

α -RANK: UNIFIED ITEM-FAIR RANKING FROM A CO-OPERATIVE GAME THEORY VIEW

Anonymous authors

Paper under double-blind review

ABSTRACT

Driven by economic and systematic considerations, the pursuit of item fairness in ranking has emerged as a prominent topic in recommendation and advertising applications. Prior research has suggested various fairness aspects can be aligned with the concept of distributive justice in sociology, such as utilitarianism, dealism, and egalitarianism. However, they fail to distinguish the distinctions and relationships among these fairness dimensions in ranking. In fact, item fairness can be viewed as a unified challenge of fairly allocating the constrained and fluctuating resources, from the perspective of cooperative game theory. In our work, we introduce the smooth α -fairness objective for different fairness and unify item fairness as a cooperative game problem. In such games, items are considered as the players dividing the “cake” of user attention. In such games, we analyze the α -fairness objective from a theoretical way and introduce an efficient approach called α -rank. Firstly, we re-form several important axioms in cooperative games to tell us how item fairness principles exhibit when the resource “cake” changes in ranking. Then we designed α -rank, which applies the optimal transport to conduct item fairness. Theoretical analysis provides an upper bound, showcasing the maximum total utility loss across different fairness degrees. We conducted experiments in two ranking applications: recommendation and advertising. The experimental results demonstrate that α -rank effectively and efficiently outperforms the baseline methods.

1 INTRODUCTION

Ranking techniques have found extensive application in web-based platforms, such as determining which items to display to users in recommendation and advertising scenarios with limited exposure slots (Xu et al., 2018; Baeza-Yates et al., 1999). Recently, researchers have emphasized the importance of item fairness in ranking, as it not only prevents monopolization but also contributes to the creation of a healthier ecosystem (Xu et al., 2023a; Patro et al., 2020; Do et al., 2021; Li et al., 2022; Lipani, 2016).

Different from the user fairness, which pertains to ensuring that everyone has fundamental rights and responsibilities (Matsumoto & Juang, 2016; Abdollahpouri et al., 2019). Item fairness, which aims to equitably distribute items among users, is closely aligned with the concept of distributive justice (Lamont, 2017; Matsumoto & Juang, 2016) in sociology. Previous research papers advocate item from distinct principles: utilitarianism objective (Baeza-Yates et al., 1999; Lacerda et al., 2006), which focused on maximizing summarization of all item utilities; dealism objective aligned with proportion fairness (Ben-Porat & Tennenholtz, 2018), which aims to strive to achieve an allocation where items possess resources in proportion to their respective weights or importance; egalitarianism objective aligned with max-min fairness (Xu et al., 2023a; Do et al., 2021; Patro et al., 2020), which equalizes the utilities of all item utilities involved in the decision-making process.

Although previous ranking models have introduced effective algorithms aligned with specific fairness aspects, they often lack a clear distinction between these underlying fairness principles. Inspired from cooperative game theory, the concept of item fairness in ranking closely resembles the notion of fair resource allocation (Matsumoto & Juang, 2016; Xu et al., 2023a), which primarily focuses on finding a suitable resource allocation method that caters to the utility of all involved parties

in an economic way. In a simple way, item fairness can be seen as a challenge involving resource allocation in scenarios where resources are both limited and subject to fluctuation.

In our work, we approach the issue of item fairness in ranking from a unified perspective rooted in cooperative game theory (Branzei et al., 2008; Peleg & Sudhölter, 2007). Within the framework of cooperative games, each item is viewed as a participant tasked with fairly dividing the “cake” of limited exposure slots. Inspired by cooperative games, we introduce the concept of α -fairness (Xu & Cumanan, 2017; Bertsimas et al., 2012) to achieve a well-balanced equilibrium among item fairness principles. As α approaches 0, 1, and ∞ , it corresponds to the utilitarianism, dealism, and egalitarianism solutions, respectively. Optimizing α -fairness offers a smooth and adaptable approach to achieve item fairness in accordance with varying requirements.

Then we analyze the α -fairness objective from a theoretical way and introduce an efficient approach called α -ranking. Specifically, we begin by reforming several key axioms of cooperative games designed for item fairness. The axioms describe how different item fairness principles behave when the amount of resources, such as limited exposure slots, changes.

After that, we propose an efficient approach named α -rank to efficiently tackle the α -fairness optimization objective in ranking. Firstly, we identify an upper-bound function for the target problem, which conforms to the structure of a standard cooperative game and can be efficiently resolved. Then, we utilize Sinkhorn algorithm (Swanson et al., 2020) of optimal transport (OT) (Pham et al., 2020; Peyré et al., 2019) to map the upper-bound function back to the original space, thus arriving at our ranking results efficiently. Finally, we offer theoretical insights into the maximum loss of total item utilities under various α values through the upper-bound function.

We also apply α -ranking into real-world ranking scenarios, specifically in recommendation and advertising, using two extensive public datasets. Experiment results demonstrate that α -rank can achieve better performance while maintaining the efficiency required for industrial ranking systems.

2 RELATED WORKS

Fairness principle: Cultural perspectives on fairness exhibit significant variations, as extensively explored in sociological research (Tyler & Allan Lind, 2002; Tyler & Smith, 1995). In practice, two common fairness definitions are rooted: equality and equity (Matsumoto & Juang, 2016). Equality is defined as: *everyone is treated the same and provided the same resources to succeed*, which aims to ensure the fundamental rights and responsibilities of each individual. While equity is defined as: *ensuring that resources are equally distributed based on needs*, which is close to the concept of distributive justice (Lamont, 2017). In distributive justice, there are three types of allocation principles. Utilitarianism proposed by Aristotle (Sen, 1979), which aims to maximize the summation of utilities. As for the dealism proposed by Nash (Nash Jr, 1950), it focuses on reaching an agreement point based on the deals previously made by each side. Egalitarianism (Rawls, 1971) aims to equalize the utilities of all individuals. Item fairness is more related to the distributive justice realm. In this paper, we apply the cooperative games to unify the three principles in item fairness.

Item fairness methods: Regarding item fairness, previous work often focused on two types: individual fairness (Marras et al., 2022; Li et al., 2021), which concentrates on equitable treatment for individuals, and group fairness, which categorizes items groups (Ge et al., 2021; Xu et al., 2023a). Our work primarily focuses on individual fairness as the main objective, while group fairness can be formulated in a similar manner. For different fairness aspects, Current mainstream ranking systems (Rendle et al., 2012; Xue et al., 2017; Yang et al., 2019) apply utilitarianism to optimize the summation of platform profit. Dealism often relates to the proportion fairness (Bertsimas et al., 2011). For example, Ben-Porat & Tennenholtz (2018); Patro et al. (2020); Biswas et al. (2021) proposed the Shapley algorithm to reach the point. Optimizing objective of Egalitarianism could be Gini Index (Do & Usunier, 2022), max-min fairness (Xu et al., 2023a; Do et al., 2021) and distance of different groups (Jiang et al., 2021). However, they often focused on one type of fairness and failed to distinguish the connections between different fairness principles.

Cooperative games: The field of game theory (Von Neumann & Morgenstern, 1947), is commonly divided into: cooperative games and non-cooperative games. Different from non-cooperative game (Nash, 1951), cooperative game involves players whose interests are neither completely opposed nor completely coincident, allowing them to communicate and collaborate. In cooperative game

Table 1: Detailed explanations of variable in item fairness

Symbol	Value	Application	Explanation
\mathbf{v}_i	$\mathbf{v}_i = \sum_u w_{u,i} \mathbf{x}_{u,i}$	amortized ranking (Xu et al., 2023a; Biega et al., 2018)	the utility of the item i within the user arrival times
$w_{u,i}$	$w_{u,i} = 1$	exposure-based fairness (Xu et al., 2023a; Patro et al., 2020)	charging according to one exposure of item/advertise i to user u
	$w_{u,i} = \text{ctr}_{u,i}$	CTR-based fairness (Rendle et al., 2012; Xue et al., 2017)	charging according to the user u clicked on the item i once
	$w_{u,i} = \text{ctr}_{u,i} * \text{cvr}_{u,i}$	CVR-based fairness (Yang et al., 2019; Liu et al., 2021)	charging according to the user u converted on the item i once
$\text{ctr}_{u,i}, \text{cvr}_{u,i}$	CTR/CVR value	CTR/CVR billing	CTR/CVR value of user u to item i
γ_i	$\gamma_i = \beta_i$	recommendation (Rendle et al., 2012)	β_i serves as the adjustment factor for each item
	$\gamma_i = \text{bid}_i * \beta_i$	advertising (Yang et al., 2019; Liu et al., 2021)	bid_i represents the bidding value of the advertiser.

theory, (Shapley et al., 1953) proposed the concept of Sharpley value, offering an approach for fairly allocating in cooperative games. Another approach, proposed by (Nash Jr, 1950), is rooted in the concept of bargaining. To highlight the weight of importance of players’ bargaining power, Nash (Nash Jr, 1950) introduced a generalized framework of bargaining. Kalai & Smorodinsky (1975) proposed a solution that focuses on the proportion to the ideal utility of each player. Based on this, Kalai (1977) proposed a max-min method to equalize the utility of all players involved. In our research, we approach item fairness as a problem of equitable resource allocation for items, drawing inspiration from the perspective of cooperative game theory.

3 PROBLEM FORMULATION

We first define some notations for the problem. For vector $\mathbf{x} \in \mathbb{R}^n$, let \mathbf{x}_i denote the i -th element of the vector. For vector $\mathbf{x} \in \mathbb{R}^{n \times m}$, let $\mathbf{x}_{i,j}$ denote the element of i -th row and j -th column. \mathbf{A}_i denote the i -th column vector of \mathbf{A} . $\mathbf{x} \geq \mathbf{y}$ denotes element \mathbf{x}_i should be greater or equal to $\mathbf{y}_i, \forall i$.

In this section, we will formulate the item fairness in ranking into a constrained optimization problem. In the context of ranking, we define \mathcal{U} representing the set of users, \mathcal{I} representing the set of items. When a user $u \in \mathcal{U}$ interacts with the system, the number of retrieved items is typically limited and is often defined by a constant value denoted as K . For each user u , the decision vector $\mathbf{x}_u \in \{0, 1\}^{|\mathcal{I}|}$, where $\mathbf{x}_{u,i} = 1$ denotes item i should be recommended to user u , otherwise, $\mathbf{x}_{u,i} = 0$. The utilities of items can be represented as a vector $\mathbf{v} \in \mathbb{R}_+^{|\mathcal{I}|}$, where utility of certain item \mathbf{v}_i relates to the decision vector \mathbf{x}_u .

Then, we can write the ranking problem into the following mathematical program in a general way:

$$\mathbf{v}^f = \arg \max_{\mathbf{v} \in \mathcal{D}} f(\mathbf{v}), \quad \mathcal{D} = \{\mathbf{v}(\mathbf{x}_u) | \mathbf{1}^\top \mathbf{x}_u = K, \forall u \in \mathcal{U}\}, \quad (1)$$

where $f(\cdot)$ represents the fairness optimization objective of ranking, which can vary depending on different objectives or proposals. To better understand the item fairness application, we give an illustrated example in Appendix E. Previous studies proposed distinct types of optimization objectives that correspond to different principles of fair resource allocation in terms of $f(\cdot)$:

(1) Utilitarianism (w/o fairness) (Rendle et al., 2012): $f(\mathbf{v}) = \sum_i \gamma_i \mathbf{v}_i$

(2) Dealism (proportion fairness) (Li et al., 2022): $f(\mathbf{v}) = \sum_i \gamma_i \log v_i$

(3) Egalitarianism (max-min fairness) (Xu et al., 2023a): $f(\mathbf{v}) = \min_i \gamma_i v_i$,

where the γ_i is the weight of each item. In rankings, v_i and γ_i have different forms, which are listed in the Table 1. In the Table 1, CTR and CVR is the abbreviation of click-through-rate (CTR), and conversion-through-rate (CVR) (Yang et al., 2019).

For various fairness principles within the objectives of ranking: utilitarianism (Matsumoto & Juang, 2016) strives to maximize the overall utilities of the items, seeking to optimize the collective benefit. Dealism (Bertsimas et al., 2011) aims to strive to allocate items that possess resources in proportion to their respective weights γ_i . Egalitarianism (Bertsimas et al., 2011) aims to equalize the utilities of items by enhancing the utility of the worst-off items, promoting fairness through improved distribution of benefits. Previous work also proposed to trade-off the different optimizing objectives for item fairness (Abdollahpouri & Burke, 2019; Abdollahpouri et al., 2020; Hao et al., 2021; Naghiaei et al., 2022).

In cooperative games, the α -fairness (Bertsimas et al., 2012) provides a smooth way to unify the three types of fairness principles:

$$f(\mathbf{v}; \alpha) = \begin{cases} \sum_i \frac{v_i^{1-\alpha}}{1-\alpha} & \text{if } \alpha > 0, \alpha \neq 1 \\ \sum_i \log(v_i) & \text{if } \alpha = 1 \end{cases}, \quad W(\alpha) = \max_{\mathbf{v} \in \mathcal{D}} f(\mathbf{v}; \alpha). \quad (2)$$

where α approaches 0, 1, and ∞ , it corresponds to the utilitarianism, dealism, and egalitarianism solutions, respectively.

4 OUR FRAMEWORK

In this section, we will first introduce five axioms of the α -fairness in ranking from the view of cooperative game in a theoretical way. Then, we present the α -ranking to efficiently and effectively solve the item fairness in ranking.

4.1 AXIOMS IN ITEM FAIRNESS

In this section, we will re-form several axioms in cooperative game theory (Bertsimas et al., 2011) that one might seek in an item-fair ranking system. These axioms show that how item fairness principles will behave when there are changes in available resources.

Axiom 1 (Pareto Optimality) *The utilities of items \mathbf{v}^f is Pareto optimal, that is, there does not exist another solution $\mathbf{v} \in \mathcal{U}$ so that utility vector $\mathbf{v} \geq \mathbf{v}^f$ and $\mathbf{v}^f \neq \mathbf{v}$.*

Axiom 2 (Symmetry) *When two items i and j possess equal weight and charging weight values, indicated by $\gamma_i = \gamma_j$, they are expected to yield the same utility outcome, denoted as $\mathbf{v}_i = \mathbf{v}_j$.*

Axiom 3 (Affine Invariance) *If we have an affine operator $A(\mathbf{v}_i) = c_i \mathbf{v}_i$, $c_i > 0$, then fair allocation under ranking is equal to the affine transformation of the fair allocation under the original system, i.e. $\arg \max_{\mathbf{v}} f(A(\mathbf{v})) = A(\arg \max_{\mathbf{v}} f(\mathbf{v}))$.*

Axiom 4 (Independence of Irrelevant Alternatives) *If $\mathcal{D}_1, \mathcal{D}_2$ are two utility feasible set such that $\mathcal{D}_1 \subset \mathcal{D}_2$ and $\arg \max_{\mathbf{v} \in \mathcal{D}_2} f(\mathbf{v}) \in \mathcal{D}_1$, then $\arg \max_{\mathbf{v} \in \mathcal{D}_1} f(\mathbf{v}) = \arg \max_{\mathbf{v} \in \mathcal{D}_2} f(\mathbf{v})$.*

Axiom 5 (Resource Monotonicity) *Let $\mathcal{D}_1, \mathcal{D}_2$ be two utility sets and $\mathcal{D}_1 \subset \mathcal{D}_2$ and $\mathcal{D}_1 \neq \mathcal{D}_2$, we have $\arg \max_{\mathbf{v} \in \mathcal{D}_1} f(\mathbf{v}) \leq \arg \max_{\mathbf{v} \in \mathcal{D}_2} f(\mathbf{v})$.*

Axiom 1 (Pareto optimality) ensures that no situation can arise where the utility of two items can simultaneously increase. Axiom 2 (Symmetry) ensures ranking model cannot differentiate the items by their attributes. Axiom 3 (Affine Invariance) guarantees that the ranking outcome remains unchanged regardless of the choice of utility numeraire. Axiom 4 (Independence of Irrelevant Alternatives) illustrates that if resources are decreased for an item and the original solution lies within the feasible region, then the solution remains the same as the original. Axiom 5 (Resource Monotonicity) illustrates that increasing the feasible set will give each item equal or greater utility. Detailed analysis can be seen in Axiom 1-4 (Nash Jr, 1950) and Axiom 5 (Kalai & Smorodinsky, 1975).

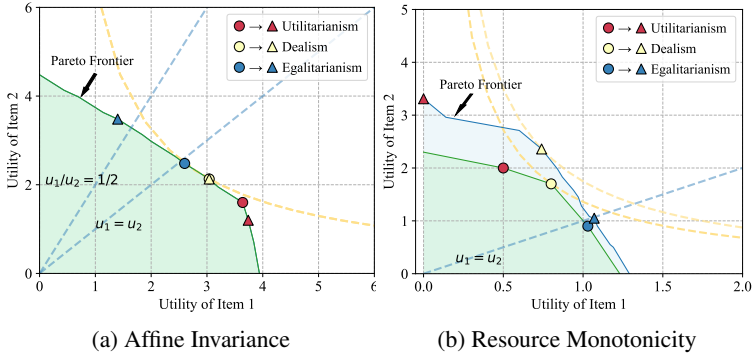


Figure 1: Toy examples to illustrate the axioms of item fairness. Two items (item 1 and item 2) are recommended to 30 users, with the constraint that each user can only be exposed to one item (i.e. $K = 1$). Circles and triangles are utilized to visually depict the shifts in optimal solutions for each fairness criterion when faced with changes in available resources.

Theorem 1 *Utilitarianism, dealism, and egalitarianism all adhere to Axioms 1, 2, and 4. However, utilitarianism and dealism fail to meet Axiom 5 (Resource Monotonicity), while utilitarianism and egalitarianism do not conform to Axiom 3.*

Remark 1 *The axioms indicate that as α varies from 0 to 1, the system disregards the numeraire of utilities and allocates resources in proportion to their weight γ_i more often.*

Remark 2 *The axioms suggest that as α increases from 1 to ∞ , the system instructs the platform to enhance the utility of the worst-off item and simultaneously improve the utility of all items when the resource increases more often.*

The proof of Theorem 1 can be seen in Appendix A. The theorem indicates that various fairness principles exhibit varying performance as the available resources change.

In order to better understand the axioms from the view of cooperative games, we conducted a simulation to analysis of two axioms. Figure 1 illustrates a real ranking scenario where the resource “cake” to explore how different fairness principles perform in ranking tasks. We set each item’s weight γ_i to 2 and 3, respectively.

Figure 1 (a) illustrates that optimal points corresponding to various fairness principles if change the numeraire of resource “cake”. In this simulation, we have doubled the utility of item 2 while keeping the utility of item 1 unchanged. The experimental findings demonstrate that utilitarianism and egalitarianism do not adhere to Axiom 3 (Affine Invariance). Conversely, dealism maintains its behavior of allocating resources in proportion to their weight γ_i for items 2 and 3, respectively, regardless of the choice of resource numeraire.

Figure 1 (b) depicts the optimal points corresponding to various fairness principles if we increase the size of the resource “cake”. In this illustration, the green regions represent the original feasible region, while the blue regions indicate the expanded region where additional resources are allocated (assigning to more 10 users than the initial 30 users). The experimental findings demonstrate that utilitarianism and dealism do not adhere to Axiom 5 (Resource Monotonicity). Conversely, egalitarianism exhibits the ability to enhance the utility of both items in the presence of resource changes.

4.2 α -RANK ALGORITHM

In this section, we will introduce α -rank approach to efficiently handle the α -fairness objective optimization in equation 2. The overall algorithm workflow can be seen in Algorithm 1.

We observe that directly optimizing the equation 2 requires huge computational costs since it is a non-linear, large-scale, and integral programming (Bertsekas, 1997). Therefore, firstly, we construct an easy-solved standard cooperative game programming (equation 3), which is the upper bound function of equation 2. Then we apply the transport optimal (OT) projection method to obtain the final ranking result (equation 4). Finally, we prove a theoretical result to show the maximum social

Algorithm 1: Algorithm of α -rank

Input: User set \mathcal{U} , item set \mathcal{I} , ranking size K , fairness coefficient α , OT coefficient λ , item weight $\gamma_i, \forall i \in \mathcal{I}$, user-item 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 e^* from equation 3.
- 2: Initialize $\mathbf{m} = K\mathbf{1}$, $\mathbf{n} = e^*$, $\mathbf{C}_{u,i} = \gamma_i w_{u,i}, \forall u \in \mathcal{U}, \forall i \in \mathcal{I}$, $\mathbf{B} = e^{-\frac{c}{\lambda}}$
- 3: **for** $t = 1, \dots, T$ **do**
- 4: $\mathbf{m} = K\mathbf{1} \odot \mathbf{B}\mathbf{n}$
- 5: $\mathbf{n} = e^* \odot \mathbf{B}\mathbf{m}$
- 6: **end for**
- 7: $\tilde{\mathbf{x}} = \text{diag}(\mathbf{m})\mathbf{B}\text{diag}(\mathbf{n})$
- 8: $L_K(u) = \arg \max_{S \subset \{1,2,\dots,|\mathcal{I}|\}, |S|=K} \sum_{i \in S} \tilde{\mathbf{x}}_{u,i}, \quad \forall u \in \mathcal{U}$

utility loss across different fairness degree, named price of fairness (POF) of ranking (Bertsimas et al., 2011; 2012).

4.2.1 UPPER BOUND FUNCTION CONSTRUCTION

Theorem 2 *There exists $\tau > 0$, s.t. we have the following function*

$$\begin{aligned} \hat{W}(\alpha) &= \max_{\mathbf{e}} \sum_i \gamma_i \eta_i g(\mathbf{e}; \alpha) \\ \text{s.t. } \sum_{i \in \mathcal{I}} \mathbf{e}_i &= K, \quad 0 \leq \mathbf{e}_i \leq 1, \quad \eta_i = \tau \sum_{u \in \mathcal{U}} w_{u,i}, \forall i \in \mathcal{I} \\ g(\mathbf{e}; \alpha) &= \begin{cases} \sum_i \frac{\mathbf{e}_i^{1-\alpha}}{1-\alpha} & \text{if } \alpha > 0, \alpha \neq 1 \\ \sum_i \log(\mathbf{e}_i) & \text{if } \alpha = 1 \end{cases}, \end{aligned} \quad (3)$$

where $\hat{W}(\alpha) \geq \max_{\mathbf{v} \in \mathcal{U}} f(\mathbf{v}; \alpha)$ and the variable $\mathbf{e}_i = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \mathbf{x}_{u,i}$, which is the averaged exposure of certain item i within a period of time.

The proof of Theorem 2 can be seen in Appendix B. The optimal value e^* represents the average exposure of items achieved under the α -fairness optimization objective. Then we will apply the Sinkhorn algorithm (Pham et al., 2020) to project the averaged exposure e^* to recommendation list $\mathbf{x} \in \{0, 1\}^{|\mathcal{U}| \times |\mathcal{I}|}$ discussed in Section 3.

4.2.2 OPTIMAL TRANSPORT PROJECTION

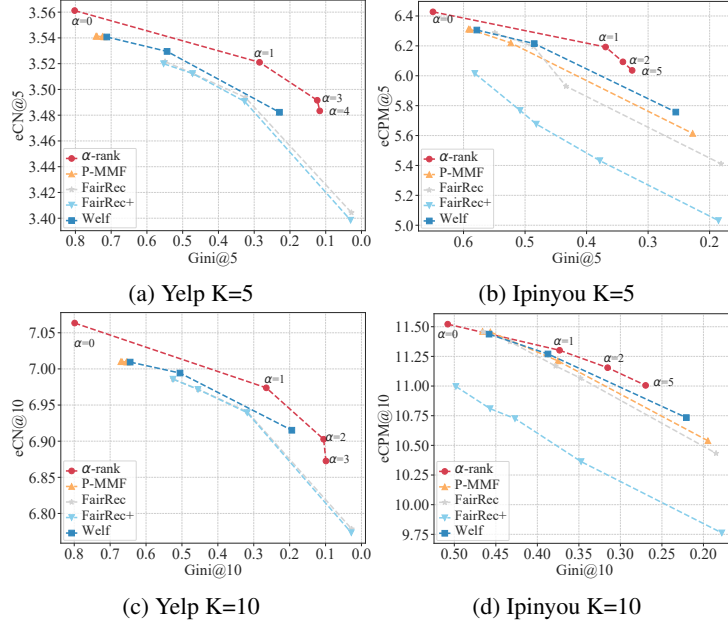
We obtain the final ranking result by utilizing the following sample process, where $\tilde{\mathbf{x}}$ (i.e. ranking score distribution) is derived from the OT projection process.

$$L_K(u) = \arg \max_{S \subset \{1,2,\dots,|\mathcal{I}|\}, |S|=K} \sum_{i \in S} \tilde{\mathbf{x}}_{u,i}, \quad \forall u \in \mathcal{U}. \quad (4)$$

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

$$\tilde{\mathbf{x}} = \arg \min_{\mathbf{x} \geq 0} \langle \mathbf{x}, -\mathbf{C} \rangle + \lambda H(\mathbf{x}) \quad \text{s.t.} \quad \mathbf{x}\mathbf{1} = K\mathbf{1}, \quad \mathbf{1}^\top \mathbf{x} = e^*, \quad (5)$$

where $\mathbf{1}$ denotes a vector of ones, e^* denotes the optimal value of equation 3 and λ is the coefficient of entropy regularizer. $\langle \mathbf{x}, -\mathbf{C} \rangle$ results transport plan lies on the Pareto frontier. $H(\mathbf{x}) = \sum_u \sum_i \mathbf{x}_{u,i} \log(\mathbf{x}_{u,i})$, 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 e^* .

Figure 2: Pareto frontier of four different datasets with different top- K ranking.

This problem can be efficiently solved by the Sinkhorn algorithm (Swanson et al., 2020), 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, $\mathbf{B} = e^{-\frac{C}{x}}$, and $\mathbf{m} \in \mathbb{R}^{|\mathcal{U}|}$, $\mathbf{n} \in \mathbb{R}^{|\mathcal{I}|}$, which iteratively computes

$$\mathbf{m} \leftarrow K\mathbf{1} \oslash \mathbf{B}\mathbf{n}, \quad \mathbf{n} \leftarrow \mathbf{e}^* \oslash \mathbf{B}\mathbf{m},$$

where \oslash denotes element-wise division.

4.2.3 PRICE OF ITEM FAIRNESS

Typically, when conducting fairness adjustments in ranking, it may result in the redistribution of resources that can lead to a reduction in the total utilities ($\sum_i v_i$) of the system. In this section, we aim to bound the price of item fairness (POF) (Bertsimas et al., 2011), which measures the maximum social utility loss across different fairness degrees, i.e. different α values.

Theorem 3 *The price of item fairness is quantified as the relative reduction in the sum of utilities when comparing the fair solution to the utilitarian solution, represented as:*

$$POF = \frac{W(0) - W(\alpha)}{W(0)} \leq 1 - O(|\mathcal{U}|^{-\frac{\alpha}{1+\alpha}}), \quad (6)$$

Remark 3 *The Theorem 3 holds that when increasing the item fairness degrees (α becomes larger) in a ranking system, there is a bound on the rate $1 - O(|\mathcal{U}|^{-\frac{\alpha}{1+\alpha}})$ at which utilities will decrease.*

5 EXPERIMENT

We evaluate the performance of α -rank. In the experiment, we mainly conduct the CTR/CVR-based fairness discussed in Table 1. For the exposure-based fairness, please see the Appendix I. The source code and experiments have been shared in supplementary file.

5.1 EXPERIMENTAL SETTINGS

Dataset. The experiments were based on two large-scale, publicly available ranking applications, including: **Yelp**¹: a large-scale businesses recommendation dataset. It has 154543 samples, which

¹<https://www.yelp.com/dataset>

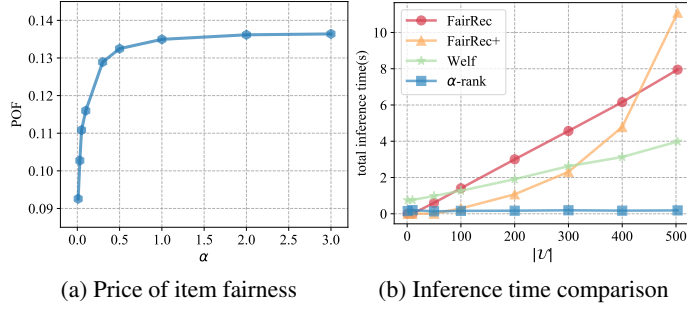


Figure 3: Sub-figure (a) illustrates the price of item fairness (POF) change w.r.t fairness degree α . Sub-figure (b) describes online inference items for α -rank and other baselines w.r.t user size $|\mathcal{U}|$.

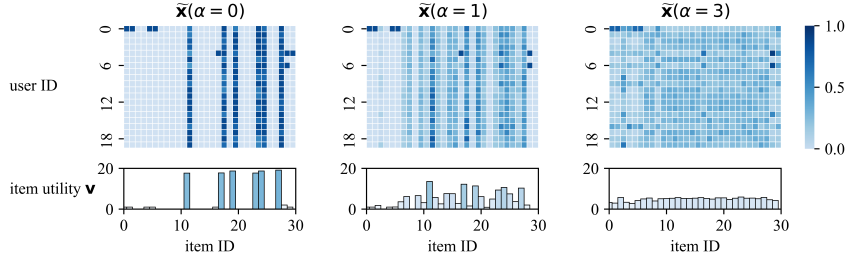


Figure 4: Visualization of α -rank result.

contains 17034 users, 11821 items. **Ipinyou** (Liao et al., 2014)²: a large-scale advertising dataset. We only used the clicked data, which contains 18588 samples, which contains 18565 users, 149 advertisements. Every advertisement has a bidding price.

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. Following (Zhang et al., 2022; Xu et al., 2023b), we used BPR (Rendle et al., 2012) model to compute the CTR-CVR value of user-item pair. For each user-item pair (u, i) , the model will output the CTR-CVR value $w_{u,i}$. For the item weight γ_i , a value of 1 is assigned for recommendation applications, while for advertising applications, $\gamma_i = \log(\text{bid}_i)$, where bid_i represents the bidding price of an advertisement.

Evaluation. As for the evaluation metrics, the performances of the models were evaluated from two aspects: social welfare, and fairness degree. As for the social welfare, following the practices in (Wu et al., 2021; Xu et al., 2023b; Yang et al., 2019), we utilized expected Click/Conversion Number (eCN) for recommendation application and expected Cost Per Mile (eCPM) for advertising application under top- K ranking.:

$$\text{eCN@K} = \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{I}} \mathbf{v}_i, \quad \text{eCPM@K} = \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{I}} \text{bid}_i \mathbf{v}_i. \quad (7)$$

As for the fairness degree, we utilized the Gini Index (Do & Usunier, 2022; Do et al., 2021), which is the most common measure of item utility inequality under top- K ranking. Formally, it defines as:

$$\text{Gini@K} = \frac{\sum_i \sum_j |\gamma_i \mathbf{v}_i - \gamma_j \mathbf{v}_j|}{2|\mathcal{I}| \sum_i \gamma_i \mathbf{v}_i}, \quad (8)$$

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 (Patro et al., 2020) and **FairRec+** (Biswas et al., 2021) proposed to ensure Max-Min Share (α -MMS) of exposure for the items. **Welf** (Do et al., 2021) use the Frank-Wolfe algorithm to maximize the Welfare functions of worst-off items. **P-MMF** (Xu et al., 2023a) utilized the mirror descent method to improve the worst-off item’s utility.

²<http://contest.ipinyou.com/>

5.2 EXPERIMENT RESULTS

Figure 2 shows the Pareto frontiers Xu et al. (2023a) of Gini Index (abbreviated as Gini.) and eCN/eCPM on two application datasets with different ranking size K . The Pareto frontiers were 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.

Analyzing the Pareto frontiers, it becomes evident that the proposed α -rank method consistently outperforms the baseline methods (as indicated by the α -rank curves occupying the upper right corner). This Pareto dominance signifies that, for a given eCN@K/eCPM@K level, α -rank achieves superior Gini@K values, and for a given Gini@K level, it attains better eCN@K/eCPM@K performance. These results highlight that α -rank significantly outperforms the baseline methods.

5.3 EXPERIMENT ANALYSIS

We also conducted experiments to analyze α -rank on Yelp for Top-10 ranking. For ablation studies and Lorenz curve Gastwirth (1971) analysis, please see Appendix H and Appendix G, respectively.

Price of item fairness. Firstly, we conducted an experiment to demonstrate how the price fairness of the item in Figure 3 (a) changes with respect to variations in the fairness degree α , ranging from 0.0 to 3.0. we directly compute the POF based on equation 6. From the curve, it is evident that as we increase the fairness degree α , the α -rank approach leads to a reduction in the total utilities of items. The experiment verified the theoretical analysis results in Theorem 3.

Inference time. We conducted experiments to investigate the total inference time of the α -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 relation to the varying number of users, all while keeping the number of items constant.

Figure 3 (b) reports the curves of total inference time (s) w.r.t. user size $|\mathcal{U}|$. It’s worth noting that the α -rank method exhibits a remarkably low inference time, 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 α -rank consistently maintains a low inference time. The α -rank method involves matrix operations with limited sensitivity to changes in user size.

Visualizing ranking results. In Figure 4, we visualize the ranking result matrix \tilde{x} and the utility vector v of items generated by the α -rank method for different values of α (0, 1, and 3), where these values correspond to utilitarianism, dealism, and egalitarianism, respectively. The distribution of vector v reflects the fairness degree of items. The detailed histogram of utility level of items under different α can be seen in Appendix F.

The results clearly demonstrate that the utilitarianism 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 dealism, α -rank approach tends to distribute rankings to items in proportion to their contribution to the market. For egalitarianism, α -rank method strives to provide equal exposure and similar utilities to every item in the ranking. The experiment also served as validation that α -rank method can effectively adapt to various fairness principles as intended.

6 CONCLUSION

This paper proposes α -rank model that aims to unify the item fairness in ranking from the cooperative game theory view. Firstly, we conducted an analysis of various fairness principles in ranking and unified these principles within the framework of cooperative game theory. Then we introduced the approach of α -rank can well-balance different fairness principles. Theoretical results to establish the maximum total utility loss for different values of α . Finally, Experiment results show that α -rank can outperform the state-of-the-art baselines efficiently and effectively.

REFERENCES

- Himan Abdollahpouri and Robin Burke. Multi-stakeholder recommendation and its connection to multi-sided fairness. *arXiv preprint arXiv:1907.13158*, 2019.
- Himan Abdollahpouri, Masoud Mansoury, Robin Burke, and Bamshad Mobasher. The unfairness of popularity bias in recommendation. *arXiv preprint arXiv:1907.13286*, 2019.
- Himan Abdollahpouri, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Pizzato. Multistakeholder recommendation: Survey and research directions. *User Modeling and User-Adapted Interaction*, 30(1):127–158, 2020.
- Ricardo Baeza-Yates, Berthier Ribeiro-Neto, et al. *Modern information retrieval*, volume 463. ACM press New York, 1999.
- Omer Ben-Porat and Moshe Tennenholtz. A game-theoretic approach to recommendation systems with strategic content providers. *Advances in Neural Information Processing Systems*, 31, 2018.
- Dimitri P Bertsekas. Nonlinear programming. *Journal of the Operational Research Society*, 48(3): 334–334, 1997.
- Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. The price of fairness. *Operations research*, 59(1):17–31, 2011.
- Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. On the efficiency-fairness trade-off. *Management Science*, 58(12):2234–2250, 2012.
- Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international acm sigir conference on research & development in information retrieval*, pp. 405–414, 2018.
- Arpita Biswas, Gourab K Patro, Niloy Ganguly, Krishna P Gummadi, and Abhijnan Chakraborty. Toward fair recommendation in two-sided platforms. *ACM Transactions on the Web (TWEB)*, 16(2):1–34, 2021.
- Rodica Branzei, Dinko Dimitrov, and Stef Tijs. *Models in cooperative game theory*, volume 556. Springer Science & Business Media, 2008.
- Virginie Do and Nicolas Usunier. Optimizing generalized gini indices for fairness in rankings. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 737–747, 2022.
- Virginie Do, Sam Corbett-Davies, Jamal Atif, and Nicolas Usunier. Two-sided fairness in rankings via lorenz dominance. *Advances in Neural Information Processing Systems*, 34:8596–8608, 2021.
- Joseph L Gastwirth. A general definition of the lorenz curve. *Econometrica: Journal of the Econometric Society*, pp. 1037–1039, 1971.
- Yingqiang Ge, Shuchang Liu, Ruoyuan Gao, Yikun Xian, Yunqi Li, Xiangyu Zhao, Changhua Pei, Fei Sun, Junfeng Ge, Wenwu Ou, et al. Towards long-term fairness in recommendation. In *Proceedings of the 14th ACM international conference on web search and data mining*, pp. 445–453, 2021.
- Qianxiu Hao, Qianqian Xu, Zhiyong Yang, and Qingming Huang. Pareto optimality for fairness-constrained collaborative filtering. In *Proceedings of the 29th ACM International Conference on Multimedia*, MM '21, pp. 5619–5627, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450386517. doi: 10.1145/3474085.3475706.
- Zhimeng Jiang, Xiaotian Han, Chao Fan, Fan Yang, Ali Mostafavi, and Xia Hu. Generalized demographic parity for group fairness. In *International Conference on Learning Representations*, 2021.
- Ehud Kalai. Proportional solutions to bargaining situations: interpersonal utility comparisons. *Econometrica: Journal of the Econometric Society*, pp. 1623–1630, 1977.

- Ehud Kalai and Meir Smorodinsky. Other solutions to nash’s bargaining problem. *Econometrica: Journal of the Econometric Society*, pp. 513–518, 1975.
- Anisio Lacerda, Marco Cristo, Marcos André Gonçalves, Weiguo Fan, Nivio Ziviani, and Berthier Ribeiro-Neto. Learning to advertise. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 549–556, 2006.
- Julian Lamont. *Distributive justice*. Routledge, 2017.
- Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, and Yongfeng Zhang. Towards personalized fairness based on causal notion. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1054–1063, 2021.
- Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, Juntao Tan, Shuchang Liu, and Yongfeng Zhang. Fairness in recommendation: A survey. *arXiv preprint arXiv:2205.13619*, 2022.
- Hairen Liao, Lingxiao Peng, Zhenchuan Liu, and Xuehua Shen. ipinyou global rtb bidding algorithm competition dataset. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*, pp. 1–6, 2014.
- Aldo Lipani. Fairness in information retrieval. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pp. 1171–1171, 2016.
- Xiangyu Liu, Chuan Yu, Zhilin Zhang, Zhenzhe Zheng, Yu Rong, Hongtao Lv, Da Huo, Yiqing Wang, Dagui Chen, Jian Xu, et al. Neural auction: End-to-end learning of auction mechanisms for e-commerce advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 3354–3364, 2021.
- Mirko Marras, Ludovico Boratto, Guilherme Ramos, and Gianni Fenu. Equality of learning opportunity via individual fairness in personalized recommendations. *International Journal of Artificial Intelligence in Education*, 32(3):636–684, 2022.
- David Matsumoto and Linda Juang. *Culture and psychology*. Cengage Learning, 2016.
- Mohammadmehdi Naghiaei, Hossein A Rahmani, and Yashar Deldjoo. Cpfair: Personalized consumer and producer fairness re-ranking for recommender systems. *arXiv preprint arXiv:2204.08085*, 2022.
- John Nash. Non-cooperative games. *Annals of mathematics*, pp. 286–295, 1951.
- John F Nash Jr. The bargaining problem. *Econometrica: Journal of the econometric society*, pp. 155–162, 1950.
- Gourab K Patro, Arpita Biswas, Niloy Ganguly, Krishna P Gummadi, and Abhijnan Chakraborty. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *Proceedings of The Web Conference 2020*, pp. 1194–1204, 2020.
- Bezalel Peleg and Peter Sudhölter. *Introduction to the theory of cooperative games*, volume 34. Springer Science & Business Media, 2007.
- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Khiem Pham, Khang Le, Nhat Ho, Tung Pham, and Hung Bui. On unbalanced optimal transport: An analysis of sinkhorn algorithm. In *International Conference on Machine Learning*, pp. 7673–7682. PMLR, 2020.
- John Rawls. A theory of justice. *Cambridge (Mass.)*, 1971.
- Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618*, 2012.
- Harvey M Salkin and Cornelis A De Kluyver. The knapsack problem: a survey. *Naval Research Logistics Quarterly*, 22(1):127–144, 1975.

- Amartya Sen. Utilitarianism and welfarism. *The journal of Philosophy*, 76(9):463–489, 1979.
- Lloyd S Shapley et al. A value for n-person games. 1953.
- Kyle Swanson, Lili Yu, and Tao Lei. Rationalizing text matching: Learning sparse alignments via optimal transport. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5609–5626, 2020.
- Tom R Tyler and E Allan Lind. Procedural justice. In *Handbook of justice research in law*, pp. 65–92. Springer, 2002.
- Tom R Tyler and Heather J Smith. Social justice and social movements. 1995.
- John Von Neumann and Oskar Morgenstern. Theory of games and economic behavior, 2nd rev. 1947.
- Yao Wu, Jian Cao, Guandong Xu, and Yudong Tan. 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*, pp. 1013–1022, 2021.
- Chen Xu, Sirui Chen, Jun Xu, Weiran Shen, Xiao Zhang, Gang Wang, and Zhenhua Dong. P-mmf: Provider max-min fairness re-ranking in recommender system. In *Proceedings of the ACM Web Conference 2023*, pp. 3701–3711, 2023a.
- Chen Xu, Xiaopeng Ye, Jun Xu, Xiao Zhang, Weiran Shen, and Ji-Rong Wen. Ltp-mmf: Towards long-term provider max-min fairness under recommendation feedback loops. *arXiv preprint arXiv:2308.05902*, 2023b.
- Jun Xu, Xiangnan He, and Hang Li. Deep learning for matching in search and recommendation. In *The 41st International ACM SIGIR Conference on Research Development in Information Retrieval*, SIGIR '18, pp. 1365–1368, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450356572. doi: 10.1145/3209978.3210181.
- Peng Xu and Kanapathippillai Cumanan. Optimal power allocation scheme for non-orthogonal multiple access with α -fairness. *IEEE Journal on Selected Areas in Communications*, 35(10):2357–2369, 2017.
- Hong-Jian Xue, Xinyu Dai, Jianbing Zhang, Shujian Huang, and Jiajun Chen. Deep matrix factorization models for recommender systems. In *IJCAI*, volume 17, pp. 3203–3209. Melbourne, Australia, 2017.
- Xun Yang, Yasong Li, Hao Wang, Di Wu, Qing Tan, Jian Xu, and Kun Gai. Bid optimization by multivariable control in display advertising. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 1966–1974, 2019.
- Xiao Zhang, Sunhao Dai, Jun Xu, Zhenhua Dong, Quanyu Dai, and Ji-Rong Wen. Counteracting user attention bias in music streaming recommendation via reward modification. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 2504–2514, 2022.

A PROOF OF THEOREM 1

A.1 AXIOM 1

Suppose exist another solution $\mathbf{v} \in \mathcal{U}$ so that utility vector $\mathbf{v} \geq \mathbf{v}^f$ and $\mathbf{v}^f \neq \mathbf{v}$.

For the utilitarianism, we have

$$\mathbf{1}^\top \mathbf{v} \geq \mathbf{1}^\top \mathbf{v}^f, \quad \mathbf{v}^f \neq \mathbf{v},$$

which verifies that \mathbf{v} is a better solution compared to \mathbf{v}^f . There is a contradiction. Therefore, utilitarianism satisfies the Axiom 1.

For the dealism, we have

$$\mathbf{1}^\top \log(\mathbf{v}) \geq \mathbf{1}^\top \log(\mathbf{v}^f), \quad \mathbf{v}^f \neq \mathbf{v},$$

which also verifies that \mathbf{v} is a better solution compared to \mathbf{v}^f . There is a contradiction. Therefore, dealism satisfies the Axiom 1.

For the egalitarianism, we have

$$\mathbf{v}_i = \mathbf{v}_j, \forall i, j \in \mathcal{I}, \mathbf{v} \in \mathcal{U}.$$

Therefore, if $\mathbf{v} \geq \mathbf{v}^f$ and $\mathbf{v}^f \neq \mathbf{v}$, there must be

$$\mathbf{v}_i > \mathbf{v}_i^f, \forall i \in \mathcal{I},$$

and therefore,

$$\min_i \mathbf{v}_i > \min_i \mathbf{v}_i^f,$$

which is a contradiction with the optimization objective. Therefore, egalitarianism satisfy the Axiom 1.

Q.E.D.

A.2 AXIOM 2

As the formulation doesn't take item attributes into account and applies the same operation to all items, it becomes evident that the Axiom holds.

A.3 AXIOM 3

For utilitarianism, consider the following CTR matrix for recommendation,

$$\begin{bmatrix} 0.1 & 0.2 & 0.3 \\ 0.5 & 0.4 & 0.1 \\ 0.7 & 0.5 & 0.3 \end{bmatrix},$$

where the element in i-th row and j-column is $w_{u,i}$ and ranking size is 1. In such a matrix, the utilitarianism solution is $0.3+0.5+0.7 = 1.5$ since it always chooses the highest value among the items of certain user (i.e. selects the highest value of one row). However, we apply the transformation A , such that the matrix becomes

$$\begin{bmatrix} 0.1 * 4 & 0.2 & 0.3 \\ 0.5 * 4 & 0.4 & 0.1 \\ 0.7 * 4 & 0.5 & 0.3 \end{bmatrix},$$

then the utilitarianism solution is $0.4 + 2 + 2.8 \neq 4 * 1.5$. Therefore, utilitarianism does not satisfy Axiom 3.

For dealism, for the transformation A , we have

$$f(A(\mathbf{v})) = \max_i \sum \log(c_i \mathbf{v}_i) = \max_i \sum \log(\mathbf{v}_i) = f(\mathbf{v}).$$

Therefore, dealism satisfies Axiom 3.

For egalitarianism, it is easy to observe if could happen

$$\arg \min_i \mathbf{v}_i \neq \arg \min_i c_i \mathbf{v}_i,$$

therefore, egalitarianism does not satisfy Axiom 3.

A.4 AXIOM 4

It is evident that all three optimization objectives do not take resource available levels into consideration. In simpler terms, if the resources allocated to an item are reduced, and the initial solution still falls within the feasible region, the solution remains unchanged from the original.

A.5 AXIOM 5

For the utilitarianism, consider the matrix

$$\begin{bmatrix} 0.1 & 0.2 & 0.3 \\ 0.5 & 0.4 & 0.1 \end{bmatrix},$$

where the element in i -th row and j -column is $w_{u,i}$ and ranking size is 1. In the context of ranking applications, we assume that each item is exposed at most once. Therefore, the ranking model will choose 0.3 in the first row and 0.5 in the second row. However, if we increase the resource to

$$\begin{bmatrix} 0.1 & 0.2 & 0.3 \\ 0.5 & 0.3 & 0.1 \\ 0.7 & 0.1 & 0.1 \end{bmatrix},$$

the ranking model will choose 0.3 in the first row, 0.3 in the second row and 0.7 in the third row. The first item's utility decreases from 0.4 to 0.3. Therefore, utilitarianism does not satisfy Axiom 5.

Similarly, dealism does not satisfy Axiom 5.

For egalitarianism, $\mathcal{D}_1 \subset \mathcal{D}_2$ and $\mathcal{D}_1 \neq \mathcal{D}_2$, we have

$$\mathbf{v}_i = \mathbf{v}_j, \forall i, j \in \mathcal{I}, \mathbf{v} \in \mathcal{U},$$

and since it satisfies the Axiom 1, we therefore have

$$\arg \max_{\mathbf{v} \in \mathcal{D}_1} f(\mathbf{v}) \leq \arg \max_{\mathbf{v} \in \mathcal{D}_2} f(\mathbf{v}).$$

Q.E.D.

B PROOF OF THEOREM 2

According to the equation 1, we have

$$\begin{aligned} W(\alpha) &= \max f(\mathbf{v}; \alpha), \\ \text{s.t. } \mathbf{1}^\top \mathbf{x}_u &= K, \quad \forall u \in \mathcal{U} \\ \mathbf{x}_{u,i} &= \{0, 1\}, \quad \forall i \in \mathcal{I}, u \in \mathcal{U} \\ \mathbf{v}_i &= \gamma_i \sum_{u \in \mathcal{U}} w_{u,i} \mathbf{x}_{u,i}, \quad \forall i \in \mathcal{I} \end{aligned}$$

Let abbreviate $|\mathcal{U}| = N$ and $\mathbf{e}_i = \frac{1}{N} \sum_{u \in \mathcal{U}} \mathbf{x}_{u,i}$.

Therefore, we can relax the first condition as

$$\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} \mathbf{x}_{u,i} = NK,$$

i.e.

$$\sum_{i \in \mathcal{I}} \mathbf{e}_i = K \tag{9}$$

We also relax the second condition as

$$0 \leq \sum_{u \in \mathcal{U}} \mathbf{x}_{u,i} \leq N, \forall i \in \mathcal{I}$$

i.e.

$$0 \leq \mathbf{e}_i \leq 1, \forall i \in \mathcal{I}. \tag{10}$$

Finally, there exists $\tau > 0$, we can relax the third condition as

$$\sum_{u \in \mathcal{U}} w_{u,i} \mathbf{x}_{u,i} \leq \tau \sum_{u \in \mathcal{U}} w_{u,i} \sum_{u \in \mathcal{U}} \mathbf{x}_{u,i},$$

i.e.

$$\sum_{u \in \mathcal{U}} w_{u,i} \mathbf{x}_{u,i} \leq N \eta_i \mathbf{e}_i, \quad (11)$$

where $\eta_i = \tau \sum_{u \in \mathcal{U}} w_{u,i}$.

Combining Equation 9, 10, 11, we have $\hat{W}(\alpha) \geq W(\alpha)$.

Q.E.D

C LEMMA 1

Lemma 1 Given the vector $\mathbf{x} \in \mathbb{R}^N$, $\sum_{i=1}^N w_i \mathbf{x}_i^{1-\alpha} \geq C$, for any $\alpha > 0$, $w_i > 0$, $\mathbf{x}_i > 0$, we have

$$\sum_{i=1}^N w_i \mathbf{x}_i \geq C \sigma^\alpha N^{-\frac{\alpha}{1+\alpha}}, \quad (12)$$

where $\sigma = \min_i \mathbf{x}_i$.

C.1 PROOF OF LEMMA 1

We can observe that $f(x) = x^{1+\alpha}$ is a concave function, therefore, we apply the Hölder inequality, we have when $\frac{1}{1+\alpha} + \frac{\alpha}{1+\alpha} = 1$,

$$C \leq \sum_{i=1}^N w_i \mathbf{x}_i * \frac{1}{\mathbf{x}_i^\alpha} \leq \left(\sum_{i=1}^N w_i^{1+\alpha} \mathbf{x}_i^{1+\alpha} \right)^{\frac{1}{1+\alpha}} \left(\sum_{i=1}^N \frac{1}{\mathbf{x}_i^{1+\alpha}} \right)^{\frac{\alpha}{1+\alpha}} \quad (13)$$

From concave attribute, we have

$$\left(\sum_{i=1}^N w_i \mathbf{x}_i \right)^{1+\alpha} \geq \sum_{i=1}^N w_i^{1+\alpha} \mathbf{x}_i^{1+\alpha},$$

therefore, we have

$$\left(\sum_{i=1}^N w_i^{1+\alpha} \mathbf{x}_i^{1+\alpha} \right)^{\frac{1}{1+\alpha}} \leq \sum_{i=1}^N w_i \mathbf{x}_i \quad (14)$$

Let $\sigma = \min_i \mathbf{x}_i$, we have

$$\left(\sum_{i=1}^N \frac{1}{\mathbf{x}_i^{1+\alpha}} \right)^{\frac{\alpha}{1+\alpha}} \leq \frac{N^{\frac{\alpha}{1+\alpha}}}{\sigma^\alpha} \quad (15)$$

Then combining equation 14 and equation 15 the equation 13 becomes:

$$\sum_{i=1}^N w_i \mathbf{x}_i \geq C \sigma^\alpha N^{-\frac{\alpha}{1+\alpha}}$$

Q.E.D

D PROOF OF THEOREM 3

To simply the symbol, let $m_i = \gamma_i \eta_i$. The objective becomes:

$$\begin{aligned} \hat{W}(\alpha) &= \max_e \sum_i m_i \frac{e_i^{1-\alpha}}{1-\alpha} \\ \text{s.t.} \quad &\sum_i e_i = K \\ &0 \leq e_i \leq 1 \end{aligned} \quad (16)$$

, where the input e optimal value is \mathbf{z} , which represents be the best allocation under the α -fairness criterion.

Firstly, we will bound the \hat{W}^α : without generality, we assume that:

$$m_1 \mathbf{z}_1 \geq m_2 \mathbf{z}_2 \geq \dots, \geq m_{|\mathcal{I}|} \mathbf{z}_{|\mathcal{I}|}. \quad (17)$$

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

$$\nabla \hat{W}^\alpha(\mathbf{z})(e - \mathbf{z}) \leq 0, \forall e \in \mathcal{E},$$

where

$$\mathcal{E} = \{e \mid \sum_i e_i = K, \quad 0 \leq e_i \leq 1\}.$$

The equation can be equivalently written as:

$$\mathbf{g}^\top \mathbf{e} \leq 1, \forall e \in \mathcal{E}, \quad (18)$$

where

$$\mathbf{g}_i = \frac{m_i \mathbf{z}_i^{-\alpha}}{\sum_i m_i \mathbf{z}_i^{1-\alpha}}. \quad (19)$$

We observe the Equation equation 18, which is a well-studied knapsack problem (Salkin & De Kluyver, 1975), with the best solution:

$$\frac{\sum_{k=1}^K m_{|\mathcal{I}|-k+1} \mathbf{z}_{|\mathcal{I}|-k+1}^{-\alpha}}{\sum_i m_i \mathbf{z}_i^{1-\alpha}} \leq 1, \quad (20)$$

since according the equation 17, we have

$$m_1 \mathbf{z}_1^{-\alpha} \leq m_2 \mathbf{z}_2^{-\alpha} \leq \dots, \leq m_{|\mathcal{I}|} \mathbf{z}_{|\mathcal{I}|}^{-\alpha}.$$

From the equation 20, exists $0 < \lambda < 1$, we have:

$$\sum_i m_i \mathbf{z}_i^{1-\alpha} \geq \lambda W^0 \mathbf{z}_1^{-\alpha}, \quad (21)$$

since

$$\sum_{k=1}^K m_{|\mathcal{I}|-k+1} \mathbf{z}_{|\mathcal{I}|-k+1}^{-\alpha} \geq \mathbf{z}_1^{-\alpha} \lambda \sum_{k=1}^K m_{[k]} = \lambda W^0 \mathbf{z}_1^{-\alpha},$$

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

Taking the equation 21 into Lemma 1, we have

$$\sum_i m_i \mathbf{z}_i \geq \lambda W^0 \left(\frac{\mathbf{z}_{|\mathcal{I}|}}{\mathbf{z}_1} \right)^\alpha N^{-\frac{\alpha}{1+\alpha}}. \quad (22)$$

Therefore,

$$\begin{aligned} \text{POF} &= \frac{W(0) - W(\alpha)}{W(0)} \leq \frac{W(0) - \hat{W}(\alpha)}{W(0)} \\ &\leq 1 - \lambda \left(\frac{\mathbf{z}_{|\mathcal{I}|}}{\mathbf{z}_1} \right)^\alpha N^{-\frac{\alpha}{1+\alpha}} \\ &= 1 - O(N^{-\frac{\alpha}{1+\alpha}}). \end{aligned}$$

Q.E.D.

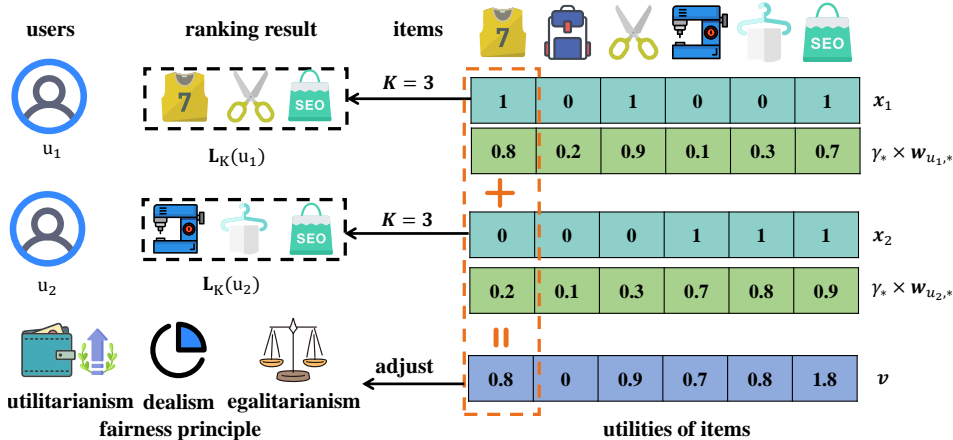


Figure 5: An illustrative example to visualize item fairness in ranking applications.




Figure 6: Visualization of items utilities for different α .

E AN ILLUSTRATIVE EXAMPLE OF ITEM FAIRNESS

In this section, we will give an illustrative example of item fairness.

Figure 5 provides an illustrative example of item fairness in the ranking process. In this scenario, we assume the recommender system includes users u_1 and u_2 , and there are six items in the RS available for recommendation to these users.

In a simple way, we suppose u_1 and u_2 , and the system ranks a list of $K = 3$ items from six item candidates. The ranking model will make the ranking list $L_K(u_1), L_K(u_2)$ based on the decision variable x . For example, in the case of user u_1 , the recommendation will expose items with a value of 1 for the decision vector x .

We also set the weight $\gamma_i \times w_{u,i}$ of item i . The utilities v of items are determined by aggregating the scores of the exposed items across the entire user set. For example, for the first item i_1 (yellow cloth), the utility v_1 is computed as $1 * 0.8 + 0 * 0.2 = 0.8$. We aim to balance the utilities of items from utilitarianism, dealism, and egalitarianism, respectively.

F VISUALIZATION OF ITEM UTILITY

To give a better understanding of how the hyperparameter α influences fairness under different ethical principles (utilitarianism, dealism, and egalitarianism), we plotted histograms of item utility for each of the different values of α (0, 1, and 3).

As illustrated in Figure 6, when α increases from 0 to 3, the maximum utility of items decreases, while the minimum utility increases. This indicates a shift toward a more equitable and fair distribution of utility among items.

When $\alpha = 0$ (corresponding to utilitarianism), the difference in utility level between items is very large (more than 56.9). When $\alpha = 3$ (corresponding to egalitarianism), utility levels converge, and the difference between the highest and lowest utility narrows to a mere 8.34.

Obviously, when α is small, α -rank focuses more on maximizing the overall utility, allocating most exposure toward the item with the largest $w_{u,i}$. On the contrary, for larger values of α , α -rank places greater emphasis on fairness among items. When α falls in between, α -rank strives to achieve a balance between accuracy and fairness, ensuring that more items receive adequate exposure while enhancing the overall utility of the item group.

G LORENZ CURVE OF DIFFERENT FAIRNESS MODELS

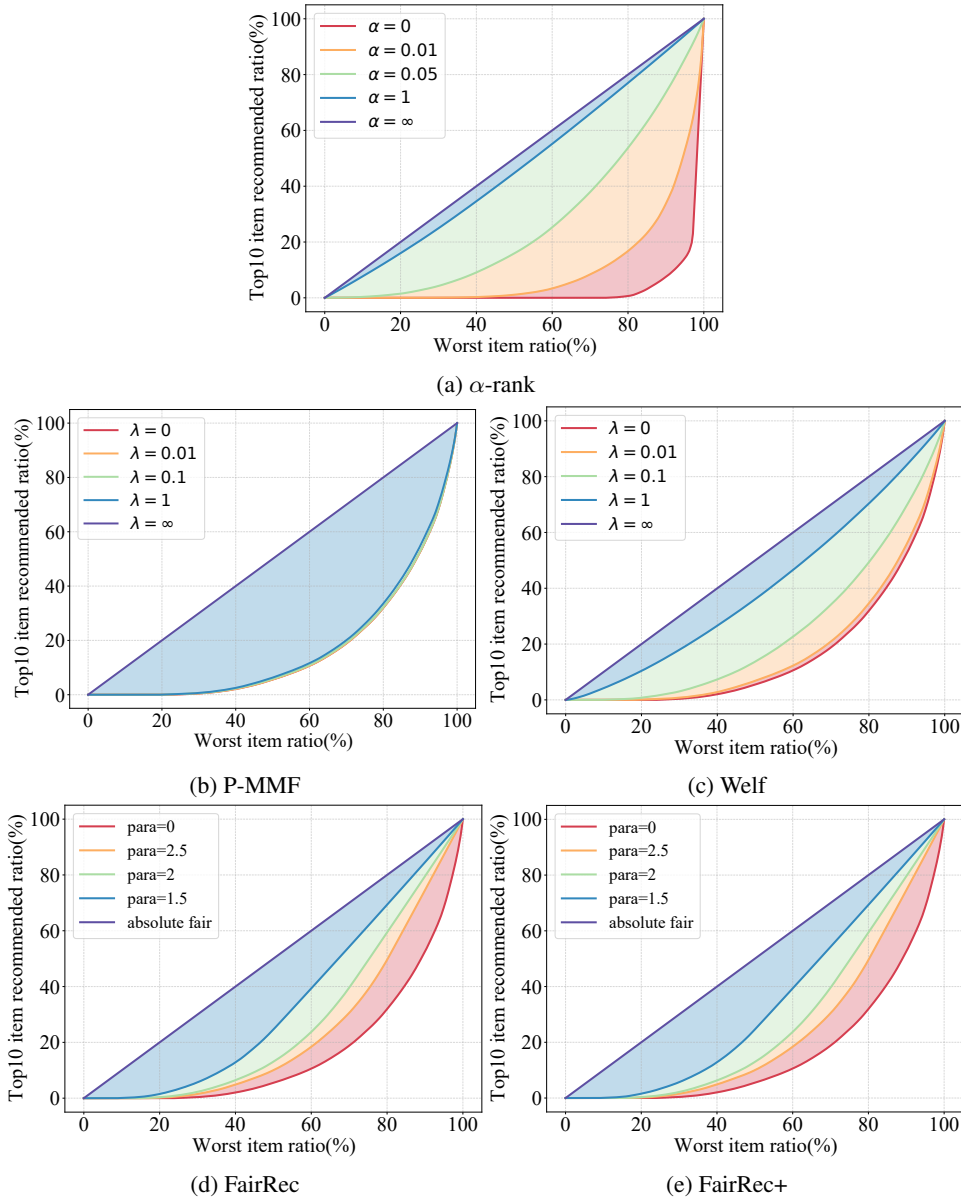


Figure 7: The Lorenz curves for different models on item exposure.

To comprehensively analyze the item fairness of α -rank and other baselines, we conducted an analysis employing Lorenz curves Gastwirth (1971), as shown in Figure 7, to visualize the distribution of item utility. Lorenz curve is a well-established tool for assessing distributional inequality in dif-

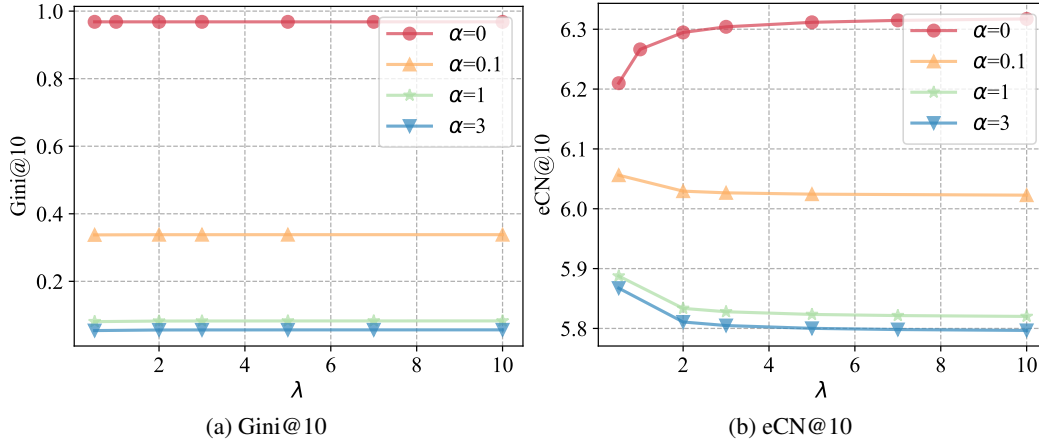


Figure 8: Gini and eCN change w.r.t the coefficient of entropy regularizer λ in OT under different α

ferent scenarios. The gini coefficient is the area between the Lorenz curve and the absolute fairness line.

Here are the steps we took to plot the Lorenz curve. Firstly, we sort the items according to their utility level, from low to high. Then, we calculate the percentage of total utility for each percentage of items separately. Finally, we plot these percentage data as a curve that shows the proportion of utility ($y\%$) occupied by the worst $x\%$ of items. Note that the x -axis represents the cumulative percentage of item numbers, while the y -axis represents the cumulative utility percentage. It is crucial to highlight that the diagonal line extending towards the upper right corner signifies a state of absolute fairness.

We can see that in Figure 7 (a), when $\alpha \rightarrow 0$, α -rank predominantly emphasizes total utility level, and the fairness shows poor performance. In contrast, as $\alpha \rightarrow \infty$, α -rank increasingly prioritizes fairness, resulting in a distribution of item utility that tends towards egalitarianism.

We can show that influence can be exerted on the equity of utility distribution through parameter tuning of other baselines P-MMF, Welf, FairRec, and FairRec+. Compared with other baseline models, α -rank has a greater range of control over the distribution through hyperparameter adjustments. As the parameter α varies from 0 to infinity, it causes the Lorenz curve to shift in a manner where all item utilities can be taken into consideration. However, other baseline methods are sensitive to their parameters, which can make them challenging to adjust with respect to the requirements of practical applications.


H ABLATION EXPERIMENT OF α -RANK

To better investigate the performance of our model under different parameters settings, we conducted a series of ablation experiments on the Yelp dataset under ranking size $K = 10$. Similar experiment results are also observed on other dataset and other ranking size K .

H.1 ABLATION STUDY ON OT COEFFICIENT λ

In this section, we conducted experiments using two metrics, Gini@10 and eCN@10, to investigate how the OT regularizer coefficient λ influences model performance. We adjusted λ value within the range of $[0, 10]$.

Figure 8 (a) illustrates how the Gini metric changes with λ . We can observe that Gini remains relatively stable across the λ values ranging from 0 to 10. Therefore, we can conclude that λ value is not sensitive to the fairness metric Gini. Moreover, as the parameter α increases, α -rank begins to prioritize fairness among items more prominently, resulting in a decrease in Gini.

Figure 9: Gini and eCN change w.r.t user size $|\mathcal{U}|$

Curves in Figure 8 (b) report the change in eCN with respect to λ , ranging from 0 to 10. We can see that the eCN value shows a drastic change when λ transitions from 0 to 4. Specifically, when $\alpha = 0$, eCN increases by 1.6% (from 6.21 to 6.31), and when $\alpha = 3$, eCN decreases by 1.3% (from 5.87 to 5.80). As α increases, the decreasing trend of eCN@10 becomes more pronounced.

In the OT problem, λ controls the size of the entropy regularizer. By tuning the parameter λ , we can trade-off between smoothness and convexity of the exposure probability distribution among items. A higher λ will emphasize more on the smoothness of the probability distribution, implying a reduced disparity in exposure among items. When α is relatively small ($\alpha \leq 0.01$), 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 α becomes larger, the increased smoothness of probability distributions leads to a decrease in exposure for items with higher $w_{u,i}$, consequently resulting in a decrease in eCN@10.

H.2 ABLATION STUDY ON USER SIZE $|\mathcal{U}|$

In this section, we conducted experiments on the performance of α -rank with respect to user size $|\mathcal{U}|$, which is significant for understanding how α -rank adapts to different user sizes. For the experimental setting, we randomly sampled $|\mathcal{U}| \in [5, 500]$ numbers of users. Finally, we plotted the curve based on the mean and variance of Gini@10 and eCN@10.

Figure 9 (a) illustrates that Gini@10 remains nearly constant regardless of changes in user size. This shows that α -rank demonstrates consistent fairness performance even when confronted with varying user sizes.

However, Figure 9 (b) reveals a decline in eCN@10 as user size increases. Furthermore, the extent of eCN@10 reduction diminishes with higher α values. This is because when the number of users is relatively small, their interests tend to be more focused, making it easier for α -rank to allocate exposure. As α increases from 0 to 3, α -rank transitions from prioritizing accuracy to equalizing exposure across items.

H.3 ABLATION STUDY ON ITEM SIZE $|\mathcal{I}|$

Similarly, we study the influence of item size $|\mathcal{I}|$ on the performance of α -rank. To conduct our analysis, we randomly sampled $|\mathcal{I}| \in [20, 300]$ numbers of items for three times and plotted the figure using the mean and variance of Gini@10 and eCN@10.

The curves, drawn in Figure 10 (a), reveal a notable trend: Gini@10 increases as item size grows. The trend is particularly pronounced when α is relatively small ($\alpha < 0.01$). In contrast, when α is relatively large (i.e. $\alpha \geq 1$), we can barely see the trend of increase, and the curves nearly overlap, where Gini@10 is extremely close to 0. This observation highlights the role of the parameter α in




Figure 10: Gini and eCN change w.r.t the item size $|Z|$

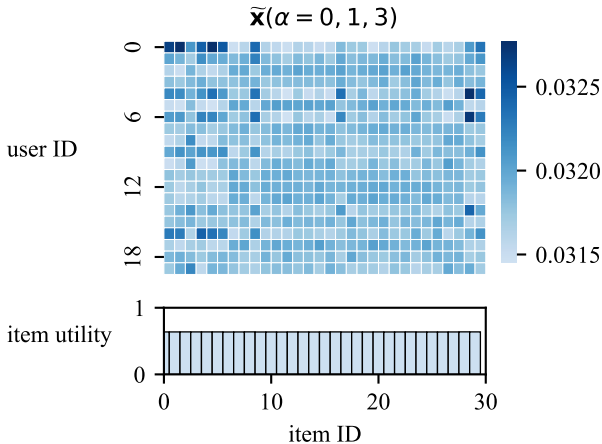


Figure 11: Visualization of exposure-based α -rank result

stabilizing item fairness, even when dealing with large item sizes. Moreover, it demonstrates that by adjusting α , we can maintain item fairness even if the item size is high.

In Figure 10 (b), we observed a similar pattern to Figure 10 (a). With item size $|Z|$ increases, eCN@10 also rises and such increase becomes more and more apparent as α increases. This suggests that α -rank can achieve a more accurate and user-satisfying result when applied to larger item datasets.

I EXPOSURE-BASED FAIRNESS

In this section, we aim to analyze the exposure-based fairness Xu et al. (2023a); Patro et al. (2020) performance of α -fairness and other baselines. In exposure-based fairness, we set $w_{u,i} = 1$ discussed in Table 1. The exposure-based fairness focuses on the exposure of each item as their utilities.

I.1 VISUALIZING RANKING RESULT

In Figure 11, we visualized the ranking result matrix \tilde{x} and the utility vector v of exposure-based α -rank at different values of alpha (0, 1, and 3). As we can see in Figure 11, no matter what value of α we choose, exposure-based α -rank tends to equalize item exposure and allocate utility more fairly among items.

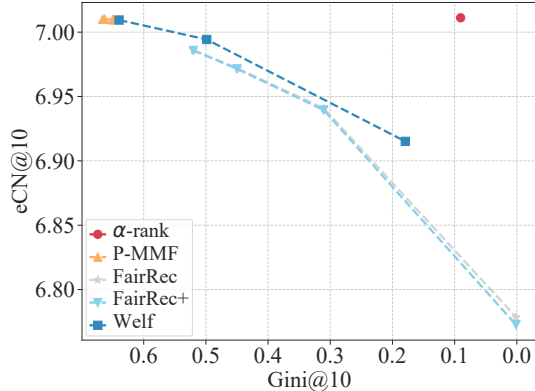


Figure 12: Pareto Frontier of exposure-based α -rank and other baselines




Figure 13: Gini and eCN change w.r.t the coefficient of OT regularizer λ in OT under different α

I.2 PARETO FRONTIER

To verify the effectiveness of α -rank in exposure-based fairness, we conducted a comparative analysis by plotting the Pareto frontiers of α -rank against all baseline methods. Figure 12 illustrates the performance of α -rank in comparison to other baseline algorithms. From the figure, we can observe that α -rank also outperforms the baselines in terms of both Gini@10 and eCN@10 metrics in the realm of exposure-based fairness. These results also verify the effectiveness of α -rank over the baselines.

I.3 ABLATION STUDY ON OT COEFFICIENT λ

We also study the effect of the hyperparameter λ in the performance of α -rank in exposure-based fairness. Specifically, we explore how varying values of λ influence the quality of recommendations when α is set to different values within the range of 0 to 3.

Figure 13 presents that both Gini@10 and eCN@10 curves show a consistent decreasing trend as λ increases. This trend becomes less pronounced as λ reaches larger values. The experiment also verifies the parameter λ in the OT projection controls the smoothness and the convexity of the exposure probability distribution between items.

Furthermore, it is noteworthy that curves for different α are basically the same. In the exposure-based α -rank algorithm, the weight η_i of items in equation 3 is the same for all items. In simpler terms, every item holds the same weight within the system, resulting in the same underlying principle regardless of the specific value of α .