

The Bandwagon Effect: Not Just Another Bias

Norman Knyazev
Radboud University
Nijmegen, The Netherlands
norman.knyazev@ru.nl

Harrie Oosterhuis
Radboud University
Nijmegen, The Netherlands
harrie.oosterhuis@ru.nl

ABSTRACT

Optimizing recommender systems based on user interaction data is mainly seen as a problem of dealing with selection bias, where most existing work assumes that interactions from different users are independent. However, it has been shown that in reality user feedback is often influenced by earlier interactions of other users, e.g. via average ratings, number of views or sales per item, etc. This phenomenon is known as the *bandwagon effect*.

In contrast with previous literature, we argue that the bandwagon effect should not be seen as a problem of statistical bias. In fact, we prove that this effect leaves both individual interactions and their sample mean unbiased. Nevertheless, we show that it can make estimators inconsistent, introducing a distinct set of problems for convergence in relevance estimation. Our theoretical analysis investigates the conditions under which the bandwagon effect poses a consistency problem and explores several approaches for mitigating these issues. This work aims to show that the bandwagon effect poses an underinvestigated open problem that is fundamentally distinct from the well-studied selection bias in recommendation.

CCS CONCEPTS

• Information systems → Social recommendation; Learning to rank; • Mathematics of computing → Time series analysis.

KEYWORDS

Bandwagon Effect, Consistency, Statistical Bias, Recommendation

ACM Reference Format:

Norman Knyazev and Harrie Oosterhuis. 2022. The Bandwagon Effect: Not Just Another Bias. In *Proceedings of the 2022 ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR '22)*, July 11–12, 2022, Madrid, Spain. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3539813.3545142>

1 INTRODUCTION

One of the main challenges in optimizing Recommender Systems (RecSys) is correctly capturing user preferences from historical interaction data [17, 20, 21, 38, 50]. A widely adopted approach is to train a model to predict known ratings and then apply it to previously unobserved user-item pairs [35, 48, 54]. However, in the last decade it has become apparent that observed interactions are often a very skewed representation of the actual user preferences [4, 35, 47]. In particular, ratings are known to be affected by various forms of

selection bias, such as *popularity bias* [47], *positivity bias* [35] or *position bias* [34]. Popularity bias, for instance, occurs because as most interactions involve only a small number of very popular items, the patterns in observed ratings may not necessarily extrapolate to the data as a whole [4, 35, 45, 47]. It is now well understood that ignoring statistical biases during optimization is likely to lead to suboptimal recommendation performance [4, 18, 45] and numerous solutions have been proposed to address various forms of selection bias in recommendation data [18, 22, 39, 40, 40, 43, 45, 57, 63].

In this work, we will investigate the related but fundamentally distinct *bandwagon effect*, also referred to as *herding* or *conformity bias*. Despite being observed by multiple studies, this often-overlooked effect occurs because users tend to follow the behavior and opinions of others [12, 46]. That is, if a user is told that many other users have liked an item, they themselves become more likely to give a positive rating. Conversely, when a user is told that many other users dislike an item, their own rating is likely to become lower. Several studies have observed this behavior when exposing the average rating from previous users [1, 7, 12, 51] and analogous trends have also been reported for implicit feedback [15, 36]. Consequently, the bandwagon effect can pose a serious problem when trying to infer user preferences from interaction data.

Despite its well-known existence, there have only been limited attempts at reducing the influence of the bandwagon effect [30, 56, 65] and few authors have theoretically analyzed it in statistical terms [60, 61]. Our work will build on a mathematical model proposed by Xie *et al.* [60, 61], who in line with earlier literature [56, 65] interpret the bandwagon effect as a problem of statistical bias.

In this work, we provide a rigorous theoretical analysis of the bandwagon effect and its statistical impact on rating and Click-Through Rate (CTR) estimation. In stark contrast with previous work, we prove that a simple form of bandwagon effect of Xie *et al.* does *not* introduce statistical bias, as both individual ratings and their sample mean yield unbiased estimates under the effect. Nonetheless, our analysis also reveals that the bandwagon effect can create serious problems for convergence by making the sample mean inconsistent. Therefore, despite not being a bias issue, the bandwagon effect can still introduce significant errors during optimization. To understand when this may pose a problem, we theoretically prove a condition under which inconsistency is guaranteed to arise for the sample mean. Furthermore, our theoretical and simulated results show that even when an estimator is consistent, the convergence can still be so slow that finding the true user preferences becomes practically infeasible. We therefore investigate whether existing debiasing methods can be adapted to handle these convergence issues. However, as none of our methods are able to combine theoretical guarantees with robust empirical performance, we conclude that the bandwagon effect currently remains an open problem that needs more attention in future work.



This work is licensed under a Creative Commons Attribution International 4.0 License.

ICTIR '22, July 11–12, 2022, Madrid, Spain

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9412-3/22/07.

<https://doi.org/10.1145/3539813.3545142>




Figure 1: Impact of selection bias and bandwagon effect on average rating estimation. Under popularity bias, users are less likely to interact with unpopular items, leaving most of their ratings unobserved. Under the bandwagon effect, new ratings are skewed in the direction of previous users’ opinion, with larger group size eliciting a stronger effect.

2 RELATED WORK

2.1 Debiasing Interactions for Recommendation

Learning from user interactions has a long tradition in the field of recommender systems [16, 25, 32, 35, 45]. A common approach is to cast the recommendation problem as a rating prediction task where the goal is to predict the ratings a user would give to items [16, 25, 32]. However, one should not straightforwardly base the optimization on the ratings observed in interaction logs as they are often not representative of actual user preferences due to selection bias [16, 32]. The most well-known and well-studied form of bias in rating data is popularity bias: a small selection of popular items receives the large majority of user ratings and is thus grossly overrepresented [4, 35, 47]. Positivity bias is also prevalent: users may be more likely to leave a rating on items they like; consequently, low ratings become underrepresented in the resulting interaction data [35]. As seen in Figure 1 (top), if popular items that attract the majority of ratings also happen to be more liked by users (i.e. receive higher ratings), the average rating estimate from observed ratings becomes skewed towards a higher value. Schnabel et al. [45] showed that not accounting for selection bias makes the optimization procedure focus on statistically overrepresented user-item combinations. As a solution, they propose to use Inverse Propensity Scoring (IPS) to unbiasedly estimate and optimize the performance of a system [29]. Debiasing for recommendation has become a very popular topic in recent years: for example, Wang et al. [57] propose a doubly robust estimator that combines IPS with regression estimates; Saito et al. [40] adapted the IPS matrix factorization approach to debias click data instead of ratings; and Jeunen and Goethals [22] incorporate a prior distribution to avoid overestimation during prediction.

Selection bias is also widespread in clicks. Its most well-studied form is position bias, which occurs when the position at which an item is displayed heavily impacts the attention the item will receive, and thus how many users will click on it [8]. Similar to ratings, IPS is a popular method to correct for selection bias in clicks [23, 33].

The problem of selection bias in explicit feedback, i.e. ratings, and implicit feedback, e.g. clicks, is thus well recognized and has been addressed extensively by the existing literature [16, 22, 23, 25, 32, 33, 35, 40, 45, 57]. This can lead to the impression that learning from interactions is mainly a selection bias problem. Nevertheless,

our current work on the bandwagon effect will show that this is definitely not the case.

2.2 The Bandwagon Effect

The *bandwagon effect*, or *herding*, describes a phenomenon whereby individuals tend to exhibit a greater affinity towards an item when they believe other individuals also like it [37, 46]. The decisions of earlier individuals can propagate through *information cascades* to individuals also taking the same decision later, leading to snowballing effects, as well documented in social psychology, economics and finance, and, more recently, web information systems [3, 5, 7, 37].

Online platforms frequently include various social cues next to items to assist users with their choice, e.g. the number and the average value of previous user ratings. Numerous studies have shown that users actively leverage such information when making purchasing or consumption decisions [1, 10, 12, 13, 15, 19, 24, 36, 41, 42, 51]. Hanson and Putler [15] found that artificially inflating download counts for some of the products on a software download service led to those products receiving more subsequent downloads than their unaltered counterparts. Similarly, Qin et al. [36] report that Chrome Web Store users rarely click on items displayed with a low download count. Flanagin and Metzger [12] noted that users exposed to the average rating also rated movies closer to that average. Analogously, Sundar et al. [51] found that participants reported higher purchasing intent for high-rated items over low-rated items and that this effect arose from an improved perception of quality and value. Both studies also reported that the strength of the effect was positively correlated with the item’s number of ratings, consistent with earlier works associating larger group size with stronger herding [2, 19, 27]. Altogether, these observations suggest that users incorporate the experience of others into their own judgements and the strength of this effect is correlated with the size of the group.

This phenomenon can have a number of negative consequences. On an individual level, the bandwagon effect can be so strong that users may act against their own interests: Dholakia and Soltysinski [10] highlight that eBay users frequently bid on items that already have many bids, whilst ignoring comparable or superior listings with fewer bids. However, another consequence of the bandwagon effect – and the focus of this paper – is the difficulty it brings to estimating item relevance from user interactions. In the pioneering experiment of Salganik et al. [41], a large number of sequentially

arriving participants browsed a randomly shuffled grid of the same 48 songs and could interact with songs of their choice. The participants were randomly assigned to one of nine groups, with users in groups 2 – 9 additionally presented with the download count from previous users *from their group* alongside each song. The authors observed that songs had a higher probability of being downloaded in a group where they already had a high number of downloads. In contrast, a song shown next to a low download count was less likely to get downloaded compared to the same song in groups where it had a high or no presented download count (group 1). Moreover, with the exception of the very best and worst songs, as determined by the control group 1, the final ranking according to download counts varied wildly across different groups. In their follow-up work, Salganik and Watts [42] further showed that the bandwagon effect may be so strong that artificially altering the presented download count can make unpopular songs popular and vice-versa, with only the best songs eventually recovering their true relevance ranking. Taken together, the above results suggest that under the bandwagon effect even a large amount of user feedback is not guaranteed to yield an accurate estimate of true user preference. This is also demonstrated in Figure 1 (bottom): users' true preferences are intertwined with the ratings of earlier users and this influence grows as the number of ratings increases. As a consequence, even when assuming no selection bias, true average rating cannot be determined via the sample mean.

Recently, there has been some interest in modelling herding and disentangling the effects of the crowd on the individual. Wang et al. [56] propose an estimator maximizing a surrogate function of the likelihood of the data and demonstrate the effectiveness of their approach for prediction of subsequent ratings. Zhang et al. [65] show that subsequent rating estimation can be further improved by considering that some users may actually choose to distance their opinion from the crowd. Liu et al. [30], Wei et al. [59], Zhang et al. [66] all propose various extensions to classical recommender models by using the rating mean and/or count as additional inputs. Finally, Xie, Zhong and colleagues formulate the bandwagon process as either a rating reweighting or hidden state prediction problem and show theoretically that the bandwagon effect slows down the convergence of the sample mean [60, 61, 67].

Notably, most of the above works treat the bandwagon effect as a problem of statistical bias. Furthermore, to the best of our knowledge, only the works of Xie et al. [60, 61] and Zhong et al. [67] attempt to establish a link between the bandwagon effect and statistical theory. Our theoretical analysis will contrast with the existing literature by showing that the bandwagon effect can affect relevance estimation without introducing statistical bias.

3 BACKGROUND

3.1 Preliminaries

Before we address the bandwagon effect specifically, this section will introduce our problem setting and the statistical properties that our analysis will consider. For the sake of generality, we focus on *unpersonalized* rating estimation, as we expect the issues observed for this simple task to also apply to the more involved forms of recommendation. Furthermore, our findings are not specific to




Figure 2: Three estimators with varying bias and consistency; bold lines show the mean estimate over 1,000 independent runs, shaded areas indicate 90% confidence intervals. Top - unbiased and consistent: its expected value and point of convergence are both correct. Middle - biased and inconsistent: its expected value and point of convergence are both incorrect. Bottom - unbiased and inconsistent: its expected value is correct but its confidence interval will never converge.

recommendation and can be readily applied to other problems involving relevance estimation such as CTR estimation.

In our generic unpersonalized recommendation setting, the goal is to estimate the average preference of all users for each item. We assume that this value $p \in (0, 1)$ exists for each item and represents the percentage of users that truly like the item. Furthermore, users interact with the items by assigning them binary ratings s.t. $r_i \in \{0, 1\}$ is the i 'th rating for an item. While we refer to r_i as a rating, it can equally well represent many other binary interactions such as *clicks*, *purchases*, *shares*, etc. As discussed in Section 2.2, in real systems users are often able to see how many other users have rated an item and what their average rating was, such that past ratings can affect the behaviour of the current user. We define p to be the *true preference* or the percentage of users who like an item when they have no knowledge of the preferences of any other users. Our goal is then to estimate p for a *single* item from the given set of observed ratings $\{r_1, r_2, r_3 \dots\}$.

Let \hat{p}_n indicate an estimate of p based on n observed ratings. Our analysis will investigate three important beneficial properties that we generally desire from such an estimate. To start, we normally want the estimate to be unbiased. In other words, in expectation it should be equal to the true value of p :

$$\text{Unbiased}(\hat{p}_n) \iff \mathbb{E}[\hat{p}_n] = p. \quad (1)$$

Understandably, bias and unbiasedness have received a lot of attention in previous selection bias literature, as it is desirable to avoid systematic errors in rating estimation [23, 34, 44, 45, 58]. However, unbiasedness on its own does not guarantee that an estimator is actually usable in practice.

The second important and often overlooked property that we may want from an estimator is consistency. An estimator is considered consistent if the estimate converges in probability to the true value of p as the size of available data approaches infinity:

$$\text{Consistent}(\hat{p}_n) \iff \lim_{n \rightarrow \infty} \hat{p}_n = p. \quad (2)$$

Whilst at a first glance unbiasedness and consistency may appear to refer to the same property, this is certainly *not* the case. For example, estimators can be biased yet consistent [53] or unbiased and inconsistent. As illustrated in Figure 2, the problem with an unbiased inconsistent estimator is that its confidence intervals do not converge to a zero error even as the size of the data increases indefinitely. This means its estimation errors will not disappear even after gathering more data, thus preventing accurate convergence.

Nevertheless, whilst consistency guarantees that eventually the estimate will be correct, the amount of data required for convergence may still be unattainable in practice. Therefore, the final property we will consider is (data) efficiency, measured by the expected squared error (equivalent to variance for unbiased estimators):

$$\text{Efficiency}(\hat{p}_n) = \mathbb{E}[(\hat{p}_n - p)^2]. \quad (3)$$

Overall, efficiency represents a more practical side of an estimator: the error we can expect it to have w.r.t. the true value p after n samples. The effect the number of samples has on this measure can be seen as the speed of convergence. While most existing work on learning from user interactions focuses on correcting for bias, they often also apply propensity clipping or self-normalization which introduce a little bias to greatly reduce variance [23, 40, 44, 52, 62].

Our theoretical investigation will consider how the bandwagon effect impacts all three properties: bias, consistency and efficiency.

3.2 A Formal Model of the Bandwagon Effect

The mathematical model of the bandwagon effect we will use in this paper is a simplified version of that of Xie et al. [60, 61]. This model represents a sequential unpersonalized recommendation scenario where users rate the item in succession and one at a time. It assumes there are two factors that influence the probability of a positive value for the n 'th rating ($r_n = 1$): (i) p , the true proportion of users who would like the item given that they do not know anything of the other users' preferences (no bandwagon effect); and (ii) \bar{p}_{n-1} , the mean of the ratings that have taken place previously. The rating probability is modelled as a linear combination of these two factors:

$$P(r_n = 1) = \lambda_n p + (1 - \lambda_n) \bar{p}_{n-1}, \quad \bar{p}_{n-1} = \frac{1}{n-1} \sum_{i=1}^{n-1} r_i. \quad (4)$$

$\lambda_n \in [0, 1]$ captures the degree to which users adjust their behavior to match that of the crowd. For example, $\forall \lambda_n = 1$ indicates a process with no bandwagon effect and $\forall n > 1 : \lambda_n = 0$ a process where true preferences play no role except for the very first user. As previously mentioned, earlier works have shown a positive correlation with the strength of herding and the size of the group [2, 19, 27]. We capture this increase in effect by assuming that the influence of the sample mean cannot decrease as the number of ratings increases:

$$\lambda_1 = 1 \wedge (\forall n > 1 : 0 \leq \lambda_n \leq \lambda_{n-1}). \quad (5)$$

In addition, this assumption states that the first rating is completely unaffected by the bandwagon effect (i.e. $\lambda_1 = 1$), as the first user has

no information about the preferences of others. We have chosen the above formulation over others (e.g. [56, 65]) due to its amenability to statistical analysis whilst still capturing the key aspects of the bandwagon effect.

Before we begin our theoretical analysis, we will briefly discuss two claims from earlier work by Xie and Zhong [60] so that they can later be compared with our findings in Section 4. First, Xie and Zhong refer to the bandwagon effect as *persuasion bias*, by which they mean that the bandwagon effect biases the rating distribution such that it no longer represents the true underlying preference. Second, they show that as long as the lower bound of λ_n is above 0, the bandwagon effect will eventually be eliminated since the sample mean \bar{p}_n is guaranteed to converge to the true parameter value:

$$\inf_{n \in \mathbb{N}^+} \lambda_n > 0 \implies \lim_{n \rightarrow \infty} \bar{p}_n = p. \quad (6)$$

Our theoretical analysis will take another critical look at the bias and consistency of \bar{p}_n . Surprisingly, our findings regarding bias heavily contrast with previous work on the bandwagon effect.

4 THEORETICAL ANALYSIS: THE BANDWAGON EFFECT

Now that Section 3.2 has introduced our mathematical model of the bandwagon effect, we will begin our theoretical analysis of its effects on bias, consistency and efficiency. First, we want to again consider the claim of previous work that the bandwagon effect introduces bias to the rating distribution [60, 61]. From our bandwagon model (Equation 4) it appears that the probability of a positive rating is indeed a distortion of the true preference p . However, Theorem A.1 proves that the expected values of all sample means \bar{p}_n and all ratings r_n are actually equal to the true preference:

$$\forall n : \mathbb{E}[\bar{p}_n] = \mathbb{E}[r_n] = p. \quad (7)$$

In other words, under our model, both the sample means and actual ratings provide *unbiased* estimates of the true preference, corresponding to either top or bottom, but not center of Figure 2. This finding contrasts heavily with how previous work has approached the bandwagon effect as a bias problem [30, 56, 59–61, 64, 65, 67] and similarly, how learning from user interactions is generally also framed as a bias problem [16, 22, 23, 25, 32, 33, 35, 40, 45, 57]. The bandwagon effect *may* be seen as an increased impact of earlier interactions on the estimate of the item's relevance. However, in contrast with selection bias, there is no predetermined rating value that is disproportionately represented for the purposes of average relevance estimation or loss minimization.

Nevertheless, the lack of bias does not mean that the bandwagon effect does not introduce other problems. As discussed in Section 3.1, there are other statistical properties besides unbiasedness that one should consider. Theorem A.4 proves the following efficiency of the sample mean, i.e. its expected (squared) error:

$$\mathbb{E}[(\bar{p}_n - p)^2] = p(1-p) \left(\frac{1}{n^2} + \sum_{i=1}^{n-1} \frac{1}{i^2} \prod_{j=i+1}^n \frac{(j-1)(j+1-2\lambda_j)}{j^2} \right).$$

Consistent with our prior discussion, the above formulation reveals that when $\forall \lambda_n = 1$, the variance of \bar{p}_n is equivalent to the variance of sampling from the true preference p . Conversely, when $\forall n > 1 : \lambda_n = 0$, the expected error does not decrease after the first

interaction i.e. $\forall n : \mathbb{E}[(\bar{p}_n - p)^2] = \mathbb{V}[\bar{p}_1] = p(1-p)$. Obviously, an error that never decreases provides a serious convergence problem.

To investigate whether the sample mean is consistent, we look at the value of the expected error in the limit of infinite samples. Theorem A.5 proves the following expected error in the limit:

$$\lim_{n \rightarrow \infty} \mathbb{E}[(\bar{p}_n - p)^2] = \lim_{n \rightarrow \infty} p(1-p) \left(\sum_{i=1}^{n-1} \frac{1}{i(i+1)} \prod_{j=i+1}^{n-1} \left(1 - \frac{2\lambda_j}{j+1}\right) \right).$$

Importantly, if the expected error is not zero in the limit as in Figure 2 (bottom), the estimator is not consistent and vice-versa. Based on this fact, we prove a novel necessary condition for the consistency of \bar{p}_n in Theorem A.6:

$$\forall \epsilon : \lim_{n \rightarrow \infty} \Pr(|\bar{p}_n - p| > \epsilon) = 0 \longrightarrow \lim_{n \rightarrow \infty} \sum_{i=1}^n \lambda_i^2 = \infty. \quad (8)$$

In other words, if the sample mean is consistent (corresponds to Figure 2, top), the sum of λ_i^2 terms is unbounded. Equivalently, convergence of $\lim_{n \rightarrow \infty} \sum_{i=1}^n \lambda_i^2$ leads to inconsistency. We note that whilst the sufficient condition of Xie and Zhong [60] (Equation 6) highlights the situations where we can guarantee consistency, it cannot be used to guarantee inconsistency. For instance, one cannot use it to show whether $\lambda_n = n^{-1}$ leads to issues with convergence as $\inf_{n \in \mathbb{N}^+} \lambda_n = 0$. On the other hand, as $\lim_{n \rightarrow \infty} \sum_{i=1}^n i^{-2}$ is finite, Theorem A.6 shows that \bar{p}_n is inconsistent. Our condition can thus be seen as complementary to the sufficient condition of earlier work, qualifying the set of λ sequences where inconsistency is certain.

We can also leverage Theorem A.5 to quantify the asymptotic expected absolute error for any specific λ values. For instance, Proposition A.8 proves that for $\lambda_i = c^{i-1}$ with $c \in (0, 1)$:

$$\lim_{n \rightarrow \infty} \mathbb{E}[|\bar{p}_n - p|] \geq 2p(1-p) \exp(-c\phi(c, 1, 2) - c^2\phi(c^2, 2, 2))$$

$$\text{s.t. } \phi(z, s, a) \equiv \sum_{k=0}^{\infty} \frac{z^k}{(a+k)^s}. \quad (9)$$

Similar bounds may also be derived for other λ formulations.

Our analysis proved that neither the sample mean \bar{p}_n nor the individual ratings r_n are biased, a finding that appears to heavily contrast with earlier work that casts the bandwagon effect as a problem of bias [30, 56, 59–61, 64, 65, 67]. We speculate that the bias framing may at least in part stem from the dependency between ratings. More specifically, conditioning the expectation of \bar{p}_n and r_n on an earlier \bar{p}_m by Theorem A.1 and Proposition A.7 leads to:

$$\forall n > m > 1 : \mathbb{E}[\bar{p}_{n-1} | \bar{p}_m] - p = (\bar{p}_m - p) \prod_{i=m+1}^{n-1} \left(1 - \frac{\lambda_i}{i}\right), \quad (10)$$

$$\forall n > m > 1 : \mathbb{E}[r_n | \bar{p}_m] - p = (1 - \lambda_n) \mathbb{E}[\bar{p}_{n-1} - p | \bar{p}_m]. \quad (11)$$

This can be seen as a *conditional* bias, i.e. conditioned on any \bar{p}_m being incorrect, all subsequent ratings and sample means are biased:

$$\forall n > m \geq 1 : \mathbb{E}[\bar{p}_n | \bar{p}_m \neq p] \neq p. \quad (12)$$

We speculate that previous work has interpreted this conditional bias as bias in the rating distribution and sample mean. Nonetheless, as discussed above, our bandwagon effect model does not lead to systematic errors in relevance estimation. In Section 7, we will argue that therefore the bandwagon effect is best interpreted as a problem of consistency and not statistical bias.

To summarize the main findings of our theoretical analysis: we have proven that the bandwagon effect does not introduce bias to the rating distribution but can make the sample mean inconsistent. We have also derived a necessary condition for consistency, expanding on prior work. Lastly, our formulation of the expected error can quantify the asymptotic error for any specific λ values.

5 BANDWAGON EFFECT MITIGATION

5.1 Correcting for the Bandwagon Effect

We have established that while the bandwagon effect does not make ratings biased, it can lead to convergence problems. As possible solutions, inspired by analogous approaches in click debiasing [6, 55], we will propose an affine and a maximum likelihood estimator.

Firstly, similar to earlier unbiased Learning to Rank work [55], we notice that the model in Equation 4 is invertible. We therefore propose an intermediate affine estimate based on a single rating:

$$\hat{r}_i = \frac{r_i - (1 - \lambda_i)\bar{p}_{i-1}}{\lambda_i}. \quad (13)$$

Theorems A.9 and A.10 prove that our intermediate estimator is both unbiased and conditionally unbiased w.r.t. previous ratings:

$$\forall i > j \geq 1 : \mathbb{E}[\hat{r}_i | \hat{r}_j] = \mathbb{E}[\hat{r}_i | \bar{p}_j] = \mathbb{E}[\hat{r}_i] = p. \quad (14)$$

Therefore, using \hat{r}_i instead of r_i in an averaging estimator may allow us to speed up convergence by incorporating only the information about the user's true preference without reinforcing the previous $i-1$ ratings. With this in mind, we propose our affine estimator that uses all of the observed ratings via a weighted mean:

$$\hat{p}_n = \frac{1}{\sum_{i=1}^n \omega_i} \sum_{i=1}^n \omega_i \hat{r}_i, \quad \text{s.t. } \forall 1 \leq i \leq n : \omega_i > 0. \quad (15)$$

Clearly \hat{p}_n is also unbiased as it is simply a weighted average over the unbiased estimates \hat{r}_i . Additionally, Theorem A.14 proves it is consistent if $\sum_{i=1}^n (\omega_i / \lambda_i)^2$ grows slower in n than $(\sum_{i=1}^n \omega_i)^2$, with Corollary A.15 providing specific conditions $\forall \omega_i = 1$ and $\forall \omega_i = \lambda_i$. Furthermore, Theorem A.13 proves \hat{p}_n has the following variance:

$$\mathbb{V}[\hat{p}_n] = \frac{1}{(\sum_{i=1}^n \omega_i)^2} \sum_{i=1}^n \frac{\omega_i^2}{\lambda_i^2} (p(1-p) - (1-\lambda_i)^2 \mathbb{V}[\bar{p}_{i-1}]). \quad (16)$$

Unfortunately, interpreting $\mathbb{V}[\hat{p}_n]$ is not straightforward as in addition to the choice of ω_i it depends on both the rating values as well as their order. Analogous to its use in selection bias [23, 49], one can also apply clipping to the bandwagon effect, e.g. $\hat{\lambda}_i = \max(\lambda_i, \tau)$ s.t. $\tau > 0$. However, in contrast with IPS, \hat{r}_n and \hat{p}_n would remain unbiased, though not conditionally unbiased (Corollary A.11): clipping or misestimation s.t. $\hat{\lambda}_i \neq \lambda_i$ may still affect the rate of convergence.

Secondly, inspired by the approach of click models [6], we also propose the approximate maximum likelihood estimator p_n^* , obtained via iterative application of Newton's method [14] to maximize the log-likelihood of the data L :

$$L = \sum_{i=1}^n r_i \ln(\lambda_i p_n^* + (1-\lambda_i)\bar{p}_{i-1}) + (1-r_i) \ln(1 - (\lambda_i p_n^* + (1-\lambda_i)\bar{p}_{i-1})). \quad (17)$$

Because both the likelihood and the solution found by Newton's method appear intractable, we are unable to prove this estimator is unbiased. Nonetheless, maximum likelihood estimators have been

very effective when applied to user interactions in the past [6, 68] and can be consistent under certain feasible conditions [9].

5.2 Bandwagon Effect Estimation

Both our proposed estimators require accurate λ values to capture the bandwagon effect well. Inspired by the experiments of Salganik et al. [41], we propose the following parameter estimation scheme that can be applied when dealing with one or more items affected by the same bandwagon effect.

We will assume that λ values are shared across a subset of $K \geq 1$ items and that all users are randomly assigned to one of M bins s.t. users give n ratings per item and are only presented with rating information regarding users in the same bin. Let p^k indicate the true preference of all users w.r.t. item k (analogous to p) and $\hat{\lambda}_i$ an estimate of λ_i . The log-likelihood of this data is then:

$$L = \sum_{k=1}^K \sum_{m=1}^M \sum_{i=1}^n r_{kmi} \ln(\hat{\lambda}_i p^k + (1 - \hat{\lambda}_i) \bar{p}_{i-1}^{km}) + (1 - r_{kmi}) \ln(1 - \hat{\lambda}_i p^k - (1 - \hat{\lambda}_i) \bar{p}_{i-1}^{km}). \quad (18)$$

The maximum likelihood estimate can be found by maximizing L , for instance through the iterative application of Newton's method. In order to get $\hat{\lambda}_i$ values for $i > n$, we propose fitting them to a curve, e.g. $\hat{\lambda}_i = a + (1 - a) \cdot b^{i-1}$ where $a \in [0, 1]$ and $b \in (0, 1)$. Note that as p^k may be unknown, one may instead estimate λ_i values by using its unbiased estimate $\bar{p}^k = \frac{1}{Mn} \sum_{m=1}^M \sum_{i=1}^n r_{kmi}$.

6 EXPERIMENTS AND RESULTS

6.1 Experimental Setup

To investigate how the bandwagon can affect convergence and whether our proposed estimators can mitigate negative effects, we performed several experiments with simulated user interactions.

For an item with true relevance $p = 0.4$, we generated user interactions under three conditions: a *no bandwagon* setting where $\forall \lambda_i = 1$, a *weak bandwagon* setting where $\lambda_i = 0.6 + 0.4 \cdot 0.9^{i-1}$ and a *strong bandwagon* setting where $\lambda_i = 0.1 + 0.9 \cdot 0.95^{i-1}$. Our comparison includes the following four estimators: (i) the sample mean \bar{p}_n (Equation 4); (ii) affine mean \hat{p}_n , which is the affine estimator (Equation 15) with uniform weights $\forall i : \omega_i = 1$; (iii) affine weighted \tilde{p}_n , which is the affine estimator but with the weights $\forall i : \omega_i = \lambda_i$; and (iv) the maximum likelihood estimate p_n^* (Equation 17). For statistical significance, the entire sampling process was repeated 1,000 times independently, all our graphs include 90% confidence bounds over these independent runs.

Furthermore, to investigate how sensitive our estimators are to misestimated $\hat{\lambda}$ values, we computed the estimator values with both the actual λ values and estimated $\hat{\lambda}$ values. We followed the estimation scheme of Section 5.2, generating additional ratings with $K = 1$, $p_1 = 0.4$, $M = 20$ and $n = 100$ to represent a plausible production-level scenario [41]. The result obtained from fitting a curve: $\hat{\lambda}_i = a + (1 - a) \cdot b^{i-1}$ where $a \in [0, 1]$ and $b \in (0, 1)$ to the data was: $\hat{\lambda}_i = 0.33 + 0.67 \cdot 0.92^{i-1}$. We only report the results for misestimated $\hat{\lambda}$ values in the strong bandwagon setting, due to high similarity with the weak setting.




Figure 3: Convergence of \bar{p}_n under various levels of bandwagon. Solid and dashed lines indicate points where a 90% confidence interval over 1,000 independent runs is within 0.05 and 0.01 of p respectively, $p = 0.4$.

6.2 Estimator Bias and Convergence

Our results are visualized in Figure 3, which displays the convergence rate of the sample mean under various levels of bandwagon effect, and Figure 4, which shows the behavior of all estimators with known λ in the weak and strong bandwagon settings.

First, we consider the influence of the bandwagon effect on the sample mean (\bar{p}_n): from Figure 3 it is immediately clear that the effect's strength has a significant impact on the rate of convergence. Whilst under no bandwagon effect the confidence intervals are within 0.05 of p after fewer than 300 samples and within 0.01 after roughly 8,000 samples, under the weak bandwagon effect these numbers exceed 1,000 and 33,000 respectively. Under the strong bandwagon effect the confidence intervals increase even further, exceeding 0.05 after 10^6 samples. However, in all cases the confidence bounds still appear to decrease steadily, suggesting that \bar{p}_n will converge eventually, as also predicted by the sufficient condition for consistency (Equation 6). Nevertheless, it is also clear that the bandwagon effect introduces a significant loss in data efficiency. In practice, this could mean that an infeasibly large amount of data is required to get a decent estimate under a strong bandwagon effect.

Next, we examine the impact on the bias of the sample mean and other estimators. As shown in Figure 4 (top, middle), even prior to convergence, all estimators including \bar{p}_n make negligible errors in expectation, strongly suggesting that they are all unbiased. Whilst the unbiasedness was already theoretically proven for \bar{p}_n and the affine estimators \hat{p}_n and \tilde{p}_n , it is surprising to see that even at low numbers of samples the expected error of the maximum likelihood estimate p_n^* appears negligible. Our results thus seem to indicate that all the estimators are unbiased or have an extremely small bias, further confirming that the bandwagon effect is not a bias problem.

We then focus on the weak bandwagon condition (Figure 4 top). Here, all estimators converge to p , at a rate similar to that of \bar{p}_n . That the sample mean and affine estimators converge is not surprising, as the setting meets the sufficient condition for consistency. However, in this scenario none of the estimators actually provide a clear improvement over the sample mean, indicating that they are unable to apply significant beneficial corrections to the bandwagon effect. The lack of improvement over \bar{p}_n is particularly concerning as this indicates that our estimators are similarly affected by the loss of data efficiency induced by the bandwagon effect.




Figure 4: Convergence of \hat{p}_n (left), \tilde{p}_n (center) and p_n^* (right) compared to \bar{p}_n with $p = 0.4$ and $\lambda_i = 0.6 + 0.4 \cdot 0.9^{i-1}$ (top), $\lambda_i = 0.1 + 0.9 \cdot 0.95^{i-1}$ (center and bottom) and misestimated $\hat{\lambda}_i = 0.33 + 0.67 \cdot 0.92^{i-1}$ (bottom); bold lines show the mean estimate over 1,000 independent runs, shaded areas indicate 90% confidence intervals.

Nonetheless, one may expect the convergence to slow down even further in the strong bandwagon setting shown in the middle row of Figure 4. Here, both the affine and maximum likelihood estimators are actually able to converge on the true value much sooner than the sample mean. However, confidence bounds of the former also become wider after roughly 10 sampled ratings and stay wider until roughly 1,000 samples, with the weighting of \tilde{p}_n seemingly reducing this effect when compared with \hat{p}_n . Thus during this period, the sample mean actually appears to be a more reliable choice in terms of expected error. This observation illustrates that choosing between estimators is not always straightforward, as in this case the optimal choice between \bar{p}_n and \tilde{p}_n depends on the number of samples that are available. Importantly, the maximum likelihood estimator does not have this drawback: its confidence bounds never appear wider than the sample mean while still converging quicker.

Across both settings, we have thus observed no noticeable bias problems for any estimator. Importantly, the sample mean has serious convergence issues under the strong bandwagon, where the other estimators are able to converge much sooner. Nonetheless, despite its high efficiency, the affine estimator was also found to initially have the highest variance. It appears that when λ values are known, the maximum likelihood estimator is the most reliable in spite of the lack of proven theoretical guarantees for this estimator. However, even the maximum likelihood estimator is unable to fully undo the slowdown introduced by either level of the bandwagon effect, highlighting the substantial impact this effect still has on effective sample size and relevance estimation.

6.3 Bandwagon Misestimation

To investigate the effect of misestimated $\hat{\lambda}$ values, we look at the bottom row in Figure 4. In particular, it shows what happens when we underestimate a strong bandwagon effect. Clearly, we see that no noticeable bias has been introduced to any estimator. Interestingly, the affine estimators actually suffer less from initial variance: however, we speculate that overestimating the bandwagon effect could instead enlarge the variance. Nevertheless, all our estimators converge substantially slower than when λ was known (cf. middle row), with their confidence bounds now only moderately smaller than those of \bar{p}_n . We thus conclude that the ability of our estimators to accelerate convergence is very sensitive to correct estimation of λ , where we do not observe an advantage for any of the estimators.

The results with misestimated $\hat{\lambda}$ values demonstrate that λ estimation is a crucial part of effective bandwagon correction. Researching bandwagon estimation methods is therefore an important task for future work: since without robust estimation, correcting for the bandwagon effect appears to remain an open problem.

7 DISCUSSION

So far, we have investigated the bandwagon effect in both theoretical and empirical terms. In contrast with the well-studied selection bias [22, 43, 45, 63], our experimental results indicate that the bandwagon effect mainly poses a problem for convergence without introducing noticeable statistical bias. Yet the effect is generally referred to as a bias by previous work [56, 60, 61, 67]. This raises a

higher-level conceptual question: *should the bandwagon effect be interpreted as a bias?* The remainder of this section will contrast three possible perspectives on this question that we deem valuable for the understanding of the field.

On one hand, both our theoretical and empirical results clearly indicate that statistical bias is not the issue. The expected value of both the sample mean and individual ratings is equal to the true relevance, with no noticeable error in the mean value of each estimator in our experiments. Therefore, in statistical terms, it seems that we are purely dealing with convergence and consistency issues and not with bias. This distinction is important as solutions to inconsistency may be very different than those for bias: i.e. our experimental results suggest that methods based on debiasing techniques are not particularly effective at solving convergence issues.

On the other hand, our theoretical analysis also revealed that the bandwagon effect introduces conditional bias: given that there is an error in the sample mean after m ratings, the expected value of any subsequent rating or sample mean will be incorrect. In other words, under the bandwagon effect any errors in the sample mean, which are unavoidable in practice, will subsequently introduce (conditional) bias. A proponent of the first perspective could argue that this view is mistaking convergence issues for bias. Nevertheless, we argue that as long as both parties understand the difference between *bias* and *conditional bias*, their disagreement is mainly about semantics and not on the underlying nature of the issue.

Alternatively, one could consider the bandwagon effect as actually changing the *true* preferences of users. This view poses that the observed ratings are not a distortion of the true preferences but that the actual preferences of users have changed. One could argue that this is in line with the users' experience as they do not make a conscious choice to alter their ratings in response to others' feedback but aim to assign ratings according to their satisfaction. Instead, their enjoyment of a product may still be subconsciously influenced by the enjoyment of others. Thus the advantage of this view is the apparent alignment with the users' own experience. However, their experience is precisely unaware of any subconscious factors such as the bandwagon effect. The danger is that by not viewing the bandwagon effect as a misrepresentation of preferences, one may be less inclined to correct for it. However, when training a recommender system, it surely seems critical to know whether high ratings were given because a user was affected by previous users, or whether they made this decision solely based on the item itself.

Apart from these conceptual perspectives, we strongly argue that in statistical terms the bandwagon effect is a problem of convergence and consistency and not of statistical bias. Nevertheless, outside of pure statistical terms, we think that each of the three views brings a valuable perspective to the field. Most importantly, we think the field should be aware of the difference between the bandwagon effect and statistical bias. As long as this distinction is understood, we think it is acceptable if one chooses to conceptually categorize the bandwagon effect as another form of bias, although we would advise against it. Overall, our work aims to promote a more complete statistical understanding of the issues surrounding user interactions that goes beyond bias to include variance, consistency and convergence.

8 CONCLUSION

In this paper we examined the bandwagon effect: an underexplored phenomenon where users' interactions are affected by interactions of earlier users. While previous work has approached the bandwagon effect as problem of bias [60, 61], our theoretical analysis revealed that it does not add bias to the rating distribution nor to the sample mean. Instead, we have shown that it is better viewed as a problem of convergence and consistency. We derived a necessary condition for consistency under the effect, thereby also adding to earlier consistency claims [60]. We then proposed several novel estimators to mitigate the convergence issues. While our empirical analysis revealed that our novel estimators can substantially increase the convergence rate under a strong bandwagon effect, all estimators also carry their own limitations. Moreover, even when using the empirically best-performing estimator, our purposely simple model of the bandwagon effect still significantly increases the number of samples required for accurate relevance estimation. Future work may wish to examine the impact of more complex bandwagon models (e.g. [65, 67]) as well as parameter and rating scale choices on bias and convergence, both in unpersonalized as well as *personalized* settings. For the latter case in particular, user preference may be further associated with the time of arrival (e.g. item's fans are both more likely to like it and to rate it early) or λ values for one timestep may vary across users. Nevertheless, even under our simple model, our findings lead us to conclude that there is no clear and well-understood solution for the bandwagon effect.

Acknowledgments

This work was supported by the Google Research Scholar Program. All content represents the opinion of the author, which is not necessarily shared or endorsed by their employers and/or sponsors.

Code and data

To facilitate the reproducibility of the reported results, our experimental implementation is publicly available at <https://github.com/NKNY/bandwagonictir2022>.

A APPENDIX

THEOREM A.1 (BIAS OF SAMPLE MEAN AND SINGLE RATING). \bar{p}_n and r_n (Equation 4) are unbiased estimators of parameter p for all n :

$$\forall n : \mathbb{E}[\bar{p}_n] = \mathbb{E}[r_n] = p. \quad (19)$$

PROOF. $\mathbb{E}[\bar{p}_1] = \mathbb{E}[r_1] = p$ by Equations 4-5. Then by induction:

$$\mathbb{E}[\bar{p}_n] = \frac{(n-1)\mathbb{E}[\bar{p}_{n-1}] + \mathbb{E}[r_n]}{n} = \frac{\lambda_n p + (n-\lambda_n)\mathbb{E}[\bar{p}_{n-1}]}{n} = p.$$

r_n is unbiased as it is a weighted sum of p and unbiased $\mathbb{E}[\bar{p}_{n-1}]$:

$$\mathbb{E}[r_n] = \mathbb{E}[\lambda_n p + (1-\lambda_n)\bar{p}_{n-1}] = \lambda_n p + (1-\lambda_n)\mathbb{E}[\bar{p}_{n-1}] = p. \quad \square$$

LEMMA A.2 (VARIANCE OF r_i). Under the bandwagon effect the variance of any rating r_i is $p(1-p)$.

PROOF. By Theorem A.1:

$$\forall i : \mathbb{V}[r_i] = \mathbb{E}[r_i^2] - \mathbb{E}^2[r_i] = \mathbb{E}[r_i] - \mathbb{E}^2[r_i] = p - p^2. \quad \square$$

LEMMA A.3. Covariance of \bar{p}_{n-1} and r_n is $(1-\lambda_n)^2 \mathbb{V}[\bar{p}_{n-1}]$.

PROOF.

$$\begin{aligned} \text{Cov}[\bar{p}_{n-1}, r_n] &= \mathbb{E}[(\bar{p}_{n-1} - p)(\lambda_n p + (1 - \lambda_n)\bar{p}_{n-1} - p)] \quad (20) \\ &= (1 - \lambda_n)^2 \mathbb{V}[\bar{p}_{n-1}]. \quad \square \end{aligned}$$

THEOREM A.4 (EFFICIENCY OF SAMPLE MEAN). *The efficiency of the sample mean under the bandwagon effect defined in Equation 4 is*

$$\mathbb{E}[(\bar{p}_n - p)^2] = p(1-p) \left(\frac{1}{n^2} + \sum_{i=1}^{n-1} \frac{1}{i^2} \prod_{j=i+1}^n \frac{(j-1)(j+1-2\lambda_j)}{j^2} \right).$$

PROOF. By Theorems A.1 and Lemmas A.2, A.3:

$$\begin{aligned} \mathbb{E}[(\bar{p}_n - p)^2] &= \frac{(n-1)^2}{n^2} \mathbb{V}[\bar{p}_{n-1}] + \frac{2(n-1)}{n^2} \text{Cov}[r_n, \bar{p}_{n-1}] \\ &+ \frac{\mathbb{V}[r_n]}{n^2} = \frac{(n-1)(n+1-2\lambda_n)}{n^2} \mathbb{V}[\bar{p}_{n-1}] + \frac{p(1-p)}{n^2}. \quad (21) \end{aligned}$$

The proof is completed by unrolling the process to $\mathbb{V}[\bar{p}_1] = p(1-p)$ and collecting all the terms. \square

THEOREM A.5. *The asymptotic efficiency of the sample mean is:*

$$\lim_{n \rightarrow \infty} \mathbb{E}[(\bar{p}_n - p)^2] = \lim_{n \rightarrow \infty} p(1-p) \left(\sum_{i=1}^{n-1} \frac{1}{i(i+1)} \prod_{j=i+1}^{n-1} \left(1 - \frac{2\lambda_j}{j+1} \right) \right).$$

PROOF. By re-grouping numerators and denominators in the product of Theorem A.4 and cancelling out extra terms:

$$\mathbb{E}[(\bar{p}_n - p)^2] = p(1-p) \left(\frac{1}{n^2} + \sum_{i=1}^{n-1} \frac{n+1-2\lambda_n}{i(i+1)n} \prod_{j=i+1}^{n-1} \left(1 - \frac{2\lambda_j}{j+1} \right) \right).$$

The proof is completed by taking $\lim_{n \rightarrow \infty}$ of above expression. \square

THEOREM A.6 (SAMPLE MEAN CONSISTENCY). $\sum \lambda_i^2$ diverging is a necessary condition for the consistency of \bar{p}_n (Equation 4):

$$\forall \epsilon > 0 : \lim_{n \rightarrow \infty} \Pr(|\bar{p}_n - p| > \epsilon) = 0 \longrightarrow \lim_{n \rightarrow \infty} \sum_{i=1}^n \lambda_i^2 = \infty. \quad (22)$$

PROOF. By Loève [31, p. 159], consistency of \bar{p} implies its variance vanishing: $\forall \epsilon > 0 : \Pr(|\bar{p}_n - p| \geq \epsilon) \geq \mathbb{V}[\bar{p}_n] - \epsilon^2$. Then: $\forall \epsilon > 0 : \lim_{n \rightarrow \infty} \Pr(|\bar{p}_n - p| > \epsilon) = 0 \rightarrow \lim_{n \rightarrow \infty} \text{Var}(\bar{p}_n) = 0 \rightarrow \forall i : \lim_{n \rightarrow \infty} \prod_{j=i+1}^{n-1} \left(1 - \frac{2\lambda_j}{j+1} \right) = 0 \rightarrow \forall i : \lim_{n \rightarrow \infty} \sum_{j=i+1}^{n-1} \frac{\lambda_j}{j+1} = \infty \rightarrow \forall i : \lim_{n \rightarrow \infty} \sqrt{\sum_{j=i+1}^{n-1} \lambda_j^2} \sqrt{\sum_{j=i+1}^{n-1} \frac{1}{(j+1)^2}} = \infty \rightarrow \lim_{n \rightarrow \infty} \sum_{i=1}^n \lambda_i^2 = \infty$. The second step is due to partial products in Theorem A.5, the third by Leonard [28, p. 6], the fourth by Cauchy-Schwarz inequality and the fifth by convergence of p-series with $p = 2$. \square

PROPOSITION A.7 (CONDITIONAL BIAS OF SAMPLE MEAN).

$$\mathbb{E}[\bar{p}_n - p \mid \bar{p}_m] = (\bar{p}_m - p) \prod_{i=m+1}^n \left(1 - \frac{\lambda_i}{i} \right) \text{ where } m < n.$$

PROOF. By repeatedly applying the induction step in Theorem A.1, replacing $\mathbb{E}[\bar{p}_n]$ with $\mathbb{E}[\bar{p}_n \mid \bar{p}_m]$ and noting that $\mathbb{E}[\bar{p}_m \mid \bar{p}_m] = \bar{p}_m$:

$$\mathbb{E}[\bar{p}_n - p \mid \bar{p}_m] = \mathbb{E}[\bar{p}_{n-1} - p \mid \bar{p}_m] \left(1 - \frac{\lambda_n}{n} \right) = (\bar{p}_m - p) \prod_{i=m+1}^n \left(1 - \frac{\lambda_i}{i} \right). \quad \square$$

PROPOSITION A.8 (ERROR BOUND OF \bar{p}_n). *If $\lambda_i = c^{i-1}$ s.t. $c \in (0, 1)$, the asymptotic expected absolute error of the sample mean is at least $2p(1-p) \exp(-c\phi(c, 1, 2) - c^2\phi(c^2, 2, 2))$, where $\phi(z, s, a) \equiv \sum_{k=0}^{\infty} \frac{z^k}{(a+k)^s}$ is Lerch transcendent and converges to a fixed value $\forall z < 1$ [11].*

PROOF.

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{E}[|\bar{p}_n - p|] &\geq \lim_{n \rightarrow \infty} \mathbb{E}[|\mathbb{E}[\bar{p}_n - p \mid r_1]|] = 2p(1-p) \\ &\cdot \exp\left(\sum_{i=2}^{\infty} \ln\left(1 - \frac{\lambda_i}{i} \right) \right) \geq 2p(1-p) \exp\left(\sum_{i=2}^{\infty} \left(-\frac{\lambda_i}{i} - \frac{\lambda_i^2}{i^2} \right) \right) \\ &= 2p(1-p) \exp\left(-c \sum_{i=0}^{\infty} \frac{c^i}{(2+i)} - c^2 \sum_{i=0}^{\infty} \frac{c^{2i}}{(2+i)^2} \right). \end{aligned}$$

The first step is due to the law of total expectation and Jensen's inequality for conditional expectation. The second is by Proposition A.7 with $m = 1$. The third is due to a known inequality $\forall x > -0.68 : \ln(1+x) \geq x - x^2$ [26]. \square

THEOREM A.9. *Each \hat{r}_i (and thus also \hat{p}_n) is unbiased for all $\hat{\lambda}_i$.*

PROOF. By Theorem A.1, for any choice of i and $\hat{\lambda}_i$:

$$\mathbb{E}[\hat{r}_i \mid \hat{\lambda}_i] = \mathbb{E}\left[\frac{r_i - (1 - \hat{\lambda}_i)\bar{p}_{i-1}}{\hat{\lambda}_i} \right] = \frac{\mathbb{E}[r_i - \bar{p}_{i-1}]}{\hat{\lambda}_i} + \mathbb{E}[\bar{p}_{i-1}] = p. \quad \square$$

THEOREM A.10 (CONDITIONAL BIAS OF $\hat{r}_i \mid \hat{\lambda}_i = \lambda_i$). *If $\hat{\lambda}_i = \lambda_i$, the affine mean estimator \hat{r}_i is an unbiased estimator of p when conditioned on \hat{r}_j or \bar{p}_j s.t. $\forall j < i$.*

PROOF. For any function of previous ratings $f_i = f(r_1, \dots, r_{i-1})$:

$$\begin{aligned} \mathbb{E}[\hat{r}_i \mid f_i, \hat{\lambda}_i] &= \mathbb{E}\left[\frac{\lambda_i p + (1 - \lambda_i)\bar{p}_{i-1} - (1 - \hat{\lambda}_i)\bar{p}_{i-1}}{\hat{\lambda}_i} \mid f_i \right] \\ &= \frac{\lambda_i}{\hat{\lambda}_i} p + \left(\frac{\hat{\lambda}_i - \lambda_i}{\hat{\lambda}_i} \right) \mathbb{E}[\bar{p}_{i-1} \mid f_i] = p. \quad (23) \end{aligned}$$

The proof is completed by setting $f_i = \hat{r}_j$ or $f_i = \bar{p}_j$. \square

COROLLARY A.11 (CONDITIONAL BIAS OF $\hat{r}_i \mid \hat{\lambda}_i \neq \lambda_i$). *Affine estimator \hat{r}_i is not unbiased when conditioned on a function of previous interactions as the last equality of Equation 23 may not hold.*

LEMMA A.12 (COVARIANCE OF AFFINE ESTIMATORS). *\hat{r}_i and $\hat{r}_{j < i}$ are uncorrelated as long as $\hat{\lambda}_i = \lambda_i$.*

PROOF. By Theorem A.10 and the law of total expectation:

$$\begin{aligned} \mathbb{E}[\hat{r}_i \hat{r}_j] &= \mathbb{E}[\mathbb{E}[\hat{r}_i \hat{r}_j \mid \hat{r}_j]] = \mathbb{E}[\hat{r}_j \mathbb{E}[\hat{r}_i \mid \hat{r}_j]] \\ &= \mathbb{E}[\hat{r}_j \mathbb{E}[\hat{r}_i]] = \mathbb{E}[\hat{r}_j] \mathbb{E}[\hat{r}_i]. \quad \square \end{aligned}$$

THEOREM A.13. *The variance of affine estimator in Equation 15 is:*

$$\mathbb{V}[\hat{p}_n] = \frac{1}{(\sum_{i=1}^n \omega_i)^2} \sum_{i=1}^n \frac{\omega_i^2}{\lambda_i^2} \left(p(1-p) - (1 - \lambda_i)^2 \mathbb{V}[\bar{p}_{i-1}] \right).$$

PROOF. By Lemmas A.12, A.3 and A.2:

$$\begin{aligned} \mathbb{V}[\hat{p}_n] &= \frac{1}{\left(\sum_{i=1}^n \omega_i\right)^2} \left(\sum_{i=1}^n \omega_i^2 \mathbb{V}[\hat{r}_i] + 2 \sum_{i=1}^{n-1} \sum_{j=i+1}^n \omega_i \omega_j \text{Cov}[\hat{r}_i, \hat{r}_j] \right) \\ &= \frac{1}{\left(\sum_{i=1}^n \omega_i\right)^2} \sum_{i=1}^n \frac{\omega_i^2}{\lambda_i^2} \left(\mathbb{V}[r_i] + (1 - \lambda_i)^2 \mathbb{V}[\bar{p}_{i-1}] - 2(1 - \lambda_i) \right. \\ &\quad \cdot \text{Cov}[r_i, \bar{p}_{i-1}] \left. \right) = \left(\sum_{i=1}^n \omega_i \right)^{-2} \sum_{i=1}^n \frac{\omega_i^2}{\lambda_i^2} \left(p - p^2 - (1 - \lambda_i)^2 \mathbb{V}[\bar{p}_{i-1}] \right). \quad \square \end{aligned}$$

THEOREM A.14 (CONVERGENCE OF AFFINE ESTIMATOR). *When $\forall \hat{\lambda}_i = \lambda_i$, $\hat{p}_n = \frac{\sum_{i=1}^n \omega_i \hat{r}_i}{\sum_{i=1}^n \omega_i}$ converges within ϵ of the true value p with probability of at least $1 - \alpha$ if $\left(\sum_{i=1}^n \omega_i\right)^2 \geq \frac{1}{2\epsilon^2} \log \frac{2}{\alpha} \sum_{i=1}^n \left(\frac{\omega_i}{\lambda_i}\right)^2$.*

PROOF. Define X_n as $\sum_{i=1}^n \omega_i (\hat{r}_i - p)$, $X_0 = 0$. Then X_n is a martingale as $\mathbb{E}[X_n | X_0, X_1, \dots, X_{n-1}] = X_{n-1} + \mathbb{E}[w_i (\hat{r}_i - p) | \bar{p}_{n-1}] = X_{n-1}$. For each n : $X_n - X_{n-1} = \omega_n (\hat{r}_n - p)$, whose values for $r_n = 0$ and $r_n = 1$ are separated by $d_n = \frac{\omega_n}{\lambda_n}$. By Azuma's inequality:

$$\begin{aligned} \Pr\left(\left|\frac{\sum_{i=1}^n \omega_i (\hat{r}_i - p)}{\sum_{i=1}^n \omega_i}\right| > \epsilon\right) &= \Pr\left(|X_n - X_0| > \epsilon \sum_{i=1}^n \omega_i\right) \\ &\leq 2 \exp\left(-\frac{2\epsilon^2 \left(\sum_{i=1}^n \omega_i\right)^2}{\sum_{i=1}^n d_i^2}\right). \end{aligned} \quad (24)$$

The proof is completed by taking the logarithm of both sides and redistributing the terms. \square

COROLLARY A.15 (CONVERGENCE OF \hat{p}_n FOR SPECIFIC ω_i). *By Theorem A.14, when $\forall i$: $\hat{\lambda}_i = \lambda_i$, $\omega_i = 1$, \hat{p}_n is a consistent estimator of p if $\sum_{i=1}^n \frac{1}{\lambda_i^2}$ grows slower than n^2 . Similarly, when $\forall \hat{\lambda}_i = \lambda_i = \omega_i$, \hat{p}_n is a consistent estimator of p if $\left(\sum_{i=1}^n \lambda_i\right)^2$ grows quicker than n .*

REFERENCES

- [1] Gediminas Adomavicius, Jesse C. Bockstedt, Shawn Curley, and Jingjing Zhang. 2016. Understanding Effects of Personalized vs. Aggregate Ratings on User Preferences. *CEUR Workshop Proceedings* 1679 (2016), 14–21.
- [2] Solomon E. Asch. 1951. Effects of Group Pressure upon the Modification and Distortion of Judgments. In *Groups, Leadership and Men; Research in Human Relations*. Carnegie Press, 177–190.
- [3] Christian Borgheci and Jean-Philippe Bouchaud. 2010. Spatial Correlations in Vote Statistics: A Diffusive Field Model for Decision-Making. *The European Physical Journal B* 75, 3 (2010), 395–404.
- [4] Rocío Cañamares and Pablo Castells. 2018. Should I Follow The Crowd? A Probabilistic Analysis of the Effectiveness of Popularity in Recommender Systems. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 415–424.
- [5] Franco Caparrelli, Anna M. D'Arcangelis, and Alexander Cassuto. 2004. Herding in the Italian Stock Market: A Case of Behavioral Finance. *Journal of Behavioral Finance* 5, 4 (2004), 222–230.
- [6] Aleksandr Chuklin, Ilya Markov, and Maarten de Rijke. 2015. *Click Models for Web Search*. Morgan & Claypool Publishers.
- [7] Dan Cosley, Shyong K. Lam, Istvan Albert, Joseph A. Konstan, and John Riedl. 2003. Is Seeing Believing? How Recommender Interfaces Affect Users' Opinions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 585–592.
- [8] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An Experimental Comparison of Click Position-Bias Models. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*. 87–94.
- [9] Martin J. Crowder. 1976. Maximum Likelihood Estimation for Dependent Observations. *Journal of the Royal Statistical Society. Series B (Methodological)* 38, 1 (1976), 45–53.
- [10] Utpal M. Dholakia and Kerry Soltysinski. 2001. Coveted or Overlooked? The Psychology of Bidding for Comparable Listings in Digital Auctions. *Marketing Letters* 12, 3 (2001), 225–237.
- [11] Erasmo M. Ferreira, Anderson K. Kohara, and Javier Sesma. 2017. New Properties of the Lerch's Transcendent. *Journal of Number Theory* 172 (2017), 21–31.
- [12] Andrew J. Flanagan and Miriam J. Metzger. 2013. Trusting Expert- versus User-Generated Ratings Online: The Role of Information Volume, Valence, and Consumer Characteristics. *Computers in Human Behavior* 29, 4 (2013), 1626–1634.
- [13] W. Wayne Fu and Clarice C. Sim. 2011. Aggregate Bandwagon Effect on Online Videos' Viewership: Value Uncertainty, Popularity Cues, and Heuristics. *Journal of the American Society for Information Science and Technology* 62, 12 (2011), 2382–2395.
- [14] Aurel Galántai. 2000. The Theory of Newton's Method. *J. Comput. Appl. Math.* 124, 1 (2000), 25–44.
- [15] Ward A. Hanson and Daniel S. Putler. 1996. Hits and Misses: Herd Behavior and Online Product Popularity. *Marketing Letters* 7, 4 (1996), 297–305.
- [16] José M. Hernández-Lobato, Neil Houlsby, and Zoubin Ghahramani. 2014. Probabilistic Matrix Factorization with Non-random Missing Data. In *Proceedings of the 31st International Conference on Machine Learning*. 1512–1520.
- [17] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *2008 Eighth IEEE International Conference on Data Mining*. 263–272.
- [18] Jin Huang, Harrie Oosterhuis, Maarten de Rijke, and Herke van Hoof. 2020. Keeping Dataset Biases out of the Simulation: A Debaised Simulator for Reinforcement Learning based Recommender Systems. In *Fourteenth ACM Conference on Recommender Systems*. 190–199.
- [19] Jen-Hung Huang and Yi-Fen Chen. 2006. Herding in Online Product Choice. *Psychology & Marketing* 23, 5 (2006), 413–428.
- [20] Ilija Ilievski and Sujoy Roy. 2013. Personalized News Recommendation Based On Implicit Feedback. In *Proceedings of the 2013 International News Recommender Systems Workshop and Challenge*. 10–15.
- [21] Dietmar Jannach, Lukas Lerche, and Markus Zanker. 2018. Recommending Based on Implicit Feedback. In *Social Information Access*. Springer, 510–569.
- [22] Olivier Jeunen and Bart Goethals. 2021. Pessimistic Reward Models for Off-Policy Learning in Recommendation. In *Fifteenth ACM Conference on Recommender Systems*. 63–74.
- [23] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. 2017. Unbiased Learning-to-Rank with Biased Feedback. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. 781–789.
- [24] Jinyoung Kim and Andrew Gambino. 2016. Do We Trust the Crowd or Information System? Effects of Personalization and Bandwagon Cues on Users' Attitudes and Behavioral Intentions toward a Restaurant Recommendation Website. *Computers in Human Behavior* 65 (2016), 369–379.
- [25] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (2009), 30–37.
- [26] László Kozma. 2021. Useful Inequalities Cheat Sheet. Retrieved February 14, 2022 from http://www.lkozma.net/inequalities_cheat_sheet/ineq.pdf
- [27] Bibb Latané and Sharon Wolf. 1981. The Social Impact of Majorities and Minorities. *Psychological Review* 88, 5 (1981), 438–453.
- [28] Isaac Leonard. 2012. Notes on the Riemann Zeta Function. MATH 324, Summer 2012, Elementary Number Theory, University of Alberta. Retrieved February 14, 2022 from <http://www.math.ualberta.ca/~isaac/math324/s12/zeta.pdf>
- [29] Roderick J. A. Little and Donald B. Rubin. 2019. *Statistical Analysis with Missing Data*. John Wiley & Sons.
- [30] Yiming Liu, Xuezhi Cao, and Yong Yu. 2016. Are You Influenced by Others When Rating?: Improve Rating Prediction by Conformity Modeling. In *Proceedings of the 10th ACM Conference on Recommender Systems*. 269–272.
- [31] M. Loève. 1977. Probability Concepts. In *Probability Theory I*. Springer, 151–176.
- [32] Benjamin M. Marlin and Richard S. Zemel. 2009. Collaborative Prediction and Ranking with Non-Random Missing Data. In *Proceedings of the Third ACM Conference on Recommender Systems*. 5–12.
- [33] Harrie Oosterhuis. 2020. *Learning from User Interactions with Rankings: A Unification of the Field*. Ph.D. Dissertation. Informatics Institute, University of Amsterdam.
- [34] Harrie Oosterhuis and Maarten de Rijke. 2020. Policy-Aware Unbiased Learning to Rank for Top-k Rankings. *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (2020)*, 489–498.
- [35] Bruno Pradel, Nicolas Usunier, and Patrick Gallinari. 2012. Ranking with Non-Random Missing Satings: Influence of Popularity and Positivity on Evaluation Metrics. In *Proceedings of the Sixth ACM Conference on Recommender Systems*. 147–154.
- [36] Zhen Qin, Honglei Zhuang, Rolf Jagerman, Xinyu Qian, Po Hu, Dan Chary Chen, Xuanhui Wang, Michael Bendersky, and Marc Najork. 2021. Bootstrapping Recommendations at Chrome Web Store. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3483–3491.
- [37] Ramsey M. Raafat, Nick Chater, and Chris Frith. 2009. Herding in Humans. *Trends in Cognitive Sciences* 13, 10 (2009), 420–428.
- [38] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2015. Recommender Systems: Introduction and Challenges. In *Recommender Systems Handbook*. Springer, 1–34.
- [39] Yuta Saito. 2020. Doubly Robust Estimator for Ranking Metrics with Post-Click Conversions. In *Fourteenth ACM Conference on Recommender Systems*. 92–100.

- [40] Yuta Saito, Suguru Yaginuma, Yuta Nishino, Hayato Sakata, and Kazuhide Nakata. 2020. Unbiased Recommender Learning from Missing-Not-At-Random Implicit Feedback. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 501–509.
- [41] Matthew J. Salganik, Peter S. Dodds, and Duncan J. Watts. 2006. Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science* 311, 5762 (2006), 854–856.
- [42] Matthew J. Salganik and Duncan J. Watts. 2008. Leading the Herd Astray: An Experimental Study of Self-fulfilling Prophecies in an Artificial Cultural Market. *Social Psychology Quarterly* 71, 4 (2008), 338–355.
- [43] Masahiro Sato, Janmajay Singh, Sho Takemori, and Qian Zhang. 2021. Causality-Aware Neighborhood Methods for Recommender Systems. In *European Conference on Information Retrieval*. 603–618.
- [44] Masahiro Sato, Sho Takemori, Janmajay Singh, and Tomoko Ohkuma. 2020. Unbiased Learning for the Causal Effect of Recommendation. *Fourteenth ACM Conference on Recommender Systems (2020)*, 378–387.
- [45] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. 2016. Recommendations As Treatments: Debiasing Learning and Evaluation. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning*. 1670–1679.
- [46] Herbert A. Simon. 1954. Bandwagon and Underdog Effects and the Possibility of Election Predictions. *The Public Opinion Quarterly* 18, 3 (1954), 245–253.
- [47] Harald Steck. 2011. Item Popularity and Recommendation Accuracy. In *Proceedings of the Fifth ACM Conference on Recommender Systems*. 125–132.
- [48] Harald Steck. 2013. Evaluation of Recommendations: Rating-Prediction and Ranking. In *Proceedings of the Seventh ACM Conference on Recommender Systems*. 213–220.
- [49] Alexander L. Strehl, John Langford, Lihong Li, and Sham M. Kakade. 2010. Learning from Logged Implicit Exploration Data. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems*. 2217–2225.
- [50] Xiaoyuan Su and Taghi M. Khoshgoftaar. 2009. A Survey of Collaborative Filtering Techniques. *Advances in Artificial Intelligence* (2009).
- [51] S. Shyam Sundar, Anne Oeldorf-Hirsch, and Qian Xu. 2008. The Bandwagon Effect of Collaborative Filtering Technology. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems*. 3453–3458.
- [52] Adith Swaminathan and Thorsten Joachims. 2015. Counterfactual Risk Minimization: Learning from Logged Bandit Feedback. In *Proceedings of the 32nd International Conference on Machine Learning*. 814–823.
- [53] Adith Swaminathan and Thorsten Joachims. 2015. The Self-Normalized Estimator for Counterfactual Learning. In *Advances in Neural Information Processing Systems*, Vol. 28. 3231–3239.
- [54] Duyu Tang, Bing Qin, Ting Liu, and Yuekui Yang. 2015. User Modeling with Neural Network for Review Rating Prediction. In *Proceedings of the 24th International Conference on Artificial Intelligence*. 1340–1346.
- [55] Ali Vardasbi, Maarten de Rijke, and Ilya Markov. 2020. Cascade Model-Based Propensity Estimation for Counterfactual Learning to Rank. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2089–2092.
- [56] 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.
- [57] Xiaojie Wang, Rui Zhang, Yu Sun, and Jianzhong Qi. 2019. Doubly Robust Joint Learning for Recommendation on Data Missing Not at Random. In *International Conference on Machine Learning*. 6638–6647.
- [58] Yixin Wang, Dawen Liang, Laurent Charlin, and David M. Blei. 2020. Causal Inference for Recommender Systems. In *Fourteenth ACM Conference on Recommender Systems*. 426–431.
- [59] Tianxin Wei, Fuli Feng, Jiawei Chen, Ziwei Wu, Jinfeng Yi, and Xiangnan He. 2021. Model-Agnostic Counterfactual Reasoning for Eliminating Popularity Bias in Recommender System. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining* (2021), 1791–1800.
- [60] Hong Xie and Mingze Zhong. 2020. Robust Product Rating Rules Against Herding Effects: Theory and Applications. In *2020 IEEE International Conference on Data Mining*. 1352–1357.
- [61] 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 Transactions on Knowledge Discovery from Data* 15, 3 (2021), 1–29.
- [62] Longqi Yang, Yin Cui, Yuan Xuan, Chenyang Wang, Serge Belongie, and Deborah Estrin. 2018. Unbiased Offline Recommender Evaluation for Missing-Not-at-Random Implicit Feedback. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 279–287.
- [63] Bowen Yuan, Yaxu Liu, Jui-Yang Hsia, Zhenhua Dong, and Chih-Jen Lin. 2020. Unbiased Ad Click Prediction for Position-aware Advertising Systems. In *Fourteenth ACM Conference on Recommender Systems*. 368–377.
- [64] Xiaoying Zhang, Hong Xie, Junzhou Zhao, and John C. S. Lui. 2020. Understanding Assimilation-Contrast Effects in Online Rating Systems: Modelling, Debiasing, and Applications. *ACM Transactions on Information Systems* 38, 1 (2020), 1–25.
- [65] 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 Eleventh ACM Conference on Recommender Systems*. 98–106.
- [66] Yang Zhang, Fuli Feng, Xiangnan He, Tianxin Wei, Chonggang Song, Guohui Ling, and Yongdong Zhang. 2021. Causal Intervention for Leveraging Popularity Bias in Recommendation. *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (2021)*, 11–20.
- [67] Mingze Zhong, Hong Xie, and Qingsheng Zhu. 2021. Quantifying Assimilate-Contrast Effects in Online Rating Systems: Modeling, Analysis and Application. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2369–2377.
- [68] Honglei Zhuang, Zhen Qin, Xuanhui Wang, Michael Bendersky, Xinyu Qian, Po Hu, and Dan Chary Chen. 2021. Cross-Positional Attention for Debiasing Clicks. In *Proceedings of the Web Conference 2021*. 788–797.