



Guaranteeing Accuracy and Fairness under Fluctuating User Traffic: A Bankruptcy-Inspired Re-ranking Approach

Xiaopeng Ye
Gaoling School of Artificial
Intelligence
Renmin University of China
Beijing, China
xpye@ruc.edu.cn

Chen Xu
Gaoling School of Artificial
Intelligence
Renmin University of China
Beijing, China
xc_chen@ruc.edu.cn

Jun Xu*
Gaoling School of Artificial
Intelligence
Renmin University of China
Beijing, China
junxu@ruc.edu.cn

Xuyang Xie
Huawei Noah's Ark Lab
Shenzhen, China
xiexuyang@huawei.com

Gang Wang
Huawei Noah's Ark Lab
Shenzhen, China
wanggang110@huawei.com

Zhenhua Dong
Huawei Noah's Ark Lab
Shenzhen, China
dongzhenhua@huawei.com

Abstract

Out of sustainable and economical considerations, two-sided recommendation platforms must satisfy the needs of both users and providers. Previous studies often show that the two sides' needs show different urgency: providers need a relatively long-term exposure demand while users want more short-term and accurate service. However, our empirical study reveals that previous methods for trading off fairness-accuracy often fail to guarantee long-term fairness and short-term accuracy simultaneously in real applications of fluctuating user traffic. Especially, when user traffic is low, the user experience often drops a lot. Our theoretical analysis also confirms that user traffic is a key factor in such a trade-off problem. How to guarantee accuracy and fairness under fluctuating user traffic remains a problem. Inspired by the bankruptcy problem in economics, we propose a novel fairness-aware re-ranking approach named BankFair. Intuitively, BankFair employs the Talmud rule to leverage periods of abundant user traffic to offset periods of user traffic scarcity, ensuring consistent user service at every period while upholding long-term fairness. Specifically, BankFair consists of two modules: (1) employing the Talmud rule to determine the required fairness degree under varying periods of user traffic; and (2) conducting an online re-ranking algorithm based on the fairness degree determined by the Talmud rule. Experiments on two real-world recommendation datasets show that BankFair outperforms all baselines regarding accuracy and provider fairness.

CCS Concepts

• **Information systems** → **Recommender systems.**

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '24, October 21–25, 2024, Boise, ID, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0436-9/24/10

<https://doi.org/10.1145/3627673.3679590>

Keywords

Provider Fairness, Two-sided Platform, Recommender System

ACM Reference Format:

Xiaopeng Ye, Chen Xu, Jun Xu, Xuyang Xie, Gang Wang, and Zhenhua Dong. 2024. Guaranteeing Accuracy and Fairness under Fluctuating User Traffic: A Bankruptcy-Inspired Re-ranking Approach. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*, October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3627673.3679590>

1 Introduction

To create a more equitable and sustainable two-sided platform [11, 17, 40, 50] (e.g., Amazon and Youtube), recommendation systems (RS) should ensure the needs of both the provider side and the user side [35, 40, 50]. On the provider side, prior studies [35, 42, 52] have indicated that providers are required to a minimum level of exposure over a period, which aligns with the “minimum wage” policy [43, 51, 52]. On the user side, established behavior economic theory [21, 26, 36] suggests that users tend to experience increased frustration and dissatisfaction when receiving poor-quality recommendations (i.e., the loss aversion effect [26, 27, 36]). Therefore, platforms should maintain a consistently high standard of recommendation quality to mitigate these negative user experiences [10, 19].

Although the needs of users and providers should be both satisfied, the urgency of the two sides' needs differs. Providers typically have relatively long-term exposure demands, which do not require immediate fulfillment but rather over an extended period [48, 52, 54]. Conversely, on the user side, the need is more short-term and immediate, because once users receive a poor recommendation, they tend to remember and be impacted by that negative experience for much longer [26, 36]. To balance the needs of both sides, some heuristic methods [33, 35, 50] and online gradient descent methods [51, 52] have been proposed to trade-off accuracy and fairness.

However, the existing trade-off methods could fail in some real applications where the user traffic inevitably fluctuates [4, 37]. For example, we conduct a simulation study shown in Figure 1(a) to observe the daily accuracy performances (black line) of a theoretical oracle method under a real industrial dataset KuaiRand¹, which

¹<https://kuairand.com/>




Figure 1: Illustrative experiments based on KuaiRand dataset. (a) Lower user traffic leads to more accuracy loss; (b) Applying distinct fairness strategies based on different traffic levels.

contains 175K user traffic (represented as the blue bars) within half a month. The red line indicates the minimum daily average accuracy level that the platform can tolerate, below which performance is considered unacceptable. The oracle method guarantees that each provider receives a predefined minimum exposure each day while maximizing accuracy. The result reported in Figure 1(a) shows that the short-term accuracy is more prone to compromise to fulfill long-term fairness requirements, particularly when there are fewer users within a day (red dots). Similar phenomena have also been observed in most existing fairness-aware models [35, 46, 50–52].

To investigate the fundamental reasons behind how user traffic impacts the accuracy-fairness trade-off performances, we provide a theoretical analysis from a constrained optimization problem, as detailed in Section 4. Our analysis reveals that user traffic is a key factor affecting the fairness-accuracy trade-off. Specifically, lower user traffic will result in a more serious accuracy loss when maintaining the same fairness degree (i.e., required minimum exposure) across different periods. Since fluctuations in user traffic are not rare in real-world scenarios [4, 37], a two-sided re-ranking algorithm adaptable to user traffic is needed to balance the needs of both users and providers.

To well incorporate the user traffic into consideration, we utilize the Talmud rule in bankruptcy problem [12, 44] to allocate different fairness requirements across periods of different user traffic. Specifically, the Talmud rule utilizes surplus fairness requirements during periods of high user traffic to compensate for the deficit in fairness requirements during periods of low user traffic, ensuring long-term fairness while enhancing accuracy during periods of low user traffic. Figure 1(b) shows an illustrative example, where we use the depth of color to represent the fairness degree on that day (i.e., a deeper color of the bar indicates that we will ensure a higher minimum exposure to providers). Across different periods, we utilize the resource-abundant phase (deep color) to compensate for the resource-scarce phase (light color). For example, we ensure more minimum exposures to each provider during user-abundant periods (April 26th to April 30th), while reducing the minimum exposure guarantee during user-scarce periods (April 24th, May 1st to May 04th). In such a way, we can amortize the accuracy loss to each day and keep high accuracy levels across all days (i.e., all above the red line), ensuring the needs of both sides.

To implement the above idea, we propose an online re-ranking model called BankFair which can ensure both long-term provider fairness and short-term user accuracy under fluctuating traffic.

Specifically, BankFair consists of two modules: (1) Module 1: we formulate the exposure allocation process as a sequential bankruptcy problem and utilize the Talmud rule to decide the fairness degree of each period; and (2) Module 2: we utilize the fairness degree obtained from the Talmud rule to guide our online recommendation algorithm, which utilizes the dual method to efficiently solve the fair re-ranking problem in an online style. By combining two modules, we can achieve user-friendly and provider-fair re-ranking effectively and efficiently.

Our main contributions can be summarized as follows:

- We emphasize the significance of guaranteeing both short-term user accuracy and long-term provider fairness under fluctuating user traffic in two-sided platforms.
- We propose a two-sided re-ranking model named BankFair, which formulates the exposure allocation process as a sequential bankruptcy problem and utilizes the Talmud rule to solve it.
- The extensive experiments demonstrate that our approach outperforms existing baselines on two datasets in terms of both fairness and accuracy.

2 Related Work

Fairness-aware re-ranking in two-sided platforms has drawn much attention in recent years. Depending on the target audience, it can be divided into three types: user fairness [28, 29], provider fairness [35, 38, 51, 53], and two-sided fairness [33, 35]. For the provider fairness studied in this paper, there are different forms: (1) max-min fairness [51], which aims to ensure the interests of the worst-off providers; (2) equity of attention [6, 32, 33, 50] which lets the exposure received by providers be proportional to their utility; (3) minimum exposure guarantee [5, 7, 30, 35, 42, 52, 54], which tends to ensure that the exposure for providers over a period exceeds a minimum threshold. Within the research line, some heuristic methods [7, 35, 54] were proposed to ensure the minimum exposure in several recommendation lists and amortize the accuracy loss among each user (e.g., greedy round robin [7, 35]). Also, some other works [30, 42, 52] formulated the re-ranking problem with exposure constraint as Integer-Programming (IP) and adopted online optimization methods (e.g., sub-gradient descent [16]) to solve it. In most existing works [7, 35, 42, 54], the minimum exposure value cannot be directly specified but is controlled by a hyperparameter, which is hard to adjust and control. Thus, FairSync [52] considered a more industrially practical setting, which guarantees an arbitrarily specified minimum exposure, and we also consider this scenario in our paper. Nonetheless, all these works overlooked the accuracy loss caused by the minimum exposure fairness under fluctuating user traffic, making it impractical in real-world industrial scenarios.

Bankruptcy problem were first introduced by O’Neill [34] and have been widely used in scarce resource allocation scenarios, such as network traffic allocation [1], water resource allocation [14, 56], vaccines resource allocation [25]. Curiel et al. [12], Dagan and Volij [13] and Thomson [45] modeled bankruptcy problem as a coalition game, and Curiel et al. [12] further proved that Talmud solution corresponds to the core of coalition game, which is stable and impregnable. Traditional bankruptcy problem is a one-time allocation process. To adapt to the time-varying nature of RS, we




Figure 2: Illustration of algorithm updates over time. The shadow area represents the user traffic of each interval. The minimum exposure needs to be guaranteed at the end of the N -th interval.

transformed the bankruptcy problem into a sequential version that involves multiple time allocations.

3 Formulation

We first define some notations for the problem. For vector $\mathbf{x} \in \mathbb{R}^n$, x_i denote the i -th element of the vector. For matrix $X \in \mathbb{R}^{n \times m}$, let $X_{i,j}$ denote the element of i -th row and j -th column. A_i denote the i -th column vector of A . $\mathbf{x} \geq \mathbf{y}$ denotes element x_i should be greater or equal to $y_i, \forall i$. kx denotes that every element x_i in \mathbf{x} will become kx_i . Next, we give a formal formulation of two-sided re-ranking in RS.

3.1 Two-sided Re-ranking in Recommender System under Fluctuating User Traffic

In RS, let $\mathcal{U}, \mathcal{I}, \mathcal{P}$ denote the set of users, items, and providers, respectively. For a provider p , there are multiple items in the set \mathcal{I}_p belonging to p . When a user $u_t \in \mathcal{U}$ arrives at time t , RS will generate a ranking list $L_K^{\text{ri}}(u_t) \in \mathcal{I}^K$ of size K . Typically, the ranking list can be generated using a user-item relevance vector $\mathbf{s}_{u_t} \in \mathbb{R}^{|\mathcal{I}|}$, where the relevance score $s_{u_t,i} \in [0, 1]$ is estimated based on the historical information, user profile, etc. The goal of two-sided re-ranking is to compute a new list $L_K(u_t)$ which well balances the short-term user accuracy and long-term provider fairness. When an item i is recommended in $L_K(u_t)$, then the corresponding provider p with such an item (i.e., $i \in \mathcal{I}_p$) can have one exposure.

3.1.1 Two-sided Re-ranking across Multiple Time Intervals. In real scenarios, different users $\mathcal{U}_n = \{u_t | t_n \leq t < t_{n+1}, n \in [1, 2, \dots, N]\}$ will arrive in RS within time interval n , where t_n is the start time of interval n . We use $r_n = |\mathcal{U}_n|$ to denote the user traffic (i.e., user number) arriving within time interval n (e.g. n means n -th hour or n -th day during the recommendation process).

Let $a(n)$ be the average re-ranking accuracy within time interval n . Following literature convention [6, 50, 51], we define the re-ranking accuracy as the ratio between the sum of position-based relevance scores in the re-ranked list $L_K(u_t)$ and that in the original ranking list $L_K^{\text{ri}}(u_t)$, which is detailed in Equation (9). Let $E \in \mathbb{N}^{|\mathcal{P}| \times N}$ be the provider earned exposure matrix, where the element $E_{p,n}$ denotes the earned exposure for provider p at interval n .

Table 1: Correspondence between elements in the bankruptcy problem and two-sided re-ranking.

Bankruptcy problem	Two-sided re-ranking
agent set \mathcal{A}	time intervals \mathcal{N}
estate E	required minimum exposure \mathbf{m}_p
claim vector \mathbf{d}	demanded minimum exposure vector \mathbf{D}_p
result vector \mathbf{x}	predicted minimum exposure vector \mathbf{M}_p

Formally, the two-sided re-ranking task can be written as a linear programming:

$$f(\phi, \mathbf{m}) = \max_E \sum_{n=1}^N a(n; E) \quad (1)$$

$$\text{s.t. } a(n; E) \geq \phi, \quad \forall n \in [1, 2, \dots, N] \quad (2)$$

$$\sum_{n=1}^N E_{p,n} \geq \mathbf{m}_p, \quad \forall p \in \mathcal{P}, \quad (3)$$

where $\mathbf{m} \in \mathbb{R}^{|\mathcal{P}|}$ denote the required minimum exposure and $\phi \in [0, 1]$ is the required minimum accuracy. The model objective (1) aims to maximize the overall re-ranking accuracy among all intervals. The two constraints represent the needs of users and providers respectively:

- **Short-term Accuracy Need:** Constraint (2) requires the accuracy $a(n)$ at each interval n should be no less than the required minimum accuracy ϕ .
- **Long-term Exposure Need:** Constraint (3) demands the cumulative exposure $\sum_{n=1}^N E_{p,n}$ of provider p within N intervals should be no less than the required minimum exposure \mathbf{m}_p .

3.1.2 Re-ranking with Algorithm Updates. Under real industrial settings as shown in Figure 2, re-ranking model $f(\phi, \mathbf{m})$ updates its parameters regularly [49, 55]. Under this condition, the ideal re-ranking problem $f(\phi, \mathbf{m})$ cannot be globally optimized for all intervals. Hence, decomposed sub-problems within each interval will be solved.

Let $\mathbf{M}_{p,n}$ be the required minimum exposure for provider p at interval n ($\mathbf{M} \in \mathbb{R}^{|\mathcal{P}| \times N}$), which will be updated at the beginning of interval $n, \forall n = 1, 2, \dots, N$. In the sub-problem, Constraint (3) is decomposed into two constraint ($E_{p,n} \geq \mathbf{M}_{p,n}, \sum_{n=1}^N \mathbf{M}_{p,n} \geq \mathbf{m}_p$). This requires the summation of the required exposure $\mathbf{M}_{p,n}$ at each interval should exceed the final required minimum exposure \mathbf{m}_p . Hence, an optimal $\mathbf{M}_{p,n}$ should be determined to make $f(\phi, \mathbf{m})$ is maximized and the accuracy constraint (2) is satisfied.

3.2 Bankruptcy Problem for Exposure Allocation

In this section, we will first introduce the definition and elements of bankruptcy problem. Then we will formulate the exposure allocation task as a bankruptcy problem.

3.2.1 Bankruptcy Problem. Bankruptcy problem provides a resource allocation solution that is suitable for scenarios with fluctuating resources and demands [1, 44, 56]. The input of a bankruptcy problem can be defined as a triplet: $(\mathcal{A}, E, \mathbf{d})$ and the output can be defined as a result vector \mathbf{x} :

Agent Set \mathcal{A} : the resource demanders which resources need to be allocated to.

Estate E : the total resource waited to be allocated $E \in \mathbb{R}$.

Claim vector \mathbf{d} : the demanded resource of agents $\mathbf{d} \in \mathbb{R}^{|\mathcal{A}|}$.

Result vector \mathbf{x} : the allocation result $\mathbf{x} \in \mathbb{R}^{|\mathcal{A}|}$, where each x_a denotes the allocated resource of agent a . The sum of allocated resources should equal the total resource, i.e., $\sum_{a \in \mathcal{A}} x_a = E$.

To solve the bankruptcy problem, for each agent $a \in \mathcal{A}$, the objective is to satisfy the fluctuating agent claims \mathbf{d} as much as possible when the estate E undergoes fluctuations. In economics, several methods have been developed to solve the bankruptcy problem, including the proportional rule [18], Talmud rule [44], etc.

3.2.2 Minimum Exposure Allocation as Bankruptcy. We correspond the elements of minimum exposure allocation with those in the bankruptcy problem, as shown in Table 1. Specifically,

- Time interval $n \in \mathcal{N}$ is regarded as the agent $a \in \mathcal{A}$ in bankruptcy problem. This is because \mathcal{N} is the set of intervals we want to allocate the required minimum exposures \mathbf{m} to.
- The minimum exposure for providers $\mathbf{m}_p \in \mathbb{R}$ can be regarded as the estate E of the bankruptcy problem, which is the total minimum exposure awaited to be allocated to intervals \mathcal{N} .
- The demanded minimum exposure $D_p \in \mathbb{R}^{|\mathcal{N}|}$ can be regarded as the claim vector \mathbf{d} of the bankruptcy problem, which stands for the demands of exposure in time intervals.
- The predicted minimum exposure vector $\mathbf{M}_p \in \mathbb{R}^{|\mathcal{N}|}$ corresponds to the result vector \mathbf{x} in the bankruptcy problem, where $\mathbf{M}_{p,n}$ denotes the actual minimum exposure provider p should receive during time interval n .

Based on our analysis (Section 4), to maintain a high level of accuracy at every interval, we need to set different demands D_p (i.e., the claim \mathbf{d}) based on fluctuating user traffic in real applications [4, 37]. Meanwhile, the platform can continuously change its provider-fair policy \mathbf{m}_p (i.e., the estate E) due to the development stage or incentive policy [3]. Therefore, due to the beneficial characteristics of dynamic resource allocation, we can formulate the exposure allocation process as a bankruptcy problem.

4 Analysis for Accuracy-Fairness Trade-off

In this section, we will analyze the accuracy loss to answer two questions: (1) Why does provider fairness hurt accuracy? (2) Why does the lower user traffic result in accuracy loss?

4.1 Why Provider Fairness Hurts Accuracy?

From Problem (1), we can draw Theorem 1:

THEOREM 1. Let $\bar{f}(\phi, \mathbf{m})$ be the maximum re-ranking accuracy (i.e., accuracy solution) and $\bar{E}_{p,n}$ be the exposure of provider p at interval n without considering provider fairness. When considering fairness in Equation (1), if \bar{E}_p do not satisfy the constraint, i.e., $\sum_{n=1}^N \bar{E}_{p,n} < \mathbf{m}_p$, then $f(\phi, \mathbf{m}) < \bar{f}(\phi, \mathbf{m})$; otherwise, $f(\phi, \mathbf{m}) = \bar{f}(\phi, \mathbf{m})$.

Proof of Theorem 1 is given in Appendix A.1. In Theorem 1, the accuracy loss brought from provider fairness is through the fairness constraint (i.e., Constraints (3) from Equation (1)). If the accuracy-centric solution fails to meet the provider-fair constraint, the compromise will lead to a solution within the feasible region




Figure 3: A toy example with two providers to illustrate the optimization process at interval i of Equation (1). Provider 1 (x-axis) requires a minimum exposure guarantee of $M_{1,i} = 4$ and provider 2 (y-axis) has no requirement ($M_{2,i} = 0$). The green area is the feasible region constructed by the fairness constraint. (a) Suppose 3 users arrive at interval i and each is recommended 5 items ($3 \times 5 = 15$ total exposures). The grey point denotes the accuracy solution $(\bar{E}_{1,i}, \bar{E}_{2,i})$ and the orange point denotes the provider-fair solution $(E_{1,i}^{\text{fair}}, E_{2,i}^{\text{fair}})$, verifying Theorem 1; (b) Suppose the number of users reduced to 2 at interval $i + 1$ ($2 \times 5 = 10$ total exposures). The red point denotes the provider-fair solution, verifying Theorem 2.

defined by the provider-fair constraint, where accuracy will be compromised.

Figure 3 (a) presents an illustrative toy example to explain the intuition of Theorem 1 geometrically. From Figure 3 (a), the accuracy solution is not in the feasible region. Therefore, to satisfy the fairness constraint, the accuracy solution (the grey point) will move to the provider-fair solution (the orange point), resulting in inevitable accuracy loss. Meanwhile, the moved distance of two points reflects the degree of accuracy loss [8].

In summary, from the perspective of optimization, provider-fair constraint brings the accuracy loss. When the accuracy-maximized solution does not satisfy the fairness constraints, it will inevitably result in a loss of accuracy.

4.2 Why Lower Traffic Cause Accuracy Loss?

Then we will answer this question through Theorem 2

THEOREM 2. When the preferences of users arriving at different times are typically random and independent [22], the expectation of accuracy loss $L = \bar{f}(\phi, \mathbf{m}) - f(\phi, \mathbf{m})$ is negatively correlated to user traffic r_n at each interval n , i.e., $\mathbb{E}[L] \propto \frac{1}{r_n}$.

Proof of Theorem 1 is given in Appendix A.2. Theorem 2 states that lower user traffic will lead to more accuracy loss by affecting the feasible region of the constrained optimization problem.

To better explain Theorem 2, Figure 3 (b) presents the example of holding the same fairness requirement across interval i and $i + 1$ (i.e., $M_{1,i} = M_{1,i+1}$). If intervals i and $i + 1$ share the same fairness constraints and the traffic at interval $i + 1$ drops to two users, the feasible region (green area) transits from Figure 3 (a) to Figure 3 (b). Consequently, the area of the feasible region reduces from 73% to 60% of the area with no fairness constraints (the area formed by the coordinate axes and the black diagonal line). As a result, the provider-fair solution at interval i (represented by the orange point) with an accuracy of $f(\phi, \mathbf{m}) = 1.5$ no longer complies with the

Algorithm 1: Online learning of BankFair

Input: Initial dual variable μ , step size η_t , penalty vector λ , required minimum exposure \mathbf{m}_p , item-provider matrix \mathbf{A} , ranking score $s_{u,i}, \forall u, i$.

Output: The re-ranking decision variable $\{x_t, t = 1, 2, \dots, \infty\}$.

- 1: **for** $n = 1$ to N **do**
- 2: // Module 1.
- 3: $\mathcal{N} = \{n, n+1, \dots, N\}$.
- 4: $\hat{\mathbf{r}}(n) = g([\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_{n-1}])$.
- 5: $D_p(n) = \alpha K \hat{\mathbf{r}}(n)$.
- 6: $\hat{\mathbf{m}}_p(n) = [\hat{\mathbf{m}}_p(n-1) - M_{p,n-1} + \beta_p]_+, \forall p \in \mathcal{P}$.
- 7: $\hat{\mathbf{M}}_p(n) = \text{TAL}(\mathcal{N}, \hat{\mathbf{m}}_p(n), D_p(n)), \forall p \in \mathcal{P}$.
- 8: $\mathbf{M}_n^\top = \hat{\mathbf{M}}(n)_1$.
- 9: // Module 2.
- 10: Initialize dual variable $\mu = 0$.
- 11: Update remaining unearned exposure $\beta_n^\top = \mathbf{M}_n^\top$.
- 12: **for** $t = 1$ to r_n **do**
- 13: $\mathbf{x}_t = \arg \max_{\mathbf{x}_t \in \mathcal{X}} \{s_{u_t}^\top \mathbf{x}_t / \hat{\mathbf{r}}_n - \mu_t^\top \mathbf{A}^\top \mathbf{x}_t\}$
- 14: $E_{n,t}^\top = \arg \max_{E_{n,t}^\top \leq \gamma} \{p_\lambda^*(-\mu_t)\}$
- 15: Update $\beta_n^\top = \beta_n^\top - \mathbf{A}_t^\top \mathbf{x}_t$.
- 16: $\tilde{\mathbf{g}}_t = -\mathbf{A}_t^\top \mathbf{x}_t + E_t$.
- 17: $\mu_{t+1} = \arg \min_{\mu \in \mathcal{D}} \langle \tilde{\mathbf{g}}_t, \mu \rangle + \frac{1}{2\eta} \|\mu - \mu_t\|_w^2$.
- 18: **end for**
- 19: **end for**

fairness constraint $E_{1,i+1} \geq 4$. To meet the constraint, the provider-fair solution at interval $i+1$ must move to the red point, accepting a lower accuracy compared to interval i , as indicated in Theorem 1.

In summary, the re-ranking accuracy loss stems from the impact of user traffic fluctuations under consistent fairness constraints. As the user traffic varies and fairness constraints remain unchanged, it becomes an important factor that affects the fairness feasible region—lower user traffic narrows the region, while higher user traffic enlarges it. Meanwhile, we also observe that the more narrow the fairness feasible region becomes, the larger the accuracy loss will be.

5 Our Approach: BankFair

In this section, we present our re-ranking algorithm BankFair to guarantee both accuracy and fairness. Figure 4 depicts the overall workflow of BankFair, and Algorithm 1 describes the overall procedure. The algorithm consists of the following two modules: (1) *Line 2-8*: module 1, after a new time interval n starts where we will get the predicted minimum exposures \mathbf{M}_n^\top as the result vectors of the bankruptcy problem. (2) *Line 9-18*: module 2, we propose an online learning algorithm to obtain the two-sided re-ranking list $L_K(u_t)$ for every user u_t according to the predicted minimum exposure \mathbf{M}_n^\top at each interval n .

5.1 Module 1: Allocating Required Exposures

In this section, we will get the predicted minimum exposure \mathbf{M}_n^\top at each interval n by transforming the traditional bankruptcy problem into a sequential problem.

In the sequential bankruptcy problem of fair re-ranking, we update the input triplet (upcoming time intervals $\mathcal{N} = \{n, \dots, N\}$, required minimum exposure $\hat{\mathbf{m}}_p(n)$, demanded minimum exposure $D_p(n)$) and get the result vector $\hat{\mathbf{M}}(n)$ at each interval n .

In this module, we first predict $D_p(n)$ in each interval n to be the demand vector using the historical data $\mathcal{H}_{n-1} := \{\mathbf{r}_s, E_s^\top\}_{s=1}^{n-1}$. Then we will use historical data \mathcal{H}_{n-1} to update the remaining required minimum exposure $\hat{\mathbf{m}}_p(n)$ of the bankruptcy problem. Finally, we use the input triplet to get the result vector $\hat{\mathbf{M}}(n)$ and output its first column as the predicted minimum exposure \mathbf{M}_n^\top .

5.1.1 Demand Vector Prediction. According to Theorem 2, we know that when the user traffic is low, adopting a larger minimum exposure will result in a greater accuracy loss. Therefore, to guarantee the accuracy level, we need to decrease the fairness degree (i.e., give a low demand for providers) at lower user traffic levels to reduce the accuracy loss. Otherwise, we need to increase the fairness degree (i.e., give a high demand for providers) to compensate for the low user traffic period, fulfilling the exposure requirements \mathbf{m} . Formally, the demand vector $D_p(n) \in \mathbb{R}^{|\mathcal{N}|}$ of each provider at interval n can be written as:

$$D_p(n) = \alpha K \hat{\mathbf{r}}(n) = \alpha K g([\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_{n-1}]), \quad (4)$$

where $\alpha > 0$ is a proportional coefficient and $\hat{\mathbf{r}}(n) = [\hat{\mathbf{r}}_n, \hat{\mathbf{r}}_{n+1}, \dots, \hat{\mathbf{r}}_N]$ is the predicted user traffic from interval n to N , which is predicted by the time series prediction model $g(\cdot)$ using historical observed data $\mathcal{H}_{n-1} := \{\mathbf{r}_s, E_s^\top\}_{s=1}^{n-1}$. In our experiments, we utilize the well-known time series prediction model Gated Recurrent Unit (GRU) [9], which can be replaced with any time series prediction model [15, 57].

5.1.2 Update Remaining Minimum Exposure. Then, we need to update the remaining required minimum exposure $\hat{\mathbf{m}}_p(n)$ at each interval n as the estate for sequential bankruptcy problem. Specifically, at the beginning of each interval n , we collect the allocation history $\mathcal{H}_{n-1} := \{\mathbf{r}_s, E_s^\top\}_{s=1}^{n-1}$ to calculate the remaining exposures $\hat{\mathbf{m}}_p(n) \in \mathbb{R}^+$ at interval n :

$$\hat{\mathbf{m}}_p(n) = [\hat{\mathbf{m}}_p(n-1) - E_{p,n-1}]_+ = \left[\mathbf{m}_p - \sum_{i=1}^{n-1} E_{p,i} \right]_+, \forall p \in \mathcal{P}, \quad (5)$$

where $[\cdot]_+ \triangleq \max\{0, \cdot\}$.

5.1.3 Obtain Output Vector. After getting the input triplet, we will utilize the Talmud rule [44] in bankruptcy problem to obtain the output vector \mathbf{M}_n^\top by getting the first column (i.e., present interval) of the matrix $\hat{\mathbf{M}}(n)$, i.e., $\mathbf{M}_n^\top = \hat{\mathbf{M}}_1(n)^\top$. $\hat{\mathbf{M}}(n) \in \mathbb{R}^{|\mathcal{P}| \times |\mathcal{N}|}$ can be viewed as the predicted future minimum exposure, which can be calculated by the Talmud rule:

$$\begin{aligned} \hat{\mathbf{M}}_{p,i}(n) &= \text{TAL}_i(\mathcal{N}, \hat{\mathbf{m}}_p(n), D_p(n)) \\ &= \begin{cases} \min\{\frac{D_{p,i}(n)}{2}, \theta\} & \text{if } \hat{\mathbf{m}}_p(n) \leq \frac{1^\top D_p(n)}{2} \\ \max\{\frac{D_{p,i}(n)}{2}, D_{p,i}(n) - \theta\} & \text{if } \hat{\mathbf{m}}_p(n) > \frac{1^\top D_p(n)}{2} \end{cases}, \end{aligned} \quad (6)$$

where θ is the parameter that makes the result vector $\hat{\mathbf{M}}_p(n)$ satisfying $\sum_{i \in \mathcal{N}} \hat{\mathbf{M}}_{p,i}(n) = \hat{\mathbf{m}}_p(n)$.

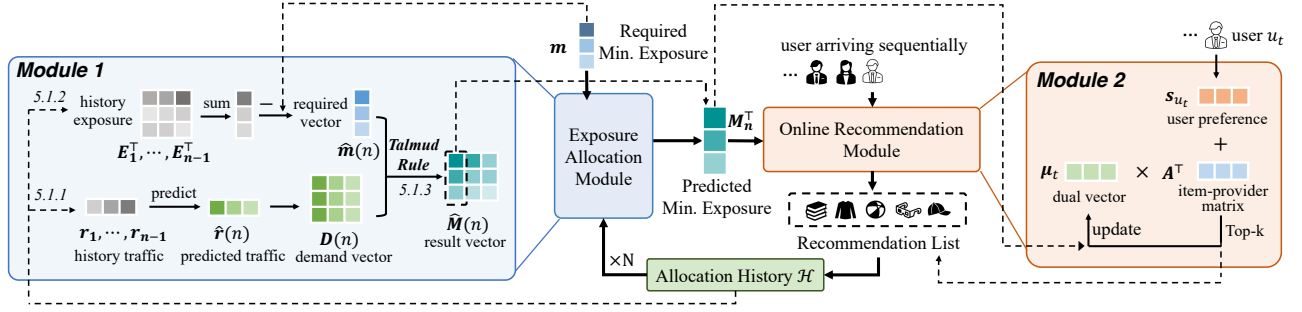


Figure 4: Workflow of the proposed BankFair model

Through the Talmud rule $TAL(\cdot)$, we can allocate the required minimum exposure $\hat{m}_p(n)$ into each interval according to the bankruptcy problem. The Talmud rule can be seen from two conditions: (1) If the minimum exposure $\hat{m}_p(n)$ does not exceed half of the total demands $\mathbf{1}^\top D_p(n)/2$: $\hat{m}_p(n)$ will be distributed equally among each time interval n until $\hat{M}_{p,n}(n)$ meets half of the demand $D_{p,n}(n)/2$. (2) If the minimum exposure $\hat{m}_p(n)$ exceed the total demands $\mathbf{1}^\top D_p(n)/2$: $\hat{m}_p(n)$ will be allocated to equalize the difference between each allocation $\hat{M}_{p,n}(n)$ and its demand $D_{p,n}(n)$.

5.2 Module 2: Online Recommendation

After getting the predicted minimum exposure M_n^\top as the provider fairness constraint at each interval n , we need to get the recommendation list $L_K(u_t)$ for each user u_t .

Following [42, 51], we first use an integer-programming (IP)-based method to formulate the two-sided re-ranking problem in an offline scenario. Then, we adapt the offline algorithm into an online version to adapt the sequential arriving characteristic of users.

5.2.1 BankFair for Offline Scenario. At each interval n , we formulate the two-sided re-ranking task as a resource allocation problem [2]:

$$\max_{\mathbf{x} \in \mathcal{X}} \frac{1}{r_n} \sum_{u_t \in \mathcal{U}_n} s_{u_t}^\top \mathbf{x}_t \quad \text{s.t.} \quad E_n^\top = \sum_t A^\top \mathbf{x}_t \geq M_n^\top, \quad (7)$$

where $\mathbf{x}_t \in \mathcal{X}$ is the decision variable that decides whether item i is recommended to user u_t . Given ranking list length K , $\mathcal{X} = \{\mathbf{x}_t \in \{0, 1\}^{|\mathcal{I}|} | \mathbf{1}^\top \mathbf{x}_t = K\}$. Specifically, for each item i , $x_{t,i} = 1$ if it is added to the re-ranking list $L_K(u_t)$ with size K , otherwise $x_{t,i} = 0$. A is the item-provider adjacent matrix, where $A_{i,p} = 1$ indicates item $i \in \mathcal{I}_p$, and 0 otherwise.

In Equation (7), the objective is to maximize the accuracy while ensuring the fairness constraint. The accuracy $a(n; \mathbf{x})$ denotes the sum of the ranking scores s_{u_t} (see Section 3) for each re-ranking list $L_K(u_t)$ on average. The fairness constraint means that provider p must receive at least $M_{p,n}$ exposures of items corresponding to p (the item-provider correspondence relation can be used A to represent) at interval n . Since we have obtain the traffic-adaptive minimum exposure M_n in the previous section, we are not imposing any minimum constraints on accuracy.

5.2.2 BankFair for Online Scenario. In real online recommendation scenarios, users usually arrive one after the other sequentially (see

Figure 4) and we should give user u_t recommendation list $L_K(u_t)$ immediately. Hence, we develop an online version of the offline algorithm at each interval n .

To efficiently solve Equation (7), we utilize Lagrangian relaxation, transforming Equation (7) into a simpler regularized optimization problem utilizing the following Theorem.

THEOREM 3. The Equation (8) is a relaxed dual problem of Equation (7):

$$W_{Dual} = \min_{\boldsymbol{\mu} \in \mathcal{D}} [h^*(A\boldsymbol{\mu})] + p_\lambda^*(-\boldsymbol{\mu}), \quad (8)$$

where $\mathcal{D} = \{\boldsymbol{\mu} \in \mathbb{R}^{|\mathcal{P}|} \mid \boldsymbol{\mu} \geq -\boldsymbol{\lambda}\}$ is the feasible region of dual variable $\boldsymbol{\mu}$ and

$$h^*(A\boldsymbol{\mu}) = \max_{\mathbf{x}_t \in \mathcal{X}} \sum_{t=1}^{r_n} [s_{u_t}^\top \mathbf{x}_t / r_n - \boldsymbol{\mu}^\top A^\top \mathbf{x}_t],$$

$$p_\lambda^*(-\boldsymbol{\mu}) = \max_{E_n^\top \leq \boldsymbol{y}} \left[-\sum_{i=1}^{|\mathcal{P}|} \lambda_i [M_{n,i}^\top - E_{n,i}^\top]_+ + \boldsymbol{\mu}^\top E_n^\top \right].$$

LEMMA 1. The conjugate function $p_\lambda^*(\cdot)$ has a close form $p_\lambda^*(-\boldsymbol{\mu}) = \boldsymbol{\mu}^\top M_n^\top + \sum_{p=1}^{|\mathcal{P}|} (\gamma_p - M_{p,n}) \max(\boldsymbol{\mu}_p, 0)$, $\forall \boldsymbol{\mu} \in \mathcal{D}$. The optimal dual variable is: $E_{p,n}^*(-\boldsymbol{\mu}) = M_{p,n}$ if $\boldsymbol{\mu}_p \in [-M_{p,n}, 0)$ and $E_{p,n}^*(-\boldsymbol{\mu}) = \gamma_p$ if $\boldsymbol{\mu}_p \geq 0$.

Due to space limitation, we put the proof of Theorem 3 and Lemma 1 in the github repository².

Within each time interval n , the algorithm computes the decision variable \mathbf{x}_t based on $h^*(A\boldsymbol{\mu})$ after a new user u_t arrives. Then, the algorithm updates the auxiliary variable E . Then, we obtain the stochastic subgradient of the dual function and use the weighted, projected subgradient method to update the dual variable $\boldsymbol{\mu}_t$.

6 Experiment

In this section, we conduct experiments on two real datasets. The source codes and experiments have been shared at GitHub².

6.1 Experimental Settings

6.1.1 Datasets. To accurately represent real user traffic, we conducted experiments on two real-world recommendation datasets. The interaction data in both datasets were randomly sampled from actual user traffic in short video applications.

²<https://github.com/shawnye2000/BankFair>

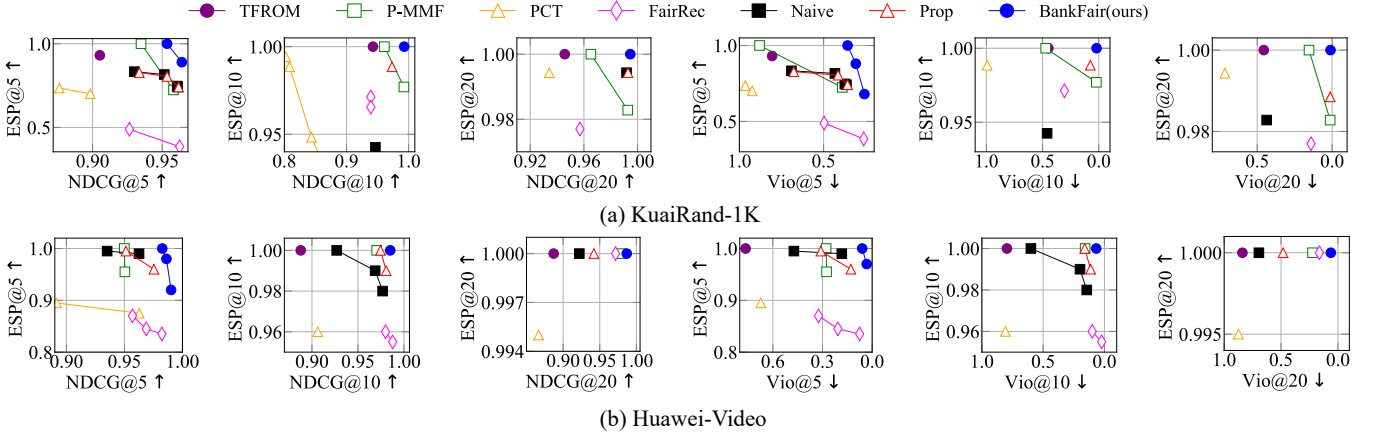


Figure 5: Pareto frontier of two different datasets with different top- K ranking. Y-axis shows ESP@ K metric, while X-axis shows NDCG@ K metric and Vio@ K metric. \uparrow means higher values are better and \downarrow favors lower values. According to the size of the dataset, for KuaiRand-1K, we set $m_p = 1000$ and for Huawei-Video, we set $m_p = 100$. For both datasets, we set $\phi = 0.95$.

KuaiRand-1K³ [20]: a dataset collected from the video-sharing mobile app, *KuaiShou*. The data are all derived from interactions from 8th April 2022 to 8th May 2022 from *KuaiShou*. We use the standard part, which contains randomly sampled user traffic of 302870 interactions from 933 users on 6825 videos with 174 providers.

Huawei-Video: a dataset collected from the short video platform in *Huawei Browser* from 2 Jan. 2024 to 8 Jan. 2024, spanning a week. It contains 19355 users, 5364 items, and 200 providers, totaling 118765 interactions.

For KuaiRand-1K, we use the interactions before 23rd April for training and 23rd April to 8th May for testing. For Huawei-Video, we use interactions before 5th Jan. for training and 6th Jan. to 8th Jan. for testing.

6.1.2 Algorithm Update Setting. We divide the datasets into a training set and a testing set chronologically. The testing set is further divided into N update intervals. To simulate the algorithm update process in real scenarios, we update the base model (e.g., LightGCN [23]) at the beginning of each interval using all the historical interactions accumulated.

6.1.3 Evaluation. For accuracy metrics, following previous practices in [6, 50, 51], we use the normalized discounted cumulative gain (NDCG) to measure the accuracy of recommendation, which is defined as the discounted cumulative relevance score between the re-ranking list $L_K(u_t)$ and the user’s original recommendation $L_K^{\text{ori}}(u_t)$. The average NDCG for all users can be defined as:

$$\text{NDCG@K} = \frac{1}{\sum_{j=1}^N |\mathcal{U}_j|} \sum_{n=1}^N \sum_{u_t \in \mathcal{U}_n} \text{NDCG}_{u_t} @K, \quad (9)$$

where $\text{NDCG}_{u_t} @K = \frac{\sum_{i \in L_K(u_t)} s_{u_t, i} / \log(\text{rank}_i + 1)}{\sum_{i \in L_K^{\text{ori}}(u_t)} s_{u_t, i} / \log(\text{rank}_i^{\text{ori}} + 1)}$, rank_i and $\text{rank}_i^{\text{ori}}$ are the ranking positions of the item i in $L_K(u_t)$ and $L_K^{\text{ori}}(u_t)$.

To evaluate the extent to which a model sacrifices user experience to achieve fairness, we use the proportion of users receiving

accuracy below the minimum requirement ϕ , which is defined as the accuracy violation (Vio):

$$\text{Vio@K} = \frac{1}{\sum_{j=1}^N |\mathcal{U}_j|} \sum_{n=1}^N \sum_{u_t \in \mathcal{U}_n} \mathbb{I}(\text{NDCG}_{u_t} @K < \phi), \quad (10)$$

where $\mathbb{I}(\cdot) = 1$ if the condition (\cdot) is satisfied otherwise 0. If all users satisfies the minimum accuracy requirement, then Vio@ K becomes 0.

For fairness evaluation, we measure the provider group’s satisfaction with their exposure over a period. Following [35, 47, 52], we use the metric of enough satisfaction group (ESP):



$$\text{ESP@K} = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \mathbb{I} \left(\left[\sum_{n=1}^N \sum_{u_t \in \mathcal{U}_n} \sum_{i \in L_K(u_t)} \mathbb{I}(i \in I_p) \right] \geq m_p \right). \quad (11)$$

The value of ESP ranges between 0 and 1. When all the providers get exposures no less than the minimum exposure guarantee m_p , then ESP becomes 1.

6.1.4 Baselines. We compare with the following two-sided fairness-aware baselines: **P-MMF** [51]: An integer-programming (IP)-based re-ranking method ensures the exposure of worst-off provider; **FairRec** [35]: A method guarantees the minimum exposure for providers and employs a greedy strategy to uphold user fairness; **TFROM** [50]: A two-sided re-ranking method which improves provider exposure and amortizes the accuracy loss among users to guarantee accuracy; **PCT** [46]: A max-margin-relevance-based method to ensures a target item exposure distribution and reduce accuracy loss for users.

We also consider two baselines that take a simple allocation rule in Module 1 (See Section 5.1) while maintaining Module 2 as same as ours. **Naive**: This algorithm sets $M_{p,n} = m_p/2$ if the predicted user traffic \hat{r}_n exceeds threshold c and $M_{p,n} = 0$ otherwise. In our experiment, we set c as the mean value of the upcoming predicted user traffic $c = \text{Mean}(\{\hat{r}_i\}_{i=n}^N)$; **Prop**: This algorithm determines the minimum exposure $M_{p,n}$ based on the proportion of predicted

³<https://kuairand.com/>


Figure 6: NDCG@10 and Vio@10 w.r.t. minimum exposure m_p .

Figure 7: (a) Vio@10 under different required minimum accuracy ϕ . (b) The effect of the fluctuation degree τ on Vio@10.

traffic for the current interval n to the overall upcoming predicted traffic: $M_{p,n} = \frac{\hat{r}_n}{\sum_{i=n}^N \hat{r}_i} m_p$.

6.1.5 Implementation Details. For the BankFair models, the proportionality coefficient α was set to $\alpha = \frac{k \sum_{p=1}^{|\mathcal{P}|} m_p}{|\mathcal{P}| \sum_{n=1}^N K r_n}$ and k was tuned among $[1, 2]$. Meanwhile, the bankruptcy problem has a requirement: the demand resource should exceed the estate, therefore, $\sum_{n=1}^N \alpha K \hat{r}_n \geq m_p, \forall p \in \mathcal{P}$. For the violation penalty vector λ , we set $\lambda_p = \beta \frac{\max_p(|I_p|)}{|I_p|} + \frac{1-\beta}{|\mathcal{P}|}$, where $\beta \in [0, 1]$ is a factor controlling the importance on small providers (providers with lower $|I_p|$). Following the practices in [50, 51], we set γ based on the number of items provided by the providers: $\gamma_p = K \hat{r}_n |I_p| / |I|, \forall p \in \mathcal{P}$.

6.2 Experimental Results

In this section, we conduct experiments to test the performance of BankFair and baselines in terms of short-term accuracy and long-term fairness, on KuaiRand-1K and Huawei-Video. In all the experiments, LightGCN [23] is chosen as the base model to generate the relevance score s . We set the algorithm update interval to 24 hours (1 day), which is the most commonly used setting. Considering different dataset sizes, we set $m_p = 1000$ on KuaiRand-1K and $m_p = 100$ on Huawei-Video. We set $\phi = 0.95$ for calculating the Vio@K metric for both datasets.

Figure 5 shows the Pareto frontiers [31] of ESP@K-NDCG@K and ESP@K-Vio@K on two datasets with different ranking sizes K . Note that, for the hyperparameter tuning for each model, we only retain the points on the Pareto front, therefore, in Figure 5, some methods will have only one point, while others have multiple.

From the Pareto frontiers, we find that BankFair Pareto dominates the baselines (i.e., the BankFair points are at the upper right corner), which means BankFair achieves higher accuracy level

(NDCG@K) and guarantee lower short-term accuracy violation (Vio@K) under the same fairness (ESP@K) level. Furthermore, we find the proposed BankFair can satisfy the required minimum exposure for all providers (i.e., ESP=100%) and ensure the minimum accuracy for nearly all users (i.e., Vio \approx 0%) on all datasets under all top-K settings.

Meanwhile, we notice that when ranking list size K is small (e.g., $K = 5$), most baselines fail to meet the short-term accuracy requirement for the majority of users (i.e., maintaining a high Vio@K) while achieving 100% ESP@K. However, BankFair guarantees the minimum accuracy requirements for nearly all users (Vio@K \approx 0%) and satisfies the required minimum exposure for all providers (i.e., ESP@K=100%) simultaneously.

All the experimental results verify that BankFair splendidly guarantees the short-term accuracy requirements and the long-term required minimum exposure of all providers.

6.3 Experimental Analysis

We conducted experiments to analyze BankFair on KuaiRand-1K. In this section, we compare the best-performed baselines in Section 6.2.

6.3.1 Influence of Required Minimum Exposure m . In real applications, the platforms often change the required minimum exposure m to fit different applications (e.g., development stage, incentive policy) [3]. In this experiment, we investigate the performance of BankFair as the required minimum exposure varies $m_p \in [100, 1000], \forall p \in \mathcal{P}$. To make fair comparisons for all selected models, we select the result with the best NDCG@K and Vio@K performance with at least 95% ESP@K under different m_p .

From Vio@10 and NDCG@10 curves in Figure 6, we can see that the proposed BankFair achieves better accuracy (higher NDCG@10 and lower Vio@10) in comparison to the baselines under different required minimum exposure m_p levels. Moreover, the accuracy metric of most baselines (except TFROM, as there are no hyperparameters to adjust) worsens by a large margin (i.e., NDCG@K decreases, Vio@K increases) when the fairness requirement becomes strict (i.e., m_p increases). However, BankFair can maintain relatively stable and high accuracy under different required minimum exposure m_p levels.

6.3.2 Influence of Required Minimum Accuracy ϕ . We conduct experiments to validate the effectiveness of BankFair under different required minimum accuracy ϕ . Sub-figure 7 (a) shows the accuracy violation (Vio@K) performance of BankFair and the best-performed baselines under different required minimum accuracy $\phi \in [0.85, 0.99]$. From Sub-figure 7 (a), we observe that BankFair guarantees a very low Vio@K under any required accuracy level and greatly outperforms all selected baselines. It is worth noting that BankFair maintains low Vio levels even at higher values of ϕ (i.e., more strict accuracy requirement).

6.3.3 Influence of User Traffic Fluctuation Degree τ . In real-world scenarios, user traffic fluctuations can vary significantly, ranging from intense to relatively steady [4, 37]. In this experiment, we validate the effectiveness of our method under varying degrees of traffic fluctuations. To realize different fluctuation degree, we set a temperature coefficient τ with which we re-sample the user traffic r_n at interval n with the probability p_n , where

Table 2: Performances of BankFair and baselines by using BPR and NCF as the base models. The experimental settings are identical to that reported in Section 6.2. * means the improvements over the baseline that achieve ESP=100% are statistically significant (t -tests and p -value < 0.05).

Base model	BPR			NCF		
	NDCG \uparrow	Vio \downarrow	ESP \uparrow	NDCG \uparrow	Vio \downarrow	ESP \uparrow
FairRec	0.9561	0.1746	0.8333	0.9429	0.2408	0.9598
PCT	0.7879	0.9957	0.9828	0.7773	0.9957	0.9828
TFROM	0.9158	0.7142	1.0000	0.9430	0.4589	1.0000
P-MMF	0.9403	0.5679	1.0000	0.9571	0.3009	1.0000
Naive	0.9377	0.6356	0.9310	0.9458	0.4926	0.9713
Prop	0.9654	0.1850	0.9540	0.9758	0.0590	1.0000
BankFair(ours)	0.9806*	0.1179*	1.0000	0.9957*	0.0024*	1.0000
Improv.	1.6% \uparrow	36.3% \downarrow		2.0% \uparrow	95.9% \downarrow	

$p_n = \text{Softmax}([r_1, r_2, \dots, r_n]/\tau)$. Intuitively, the user traffic disparities become substantial when $\tau \rightarrow 0$, and the user traffic tends to become nearly uniform when $\tau \rightarrow 1$.

Sub-figure 7 (b) shows the Vio@10 under varying fluctuation degree τ . The shaded area indicates the 95% confidence intervals of t -distribution under 5 experiments with different random seeds. From Sub-figure 7 (b), we observe that BankFair can achieve lower Vio@K under different degrees of fluctuation, especially when the fluctuation is intense (i.e., $\tau \in [0.0, 0.2]$). Meanwhile, We can observe that as τ decreases (i.e., more intense fluctuations), Vio@K of most baselines, especially Naive, increases greatly, which supports our theoretical analysis in Section 4 that lower traffic strengthens accuracy loss.

6.3.4 Ablation Study on Different Base Models. We conduct experiments to test the performances of BankFair with the two widely used ranking models of BPR [39] and NCF [24] as the base ranking model to generate the relevance score s . We select the result with the best NDCG@K and Vio@K performance with highest ESP@K. As shown in Table 2, we observe that BankFair achieves 100% ESP@K, better accuracy, and lower accuracy violation than baselines under different base models.

7 Conclusion and Future works

This paper emphasizes the significance of guaranteeing both short-term accuracy and long-term fairness under fluctuating user traffic in two-sided recommendation platforms. By formalizing the two-sided re-ranking as a constrained optimization problem, we theoretically analyze the root cause of accuracy loss under fluctuating user traffic. Inspired by the analysis and the Talmud rule in the bankruptcy problem as the exposure allocation principle, we propose a new re-ranking model called BankFair to balance accuracy-fairness well under fluctuating user traffic applications. Experiments on two real-world recommendation datasets show that BankFair can outperform all baselines in terms of both accuracy and fairness. In future work, we will also explore other forms of fairness functions, such as user fairness, under various user traffic scenarios.

Acknowledgments

This work was funded by the National Key R&D Program of China (2023YFA1008704), the National Natural Science Foundation of China (No. 62377044), Beijing Key Laboratory of Big Data Management and Analysis Methods, Major Innovation & Planning Interdisciplinary Platform for the “Double-First Class” Initiative, funds for building world-class universities (disciplines) of Renmin University of China, and PCC@RUC.

A Appendix

A.1 Proof of Theorem 1

PROOF. We first give the formal definition of $\tilde{f}(\phi, \mathbf{m})$, i.e.,

$$\tilde{f}(\phi, \mathbf{m}) = \max \sum_{n=1}^N a(n; E) \quad \text{s.t.} \quad \text{Constraint (2)}.$$

It’s important to note that Constraint (2) will be naturally satisfied without Constraint (3) in the problem $\tilde{f}(\phi, \mathbf{m})$. This is because, according to the re-ranking accuracy definition (see Section 6.1.3), the re-ranking accuracy $a(n; E)$ will be 1 without the fairness constraint, satisfying ϕ naturally.

We know that the difference between problem $\tilde{f}(\phi, \mathbf{m})$ and problem $f(\phi, \mathbf{m})$ is Constraint (3). Therefore, we can conclude that $\tilde{f}(\phi, \mathbf{m}) \geq f(\phi, \mathbf{m})$ always holds, because the feasible region of $\tilde{f}(\phi, \mathbf{m})$ contains the feasible region of $f(\phi, \mathbf{m})$.

We prove Theorem 1 in two condition:

- (1) If $\tilde{f}(\phi, \mathbf{m})$ satisfies Constraint (3), then it falls in the feasible region of $f(\phi, \mathbf{m})$. Also, since it is the solution which maximize the Objective (1), $f(\phi, \mathbf{m}) = \tilde{f}(\phi, \mathbf{m})$;
- (2) If $\tilde{f}(\phi, \mathbf{m})$ does not satisfy Constraint (3), $\tilde{f}(\phi, \mathbf{m}) \neq f(\phi, \mathbf{m})$. Since $\tilde{f}(\phi, \mathbf{m}) \geq f(\phi, \mathbf{m})$, then $\tilde{f}(\mathbf{m}) > f(\mathbf{m})$.

□

A.2 Proof of Theorem 2

PROOF. Let S_{fair} and S denote the areas of feasible regions of $f(\phi, \mathbf{m})$ and $\tilde{f}(\phi, \mathbf{m})$, respectively. First, we can know that $S \geq S_{\text{fair}}$ since the feasible region without Constraint (3) indeed contains the feasible region with Constraint (3). Then, we can compute S_{fair}/S :

$$\frac{S_{\text{fair}}}{S} = \frac{\int_{l=1}^{|\mathcal{P}|} \frac{r_n K - \sum_{p=1}^l M_{p,n}}{r_n K - \sum_{p=1}^{l-1} M_{p,n}}}{\int_{l=1}^{|\mathcal{P}|} \frac{r_n K - \sum_{p=1}^{|\mathcal{P}|} M_{p,n}}{r_n K}} = 1 - \frac{\sum_{p=1}^{|\mathcal{P}|} M_{p,n}}{r_n K}.$$

According to Theorem 1, we know $\tilde{f}(\phi, \mathbf{m}) \geq f(\phi, \mathbf{m}) > 0$ always holds. Let $f(\phi, \mathbf{m}) = k\tilde{f}(\phi, \mathbf{m})$ and $\tau = (1 - k)\tilde{f}(\phi, \mathbf{m})$, where $k \in (0, 1]$. Due to the randomness of the user’s preference, we can write the expectation of accuracy loss $\mathbb{E}[L]$:

$$\begin{aligned} \mathbb{E}[L] &= \mathbb{E}[\tilde{f}(\phi, \mathbf{m}) - f(\phi, \mathbf{m})] = \mathbb{E}[\tau] \\ &= P(k \neq 1) \mathbb{E}[\tau | k \neq 1] + P(k = 1) \mathbb{E}[\tau | k = 1] \\ &= P(k \neq 1) \mathbb{E}[\tau | k \neq 1] = P(\tilde{f}(\phi, \mathbf{m}) > f(\phi, \mathbf{m})) \mathbb{E}[\tau | k \neq 1]. \end{aligned}$$

Here, we use the Monte Carlo method [41] to approximate $P(\tilde{f}(\phi, \mathbf{m}) > f(\phi, \mathbf{m})) \approx 1 - \frac{S_{\text{fair}}}{S}$. Then we can have:

$\mathbb{E}[L] \approx \left(1 - \frac{S_{\text{fair}}}{S}\right) \mathbb{E}[\tau | k \neq 1]$, where $\mathbb{E}[\tau | k \neq 1]$ is irrelevant to the user traffic r_n . Its exact expression is out of the scope of this paper. Therefore, we can conclude that the expectation of accuracy loss L is proportional to $\frac{1}{r_n}$, i.e., $\mathbb{E}[L] \propto 1 - \frac{S_{\text{fair}}}{S} \propto \frac{1}{r_n}$. □

References

- [1] Angelos Antonopoulos. 2020. Bankruptcy problem in network sharing: Fundamentals, applications and challenges. *IEEE Wireless Communications* 27, 4 (2020), 81–87.
- [2] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. 2021. Regularized online allocation problems: Fairness and beyond. In *International Conference on Machine Learning*. PMLR, 630–639.
- [3] Amit Kumar Bardhan and Saad Ashraf. 2022. More buyers or more sellers: on marketing resource allocation strategies of competing two-sided platforms. *Electronic Commerce Research* (2022), 1–30.
- [4] Joeran Beel, Akiko Aizawa, Corinna Breiter, and Bela Gipp. 2017. Mr. DLlib: recommendations-as-a-service (RaaS) for academia. In *2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*. IEEE, 1–2.
- [5] Omer Ben-Porat and Rotem Torkan. 2023. Learning with Exposure Constraints in Recommendation Systems. In *Proceedings of the ACM Web Conference 2023*. 3456–3466.
- [6] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international acm sigir conference on research & development in information retrieval*. 405–414.
- [7] Arpita Biswas, Gourab K Patro, Niloy Ganguly, Krishna P Gummadi, and Abhijnan Chakraborty. 2021. Toward fair recommendation in two-sided platforms. *ACM Transactions on the Web (TWEB)* 16, 2 (2021), 1–34.
- [8] Stephen P Boyd and Lieven Vandenbergh. 2004. *Convex optimization*. Cambridge university press.
- [9] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* (2014).
- [10] Sierra Claridy. 2009. WHAT IS CUSTOMER SERVICE? *Consortium Journal of Hospitality & Tourism* 14, 1 (2009).
- [11] Maxime C Cohen and Renyu Zhang. 2022. Competition and cooptation for two-sided platforms. *Production and Operations Management* 31, 5 (2022), 1997–2014.
- [12] Imma J Curiel, Michael Maschler, and Stef H Tijs. 1987. Bankruptcy games. *Zeitschrift für operations research* 31 (1987), A143–A159.
- [13] Nir Dagan and Oscar Volij. 1993. The bankruptcy problem: a cooperative bargaining approach. *Mathematical Social Sciences* 26, 3 (1993), 287–297.
- [14] Dagmawi Mulugeta Degefu, Weijun He, Liang Yuan, An Min, and Qi Zhang. 2018. Bankruptcy to surplus: Sharing transboundary river basin’s water under scarcity. *Water Resources Management* 32 (2018), 2735–2751.
- [15] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. 2015. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2625–2634.
- [16] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research* 12, 7 (2011).
- [17] Thomas Eisenmann, Geoffrey Parker, Marshall W Van Alstyne, et al. 2006. Strategies for two-sided markets. *Harvard business review* 84, 10 (2006), 92.
- [18] Sinan Ertemel and Rajnish Kumar. 2018. Proportional rules for state contingent claims. *International Journal of Game Theory* 47 (2018), 229–246.
- [19] J Fleming. 2006. Why consistency is the key to profitable customer service. *CUSTOMER MANAGEMENT* 14, 4 (2006), 10.
- [20] Chongming Gao, Shijun Li, Yuan Zhang, Jiawei Chen, Biao Li, Wenqiang Lei, Peng Jiang, and Xiangnan He. 2022. KuaiRand: An Unbiased Sequential Recommendation Dataset with Randomly Exposed Videos. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 3953–3957.
- [21] Yingqiang Ge, Shuyuan Xu, Shuchang Liu, Zuohui Fu, Fei Sun, and Yongfeng Zhang. 2020. Learning personalized risk preferences for recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 409–418.
- [22] Lei Guo, Hongzhi Yin, Qinyong Wang, Tong Chen, Alexander Zhou, and Nguyen Quoc Viet Hung. 2019. Streaming session-based recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 1569–1577.
- [23] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 639–648.
- [24] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.
- [25] Mengdian Hong, Ruoying Li, and Peiran Xie. 2021. Construction of Distribution System of Coronavirus Based on Talmud Bankruptcy Distribution. In *2021 3rd International Conference on Economic Management and Cultural Industry (ICEMCI 2021)*. Atlantis Press, 2089–2097.
- [26] Daniel Kahneman and Amos Tversky. 2013. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*. World Scientific, 99–127.
- [27] Botond Köszegi and Matthew Rabin. 2006. A model of reference-dependent preferences. *The Quarterly Journal of Economics* 121, 4 (2006), 1133–1165.
- [28] Jurek Leonhardt, Avishek Anand, and Megha Khosla. 2018. User fairness in recommender systems. In *Companion Proceedings of the The Web Conference 2018*. 101–102.
- [29] Yunqi Li, Hanxiong Chen, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2021. User-oriented fairness in recommendation. In *Proceedings of the Web Conference 2021*. 624–632.
- [30] Ramon Lopes, Rodrigo Alves, Antoine Ledent, Rodrygo LT Santos, and Marius Kloft. 2024. Recommendations with minimum exposure guarantees: A post-processing framework. *Expert Systems with Applications* 236 (2024), 121164.
- [31] Alexander V Lotov and Kaisa Miettinen. 2008. Visualizing the Pareto frontier. In *Multiojective optimization: interactive and evolutionary approaches*. Springer, 213–243.
- [32] Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. 2020. Controlling fairness and bias in dynamic learning-to-rank. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*. 429–438.
- [33] Mohammadmehdi Naghiaei, Hossein A Rahmani, and Yashar Deldjoo. 2022. Cpfair: Personalized consumer and producer fairness re-ranking for recommender systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 770–779.
- [34] Barry O’Neill. 1982. A problem of rights arbitration from the Talmud. *Mathematical social sciences* 2, 4 (1982), 345–371.
- [35] Gourab K Patro, Arpita Biswas, Niloy Ganguly, Krishna P Gummadi, and Abhijnan Chakraborty. 2020. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *Proceedings of the web conference 2020*. 1194–1204.
- [36] Bibek Paudel, Sandro Luck, and Abraham Bernstein. 2018. Loss aversion in recommender systems: Utilizing negative user preference to improve recommendation quality. *arXiv preprint arXiv:1812.11422* (2018).
- [37] Karine Pires and Gwendal Simon. 2015. YouTube live and Twitch: a tour of user-generated live streaming systems. In *Proceedings of the 6th ACM multimedia systems conference*. 225–230.
- [38] Tao Qi, Fangzhao Wu, Chuhan Wu, Peijie Sun, Le Wu, Xiting Wang, Yongfeng Huang, and Xing Xie. 2022. Profairrec: Provider fairness-aware news recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1164–1173.
- [39] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [40] Jean-Charles Rochet and Jean Tirole. 2004. Two-sided markets: An overview. *Institut d’Economie Industrielle working paper* (2004), 1–44.
- [41] Reuven Y Rubinstein and Dirk P Kroese. 2016. *Simulation and the Monte Carlo method*. John Wiley & Sons.
- [42] Özge Süre, Robin Burke, and Edward C Malthouse. 2018. Multistakeholder recommendation with provider constraints. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 54–62.
- [43] Asep Suryahadi, Wenefrida Widianti, Daniel Perwira, and Sudarno Sumarto. 2003. Minimum wage policy and its impact on employment in the urban formal sector. *Bulletin of Indonesian economic studies* 39, 1 (2003), 29–50.
- [44] William Thomson. 2003. Axiomatic and game-theoretic analysis of bankruptcy and taxation problems: a survey. *Mathematical social sciences* 45, 3 (2003), 249–297.
- [45] William Thomson. 2013. Game-theoretic analysis of bankruptcy and taxation problems: Recent advances. *International Game Theory Review* 15, 03 (2013), 1340018.
- [46] Chenyang Wang, Yankai Liu, Yuanqing Yu, Weizhi Ma, Min Zhang, Yiqun Liu, Haitao Zeng, Junlan Feng, and Chao Deng. 2023. Two-sided Calibration for Quality-aware Responsible Recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 223–233.
- [47] Lequn Wang and Thorsten Joachims. 2023. Uncertainty Quantification for Fairness in Two-Stage Recommender Systems. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. 940–948.
- [48] Yifan Wang, Weizhi Ma, Min Zhang, Yiqun Liu, and Shaoping Ma. 2023. A survey on the fairness of recommender systems. *ACM Transactions on Information Systems* 41, 3 (2023), 1–43.
- [49] Dong Wei, Tao Zhou, Giulio Cimini, Pei Wu, Weiping Liu, and Yi-Cheng Zhang. 2011. Effective mechanism for social recommendation of news. *Physica A: Statistical Mechanics and its Applications* 390, 11 (2011), 2117–2126.
- [50] Yao Wu, Jian Cao, Guandong Xu, and Yudong Tan. 2021. Tfrom: A two-sided fairness-aware recommendation model for both customers and providers. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1013–1022.
- [51] Chen Xu, Sirui Chen, Jun Xu, Weiran Shen, Xiao Zhang, Gang Wang, and Zhenhua Dong. 2023. P-MMF: Provider Max-min Fairness Re-ranking in Recommender System. In *Proceedings of the ACM Web Conference 2023*. 3701–3711.

- [52] Chen Xu, Jun Xu, Yiming Ding, Xiao Zhang, and Qi Qi. 2024. FairSync: Ensuring Amortized Group Exposure in Distributed Recommendation Retrieval. *arXiv preprint arXiv:2402.10628* (2024).
- [53] Chen Xu, Xiaopeng Ye, Wenjie Wang, Liang Pang, Jun Xu, and Tat-Seng Chua. 2024. A Taxation Perspective for Fair Re-ranking. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1494–1503.
- [54] Tao Yang, Zhichao Xu, and Qingyao Ai. 2023. Vertical Allocation-based Fair Exposure Amortizing in Ranking. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*. 234–244.
- [55] Valentina Zanardi and Licia Capra. 2011. Dynamic updating of online recommender systems via feed-forward controllers. In *Proceedings of the 6th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*. 11–19.
- [56] Yang Zheng, Xuefeng Sang, Zhiwu Liu, Siqi Zhang, and Pan Liu. 2022. Water allocation management under scarcity: a bankruptcy approach. *Water Resources Management* 36, 9 (2022), 2891–2912.
- [57] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 11106–11115.