



A Taxation Perspective for Fair Re-ranking

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

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

Wenjie Wang*
wenjiewang96@gmail.com
NExT++ Research Center
National University of Singapore
Singapore

Liang Pang
pangliang@ict.ac.cn
Institute of Computing Technology
Chinese Academy of Sciences
Beijing, China

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

Tat-Seng Chua
dcscts@nus.edu.sg
NExT++ Research Center
National University of Singapore
Singapore

ABSTRACT

Fair re-ranking aims to redistribute ranking slots among items more equitably to ensure responsibility and ethics. The exploration of redistribution problems has a long history in economics, offering valuable insights for conceptualizing fair re-ranking as a taxation process. Such a formulation provides us with a fresh perspective to re-examine fair re-ranking and inspire the development of new methods. From a taxation perspective, we theoretically demonstrate that most previous fair re-ranking methods can be reformulated as an item-level tax policy. Ideally, a good tax policy should be effective and conveniently controllable to adjust ranking resources. However, both empirical and theoretical analyses indicate that the previous item-level tax policy cannot meet two ideal controllable requirements: (1) continuity, ensuring minor changes in tax rates result in small accuracy and fairness shifts; (2) controllability over accuracy loss, ensuring precise estimation of the accuracy loss under a specific tax rate. To overcome these challenges, we introduce a new fair re-ranking method named Tax-rank, which levies taxes based on the difference in utility between two items. Then, we efficiently optimize such an objective by utilizing the Sinkhorn algorithm in optimal transport. Upon a comprehensive analysis, Our model Tax-rank offers a superior tax policy for fair re-ranking, theoretically demonstrating both continuity and controllability over accuracy loss. Experimental results show that Tax-rank outperforms all state-of-the-art baselines on two ranking tasks.

CCS CONCEPTS

• **Information systems** → **Information retrieval**.

*Jun Xu and Wenjie Wang are corresponding authors. Work is partially done at Engineering Research Center of Next-Generation Intelligent Search and Recommendation, Ministry of Education.

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.

SIGIR '24, July 14–18, 2024, Washington, DC, USA

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

ACM ISBN 979-8-4007-0431-4/24/07

<https://doi.org/10.1145/3626772.3657766>

KEYWORDS

Re-ranking, Item Fairness, Taxation Process

ACM Reference Format:

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 (SIGIR '24), July 14–18, 2024, Washington, DC, USA*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3626772.3657766>

1 INTRODUCTION

Ranking tasks such as recommendation and advertising is important in personalized information filtering for users [51]. However, recent studies [12, 35] have revealed that prioritizing ranking accuracy will lead to significant unfair exposures of items, threatening the ethics and stability of ranking systems. Recently, many fair re-ranking methods [12, 35] have been developed to redistribute ranking slots among items more equitably by transforming the problem into a resource redistribution problem [49].

The investigation of resource redistribution has a long history in economics [19, 33], offering economic insights for examining the task of fair re-ranking. For example, as shown in the left part of Figure 1 (a), federal taxes are structured to redistribute subsidies to the poor while aiming to minimize the sacrifice of social welfare. Similarly, as shown in the right part of Figure 1 (a), the goal of fair re-ranking is to redistribute items, aiming for a more equitable distribution of ranking slots, without significantly compromising ranking accuracy. Since both taxing and fair re-ranking are essentially redistributing resources, we can examine fair re-ranking from the view of taxation. Intuitively, fair re-ranking can be viewed as taxing exposure from “rich” items (with more exposure) and redirects it to “poor” items (have less exposure) in Figure 1 (a).

Viewing fair re-ranking as a taxation process provides a fresh perspective to re-examine previous fair re-ranking methods and inspire new approaches. From a taxation perspective, we can theoretically show that most previous fair re-ranking methods [15, 31, 35, 48, 49] can be reformulated as an item-level tax policy, which imposes an additional tax on each item (see Section 4.1). In economics, a good tax policy should be effective and conveniently controllable to adjust ranking resources [19, 33]. However, we observe that the item-level tax policy derived from previous fair re-ranking methods

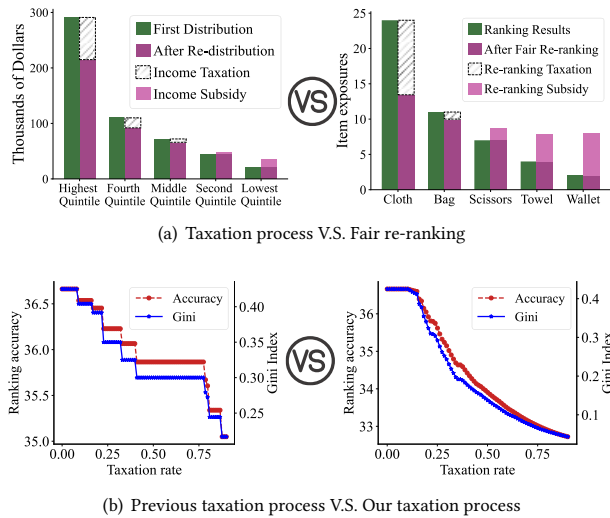


Figure 1: Sub-figure (a) illustrates that fair re-ranking can be viewed as a taxation process from an economic perspective¹. Sub-figure (b) illustrates that our taxation process exhibits better continuity compared to previous ones.

fails to meet these two ideal controllability criteria for taxation: (1) Continuity, implying that slight variations in tax rates lead to minor shifts in performances. As shown in the left part of Figure 1 (b), When tax rates vary, previous methods exhibit numerous breakpoints in fairness and accuracy performance, indicating poor continuity. (2) Controllability over accuracy loss, ensuring an accurate estimation of accuracy loss caused by a specific tax rate. However, previous methods are uncertain about such accuracy loss (refer to Section 4.2 for a more detailed theoretical analysis). These two requirements are important for a tax policy since continuity ensures the stability of the tax rate, while controllability over accuracy loss ensures better management of the trade-off between fairness and accuracy.

To meet the two aforementioned essential criteria for the taxation process for fair re-ranking, in this paper, we introduce a novel fair ranking method, termed Tax-rank. Differing from the item-level tax policy, Tax-rank introduces a distinct fair re-ranking optimization objective, which levies taxes based on the difference in utility (e.g., exposure) between two items. A detailed geometric explanation of such a taxation process is in Section 5.2.1. Upon a comprehensive analysis, our taxation process presents two key advantages: (1) Tax-rank demonstrates greater effectiveness as it adheres to continuity *w.r.t.* tax rates, as depicted in the right part of Figure 1 (b). (2) Tax-rank offers enhanced controllability over accuracy loss, as we can provide an upper bound, showcasing the maximum accuracy loss across different tax rates. More detailed theoretical analyses are in Section 5.2.2.

Meanwhile, we propose an effective algorithm to implement the taxation process of Tax-rank efficiently. We first present an easily solvable lower-bound function for the optimization objective of

Tax-rank. However, the solution of the lower-bound function does not satisfy the ranking constraint, which requires fixed-sized items (i.e. Top-K ranking) for each user. To solve this issue, we utilize the Sinkhorn algorithm [43] of optimal transport [37, 38] to project non-compliant solutions onto solutions that satisfy the constraint.

We summarize the major contributions of this paper as follows:

(1) We re-conceptualize the fair re-ranking task as a taxation process and re-examine the previous fair re-ranking methods from a novel taxation perspective.

(2) We introduce a novel fair re-ranking method named Tax-rank, incorporating a new optimization objective based on the taxation process and utilizing the Sinkhorn algorithm for efficient optimization. Theoretical evidence indicates that Tax-rank exhibits superior continuity and systematic controllability.

(3) We conduct extensive experiments on two publicly available recommendation and advertising datasets, demonstrating that Tax-rank outperforms state-of-the-art baselines in terms of both fair ranking performance and efficiency.

2 RELATED WORK

Recently, fair re-ranking tasks have become a compelling and pressing issue, driven by the need for a responsible and trustworthy ecosystem [12, 25, 27, 36]. Fairness concept in re-ranking varies for multi-stakeholders settings [1], such as user-oriented fairness [3, 24] and item-oriented fairness [35, 42, 47, 49]. The concept of user fairness suggests that ranking models should avoid delivering significantly disparate outcomes to users based on their sensitive attributes [24, 30, 46]. On the other hand, item fairness is more related to distributive justice [22, 45], which requires an equitable reallocation of ranking results for items within a healthy ecosystem. In this paper, our primary focus is on applying an economic perspective to reconsider item-fair re-ranking tasks.

Regarding item fair re-ranking, previous work can be divided into two types: one is regularized methods, which used a multi-task optimization approach with a linear combination of accuracy and fairness loss functions, incorporating a trade-off coefficient λ [15, 31, 49]. Another approach employed constraint-based methods, formulating the task as a constrained optimization problem to ensure that fairness metrics do not exceed a specified threshold [9, 35, 48, 53]. For measuring fairness, both of the methods utilize different fairness metrics: Ben-Porat and Tennenholtz [5], Biswas et al. [9], Patro et al. [35] proposed the Shapley value to optimize fairness; Do and Usunier [16] proposed to optimize Gini Index [16], Do et al. [14], Xu et al. [49] proposed to optimize max-min fairness function and Jiang et al. [20] suggested to use the variance (distance) of different item utilities to measure the item fairness. Despite achieving significant performances, existing methods show poor continuity and systematic controllability under a taxation perspective for fair re-ranking.

In the economic field, the resource is usually allocated through first distribution and re-distribution process [21]. In the process of redistribution, taxation is frequently employed as a mechanism to redistribute wealth and tackle income inequality [19, 33]. There are usually two types of taxation methods: (1) flat tax rates, such as property tax [34], which are designed with different fixed rates based on the varying amounts of property; (2) progressive tax, such

¹Federal taxes of 2016 in the U.S. <https://www.cbo.gov>.

as income taxes [39] and payroll taxes [10], typically involve progressive tax rates that increase as a taxpayer's earnings increase. Process tax often is regarded as a more useful but complex taxation method. Previous tax policies resemble a flat tax, whereas Tax-rank adopts a closer resemblance to a progressive tax. Finally, the Tax-rank policy is more closely related to α -fair optimization in cooperative games [7, 8]. However, these approaches are not suitable for fair re-ranking tasks.

3 PROBLEM FORMULATION

We first define some notations for the problem. For vector $\mathbf{x} \in \mathbb{R}^n$, let x_i denote the i -th element of the vector. For vector $\mathbf{x} \in \mathbb{R}^{n \times m}$, let $x_{i,j}$ denote the element of i -th row and j -th column.

In this section, we will first formulate the ranking task. Let \mathcal{U} representing the set of users, and \mathcal{I} representing the set of items. When a user $u \in \mathcal{U}$ arrives, the ranking system will give the user u a ranking list $L_K(u) \in \mathcal{I}^K$, which has fixed-size K items.

Then, we define the item utility in the ranking task. Let v_i be the utility of item i . In the context of ranking, v_i is typically defined as the accumulated utilities of item i across all ranking lists when $|\mathcal{U}|$ users arrive (e.g., v_i could be total exposure or click numbers of item i within a day). Formally, v_i is defined as $v_i = \sum_{u \in \mathcal{U}} w_{u,i} x_{u,i}$, where $x_{u,i} = 1$ denotes item i is in the ranking list of user u (i.e., $i \in L_K(u)$), otherwise, $x_{u,i} = 0$. If $x_{u,i} = 1$, the item i will receive an utility $w_{u,i}$. Typically, $w_{u,i}$ has two possible definitions: (1) item exposure [35, 49], which implies that if an item is exposed to a user once, it will gain one unit of utility (i.e., $w_{u,i} = 1$). (2) item click [28, 52], which implies that if an item is clicked by a user, it will gain one unit of utility. Since we do not know whether a user will click an exposed item, we utilize the estimated probability of a user clicking on an item (i.e., click-through-rate (CTR) value) as the value of $w_{u,i} \in [0, 1]$.

Finally, the goal of fair re-ranking f is to balance the various utilities of items more equitably. That is, on the one hand, f aims to maximize the weighted sum of their utilities ($\sum_{i \in \mathcal{I}} \gamma_i v_i$), where γ_i is the item weight or bidding value [52]. On the other hand, f aims to achieve as even utilities ($v_i \approx v_j, \forall i, j \in \mathcal{I}$) as possible.

4 TAXATION PERSPECTIVE FOR FAIR RE-RANKING

In this section, we utilize the taxation perspective to re-think the fair re-ranking problem.

4.1 Fair Re-ranking as Tax Policy

Firstly, the objectives of the previous fair re-ranking can be mainly divided into either regularization-based optimization [15, 31, 35, 49] or constraint-based optimization [9, 35, 48, 53]. Specifically, the objectives of regularization-based and constraint-based are formulated in Equation (1) and Equation (2), respectively:

$$W_1(\mathbf{x}) = \max_{\mathbf{x} \in \mathcal{X}} g(\mathbf{v}; \lambda_1) = \sum_{i \in \mathcal{I}} \gamma_i \sum_{u \in \mathcal{U}} w_{u,i} x_{u,i} + \lambda r(\mathbf{v}), \quad (1)$$

$$\begin{aligned} W_2(\mathbf{x}) &= \max_{\mathbf{x} \in \mathcal{X}} g(\mathbf{v}; \pi) = \sum_{i \in \mathcal{I}} \gamma_i \sum_{u \in \mathcal{U}} w_{u,i} x_{u,i}, \\ \text{s.t. } & r(\mathbf{v}) \leq \pi \end{aligned} \quad (2)$$

where $\mathbf{v}_i = \sum_{u \in \mathcal{U}} w_{u,i} x_{u,i}$, $\mathcal{X} = \{\mathbf{x} | x_{u,i} = 0/1, \sum_{i \in \mathcal{I}} x_{u,i} = K\}$ is the feasible region of variable \mathbf{x} , $\lambda_1 \in [0, \infty]$ is a trade-off coefficient, and π is the pre-defined fairness threshold. The function $r(\mathbf{v})$ is the fairness function, which aims to measure the disparity of different items' utilities. Note that in previous literature, $r(\mathbf{v})$ has many forms: $\min_{i \in \mathcal{I}} v_i$ [9, 35, 49], $-\sum_i |v_i - \bar{v}|$ [31, 44] and any other forms [25].

Then, in Theorem 1, we will illustrate that both of the objectives can be interpreted as a taxation process, which imposes an additional tax rate λ and a taxation value μ_i on each item i .

THEOREM 1. *Then optimal fair re-ranking result $\mathbf{x}(\lambda)$ with specific tax rate λ can be achieved as:*

$$\mathbf{x}^*(\lambda) = \arg \max_{\mathbf{x} \in \mathcal{X}} \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} s_{u,i} x_{u,i}, \quad (3)$$

where $s_{u,i} = \gamma_i w_{u,i} + \mu_i$, and

$$\mu = \arg \min_{\mu} \left[\sum_{u \in \mathcal{U}} \sum_{k=1}^K s_{u,[k]} + \max_{\mathbf{v}} (\lambda r(\mathbf{v}) - \mu^T \mathbf{v}) \right].$$

In simpler terms, the ranking score is like the original score $\gamma_i w_{u,i}$ but with the addition of item-level taxation μ_i with the tax rate of $\lambda \in [0, \infty]$. The tax rate λ is calculated as

$$\lambda = \begin{cases} \lambda_1, & \text{if } W_1(\mathbf{x}) \\ \arg \min_{\lambda_2} \left[\sum_{u \in \mathcal{U}} \sum_{k=1}^K c_{u,[k]} + r(\mathbf{v}) \lambda_2 - \lambda_2 \pi \right], & \text{if } W_2(\mathbf{x}) \end{cases},$$

where $c_{u,i} = \gamma_i w_{u,i}$.

The proof of the Theorem 1 can be seen in Appendix A. From Theorem 1, we can see that previous fair re-ranking methods can be regarded as an item-level tax policy with the taxation value as μ_i for ranking score $s_{u,i}$. Meanwhile, we find that the tax policy is a well-studied knapsack problem [41], where we can only choose the top K items with the highest scores $s_{u,i}$.

4.2 Re-examine Fair Re-ranking

Through examination, we theoretically demonstrate that previous item-level tax policies of fair re-ranking lack continuity and controllability over accuracy loss.

4.2.1 Non-continuity. In Theorem 2, we will establish the non Lipschitz continuous [18] of the optimization objective for the existing taxation w.r.t. the tax rate λ :

THEOREM 2. *When $r(\mathbf{v})$ is not continuous, $W_1(\mathbf{x}, \lambda)$ and $W_2(\mathbf{x}, \lambda)$ is non Lipschitz continuous wrt the tax rate λ . Formally, $\exists \epsilon > 0, \forall \delta > 0$, when $|\lambda_0 - \lambda| < \delta$, $|W_*(\mathbf{x}, \lambda) - W_*(\mathbf{x}, \lambda_0)| \geq \epsilon$, where $*$ = 1, 2.*

The detailed proof of Theorem 2 can be seen in Appendix B. Intuitively, the binary variable $x_{u,i}$ under the item-level policy and the non-differentiable property of some $r(\mathbf{v})$ lead to the non-continuity. The Theorem 2 tells us a slight change in tax rates λ leads to significant shifts in fair re-ranking performances.

4.2.2 Non-controllability over accuracy loss. we will establish the definition of the price of taxation (POT) [7] as the maximum ranking accuracy loss across various taxation levels, formulated as the ratio of ranking accuracy loss to the maximum accuracy:

$$\text{POT}(\lambda) = \frac{\text{Acc}'(0) - \text{Acc}'(\lambda)}{\text{Acc}'(0)}, \quad (4)$$

where $\text{Acc}'(\lambda)$ denotes the ranking accuracy (here, accuracy refers to the overall user-item ranking scores that these items are exposed to by users) under existing tax rate λ , *i.e.*,

$$\text{Acc}'(\lambda) = \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} \gamma_i w_{u,i} x^*(\lambda)_{u,i},$$

and $x^*(\lambda)$ is the optimal result with specific tax rate λ . Upon examining Equation (3), it becomes evident that the relationship between the ranking result $x^*(\lambda)$ and the tax rate is non-linear and non-continuous (refer to Theorem 2 and the non-linear of the fairness function $r(v)$ [35, 49]). Simultaneously, it is noteworthy that Equation (3) represents a well-studied integral knapsack problem [41], and obtaining an upper bound for such a non-linear and non-continuous knapsack problem remains a huge challenge [11]. Therefore, we are uncertain about the extent to which a specific tax rate will incur accuracy losses.

4.3 Insights from Taxation Perspective

Through the lens of taxation, the fair re-ranking mechanism operates as a dynamic adjustment tool, adapting to various systematic objectives during both periods of economic prosperity and downturn. During periods of prosperity in the ranking system, there is typically a surge in user traffic, and providers are willing to offer a greater variety of items. However, the increased competition during such periods may result in unfairness, as it can lead to a long-tail effect and the monopolization of specific items [2]. During such periods, strategically increasing the tax rate can enhance fairness, fostering a healthier ecosystem [49]. On the other hand, during economic downturns in the ranking system, user traffic tends to decrease, and providers may exhibit reduced enthusiasm for participation. In such a period, appropriately reducing tax rates can incentivize provider competition, attracting more users to the system and ensuring economic vitality. Therefore, a good fair re-ranking method not only needs to be effective but also provide a conveniently controllable way to select an appropriate hyper-parameter (*e.g.*, tax rate) for regulating fairness degree.

5 TAX-RANK

To overcome the previous tax policy issues, we propose an improved fair re-ranking method named Tax-rank.

5.1 Optimization Objective

In this section, we will first introduce the proposed taxation optimization objective of Tax-rank. Drawing inspiration from the welfare function [8], we levy taxes based on the difference in utility between two items: tax the higher utility of item i and redirect to the item j (*i.e.*, $v_i > v_j$) with the value as $\frac{\gamma_i}{\gamma_j} (\frac{v_i}{v_j})^t$. Then the overall optimization objective will be

$$\begin{aligned} \mathbf{x}^*(t) = \arg \max_{\mathbf{x} \in \mathcal{X}_s} f(\mathbf{x}; t) = & \begin{cases} \sum_i \gamma_i v_i^{1-t} / (1-t) & \text{if } t \geq 0, t \neq 1 \\ \sum_i \gamma_i \log(v_i) & \text{if } t = 1 \end{cases}, \quad (5) \\ \text{s.t. } v_i = & \sum_{u \in \mathcal{U}} w_{u,i} x_{u,i}, \quad \forall i \in \mathcal{I} \end{aligned}$$

where $\mathcal{X}_s = \{\mathbf{x} | x_{u,i} \in [0, 1], \sum_{i \in \mathcal{I}} x_{u,i} = K, \forall u \in \mathcal{U}\}$ is the feasible region of variable \mathbf{x} . Intuitively, to keep the social welfare function

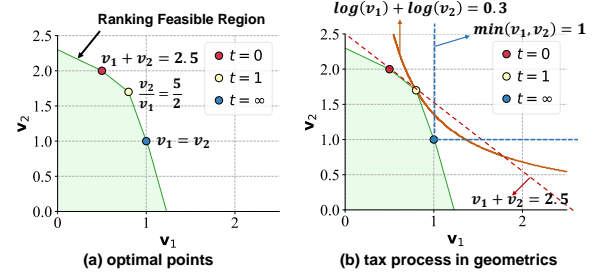


Figure 2: Geometric explanation for our taxation process, which imposes taxes based on between two items.

$f(\mathbf{x}; t)$ unchanged, for every additional unit of utility gained by the “poor” item j , the “rich” item i has to pay the tax value of $\frac{\gamma_i}{\gamma_j} (\frac{v_i}{v_j})^t$ with their weight γ_i and γ_j .

Note that, Tax-rank introduces a relaxation of the binary solution for \mathbf{x} , transitioning it into a continuous version. Here, x_t can be interpreted as the recommendation probability assigned to each item, allowing us to leverage a multi-nominal sample technique to generate ranking lists. $t \in [0, \infty)$ represents the tax rate, where a higher value results in a more fair results.

When $t = 0$, the optimization objective reduces to the accuracy-first objective; When tax rate $t = 1$, the objective can be reduced to Nash solution [32], where the tax rate will make the item utilities proportional to its weight, *i.e.*, $v_i : v_j = \gamma_i : \gamma_j, \forall i, j \in \mathcal{I}$. When t approaches ∞ , the optimization objective reduces to max-min form $\min_{i \in \mathcal{I}} v_i$, leading to even utility for every item.

5.2 Optimization Objective Analysis

In this section, we will give a formal analysis of our optimization objective in a geometric view and its controllable requirements.

5.2.1 Geometric explanation for our taxation process. To gain a better understanding of Tax-rank, we visualize the geometric to show our taxation process for taxing the difference in utility between two items. In Figure 2, the ranking system only has two items and ranking size $K = 1$. The x-axis v_1 and y-axis v_2 represent the expected click numbers of items 1 and 2 (see Section 3). We set their weight $\gamma_1 = \gamma_2 = 1$. The green area describes the feasible region of all the possible values of v_1 and v_2 under ranking constraint. The circle points are optimal solutions under different tax rates t . The line in Figure 2 (b) represents the tangent between the optimization objective and the feasible region.

In Figure 2 (a), we observe that when $t = 0$ (red points), prioritizing overall utility $v_1 + v_2 = 2.5$; at $t = 1$ (yellow points), the tax rate will make the item utilities proportional to its weight, *i.e.*, $v_1 : v_2 = 2 : 5$. At $t = \infty$ (blue points), the two items will share the same utility, *i.e.*, $v_1 = v_2$. In Figure 2 (b), we give a geometric explanation to show the optimal points change *w.r.t.* tax rate. From Figure 2, the optimal points must stay in the boundary lines (*i.e.*, Pareto frontier) of the ranking feasible region (see detailed proof in [7]). Therefore, the optimal points should be the tangent points between the contours of the optimization objective and the Pareto front. Let’s consider the taxation process in two-dimensional space.

For example, in Figure 2 (b), when $t = 0$, the red point is the tangent point between the red line $v_1 + v_2 = 2.5$ and the Pareto front. When $t = 1$, a tax of $\frac{v_2}{v_1}$ is applied to from v_2 to v_1 along the red line if $v_2 > v_1$, which means the slope $\frac{\partial v_2}{\partial v_1}$ becomes $\frac{v_1}{v_2}$. When $v_1 > v_2$, similar operation also make the slope $\frac{\partial v_2}{\partial v_1}$ becomes $\frac{v_1}{v_2}$. Then the taxation process leads to the optimization function of $\log v_1 + \log v_2$ (orange line). Similarly, $t = \infty$, a tax of $(\frac{v_2}{v_1})^\infty$ is applied at each point on the red line, resulting in a slope $\frac{\partial v_2}{\partial v_1}$ of each point become ∞ and 0 for each point when $v_2 > v_1$ and $v_1 > v_2$, respectively. Such taxation process leads to the objective becomes $\min(v_1, v_2)$ (blue line). Without loss of generality, tax the value as $\frac{Y_i}{Y_j} (\frac{v_i}{v_j})^t$ for every $v_i > v_j$ will transform the objective into the Equation (5).

In summary, the tax process of imposing taxes based on the disparity in utility between two items will determine the optimization objective of Equation (5) by influencing the slope of each point in the geometric perspective.

5.2.2 Controllable requirements. we demonstrate the optimization objective meets two ideal controllable requirements: continuous and controllability over the accuracy loss.

Continuity. Firstly, the Tax-rank policy $f(\mathbf{x}; t)$ adheres to the Lipschitz continuity property [18] of the optimization objective associated with the existing taxation concerning the tax rate t . Because we can easily observe that $f(\mathbf{x}; t)$ takes an exponential form w.r.t. t , and their first-order derivatives are also continuous [23]. Meanwhile, transforming the binary solution of previous methods to a continuous solution also ensures feasible region is continuous.

Controllability over accuracy loss. For the bound the maximum ranking accuracy loss across different taxation degrees under varying tax rates t , we will have the following Theorem:

THEOREM 3. *The price of taxation (POT) of Tax-rank is bounded:*

$$POT = \frac{\text{Acc}(0) - \text{Acc}(t)}{\text{Acc}(0)} \leq 1 - O(|\mathcal{U}|^{-\frac{t}{1+t}}), \quad (6)$$

where $\text{Acc}(t)$ denotes the accuracy under Tax-rank tax policy with tax rate t , i.e., $\text{Acc}(t) = \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} \gamma_i w_{u,i} \mathbf{x}^*(t)_{u,i}$, $\mathbf{x}^*(t)$ is the optimal fair re-ranking result with specific tax rate t in Equation (5).

Detailed proof can be seen in Appendix C. As we can see from Theorem 3: when increasing the tax rate t in a ranking system, there is a bound on the rate $1 - O(|\mathcal{U}|^{-\frac{t}{1+t}})$ at which ranking utilities will decrease. The bound offers fair re-ranking systems a controllable way to determine the appropriate tax rate t .

5.3 Algorithm

The overall algorithm workflow can be seen in Algorithm 1. We observe that directly optimizing the Equation (5) is NP-hard since it is a non-linear, large-scale, and integral programming [6]. Therefore, firstly, we construct an easy-solved standard programming (lines 1-2), which is the lower bound function of Equation (5). Then we apply the OT projection to get the final result (lines 3-9).

5.3.1 Lower Bound Function Construction.

Algorithm 1: Learning algorithm of Tax-rank

Input: User set \mathcal{U} , item set \mathcal{I} , ranking size K , tax rate t , OT coefficient λ_{ot} , item weight $\gamma_i, \forall i \in \mathcal{I}$, user-item ranking score $w_{u,i}, \forall u \in \mathcal{U}, \forall i \in \mathcal{I}$

Output: The ranking result $L_K(u), \forall u \in \mathcal{U}$

1: Get the optimal averaged exposure \mathbf{e}^* from Equation (7).

2: Initialize $\mathbf{m} = K\mathbf{1}, \mathbf{n} = \mathbf{e}^*$,

$$C_{u,i} = \gamma_i w_{u,i}, \forall u \in \mathcal{U}, \forall i \in \mathcal{I}, \mathbf{B} = e^{-\frac{C}{\lambda_{ot}}}$$

3: **for** $t = 1, \dots, T$ **do**

4: $\mathbf{m} = K\mathbf{1} \oslash \mathbf{Bn}, \mathbf{n} = \mathbf{e}^* \oslash \mathbf{Bm}$

5: **end for**

6: $\tilde{\mathbf{x}} = \text{diag}(\mathbf{m})\mathbf{B}\text{diag}(\mathbf{n})$

7: $L_K(u) = \text{Sample}_{S \subset \{1,2,\dots,|\mathcal{I}|\}, |S|=K} \tilde{\mathbf{x}}_u, \quad \forall u \in \mathcal{U}$

LEMMA 1. *There exists $\tau > 0$, s.t. we have*

$$f(\mathbf{x}; t) \geq \tilde{f}(\mathbf{x}; t) = \max_{\mathbf{e}} \sum_i \gamma_i \eta_i g(\mathbf{e}; t)$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{I}} \mathbf{e}_i = K, \quad 0 \leq \mathbf{e}_i \leq 1, \quad \eta_i = \tau \sum_{u \in \mathcal{U}} w_{u,i}, \forall i \in \mathcal{I} \quad (7)$$

$$g(\mathbf{e}; t) = \begin{cases} \sum_i \mathbf{e}_i^{1-t} / (1-t) & \text{if } t \geq 0, t \neq 1 \\ \sum_i \log(\mathbf{e}_i) & \text{if } t = 1 \end{cases},$$

This lemma can be easily achieved when $\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} w_{u,i} \mathbf{x}_{u,i} \geq \sum_{i \in \mathcal{I}} \eta_i \sum_{u \in \mathcal{U}} \mathbf{x}_{u,i}$. Let \mathbf{e}^* be the optimal value of the variable \mathbf{e} , which represents the accumulated exposure of items over all users. Then we will apply the Sinkhorn algorithm [38] to project the averaged exposure \mathbf{e}^* to recommendation decision variable $\mathbf{x} \in \mathcal{X}_S$ discussed in Section 3.

5.3.2 Optimal Transport Projection. We obtain the final ranking result by utilizing the following sample process, where $\tilde{\mathbf{x}}$ (i.e. ranking score distribution) is derived from the OT projection process.

$$L_K(u) = \text{Sample}_{S \subset \{1,2,\dots,|\mathcal{I}|\}, |S|=K} \tilde{\mathbf{x}}_u, \quad \forall u \in \mathcal{U}, \quad (8)$$

where it implies sampling K non-repeated items to user u according to the recommended probability $\tilde{\mathbf{x}}_{u,i} \in [0, 1]$ of item i .

We construct a matrix $\mathbf{C} = \mathbb{R}^{|\mathcal{U}| \times |\mathcal{I}|}$, where the element $C_{u,i} = \gamma_i w_{u,i}$. An OT problem can be formulated as:

$$\tilde{\mathbf{x}} = \arg \min_{\mathbf{x} \geq 0} \langle \mathbf{x}, -\mathbf{C} \rangle + \lambda_{ot} H(\mathbf{x}) \quad \text{s.t.} \quad \mathbf{x}\mathbf{1} = K\mathbf{1}, \quad \mathbf{1}^\top \mathbf{x} = \mathbf{e}^*, \quad (9)$$

where $\mathbf{1}$ denotes a vector of ones, \mathbf{e}^* denotes the optimal value of Equation (7) and λ_{ot} is the coefficient of entropy regularizer. $\langle \mathbf{x}, -\mathbf{C} \rangle$ results transport plan lies on the Pareto frontier and $H(\mathbf{x}) = \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} \mathbf{x}_{u,i} \log(\mathbf{x}_{u,i})$, which forces the variable $\mathbf{x}_{u,i}$ into the feasible region $[0, 1]$. The constraint condition ensures that the ranking satisfies the limitation that each user can only be ranked among the top K items, and it also guarantees that the exposure of each item aligns optimally with the predefined exposure vector \mathbf{e}^* .

This problem can be efficiently solved by the Sinkhorn algorithm [43], where the solution of the form $\tilde{\mathbf{x}} = \text{diag}(\mathbf{m})\mathbf{B}\text{diag}(\mathbf{n})$, where $\text{diag}(\cdot)$ denote the generating diagonal matrix from vector, $\mathbf{B} = e^{-\frac{\mathbf{C}}{\lambda_{ot}}}$, and $\mathbf{m} \in \mathbb{R}^{|\mathcal{U}|}, \mathbf{n} \in \mathbb{R}^{|\mathcal{I}|}$, which iteratively computes $\mathbf{m} \leftarrow K\mathbf{1} \oslash \mathbf{Bn}, \quad \mathbf{n} \leftarrow \mathbf{e}^* \oslash \mathbf{Bm}$, where \oslash denotes element-wise division.

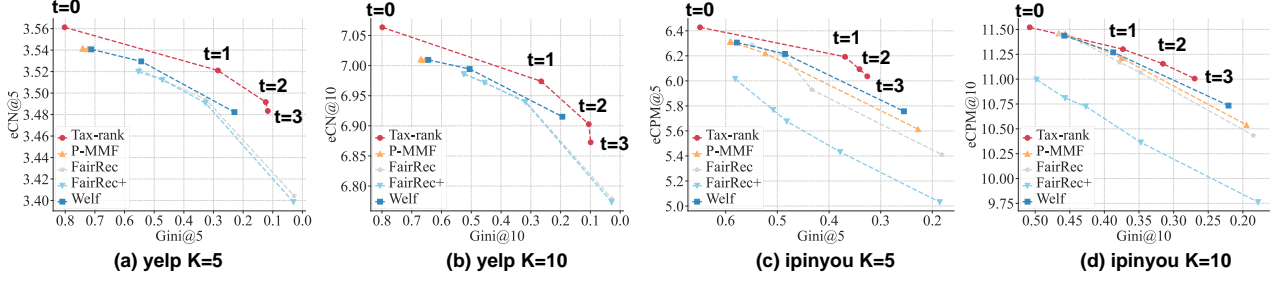


Figure 3: Pareto frontier with different top-K ranking under CTR-based settings (i.e., $w_{u,i}$ is the CTR value of user-item pair).

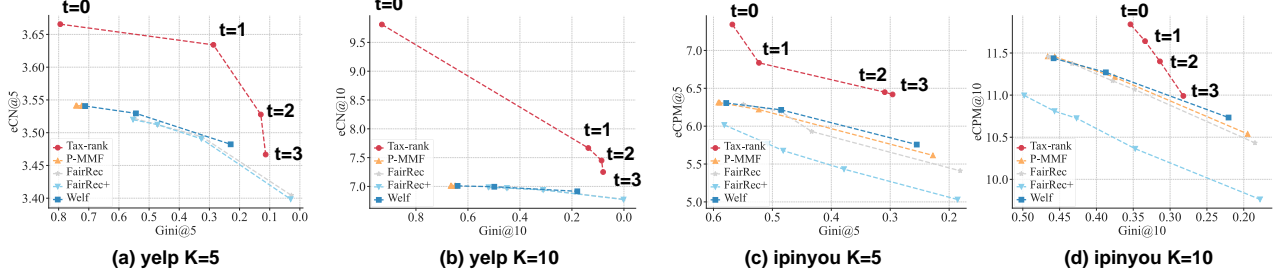


Figure 4: Pareto frontier with different top-K ranking under exposure-based settings (i.e., $w_{u,i} = 1$).

6 EXPERIMENT

We evaluate Tax-rank using two publicly available ranking datasets².

6.1 Experimental settings

Dataset. The experiments were based on two large-scale, publicly available ranking applications, including: **Yelp**³: a large-scale businesses recommendation dataset. It has 154543 samples, which contain 17034 users, 11821 items. **Ipinyou** [26]⁴: a large-scale advertising dataset. We only use the clicked data, which contains 18588 samples, which contains 18565 users, and 149 advertisements. During the pre-processing step, users and items that had interactions with fewer than 5 items or users were excluded from the entire dataset to mitigate the issue of extreme sparsity.

Evaluation. As for the evaluation metrics, the performances of the models were evaluated from two aspects: ranking accuracy, and fairness degree. As for the ranking accuracy, following the practices in [48, 50, 52], we utilize expected Click/Exposure Number (eCN) for recommendation and expected Cost Per Mile (eCPM) for advertising under top-K ranking.:

$$\text{eCN@K} = \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{I}} v_i, \quad \text{eCPM@K} = \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{I}} \text{bid}_i v_i, \quad (10)$$

where, bid_i denotes the bidding price of an advertisement, while v_i is calculated using Equation 5, which is dependent on the K .

As for the fairness degree, we utilize the Gini Index [14, 16], which is the most common measure of item utility inequality under

top-K ranking. Formally, it is defined as:

$$\text{Gini@K} = \frac{\sum_i \sum_j |\gamma_i v_i - \gamma_j v_j|}{2|I| \sum_i \gamma_i v_i}, \quad (11)$$

where it ranges from 0 to 1, with 0 representing perfect equality (every item has the same utility), and 1 representing perfect inequality (one item has all the utility, while every item else has none).

Baselines. The following representative item fairness models were chosen as the baselines: **FairRec** [35] and **FairRec+** [9] proposed to ensure Max-Min Share (t -MMS) of exposure for the items. **P-MMF** [49] utilized the mirror descent method to improve the worst-off item’s utility. Moreover, we also compare **Welf** [14], which also used the Frank-Wolfe algorithm to optimize a similar exponential form of optimization objective for two-sided fairness.

Implementation details. Following [49], CTR-based settings, we use BPR [40] model to compute the CTR value $w_{u,i} \in [0, 1]$ of each user-item pair (u, i) . For the item weight γ_i , a value of 1 is assigned for recommendation applications, while for advertising applications, $\gamma_i = \log(\text{bid}_i)$.

As for the hyper-parameters in all models, the trade-off coefficient (λ_{ot}) for OT was tuned among $[0.1, 2]$, and we implemented Tax-rank utilizing cvxpy [13] to conduct optimization.

6.2 Experimental Results

Figure 3 and Figure 4 shows the Pareto frontiers [49] of Gini Index (abbreviated as Gini.) and eCN/eCPM on two application datasets with different ranking size K . Figure 3 is generated using CTR-based settings, as described in Section 3, where the item utilities are defined as the expected number of clicks within a specified time period. On the other hand, Figure 4 is generated using exposure settings (also see Section 3), where the item utilities are defined as

²The source codes have been shared in <https://github.com/XuChen0427/Tax-rank>

³<https://www.yelp.com/dataset>

⁴<http://contest.ipinyou.com/>

the expected number of exposures. The Pareto frontiers [29] are constructed by systematically adjusting various parameters of the models and then selecting the points with the best performance in terms of both Gini@K and eCN@K/eCPM@K, resulting in an optimized trade-off between item fairness and total utilities.

Firstly, it is evident that a trade-off exists between ranking accuracy metrics (eCN@K or eCPM@K) and item fairness metric (Gini@K) *w.r.t.* the tax rate t . As the tax rate approaches 0, FairTax tends to prioritize ranking accuracy, and as the tax rate increases, FairTax tends to emphasize item fairness.

Moreover, compared with the baseline methods, it becomes evident that the proposed Tax-rank method consistently outperforms the baseline methods under both CTR-based and exposure-based settings (as indicated by the Tax-rank curves occupying the upper right corner). This Pareto dominance signifies that, for a given eCN@K or eCPM@K level, Tax-rank achieves superior Gini@K values, and for a given Gini@K level, it attains better eCN@K or eCPM@K performance. These results highlight that Tax-rank significantly outperforms the baseline methods.

6.3 Experimental Analysis

We also conducted experiments to analyze our proposed Tax-rank on Yelp dataset for top-10 ranking. Similar phenomena can be observed in other datasets and with different top-k settings.

6.3.1 Continuity. Firstly, we experiment to assess the continuity of the proposed Tax-rank and other baseline methods by adjusting the tax rate t or other parameters related to fairness. As shown in Figure 5, we draw the Lorenz Curve [17] of three best-performing baselines and our model Tax-rank under different tax rate t or other parameters. Since Lorenz Curve is a graphical representation that vividly illustrates income or wealth distribution within a population, providing insights into the level of inequality in economics.

Figure 5 shows the proportion of overall exposure percentage assumed by the bottom $x\%$ items. In simpler terms, the Lorenz Curve reveals the percentage ($y\%$) of the total exposures accumulated by the worst-off $x\%$ of items in the distribution. From Figure 5, it is evident that as the tax rate changes from 0 to ∞ , Tax-rank exhibits superior continuity. This is observed by the nearly averaged increase in the colored area between the two Lorenz Curves (can also be viewed as the fairness increasing value when increasing tax rate t). It indicates a more consistent response to a tenfold increase in the tax rate. On the contrary, as the tax rate changes, baselines, exhibit either minimal alterations (P-MMF) or an uneven increase (Welf and FairRec) in the colored area. This implies that a minor adjustment in the fair-aware parameter will result in a more significant performance alteration compared to Tax-rank.

6.3.2 Controllability over accuracy loss. Then, we experiment Figure 7 (a) to demonstrate how the price fairness of the taxation (POT) changes of Tax-rank and baseline P-MMF *w.r.t.* variations in the tax rate t , ranging from 0 to 1. Since P-MMF is a represented and best-performing tax policy in Theorem 1.

From the curve, it is evident that as we increase the tax rate t , our proposed Tax-rank and P-MMF approach leads to a reduction in the total ranking accuracy. Firstly, it is evident that the POT achieved by Tax-rank is notably lower compared to that of P-MMF,

highlighting the effectiveness of the Tax-rank approach. Secondly, we can also observe that the POT of Tax-rank obeys a smooth form, as the theoretical analysis results in Theorem 3. However, the previous ranking tax policies exhibit a more complex POT function form, resulting in poor systematic controllability.

6.3.3 Inference time. We conduct experiments to investigate the total inference time of the Tax-rank method compared to other item fairness baselines. In our analysis, our objective is to assess the total inference time across various user sizes $|\mathcal{U}|$, within real-world ranking applications. Therefore, we conducted tests to measure the total inference time of various models in terms of the varying number of users, all while keeping the number of items constant.

Figure 7 (b) reports the curves of total inference time (s) *w.r.t.* user size $|\mathcal{U}|$. It is worth noting that the Tax-rank method demonstrates remarkably low inference times, typically taking less than ten million seconds across different user sizes. Furthermore, when compared to other baseline methods, the inference time of these alternatives tends to increase either linearly or exponentially with changing user sizes, whereas Tax-rank consistently maintains a low inference time.

This is attributed to the fact that the CP solver process (Equation (7)) along with Tax-rank reduces the variable size from $|\mathcal{U}||\mathcal{I}|$ to $|\mathcal{I}|$ and employs the OT algorithm to efficiently map the solution to each user. Therefore, Tax-rank method involves matrix operations with limited sensitivity to changes in user size.

6.3.4 Visualizing fair re-ranking results. In Figure 6, we visualize the Tax-rank recommendation results under the tax rate of 0, 1, 3. In Figure 6, we visualize the ranking result matrix \tilde{x} , where the elements $\tilde{x}_{u,i}$ denotes the probability of recommending item i to user u in Equation (9). We also visualize the utility vector v , which is the amortized value for the columns value of \tilde{x} (*i.e.*, $v = \sum_{u \in \mathcal{U}} \tilde{x}_{u,i}$). The utility vector v_i represents the utility value associated with the ownership of item i .

The results clearly demonstrate that the accuracy-first solution consistently ranks the most popular items highly for users, thereby enhancing overall utility but potentially leading to market dominance by a few top items. Regarding $t = 1$, Tax-rank tends to distribute rankings to items in proportion to their contribution to the market. For $t = 3$, Tax-rank method strives to provide equal exposure and similar utilities to every item in the ranking. The experiment also served as validation that Tax-rank method can effectively adapt to various fairness principles as intended.

6.4 Ablation Experiments

In this section, we aim to conduct ablation experiments for Tax-rank. To better investigate the performance of our model under different parameter settings, we also conduct a series of ablation experiments on the Yelp dataset under ranking size $K = 10$. Similar experiment results are also observed on other datasets and other ranking sizes K .

6.4.1 OT regularizer coefficient λ_{ot} . In this section, we first conduct the ablation study for the OT regularizer coefficient λ_{ot} (Equation (9)), which trade-off the entropy loss to force the variable $x_{u,i}$ into the feasible region. Figure 8 illustrates how the ranking accuracy (eCN) and fairness metric (Gini) *w.r.t.* λ_{ot} from 0.1 – 2.0.




Figure 5: Lorenz Curve [17] of three best-performing baselines FairRec, P-MMF, Welf, and our model Tax-rank. The distinct curves in each figure are plotted by adjusting various tax rates t or different parameters.




Figure 6: Visualization of *Tax-rank* results under different tax rate $t = 0, 1, 3$. The utility vector \bar{x} is computed as the amortized value along the columns of the matrix.

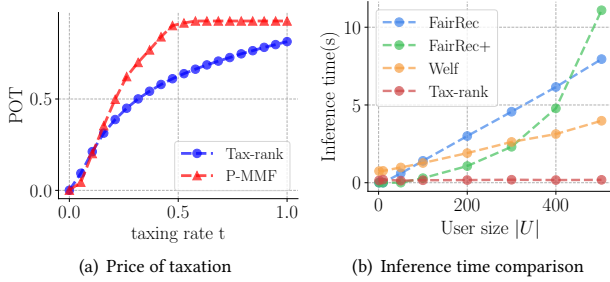


Figure 7: Sub-figure (a) illustrates the price of taxation change w.r.t tax rate t . Sub-figure (b) describes online inference items for *Tax-rank* and other baselines w.r.t user size $|U|$.




Figure 8: eCN and Gini change w.r.t. the coefficient of entropy regularizer λ_{ot} in OT under different tax rate t .

Specifically, from Figure 8 (a), we can observe that when the tax rate t is small ($t < 1$), the accuracy increases as λ_{ot} becomes larger. When the tax rate t is large ($t \geq 1$), the accuracy will decrease as λ_{ot} becomes larger. From Figure 8 (b), We can observe that Gini

remains relatively stable across the different λ_{ot} . Moreover, as the tax rate t increases, *Tax-rank* begins to prioritize fairness among items more prominently, resulting in a decrease in Gini.

The reason is the trade-off between the smoothness and convexity of the exposure probability distribution among items by tuning the parameter λ_{ot} . A higher λ_{ot} will emphasize more on the smoothness of the probability distribution, implying a reduced disparity in exposure among items. When tax rate t is relatively small ($t < 1$), a slight reduction in the difference in exposure probability among items allows some items with lower $w_{u,i}$ to gain more exposure, resulting in an increase in eCN. However, when t becomes larger, the increased smoothness of probability distributions results in a decrease in eCN consequently.

7 CONCLUSION

In this paper, we re-conceptualize the fair re-ranking task by framing it as a taxation process from an economic perspective. In theory, we reformulate previous fair re-ranking architectures as a tax process, which imposes an additional tax on each item. However, we find that previous fair re-ranking methods fail to satisfy two crucial taxation attributes: continuity and controllability over accuracy loss. To address this challenge, we introduce a novel fair re-ranking model named *Tax-rank*, whose objective levies taxes based on the difference in utility between two items. For optimization, we introduce a highly effective and efficient algorithm to optimize the objective, employing the Sinkhorn algorithm in optimal transport. Theoretical evidence shows that *Tax-rank* has superior continuity and systematic controllability compared to existing methods. Through extensive experiments on two publicly available datasets, it is evident that *Tax-rank* outperforms state-of-the-art baselines in terms of both fair ranking performance and efficiency.

ACKNOWLEDGMENTS

This work was funded by the National Key R&D Program of China (2023YFA1008704), the National Natural Science Foundation of China (No. 62377044, 62276248), the Youth Innovation Promotion Association CAS (No. 2023111), Beijing Key Laboratory of Big Data Management and Analysis Methods, Major Innovation & Planning Interdisciplinary Platform for the “Double-First Class” Initiative, Public Computing Cloud, funds for building world-class universities (disciplines) of Renmin University of China. Supported by the Outstanding Innovative Talents Cultivation Funded Programs 2024 of Renmin University of China.

APPENDIX

A PROOF OF THEOREM 1

PROOF. Let $\mathbf{a}_{[i]}$ denotes the i -th largest element of \mathbf{a} . Firstly, we will prove Equation (2) can be written as the same form of Equation (1). By utilizing the Lagrange multiplier method, we can write the objective of Equation (2) as

$$W_2(\mathbf{x}) = \min_{\lambda_2 \geq 0} \max_{\mathbf{x} \in \mathcal{X}} \sum_{i \in \mathcal{I}} \gamma_i \sum_{u \in \mathcal{U}} w_{u,i} x_{u,i} + \lambda_2 (r(\mathbf{v}) - \pi),$$

where optimal value is $\arg \min_{\lambda_2} \left[\sum_{u \in \mathcal{U}} \sum_{k=1}^K c_{n,[k]} + r(\mathbf{v}) \lambda_2 - \lambda_2 \pi \right]$, $c_{n,i} = \gamma_i w_{u,i}$. Therefore, we can only analyze the form of Equation (1). Then following Balseiro et al. [4], Xu et al. [49], we can also utilize the Lagrangian condition to break the relationship between \mathbf{v} and $\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} s_{t,i} x_{t,i}$: $W_1(\mathbf{x}) = \max_{\mathbf{x} \in \mathcal{X}} \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} s_{t,i} x_{u,i} + g(\mu^*)$, where $s_{t,i} = \gamma_i w_{u,i} + \mu_i^*$, $g(\mu) = \max_{\mathbf{v}} (\lambda r(\mathbf{v}) + \sum_{i \in \mathcal{I}} \mu_i \mathbf{v}_i)$ and

$$\mu^* = \arg \min_{\mu} \max_{\mathbf{x} \in \mathcal{X}} \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} (\gamma_i w_{u,i} + \mu_i) x_{u,i} + g(\mu).$$

□

B PROOF OF THEOREM 2

PROOF. According the Theorem 1, we can know the Equation (2) and Equation (1) is equivalent. Therefore, we mainly analyze the continuity of regularizer-based form in Equation (1): $W_1(\mathbf{x}; \lambda) = \max_{\mathbf{x} \in \mathcal{X}} \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} n_{u,i} x_{u,i} + r(\mathbf{v}) \lambda$. Then we will prove W_1 is not continuous w.r.t. tax rate λ .

We can observe that the binary ranking solution $\mathbf{x}_{u,i}$ also leads to non-continuity. We can see that if the function $W_1(\mathbf{x}; \lambda)$ changes, it implies that there exists an item i_r being removed from a u 's ranking list, while simultaneously a new item i_k is introduced and added to that u 's ranking list. Then the tax rate λ 's changing value $\delta \lambda$ rate should be at least

$$\delta \lambda > \frac{|n_{u,i_r} - n_{u,i_k}|}{|\delta r(\mathbf{v})|}$$

to result a different recommendation results, where $\delta r(\mathbf{v})$ is the change value of fairness value when item i_r is replaced with item i_k . Since the $r(\mathbf{v})$ is also often non-continuous, such as max-min form [49], therefore, we have $\exists M > 0, \delta r(\mathbf{v}) > M$.

□

C PROOF OF THEOREM 3

LEMMA 2. Given the vector $\mathbf{x} \in \mathbb{R}^N$, $\sum_{i=1}^N w_i x_i^{1+t} \geq C$, for any $t > 0, w_i > 0, x_i > 0, \sigma = \min_i x_i$ we have

$$\sum_{i=1}^N w_i x_i \geq C \sigma^t N^{-\frac{t}{1+t}}, \quad (12)$$

PROOF. We apply Hölder inequality, we have when $\frac{1}{1+t} + \frac{t}{1+t} = 1$,

$$C \leq \sum_{i=1}^N w_i x_i * \frac{1}{x_i^t} \leq \left(\sum_{i=1}^N w_i^{1+t} x_i^{1+t} \right)^{\frac{1}{1+t}} \left(\sum_{i=1}^N \frac{1}{x_i^{1+t}} \right)^{\frac{t}{1+t}}. \quad (13)$$

Also, we have $(\sum_{i=1}^N w_i x_i)^{1+t} \geq \sum_{i=1}^N w_i^{1+t} x_i^{1+t}$, therefore, we have $(\sum_{i=1}^N w_i^{1+t} x_i^{1+t})^{\frac{1}{1+t}} \leq \sum_{i=1}^N w_i x_i$. Let $\sigma = \min_i x_i$, we have

$$\left(\sum_{i=1}^N \frac{1}{x_i^{1+t}} \right)^{\frac{t}{1+t}} \leq \frac{N^{\frac{t}{1+t}}}{\sigma^t}. \quad (14)$$

Then combining Equation (14) and Equation (13) becomes:

$$\sum_{i=1}^N w_i x_i \geq C \sigma^t N^{-\frac{t}{1+t}}.$$

□

Then we will give a formal proof of Theorem 3.

PROOF. From Lemma 1, we can also write :

$$f(\mathbf{x}; t) \geq \tilde{f}(\mathbf{x}; t) \max_{\mathbf{e} \in \mathcal{E}} \sum_i m_i \frac{e_i^{1-t}}{1-t}, \quad (15)$$

where $\mathcal{E} = \{\mathbf{e} | \sum_i e_i = K, 0 \leq e_i \leq 1\}$, the input \mathbf{e} optimal value is \mathbf{z} , which represents be the best allocation under the t -fairness criterion. We briefly let $\widetilde{\text{Acc}}(t)$ denotes the accuracy value of $\tilde{f}(\mathbf{x}; t)$.

Firstly, we will bound the $\widetilde{\text{Acc}}(t)$: without generality, we assume:

$$m_1 z_1 \geq m_2 z_2 \geq \dots \geq m_{|J|} z_{|J|}. \quad (16)$$

The necessary first-order condition for the optimality of \mathbf{e} can be expressed as:

$$\nabla f(\mathbf{z}; t)(\mathbf{e} - \mathbf{z}) \leq 0, \forall \mathbf{e} \in \mathcal{E},$$

The equation can be equivalently written as:

$$\mathbf{g}^\top \mathbf{e} \leq 1, \forall \mathbf{e} \in \mathcal{E}, \quad (17)$$

where $\mathbf{g}_i = \frac{m_i z_i^{1-t}}{\sum_i m_i z_i^{1-t}}$. We observe the Equation (17), which is a well-studied knapsack problem [41], with the best solution:

$$\frac{\sum_{k=1}^K m_{|J|-k+1} z_{|J|-k+1}^{-t}}{\sum_i m_i z_i^{1-t}} \leq 1, \quad (18)$$

since according the Equation (16), we have

$$m_1 z_1^{-t} \leq m_2 z_2^{-t} \leq \dots \leq m_{|J|} z_{|J|}^{-t}.$$

From the Equation (18), exists $0 < \lambda < 1$, we have:

$$\sum_i m_i z_i^{1-t} \geq \lambda \text{Acc}(0) z_1^{-t}, \quad (19)$$

$$\sum_{k=1}^K m_{|J|-k+1} z_{|J|-k+1}^{-t} \geq z_1^{-t} \lambda \sum_{k=1}^K m_{[k]} = \lambda \text{Acc}(0) z_1^{-t},$$

where $m_{[k]}$ denotes the k -th largest element of m_i .

Taking the Equation (19) into Lemma 2, we have

$$\widetilde{\text{Acc}}(t) = \sum_i m_i z_i \geq \lambda \text{Acc}(0) \left(\frac{z_{|J|}}{z_1} \right)^t N^{-\frac{t}{1+t}}. \quad (20)$$

Therefore,

$$\begin{aligned} \text{POT} &= \frac{\text{Acc}(0) - \text{Acc}(t)}{\text{Acc}(0)} \leq \frac{\text{Acc}(0) - \widetilde{\text{Acc}}(t)}{\text{Acc}(0)} \\ &\leq 1 - \lambda \left(\frac{z_{|J|}}{z_1} \right)^t N^{-\frac{t}{1+t}} = 1 - O(N^{-\frac{t}{1+t}}). \end{aligned}$$

□

REFERENCES

- [1] Himan Abdollahpouri, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Pizzato. 2020. Multistakeholder recommendation: Survey and research directions. *User Modeling and User-Adapted Interaction* 30, 1 (2020), 127–158.
- [2] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2017. Controlling popularity bias in learning-to-rank recommendation. In *Proceedings of the eleventh ACM conference on recommender systems*. 42–46.
- [3] Himan Abdollahpouri, Masoud Mansoury, Robin Burke, and Bamshad Mobasher. 2019. The unfairness of popularity bias in recommendation. *arXiv preprint arXiv:1907.13286* (2019).
- [4] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. 2021. Regularized online allocation problems: Fairness and beyond. In *International Conference on Machine Learning*. PMLR, 630–639.
- [5] Omer Ben-Porat and Moshe Tenenholz. 2018. A game-theoretic approach to recommendation systems with strategic content providers. *Advances in Neural Information Processing Systems* 31 (2018).
- [6] Dimitri P Bertsekas. 1997. Nonlinear programming. *Journal of the Operational Research Society* 48, 3 (1997), 334–334.
- [7] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. 2011. The price of fairness. *Operations research* 59, 1 (2011), 17–31.
- [8] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. 2012. On the efficiency-fairness trade-off. *Management Science* 58, 12 (2012), 2234–2250.
- [9] 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.
- [10] John A Brittain. 1971. The incidence of social security payroll taxes. *The American Economic Review* 61, 1 (1971), 110–125.
- [11] Valentina Cacchiani, Manuel Iori, Alberto Locatelli, and Silvano Martello. 2022. Knapsack problems—An overview of recent advances. Part II: Multiple, multidimensional, and quadratic knapsack problems. *Computers & Operations Research* 143 (2022), 105693.
- [12] Yashar Deldjoo, Dietmar Jannach, Alejandro Bellogin, Alessandro Difonzo, and Dario Zanzonelli. 2022. A survey of research on fair recommender systems. *arXiv preprint arXiv:2205.11127* (2022).
- [13] Steven Diamond and Stephen Boyd. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *The Journal of Machine Learning Research* 17, 1 (2016), 2909–2913.
- [14] Virginie Do, Sam Corbett-Davies, Jamal Atif, and Nicolas Usunier. 2021. Two-sided fairness in rankings via Lorenz dominance. *Advances in Neural Information Processing Systems* 34 (2021), 8596–8608.
- [15] Virginie Do and Nicolas Usunier. 2022. Optimizing Generalized Gini Indices for Fairness in Rankings (SIGIR '22). Association for Computing Machinery, New York, NY, USA, 737–747.
- [16] Virginie Do and Nicolas Usunier. 2022. Optimizing generalized Gini indices for fairness in rankings. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 737–747.
- [17] Joseph L Gastwirth. 1971. A general definition of the Lorenz curve. *Econometrica: Journal of the Econometric Society* (1971), 1037–1039.
- [18] William W Hager. 1979. Lipschitz continuity for constrained processes. *SIAM Journal on Control and Optimization* 17, 3 (1979), 321–338.
- [19] Michelle Hanlon and Shane Heitzman. 2010. A review of tax research. *Journal of accounting and Economics* 50, 2-3 (2010), 127–178.
- [20] Zhimeng Jiang, Xiaotian Han, Chao Fan, Fan Yang, Ali Mostafavi, and Xia Hu. 2021. Generalized demographic parity for group fairness. In *International Conference on Learning Representations*.
- [21] Peter J Lambert. 1992. *The distribution and redistribution of income*. Springer.
- [22] Julian Lamont. 2017. *Distributive justice*. Routledge.
- [23] Tian Lan, David Kao, Mung Chiang, and Ashutosh Sabharwal. 2010. *An axiomatic theory of fairness in network resource allocation*. IEEE.
- [24] 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.
- [25] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, Juntao Tan, Shuchang Liu, and Yongfeng Zhang. [n. d.]. Fairness in Recommendation: Foundations, Methods and Applications. *ACM Transactions on Intelligent Systems and Technology* (n. d.).
- [26] Hairen Liao, Lingxiao Peng, Zhenchuan Liu, and Xuehua Shen. 2014. iPinYou global rtb bidding algorithm competition dataset. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*. 1–6.
- [27] Aldo Lipani. 2016. Fairness in information retrieval. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 1171–1171.
- [28] Xiangyu Liu, Chuan Yu, Zhilin Zhang, Zhenzhe Zheng, Yu Rong, Hongtao Lv, Da Huo, Yiqing Wang, Dagui Chen, Jian Xu, et al. 2021. Neural auction: End-to-end learning of auction mechanisms for e-commerce advertising. In *Proceedings of the 27th ACM Conference on Knowledge Discovery & Data Mining*. 3354–3364.
- [29] Alexander V Lotov and Kaisa Miettinen. 2008. Visualizing the Pareto frontier. In *Multiobjective optimization*. Springer, 213–243.
- [30] David Matsumoto and Linda Juang. 2016. *Culture and psychology*. Cengage Learning.
- [31] Mohammadmehdi Naghiaei, Hossein A Rahmani, and Yashar Deldjoo. 2022. Cpfair: Personalized consumer and producer fairness re-ranking for recommender systems. *arXiv preprint arXiv:2204.08085* (2022).
- [32] John F Nash Jr. 1950. The bargaining problem. *Econometrica: Journal of the econometric society* (1950), 155–162.
- [33] Birger Nerré. 2001. The concept of tax culture. In *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association*, Vol. 94. JSTOR, 288–295.
- [34] Wallace E Oates. 1969. The effects of property taxes and local public spending on property values: An empirical study of tax capitalization and the Tiebout hypothesis. *Journal of political economy* 77, 6 (1969), 957–971.
- [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] Gourab K Patro, Lorenzo Porcaro, Laura Mitchell, Qiuyue Zhang, Meike Zehlike, and Nikhil Garg. 2022. Fair ranking: a critical review, challenges, and future directions. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 1929–1942.
- [37] Gabriel Peyré, Marco Cuturi, et al. 2019. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning* 11, 5-6 (2019), 355–607.
- [38] Khien Pham, Khang Le, Nhat Ho, Tung Pham, and Hung Bui. 2020. On unbalanced optimal transport: An analysis of sinkhorn algorithm. In *International Conference on Machine Learning*. PMLR, 7673–7682.
- [39] Barry W Poulson and Jules Gordon Kaplan. 2008. State income taxes and economic growth. *Cato J.* 28 (2008), 53.
- [40] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [41] Harvey M Salkin and Cornelis A De Kluyver. 1975. The knapsack problem: a survey. *Naval Research Logistics Quarterly* 22, 1 (1975), 127–144.
- [42] Ashudeep Singh and Thorsten Joachims. 2019. Policy learning for fairness in ranking. *Advances in neural information processing systems* 32 (2019).
- [43] Kyle Swanson, Lili Yu, and Tao Lei. 2020. Rationalizing Text Matching: Learning Sparse Alignments via Optimal Transport. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 5609–5626.
- [44] Jiakai Tang, Shiqi Shen, Zhipeng Wang, Zhi Gong, Jingsen Zhang, and Xu Chen. 2023. When Fairness Meets Bias: A Debaised Framework for Fairness Aware Top-N Recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems (Singapore, Singapore) (RecSys '23)*. Association for Computing Machinery, New York, NY, USA, 200–210. <https://doi.org/10.1145/3604915.3608770>
- [45] Tom R Tyler and E Allan Lind. 2002. Procedural justice. In *Handbook of justice research in law*. Springer, 65–92.
- [46] Tom R Tyler and Heather J Smith. 1995. Social justice and social movements. (1995).
- [47] Jiayin Wang, Weizhi Ma, Jiayu Li, Hongyu Lu, Min Zhang, Biao Li, Yiqun Liu, Peng Jiang, and Shaoping Ma. 2022. Make Fairness More Fair: Fair Item Utility Estimation and Exposure Re-Distribution. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (Washington DC, USA) (KDD '22)*. Association for Computing Machinery, New York, NY, USA, 1868–1877. <https://doi.org/10.1145/3534678.3539354>
- [48] 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.
- [49] 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.
- [50] Chen Xu, Xiaopeng Ye, Jun Xu, Xiao Zhang, Weiran Shen, and Ji-Rong Wen. 2023. LTP-MMF: Towards Long-term Provider Max-min Fairness Under Recommendation Feedback Loops. *arXiv preprint arXiv:2308.05902* (2023).
- [51] Jun Xu, Xiangnan He, and Hang Li. 2018. Deep Learning for Matching in Search and Recommendation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (Ann Arbor, MI, USA) (SIGIR '18)*. Association for Computing Machinery, New York, NY, USA, 1365–1368.
- [52] Xun Yang, Yasong Li, Hao Wang, Di Wu, Qing Tan, Jian Xu, and Kun Gai. 2019. Bid optimization by multivariable control in display advertising. In *Proceedings of the 25th ACM international conference on knowledge discovery & data mining*. 1966–1974.
- [53] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P Gummadi. 2019. Fairness constraints: A flexible approach for fair classification. *The Journal of Machine Learning Research* 20, 1 (2019), 2737–2778.