systems that use these biased reviews to generate their recommendations will clearly generate biased recommendations due, in part, to this sequential bias.

We purport to fill this gap in extant literature by specifically considering sequential bias present in consumer reviews and the consequence of this bias on resulting recommendations generated. In other words, while recommender systems use user-generated data as-is, we believe there is a need to rid this data of sequential bias to provide better or less-biased recommendations. By explicitly acknowledging the existence of sequential bias and actively employing means to remove it from input data to recommender systems, one can alleviate its effect in the recommendations. We are,

therefore, interested in the scenario where a customer purchases/experiences a product and then proceeds to provide a written review of this product online. During this process, the customer also has a chance to read existing reviews on this product before writing a review. It is precisely these existing reviews that causes sequential bias in the next review that is written.

The remainder of the paper is organized as follows. We discuss related background information and literature in the next section. We study the dynamic associated with sequential bias and its effect on recommender systems in Section 3. Section 4 concludes this paper with a brief discussion.

2 Background and Related Literature

A generic framework of a recommender system is given in Figure ?? . Both implicit and explicit data are used as input to the system, which uses these to generate (i.e., learn and store explicit knowledge in) the knowledge-base. The knowledge-base essentially comprises both explicit and implicit patterns extracted from (implicit & explicit) input data. The recommender system then waits for a user to enter the system. Upon arrival of a user, who could be a potential customer, the system takes a snap-shot of this user’s characteristics and matches this with learned knowledge to generate appropriate recommendations in a timely manner. In what follows, the recommendations would be used to update system configuration either automatically or via human interference. Consequently, these new updates would eventually alter online users’ behavior towards pricing and recommendation. The closed loop assures that normal system performance can be maintained at a stabilized level according to the theory of automation.




Figure 1: A generic recommender system framework

labelframework

Although it is generally assumed that what consumers do (i.e., implicit or past behavior data) provides better information for recommender systems than what they say (explicit or consumer review data), both these data provide complementary information that are beneficially utilized in recommender systems. While both explicitly and implicitly generated data are used in recommender systems, we are interested in only explicit (i.e., user-generated reviews) input data in recommender systems. The use of user-generated explicit recommendation data has its associated challenges. We provide a brief overview of several such challenges and then some related literature in this area.

2.1 Some Challenges

Recommender systems face several challenges when dealing with explicit input data. A list of a selected few of these challenges include: (a) strategic source manipulations (e.g., profile injection attacks such as sybil, shill, random, sampling, average, bandwagon), (b) those associated with equally weighted input (e.g., user-generated reviews for a computer and a pack of chewing gum are not treated differently), (c) the endemic sparse matrix and the difficulty in generating useful patterns from such a matrix, cold start problem that arises when a user or product is new to the system and the absence of historical data on these entities and their characteristics, (d) differences in user risk tolerance levels, (e) staleness of data used in generating recommendations, (f) seasonality and trends in consumer preferences and their effect on user-generated reviews, and (g) general input data bias. We are interested in the bias which occurs due to the characteristics of the reviewed item (e.g., price, familiarity, stake to the customer, whether this item was purchased as a gift to someone else, the relationship of the purchaser to the giftee), the presence of extraneous stimulus whereby the item would not have been explicitly purchased had it not been for promotions and bundling, the highly self-selective nature of providing reviews and the sequential manner in which reviews are written. We are specifically interested in the latter - i.e., the sequence in which reviews are written and the (mostly unintended) bias that is introduced in these reviews resulting from its position in the sequence.

When perusing existing reviews on the product of interest, the intensity/magnitude of the reviewer’s sentiment as well as the positivity (or negativity) of reviews on various product features/characteristics (as well as the overall review of this product) certainly affect the reader. Several studies in the social sciences suggest that people often assign more weight to negative information than positive information of equal intensity (e.g., [2], [11]). Mizerski ([14]) found that product attributes rated unfavorably exert greater influence than those rated favorably on consumers’ attributions, beliefs and attitudes. This phenomenon has been termed the negativity effect (bias). Ahluwalia ([2]) found that highly committed consumers showed positivity effect (bias) where they weigh positive information more than negative information. I.e., there is evidence for both positivity and negativity effect depending on consumer as well as product characteristics.

In addition to the introduction of unintentional bias in user-generated product reviews, there also exists bias that are intentionally introduced due to professional relationships and friendships ([22]) and others with ulterior motives (e.g., [10], [20], [23]). Buzz marketing (e.g., [19]) is a variation on the same theme with the explicit intention of promoting a product, service, or idea.

2.2 Related Literature

The literature on recommender systems is extensive and covers a wide spectrum of related issues. We list a few from among these here. Since online product reviews are not strictly regulated, there are opportunities for the introduction of ‘fake’ or intentionally biased reviews. A popular approach to mediating this risk is through self-regulation in the form of ratings and comments. Examples of this include *x of y people found the following review helpful* on amazon.com, *x of y customers found*

this review helpful on buzzillions.com, *diggs* on digg.com, and *Karma* on slashdot.org. There are also commercial service providers such as eModeration.com where content moderation issues are evaluated in a neutral environment and KwikChex.com where negative reviews are contested.

Several studies have been reported in the literature on variations of both unintentional and intentional bias introduced in reviews. Lauw, Lim, and Wang ([12]) investigate how deviation in evaluation activities may reveal bias on the part of reviewers and controversy on the part of evaluated objects. They claim that significant deviation may help reveal bias of reviewers or controversy of objects. They define bias as the deviation in a user-assigned score from the mean score for a ‘less controversial’ object. They measure controversy as the deviation of user-assigned scores for that object.

Staddon and Chow ([22]) use association rule mining to identify bias introduced by book reviewers and authors of books they review. Specifically, they look for associations between these reviewers and authors in Web documents. When these two are present in a significant number of Web documents where either of them are present, there is a valid reason demanding closer scrutiny. They consider 300 books on amazon.com that are categorized under ‘cryptography’ and observe that this method results in high precision (few false positives)

Using data on book reviews from amazon.com, Li and Hitt ([13]) find that earlier reviews influence the item’s quality perception and have a positive bias on sales volume of the item considered. Moreover, they find that prospective customers do not take this bias into consideration when making their decisions. In practice, reviews along with other online consumer behavior are aggregated to develop effective recommendations. A common practice is to use collaborative filtering that incorporates multiple data sources to predict consumer preference and behavior, such as those that are used at Amazon and Netflix.

3 Effect of Sequential Bias in Recommender Systems

We consider several variations of bias introduced in user-generated reviews. We assume reviewers rate a product as well as its attributes. For example, a reviewer may assign an overall rating for a car as positive while assigning negative rating to its fuel efficiency, very positive rating to its cup holders, neutral rating to its maintenance cost, etc. I.e., although the overall rating is a *function* of the ratings for each of its characteristics, this function is highly likely to vary among reviewers based primarily on their individual and aggregate derived utility. For the remainder of this paper, when we refer to a review as biased or unbiased we are referring *only* to sequential bias. The absence of sequential bias does not in any way signify the presence or absence of other types of biases. However, given that we concern ourselves only with sequential bias, other types of biases are irrelevant to this work and are hence not considered.

We consider reviews written for a specific product of interest. We assume that this product is unique and that it does not have any substitutes. Since we are not interested in content analysis of written reviews, we only consider a single numerical score that is assigned by each reviewer. The modeling exercise is kept distribution-free on purpose to facilitate general applicability. Reviews

written after the first review do not assume knowledge of the existence of sequential bias. I.e., the occurrence of sequential bias is implicit and unintentional on the part of the reviewer and any given reviewer is unaware of the presence of sequential bias in earlier as well as the current review. The following notations are used in this paper:

R_i - (true) online review written by reviewer i (without sequential bias)

R'_i - online review written by reviewer i (with sequential bias)

We assume $R_1 = R'_1$

β - (bias) discount factor ($0 < \beta < 1$)

B_{ij} - bias of review R'_i on review R_j ($B_{ij} = 1$ for $i = j$)

$$s_{ij} - \text{direction (sign) of bias} = \begin{cases} 0 & \text{if } R'_i \text{ \& } R_j \text{ are both positive/negative/neutral} \\ -1 & \text{if } R'_i \text{ is positive \& } R_j \text{ is not positive} \\ -1 & \text{if } R'_i \text{ is negative \& } R_j \text{ is not negative} \\ 0 & \text{if } R'_i \text{ is neutral} \end{cases}$$

where $s_{ij} = 0$ for $R_i = 0$

$U()$ - utility function

a, b - minimum and maximum review scores

k - risk tolerance model parameter ($k > 1$)

u_i - review R'_i 's usefulness score ($0 \leq u_i \leq 1$)

The s_{ij} term directly represents sequential bias where i represents an earlier review and j represents the current review. I.e., s_{ij} determines the sign of the effect of j on i . When both reviewers i and j are positive, negative, or neutral, there is no friction between their reviews and s_{ij} equals zero. However, when reviewer i is positive or neutral and reviewer j is not positive or negative respectively reviewer i 's review exerts negative friction on reviewer j 's review due to sequential bias and this is represented by a negative s_{ij} . Similarly when reviewer i is negative or neutral and reviewer j is not negative or positive respectively, reviewer i 's review exerts positive friction on reviewer j 's review due to sequential bias and this is represented by a positive s_{ij} .

The first scenario we consider includes two entities of interest: the reviewer and the online review written by this reviewer on a given product of interest. We consider the case where the potential reviewer purchases a product, experiences this product first-hand, and then decides to write an online review of this product. Without loss of generality, we consider the reviewer to have assigned a single numerical value (for R_i) from a pre-defined range representing very negative, negative, neutral, positive, and very positive reviews on the product of interest. Clearly, any given product may have anywhere from zero through several reviews written by (presumably) different reviewers over time. We assume that each of these online reviews are independent of one another.

I.e., a reviewer is unaffected/uninfluenced by existing online reviews on the same product to abstain from intentionally/unintentionally modifying his/her (absolute) review. This is the base case where independence of each review is assumed.

The second scenario we consider is given in Figure 2. Here, we introduce sequential bias on reviews ($R_2 \cdots R_n$) that is a direct (unintentional) result of perusing reviews (previously) written by other reviewers. Of course, when there is only one existing online review for a product, the previous scenario would apply because of the sheer lack of opportunity for this review to be biased based on other (nonexistent) reviews.

In this scenario (Figure 2), the reviews written by reviewers $i = 1, 2$, and 3 are:

$$\begin{aligned} R_1, \\ R_2 + R_1(B_{12} \times s_{12}), \text{ and} \\ R_3 + R_2'(B_{23} \times s_{23}) + R_1(B_{13} \times s_{13}) \end{aligned}$$

respectively. The first reviewer is unbiased by previous reviews of this product by other reviewers. Therefore, the review provided by the first reviewer is an absolute (unbiased) review by this reader (i.e., reviewer 1 provides unbiased review R_1). The second reviewer, however, is biased by the first reviewer's review. The second reviewer's review is a composite of this reviewer's unbiased review (i.e., R_2) and the bias due to the first reviewer's review (i.e., R_1 times the bias represented by B_{12} times the direction of the bias). The magnitude of the bias (i.e., B_{12}) only provides the absolute value of the bias. When both the reviews are overall positive or both the reviews are overall negative, we assume the absence of any sequential bias since bias in this sense occurs as a compensation for friction between or negation of existing review(s) and the new review. Similarly, when the first review is neutral, we assume this review to exert no friction (and, therefore, no compensatory pressure) on the following reviews. However, when the first review is overall positive while the second review is overall negative, the second reviewer (unintentionally) compensates for existing positive review by being more negative resulting in negatively biased second review and *vice versa*. We assume additive relationship between each pair of new review and every existing review since the influence from each of the existing reviews on the new review are independent of one another. A squashing function could be used to keep the review value within a pre-specified range if necessary. We do not attempt this since it does not affect the results or resulting insights.



Figure 2: Previous reviews + addition of a new review on a product

The review by the third reviewer is similarly biased by all preceding reviews (here, reviews 1 & 2). The third reviewer's review is therefore a composite of this reviewer's unbiased review (i.e., R_3) and the bias due to the first reviewer's review (i.e., $R_1(B_{13} \times s_{13})$) and the second reviewer's

review (i.e., $R'_2(B_{23} \times s_{23})$). The general expression for the review written by the n^{th} reviewer is:

$$R_n + \sum_{i=1}^{n-1} R'_i \times B_{in} \times s_{in}$$

Assuming the bias is some form of discounting comprising a discount factor of the linear distance between the two reviews, we can represent B_{ij} as β^{j-i} and the general expression can be written as:

$$R_n + \sum_{i=1}^{n-1} R'_i \times \beta^{n-i} \times s_{in}$$

This can be justified by the generally accepted observation that (1) most perusers of existing reviews do not deal with the entire set of existing reviews but rather a manageable subset, and (2) the most recent review read is, *ceteris paribus*, likely to be more influential than the one that was read a few reviews earlier. Some form of discounting factor discounts away the reviews that are assumed to be unread and the same discounting factor assigns more weight to recent rather than earlier reviews among those that were perused.

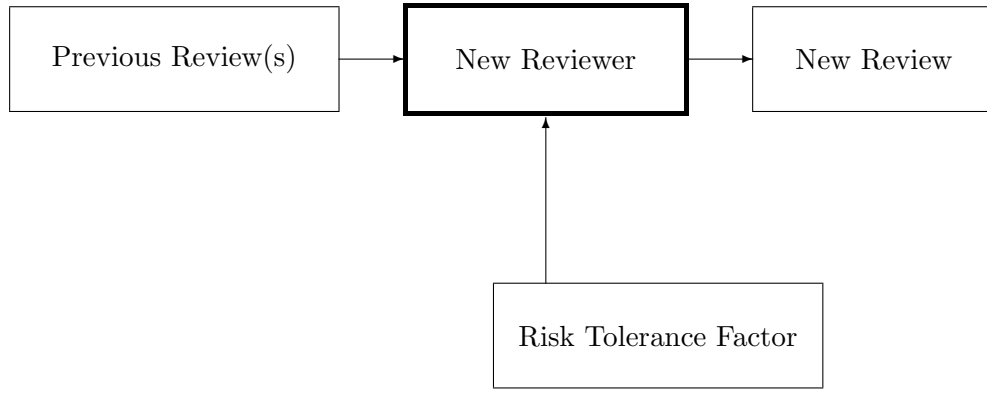


Figure 3: Previous reviews + risk tolerance + addition of a new review

The next scenario (Figure 3) incorporates the risk tolerance level or the risk attitude of the reviewer into the picture. By risk tolerance we signify or represent the extent to which this reviewer is risk-seeking, risk-neutral, or risk-averse. We model the risk tolerance level of the new reviewer using a standard power utility function:

$$U(x) = \frac{(x - a)^k}{(b - a)^k} \quad \text{where } 0 \leq U(x) \leq 1$$

Here, a and b are the minimum and maximum review scores respectively and k is a model parameter. Risk aversion is characterized by a concave utility function and risk seeking behavior is characterized by a convex function. The parameter k can be varied to obtain different levels of concavity and correspondingly different levels of risk aversion. To model risk seeking behavior using a convex utility function, we set $k > 1$. The general expression can now be written as follows:

$$R_n + \sum_{i=1}^{n-1} R'_i \times \beta^{n-i} \times s_{in} \times \frac{(R_n - a)^k}{(b - a)^k}$$

The above expression is modeled with the risk tolerance level as a multiplicative factor and not as the overall utility loss or gain in response to uncertainty. We decided on this form to enable discounting of the remaining expression by this fraction since, for example, β does not directly affect risk tolerance of any given review.

Several online stores (e.g., amazon.com) provide means to rate individual online reviews based on its usefulness to perusers of such reviews. These reviews are given a *score*, which is a ratio of the number of users who found a review useful vs. the number of users who perused that review. This is a crude measure of the belief (or, disbelief) placed by the new reviewer on any given (existing) review. In other words, previous perusers' evaluation of any given review's usefulness provides a proxy for the trust that one can place on that review. We include this usefulness score (i.e., u_i for review R_i) as a multiplicative factor in the new review where $0 \leq u_i \leq 1$. We assign $u_i = 0$ for the extreme case where *none* of the previous perusers of review R_i found this review to be useful and $u_i = 1$ for the other extreme case where *all* of the previous perusers of review R_i found it to be useful. The range in-between (i.e., $0 < u_i < 1$) represents the case where *some* (i.e., ≥ 1) of the previous perusers of review R_i found this review to be useful.

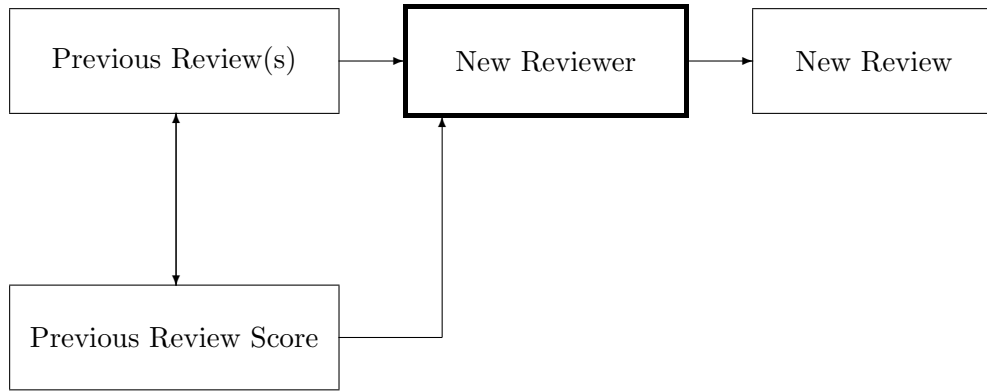


Figure 4: Previous reviews with score + addition of a new review

The next scenario considered incorporates this information (Figure 4). The new reviewer's review is now

$$R_n + \sum_{i=1}^{n-1} R'_i \times \beta^{n-i} \times s_{in} \times u_i$$

The last model (Figure 5) incorporates sequential bias based on previous reviews, previous reviews' usefulness measure scores, and the new reviewer's risk tolerance level in determining the characteristics of the new review. Here, the newest reviewer's review is given by:

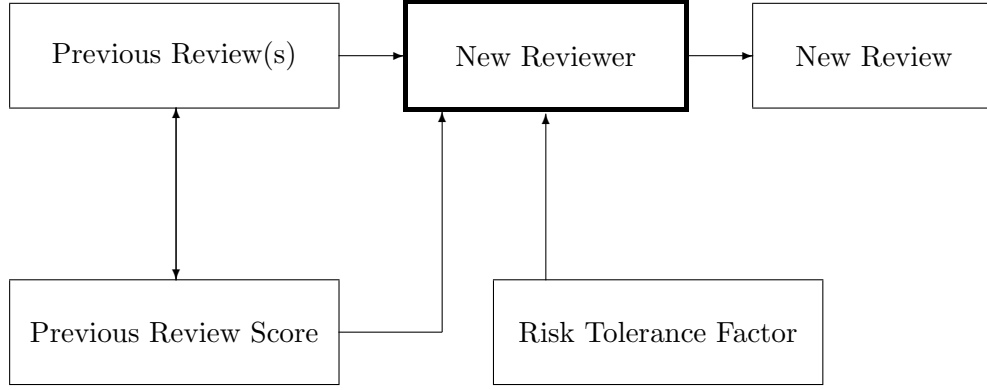


Figure 5: Previous reviews with score + risk tolerance + addition of a new review

$$R_n + \sum_{i=1}^{n-1} R'_i \times \beta^{n-i} \times s_{in} \times \frac{(R_n - a)^k}{(b - a)^k} \times u_i$$

Lemma 3.1 *Overall, in the presence of sequential bias, risk-averse reviewers are biased to write more positive reviews than risk-seeking reviewers.*

Proof Only (sequentially) biased reviews are of interest here. We have:

$$\sum_{i=1}^{n-1} R'_i \times \beta^{n-i} \times s_{in} \times \frac{(R_n - a)^k}{(b - a)^k} \times u_i \geq \sum_{i=1}^{n-1} R'_i \times \beta^{n-i} \times s_{in} \times \frac{(R_n - a)^{k'}}{(b - a)^{k'}} \times u_i$$

Here, risk-averse reviewers are represented on the left hand side and are associated with $0 < k < 1$ and risk-seeking reviewers (represented on the right hand side) are associated with $k' > 1$. Clearly, the expression on the left hand side is at least as much as that on the right hand side. \square

Lemma 3.1 refers to the cases represented in Figures 3 and 5 where the risk tolerance of the new reviewer are modeled.

Lemma 3.2 *When the reviewers are all positive (or all neutral, or all negative), the reviews written are free of sequential bias.*

Proof This is a direct consequence of the sign of the bias term being equal to zero (i.e., $s_{ij} = 0$) when the reviewers are all oriented in the same direction. I.e., when reviewer#1 is positive and reviewer#2 is also positive, reviewer#2 does not have a need to compensate for reviewer#1's review. The same applies when every review is neutral or every review is negative. In the expression we have, when $s_{ij} = 0$ the entire expression becomes zero. \square

Theorem 3.3 *The magnitude of new review provided by reviewer n (i.e., R_n) when sequential bias is present is always at least as much as that provided in the case when sequential bias is absent.*

I.e., in the last model (Figure 5) considered,

$$|\text{unbiased-}R_n| \leq |R_n + \sum_{i=1}^{n-1} R'_i \times \beta^{n-i} \times s_{in} \times \frac{(R_n - a)^k}{(b - a)^k} \times u_i|$$

Proof We represent the un(sequential)biased review by reviewer n (I.e., R_n) by ‘unbiased- R_n .’ The right hand side of this expression includes (sequential) biased R_i values.

$$|\text{unbiased-}R_n| \leq |R_n + \sum_{i=1}^{n-1} R'_i \times \beta^{n-i} \times s_{in} \times \frac{(R_n - a)^k}{(b - a)^k} \times u_i|$$

We show this proof by separately considering the three different possible cases depending on whether the new review with sequential bias (R_n) is positive, negative, or neutral.

Case 1: R_n is positive

Now, for any R'_i ($1 \leq i \leq n - 1$), when R'_i is positive, $s_{in} = 0$ (since $s_{ij} = 0$ when R_i and R_j are both positive) and the entire expression (i.e., $R'_i \times \beta^{n-i} \times s_{in} \times \frac{(R_n - a)^k}{(b - a)^k} \times u_i$) equals zero. When R'_i is negative, s_{in} is negative (i.e., -1) and the entire expression is positive. When R'_i is neutral (i.e., $R'_i = 0$), s_{in} is zero and the entire expression equals zero.

$$\therefore \sum_{i=1}^{n-1} R'_i \times \beta^{n-i} \times s_{in} \times \frac{(R_n - a)^k}{(b - a)^k} \times u_i \geq 0$$

Case 2: R_n is negative

For any R'_i ($1 \leq i \leq n - 1$), when R'_i is negative, $s_{in} = 0$ (since $s_{ij} = 0$ when R_i and R_j are both negative) and the entire expression equals zero. When R'_i is positive, s_{in} is negative (i.e., -1) and the entire expression is negative. When R'_i is neutral, s_{in} is negative (i.e., -1) and the entire expression equals zero.

$$\therefore \sum_{i=1}^{n-1} R'_i \times \beta^{n-i} \times s_{in} \times \frac{(R_n - a)^k}{(b - a)^k} \times u_i \leq 0$$

Case 3: R_n is neutral (i.e., $R_n = 0$)

For any R'_i ($1 \leq i \leq n - 1$), when R'_i is neutral, $s_{in} = 0$ (since $s_{ij} = 0$ when R_i and R_j are both neutral) and the entire expression equals zero. When R'_i is positive, s_{in} is negative (i.e., -1) and the entire expression is negative. When R'_i is negative, s_{in} is positive (i.e., +1) and the entire expression is negative.

$$\therefore \sum_{i=1}^{n-1} R'_i \times \beta^{n-i} \times s_{in} \times \frac{(R_n - a)^k}{(b - a)^k} \times u_i \leq 0 \quad \square$$

Although we specifically considered only the last model (Figure 5) in Theorem 3.3, this theorem applies to all the models considered in this paper. It should be noted that this theorem applies regardless of the degree of risk tolerance (i.e., risk-averse, risk-neutral, or risk-seeking perspective)

of the reviewers. Theorem 3.3 is significant from a recommender system's perspective since the written reviews (here, assigned scores in new reviews) must be appropriately augmented to remove sequential bias. On the other hand, the reviews can be considered to be a lower-bound for generating recommendations.

Lemma 3.4 *In the presence of sequential bias, a negative fake review has more positive effect when included earlier in the sequence.*

Proof We assume that the recommender system averages the review scores to obtain the recommendation.

Ignoring the effects of risk tolerance and scores on previous reviews to simplify notation, the reviews can be written as:

$$R_1, R_1\beta s_{1n} + R_2, R_1\beta^2 s_{1n} + R_2'\beta s_{2n} + R_3, \dots, R_1\beta^{n-1} s_{1n} + R_2'\beta^{n-2} s_{2n} + \dots + R_n$$

The difference in the average when the fake review (say, R) is placed at position p vs. q , where $p < q$, is

$$\begin{aligned} & \frac{1}{n} \times \sum_{j=n-q}^{n-p-1} R \times \beta^j \times s_{pn} \times \frac{(R-a)^k}{(b-a)^k} \times u_p \\ \text{i.e., } & \frac{1}{n} \times R \times \frac{\beta^{n-q} - \beta^{n-p}}{1 - \beta} \times s_{pn} \times \frac{(R-a)^k}{(b-a)^k} \times u_p \geq 0 \end{aligned}$$

Here, R , s_{pn} ($= s_{qn}$), and u_p ($= u_q$) are the same for this review regardless of its position in the sequence (of n reviews). The sign term (i.e., s_{pn}) is negative for negative R_i . Therefore, the earlier in the sequence of reviews the fake review is introduced, the more positive effect it has on the overall average. \square

Lemma 3.5 *In the presence of sequential bias, a positive fake review has more negative effect when included earlier in the sequence.*

Proof The proof is similar to Lemma 3.4. The difference in the average when the fake review (say, R) is placed at position p vs. q , where $p < q$, is the same as that in Lemma 3.4

$$\frac{1}{n} \times R \times \frac{\beta^{n-q} - \beta^{n-p}}{1 - \beta} \times s_{pn} \times \frac{(R-a)^k}{(b-a)^k} \times u_p \leq 0$$

Here, R , s_{pn} , and u_p are the same for this review regardless of its position in the sequence. The sign term (i.e., s_{pn}) is negative for positive R . Therefore, the earlier in the sequence of reviews the fake review is introduced, the more negative effect it has on the overall average. \square

Lemmas 3.4&3.5 signify that regardless of its type (i.e., positive or negative), a fake review has a stronger effect when placed earlier in the sequence. I.e., to enhance the recommendation of

an item, the fake review is better off earlier or later in the sequence and *vice versa* depending on whether it is negative or positive respectively. When using such data as input to a recommender system, the earlier (or later) reviews in the sequence must be scrutinized for such effects depending on the overall recommendation generated.

Theorem 3.6 *In the presence of sequential bias, assuming the magnitude of all previous reviews (i.e., R) are the same, the case where all positive reviews are stacked together and listed first results in a better (i.e., more positive and less negative) new review than otherwise when $\beta^{n-k-1}(2 - \beta^k) < 1$.*

Proof Ignoring the effects of risk tolerance and scores on previous reviews to simplify notation, the reviews can be written as:

$$R_1, R_1\beta s_{1n} + R_2, R_1\beta^2 s_{1n} + R_2'\beta s_{2n} + R_3, \dots, R_1\beta^{n-1} s_{1n} + R_2'\beta^{n-2} s_{2n} + \dots + R_n$$

Consider a scenario where there are k positive reviews followed by $(n - 1 - k)$ negative reviews. The n^{th} review can either be positive or negative.

Case 1: The n^{th} review is positive. The first k reviews are irrelevant (from a sequential bias perspective) since $s_{in} = 0$ (for $1 \leq i \leq k$) in this scenario. The n^{th} review is

$$R'_{k+1}\beta^{n-k-1} s_{(k+1)n} + R'_{k+2}\beta^{n-k-2} s_{(k+2)n} + \dots + R'_{n-1}\beta s_{(n-1)n} + R_n$$

The bias part is given by the above expression minus R_n . For the simplified case where all R values are the same, the resulting bias is:

$$R\beta \frac{1 - \beta^{n-k-1}}{1 - \beta} s \geq 0$$

Case 2: The n^{th} review is negative. The last $(n - 1 - k)$ reviews are irrelevant (from a sequential bias perspective) since $s_{in} = 0$ (for $k + 1 \leq i \leq n$) in this scenario. The n^{th} review is

$$R_1'\beta^{n-1} s_{1n} + R_2'\beta^{n-2} s_{2n} + \dots + R_k'\beta^{n-k} s_{kn} + R_n$$

Similarly, for the simplified case where all R and s values are the same, the bias is:

$$R \beta^{n-k} \frac{1 - \beta^k}{1 - \beta} s \leq 0$$

We are interested in the scenario where the positive bias is more than the negative bias. I.e., the absolute value of the bias in Case 1 is more than the absolute value of the bias in Case 2. The former case is better than the latter when the expression in Case 1 is greater than that in Case 2. I.e.,

$$\beta \frac{1 - \beta^{n-k-1}}{1 - \beta} > \beta^{n-k} \frac{1 - \beta^k}{1 - \beta}$$

which, upon simplification, results in the given expression. \square

The policy implication of Theorem 3.6 is that for an online retailer, with the facility for reviewers to write their reviews, to improve the overall recommendation of the product or service, it helps to list all the positive reviews first followed by all other reviews.

Corollary 3.7 *When sequential bias is present, the average bias in the new review when stacking all existing reviews of the same type together (all positives together or all negatives together) and listing them first (or last) is independent of the number of existing reviews of each type.*

Proof This is a direct consequence of Theorem 3.6. Consider the case where there are k positive reviews and $(n - 1 - k)$ negative reviews where $k > (n - 1 - k)$. When the k positive reviews are placed first followed by the $n - 1 - k$ negative reviews, the effect of this arrangement on the next review written from a sequential bias perspective depends on the type (i.e., positive or negative) of the new (i.e., n^{th}) review.

Case 1: The n^{th} review is positive. Here, the effect due to sequential bias is given (from Theorem 3.6) by

$$R\beta \frac{1 - \beta^{n-k-1}}{1 - \beta} s$$

Case 2: The n^{th} review is negative. Again, the effect due to sequential bias is given (from Theorem 3.6) by

$$R \beta^{n-k} \frac{1 - \beta^k}{1 - \beta} s$$

On average, the effect of sequential bias on the next review is

$$\frac{Rs}{2(1 - \beta)} (\beta - \beta^{n-k} + \beta^{n-k} - \beta^n) = \frac{Rs}{2(1 - \beta)} (\beta - \beta^n)$$

which is independent of k . □

4 Discussion & Conclusion

We considered the effect of sequential bias in online product reviews from the perspective of recommender systems. To our knowledge, none of the existing recommender systems consider the existence of sequential bias. We modeled the bias by taking into account existing reviews, the usefulness of existing reviews as determined by the number of perusers of those reviews who found them to be useful, and the risk tolerance of reviewers who write such reviews. We then studied its dynamic and formulated a model using theoretical means. Our analysis shows that influence due to sequential bias on subsequent reviews can be appreciable depending on the values of the modeled parameters. Developers or recommender systems therefore cannot ignore the existence of sequential bias in their input data. Given the popularity of recommender systems, there is an urgent need to develop means to remove influences due to sequential bias. This will alleviate its deleterious effects on the quality and accuracy of recommendations generated by recommender systems.

In our analysis, we only consider the overall review in the form of a numerical value. Considering evaluations of individual characteristics in addition to the overall review could provide additional insights and, therefore, better recommendations by a recommender system.

We assumed that a reader goes through the reviews one at a time from top to bottom of the list of existing reviews. We also assumed that the new review is written after the perusal of the entire set of existing reviews. However, when the number of reviews for an item of interest is large it is not reasonable to expect anyone to peruse the entire set (e.g., [15]). Modeling this would involve knowing (1) the exact set of reviews perused by the new reviewer, and (2) the order in which this selected set of reviews are read by this new reviewer. Some online stores (e.g., amazon.com) allow the user to sort existing reviews in several different ways (e.g., most recent review first, most useful reviews first, most negative reviews first) and this complicates the analysis even further.

While we considered the effect of sequential bias on recommender systems from a theoretical perspective, a few critical questions still remain and these can be addressed only through rigorous and strictly controlled field experiments with human subjects. Does sequential bias *significantly* influence recommendations made by recommender system by modifying the average and/or correlation values? If yes, does it matter? I.e., does it lead to significantly different recommendations? Under what conditions does this happen? And, what can the developer of a recommender system do about it? Given that sequential bias does exist in online product reviews, can we remove existing bias in reviews in such real-world databases so recommendations generated by recommender systems are more reliable and trust-worthy? We also did not explicitly model the scenario where the reviewers are mixed - i.e., some reviewers are influenced by sequential bias whereas some others are not simply because they do/did not read reviews by other reviewers. What influence do such mixed set of reviews have on the performance of recommender systems? This can be easily modeled by modifying R_i to R_i^c where $c = 1$ when review i is perused by the current reviewer and $c = 0$ otherwise. We leave these as exercises for future studies.

References

- [1] Gediminas Adomavicius and Alexander Tuzhilin. 2005. "Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions," *IEEE Transactions on Knowledge and Data Engineering*, 17(6), pp. 734-749, June.
- [2] Rohini Ahluwalia. 1996. "An Integrative Model of Market-Related Negative Information Processing," PhD Dissertation, Department of Marketing, Ohio State University.
- [3] Linden J. Ball. 2002 "Does Positivity Bias Explain Patterns of Performance on Wason's 2-4-6 task?" in Wayne D. Gray and Christian D. Schunn, *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society*, Routledge, p. 340.
- [4] Yi-Fen Chen. 2008. "Herd behavior in purchasing books online," *Computers in Human Behavior*, 24(5), pp. 1977-1992.

- [5] Guillaume Deffuant and Sylvie Huet. 2009. "Collective Increase of First Impression Bias," *Complexity*, 15(5) pp. 25-33.
- [6] Chrysanthos Dellarocas. 2005. "Reputation mechanism design in online trading environments with pure moral hazard." *Information Systems Research* 16(2) pp. 209-230.
- [7] Chrysanthos Dellarocas and Charles Wood. 2008. "The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias." *Management Science* 54(3) pp. 460-476.
- [8] Wenjing Duan, Bin Gu, and Andrew B. Whinston. 2008. "Do online reviews matter? - An empirical investigation of panel data," *Decision Support Systems*, 45 (4), pp. 1007-1016.
- [9] Ming Fan, Yong Tan, Andrew Whinston. 2005. "Evaluation and design of online cooperative feedback mechanisms for reputation management." *IEEE Transactions on Knowledge and Data Engineering* 17(2) pp. 244-254.
- [10] The Guardian. 2010 "Historian Orlando Figes admits posting Amazon reviews that trashed rivals," 23 April. <http://www.guardian.co.uk/books/2010/apr/23/historian-orlando-figes-amazon-reviews-rivals> (retrieved on 16 June 2010)
- [11] Jill G. Klein. 1991. "Negativity Effects in Impression Formation: A Test in the Political Arena," *Personality and Social Psychology Bulletin*, 17 (4), pp. 412-18.
- [12] Hady W. Lauw, Ee-Peng Lim and Ke Wang. 2006. "Bias and Controversy in Evaluation Systems" *IEEE Transactions on Knowledge and Data Engineering*, 20(11), pp. 1490-1504, November 2008.
- [13] Xinxin Li and Lorin Hitt. 2008. "Self-Selection and Information Role of Online Product Reviews," *Information Systems Research*, 19(4), pp. 456-474, December.
- [14] Richard W. Mizerski. 1982. "An Attribution Explanation of the Disproportionate Influence of Unfavorable Information," *Journal of Consumer Research*, 9 (3), pp. 301-310.
- [15] Selwyn Piramuthu. 2010. "Sequential Bias, Online Product Review, and Recommender Systems," *Proceedings of the Workshop on Data Mining in Marketing (DMM'2010) at the Industrial Conference on Data Mining (ICDM)*, pp. 5-13, IBAI Publishing.
- [16] Matthew Rabin and Joel L. Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics*, pp. 37-82.
- [17] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. 1994. "GroupLens: An Open Architecture for Collaborative Filtering of Netnews," *Proceedings of the Conference on Computer Supported Cooperative Work*, Chapel Hill, NC, pp. 175-186.

- [18] Paul Resnick, Richard Zeckhauser, John Swanson, Kate Lockwood. 2006. "The value of reputation on eBay: A controlled experiment." *Experimental Economics* 9(2) pp. 79-101.
- [19] G. Ruskin. 2005. "Commercial Alert asks FTC to Investigate Buzz Marketers for Deception," 18 October (2005) <http://www.commercialalert.org/news/news-releases/2005/10/commercial-alert-asks-ftc-to-investigate-buzz-marketers-for-deception> (retrieved on 16 June 2010)
- [20] G. Sandoval. 2006. "Digg continues to battle phony stories," *CNET.com*, December 18 (2006)
- [21] Sylvain Senecal and Jacques Nantel. 2004. "The Influence of Online Product Recommendations on Consumers' Online Choices," *Journal of Retailing*, 80, pp. 159-169.
- [22] Jessica Staddon, Richard Chow. 2008. "Detecting reviewer bias through web-based association mining," *Proceeding of the 2nd ACM workshop on Information credibility on the web*, pp. 5-10.
- [23] Robert Verkaik. 2010. "The law is catching up with those who use the internet to defame," *The Independent*, 20 April (2010)