

Fair Ranking as Fair Division: Impact-Based Individual Fairness in Ranking

Yuta Saito
Cornell University
Ithaca, NY, USA
ys552@cornell.edu

Thorsten Joachims
Cornell University
Ithaca, NY, USA
tj@cs.cornell.edu

ABSTRACT

Rankings have become the primary interface in two-sided online markets. Many have noted that the rankings not only affect the satisfaction of the users (e.g., customers, listeners, employers, travelers), but that the position in the ranking allocates exposure – and thus economic opportunity – to the ranked items (e.g., articles, products, songs, job seekers, restaurants, hotels). This has raised questions of fairness to the items, and most existing works have addressed fairness by explicitly linking item exposure to item relevance. However, we argue that any particular choice of such a link function may be difficult to defend, and we show that the resulting rankings can still be unfair. To avoid these shortcomings, we develop a new axiomatic approach that is rooted in principles of *fair division*. This not only avoids the need to choose a link function, but also more meaningfully quantifies the impact on the items beyond exposure. Our axioms of *envy-freeness* and *dominance over uniform ranking* postulate that for a fair ranking policy every item should prefer their own rank allocation over that of any other item, and that no item should be actively disadvantaged by the rankings. To compute ranking policies that are fair according to these axioms, we propose a new ranking objective related to the *Nash Social Welfare*. We show that the solution has guarantees regarding its envy-freeness, its dominance over uniform rankings for every item, and its Pareto optimality. In contrast, we show that conventional exposure-based fairness can produce large amounts of envy and have a highly disparate impact on the items. Beyond these theoretical results, we illustrate empirically how our framework controls the trade-off between impact-based individual item fairness and user utility.

CCS CONCEPTS

• **Information systems** → **Probabilistic retrieval models**; *Social recommendation*; • **Computing methodologies** → **Ranking**.

KEYWORDS

Fairness in Ranking, Fair Division, Nash Social Welfare

ACM Reference Format:

Yuta Saito and Thorsten Joachims. 2022. Fair Ranking as Fair Division: Impact-Based Individual Fairness in Ranking. In *Proceedings of the 28th*

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.

KDD '22, August 14–18, 2022, Washington, DC, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9385-0/22/08...\$15.00

<https://doi.org/10.1145/3534678.3539353>

ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22), August 14–18, 2022, Washington, DC, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3534678.3539353>

1 INTRODUCTION

Ranking interfaces are widely used to mediate online market platforms, ranging from booking a hotel to finding qualified employees. In these two-sided markets, not only the users (e.g., travelers and employers) obtain utility from the rankings, but the ranking policy also has a crucial impact on the items being ranked (e.g., hotels and job seekers). Many have noted that conventional ranking approaches [26], which exclusively optimize the utility to the users, may be unfair to the items. In particular, conventional rankings can give items with similar merit (i.e., relevance) very dissimilar exposure and thus economic opportunity [3, 27, 32].

To address this disparity of treatment, existing works on fairness-in-ranking explore how to explicitly link merit to the exposure that is allocated to individual items or groups of items [3, 23, 27, 28, 33, 35, 36]. Concrete definitions of exposure-based fairness were introduced in the seminal works by Singh and Joachims [27] and Biega et al. [3], enforcing constraints to allocate exposure proportional to relevance in expectation over groups or amortized over a sequence of rankings. This fairness-of-exposure approach has since been extended to end-to-end policy learning [24, 28, 33], dynamic ranking [23, 35], evaluation metrics [14], and contextual bandits [20, 32]. However, a key criticism of exposure-based fairness is the necessity to select a function that links merit and exposure, and there is typically no justification that this relationship should be linear – nor would any other choice be more defensible. Furthermore, we show that the fairness-of-exposure approach can still lead to rankings that violate basic fairness principles.

To overcome these shortcomings, we develop an axiomatic approach to fairness in ranking that is rooted in principles of fair division. In particular, we provide a novel formulation of fairness in ranking as a resource allocation problem, where the resources are the positions in a ranking that need to be fairly divided among the items. To concisely define fairness, we propose *envy-freeness* and *dominance over uniform ranking* as two fairness axioms that parallel widely accepted concepts from the fair-division literature [8, 18, 31]. Envy-freeness requires that no item can gain better impact by exchanging its position allocation with another item in a market. Dominance over uniform ranking requires that every item should gain better impact than under the uniform random ranking policy, ensuring that all items draw some benefit from participating in the platform. This axiomatization not only avoids the need for an arbitrary link function between relevance and exposure, it also directly captures the *impact* (e.g. clicks, conversions, revenue, streams) of

the ranking on the items, not just their visibility as in the exposure-based framework. Prior attempts to extend fairness-of-exposure to individual fairness of impact have already been shown to lead to degenerate ranking policies [27], making ours the first viable formulation of *impact-based individual item fairness*.

The key remaining problem is to find an algorithm for computing rankings that are guaranteed to fulfill the fairness axioms, and that also guarantee a notion of optimality. To this effect, we introduce a new objective for optimizing rankings that is related to the *Nash Social Welfare* (NSW). Under our formulation, maximizing the NSW is a convex program that is efficiently solvable. We prove that the ranking policy that maximizes the NSW-objective has guarantees regarding the envy-freeness of the position allocations and its dominance over the uniform ranking policy for each item. Furthermore, the rankings are Pareto optimal. We also develop an extension of the NSW called the α -NSW, which enables a steerable trade-off between user utility and item fairness through a single hyper-parameter. Beyond these conceptual and methodological contributions, we conduct extensive experiments on synthetic and real-world data. We find that conventional exposure-based fairness can produce large amounts of envy and disparate impact on the items. In contrast, the proposed NSW-maximizing policy achieves an almost envy-free impact distribution even with noisy relevance predictions, and provides more equitable benefit to all items.

2 RELATED WORK

Fair Ranking in Two-Sided Markets. There exist several notions of fairness in ranking, ranging from the ones based on the composition of a prefix of the ranking [1, 11, 12, 34] to the exposure-based item fairness. The latter argues that fairness in ranking corresponds to how ranking policies allocate exposure to individual items or groups of items based on their merit [23, 27–29, 33, 35, 36]. Within the exposure-based framework, Singh and Joachims [27] introduce a post-processing algorithm to allocate exposure for protected item groups proportional to their merit. Following Singh and Joachims [27], other fairness-of-exposure methods have been proposed such as end-to-end policy learning [28, 33] and dynamic ranking algorithms [23, 35]. Instead of guaranteeing *group* fairness, there is also a line of work aiming at achieving a fair allocation of exposure among *individual* items [3, 7]. Biega et al. [3] solve an integer program and optimize amortized share of exposure over time among each item. Bower et al. [7] propose an optimal-transport-based regularizer to enforce individual item fairness.

Although the exposure-based formulation is widely accepted for both group and individual item fairness, there is no obvious choice for how exposure should be linked to merit. Most studies require exposure to be proportional to merit — but why not proportional to merit squared or some other function? Indeed, no principled justification has been given for any choice of the link function. Our axiomatic approach overcomes this shortcoming, since it does not rely on an arbitrary link function between exposure and merit, but is solely based on the impact of the ranking on the items. Moreover, we show that the existing *exposure-based* framework can create a substantially *disparate impact* among individual items. While a disparate impact constraint has been proposed to ensure *group* item fairness [23, 27], imposing such constraints per item for individual

fairness leads to the uniform random ranking policy [27]. We are the first to formulate and enforce *impact-based individual item fairness* in a meaningful way.

Fair Division. *Fair division* has been studied extensively in algorithmic game theory [8, 18, 31]. The goal of this line of work is to allocate a set of valuable – but limited – resources or goods to the agents in a fair manner. The classical fairness desiderata considered in this field are *envy-freeness* (EF) and *proportional fair share* (PFS) [8, 18, 25]. EF requires that no agents prefer another agent’s allocation of goods to their own. PFS means that every agent receives $1/n$ of their utility over the entire set of goods, where n is the number of agents. When the goods are *divisible*, maximizing the *Nash Social Welfare*, the product of agent utilities, has been considered compelling, because its solution ensures EF and PFS as well as being Pareto optimal. In contrast, recent studies focus on a more challenging problem of fairly allocating *indivisible* goods (e.g., course seats in universities [9] or computational resources in a cloud computing environment [6]). When the goods are indivisible, no feasible allocation may satisfy EF or PFS [10]. Therefore, relaxed versions of these axioms have been considered such as *envy-freeness up to one good* (EF1) and *maximin share guarantee* (MMS) [8, 25]. For example, EF1 requires that any pairwise envy can be eliminated by removing a single good from the envied agent’s allocation. An allocation satisfying EF1 always exists for a broad class of utility functions even in the case of indivisible goods [10].

The notion of EF has been adopted to fair machine learning in binary classification [2]. This is because EF is intuitive and requires no information beyond individuals’ utility functions, which contrasts with other notions of fairness such as metric-based individual fairness [16]. Beyond binary classification, a few works have applied the axioms in fair division to recommender system applications. Patro et al. [25] map the fair recommendation problem in two-sided markets to the problem of fairly allocating indivisible goods, where the goods are the set of products. The goal of Patro et al. [25] is then to ensure EF1 among users in terms of their utility and MMS (a relaxed version of PFS) among items in terms of their exposure allocation. Do et al. [15] consider the problem of auditing the unfair behavior of recommender systems with respect to EF among users. Although these studies formulating versions of EF in the context of recommender systems are related, our contributions are unique in several ways. First, we focus on achieving fairness of *impact among individual items*. In contrast, Patro et al. [25] consider the item fairness in terms of *exposure* and Do et al. [15] consider only the *user-side* fairness. In the following sections, we argue that the exposure-based fairness can produce a substantial amount of envy and a highly disparate impact. We are the first to formulate and achieve impact-based fairness rather than fairness-of-exposure in the context of individual fairness in ranking. Second, we build on the exact notions of EF and PFS, while Patro et al. [25] relax these fairness criteria and rely on EF1 and MMS. We can formulate fairness in rankings as the problem of fairly allocating *divisible* goods by considering a class of *stochastic* ranking policies. Stochastic ranking policies have been adapted in the fair ranking literature [14, 27, 28, 33], and we rely on this stochastic formulation aiming for more desirable fairness criteria than those of previous studies. Finally, our formulation is aware of

the position bias in a ranking. Related studies [15, 25] do not take into account position bias, which is an inevitable factor in dealing with real-world ranking interfaces [21, 30].

3 AN AXIOMATIC APPROACH TO INDIVIDUAL FAIRNESS IN RANKING

We consider the following class of ranking problems, where a two-sided market platform has responsibilities to both users and items.

- On a news platform, each reader (user) receives a ranking of news stories (items) as a digest of the previous day.
- On a hiring platform, each employer (user) receives a weekly ranking of the most relevant job candidates (items) that were added to the database.
- For a scientific conference, each reviewer (user) receives a ranking of the most relevant submissions (items) during the bidding process.

All problems have in common that the platform not only aims to maximize the utility of the rankings to the users, but that it also needs to be fair to the items. Specifically, the news providers would like a fair share of the traffic for their items, the job candidates deserve to get an adequate number of connections to relevant employers, and every paper should be given an appropriate chance of finding knowledgeable reviewers.

All these problems are instances of batch-ranking problems, where $\mathcal{U} = [m]$ is a set of users, and a ranking policy π aims to optimally order a set of items $\mathcal{I} = [n]$ for each user $u \in \mathcal{U}$. We assume the full-information setting, where we have access to the true relevance labels (or their predictions) $r(u, i) \in \mathbb{R}_+$ for all user-item pairs. A ranking σ is a permutation of the set of items sampled from a ranking policy $\pi(\cdot|u)$, which is a distribution over all possible rankings of the items. We consider the general case of stochastic ranking policies, which includes deterministic ranking policies as a special case. Stochastic ranking policies have been utilized in fair ranking research to have fine-grained control over the expected exposure allocation [14, 28, 33].

3.1 Utility to Users

Like in conventional ranking frameworks, we measure the utility that a policy π provides to the users through a utility function of the following form.

$$U(\pi) := \sum_{u \in \mathcal{U}} \mathbb{E}_{\sigma \sim \pi(\cdot|u)} \left[\sum_{i \in \mathcal{I}} e(\sigma(i)) \cdot r(u, i) \right] \quad (1)$$

$\sigma(i)$ is the rank of item i in ranking σ , and $e(\cdot)$ casts the rank to the exposure probability according to the position-based model (PBM) [13]. This definition of utility captures widely-used additive ranking metrics through the choice of $e(\cdot)$. For example, when we define $e(k) := \frac{\mathbb{I}\{k \leq K\}}{\log_2(k+1)}$, the utility becomes DCG@K [19]. It is also possible to learn an application-specific $e(\cdot)$ from the data [21]. We refer to a ranking policy that maximizes the utility to the users as a utility-maximizing policy, which is defined as $\pi_{max} := \arg \max_{\pi \in \Pi} U(\pi)$.

To avoid optimizing over the exponentially large space of rankings, we utilize the fact that only the exposure probability $e(\cdot)$ depends on the rank. This allows us to more concisely express the

user utility as follows.

$$U(\pi) := \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{I}} \sum_{k=1}^n e(k) r(u, i) X_{u,i,k}^\pi \quad (2)$$

where X^π is the tensor whose (u, i, k) element $X_{u,i,k}^\pi := \mathbb{P}(\sigma(i) = k | \pi, u)$ denotes the marginal probability of item i being ranked at the k -th position for user u under policy π ,¹ and $X_{u,*,*}^\pi$ should be the doubly stochastic matrix². The benefit of using this representation is to reduce the number of parameters needed to specify the effect of policy π . Specifically, for user u , we use only $|\mathcal{I}|^2$ parameters rather than the exponential number of possible rankings, as all stochastic ranking policies with the same matrix have the same user utility. This will allow us to optimize in the space of doubly stochastic matrices. Given a doubly stochastic matrix, the Birkhoff-von Neumann (BvN) decomposition can be used to efficiently find a stochastic ranking policy that corresponds to the matrix [4, 27].

3.2 Impact on Items

While sorting the items by their probability of relevance maximizes the utility to the users [26], many existing works have noted that this naive treatment can lead to rankings that are unfair to the items [3, 28, 33]. In particular, Biega et al. [3] measure similarity between individual items by their *amortized merit*, and propose a notion of fairness requiring that *amortized exposure* should be distributed proportional to their amortized merit. This definition of fairness aims at allocating exposure similarly between items with similar merit. However, as we have already argued, this formulation lacks a clear justification for why exposure should be linked proportional to relevance – or linked via any other specific function. Furthermore, the items only indirectly care about exposure, and they more directly care about the *impact* the ranking has on them. For example, a ranking policy that predominantly shows men for high-paying jobs and shows women for low-paying jobs can perfectly obey fairness of exposure, but it clearly violates fairness of impact.

We therefore focus on fairness of impact in this work, where *impact* quantifies the effect that a ranking policy has on a specific item i . To define impact, we build on an *item-centric* version of the matrix $X_{*,i,*}^\pi$ whose (u, k) element is: $X_{u,i,k}^\pi = \mathbb{P}(\sigma(i) = k | \pi, u)$. This matrix characterizes the *allocation of positions*, i.e., how big a fraction of the k -th position in a ranking for user u goes to item i . With this notation, we define the *impact* on each item as

$$\text{Imp}_i(X_{*,i,*}^\pi) := \sum_{u \in \mathcal{U}} \sum_{k=1}^n v_i(u, k) X_{u,i,k}^\pi$$

$v_i(u, k)$ is an application-dependent impact function, which defines how much impact (e.g., expected clicks, bookings, revenue) item i receives when it is ranked at the k -th position for user u . For ease of exposition, we will use the number of relevant users who see each item under the PBM as the impact function, i.e., $v_i(u, k) := e(k) r(u, i)$. Note that even with this specialization, impact remains more meaningful than fairness of exposure, which does not differentiate between exposure to relevant and non-relevant users. Note that the

¹ $\mathbb{P}(\sigma(i) = k | \pi, u) = \mathbb{E}_{\sigma \sim \pi(\cdot|u)} [\mathbb{I}\{\sigma(i) = k\}]$ where $\mathbb{I}\{\cdot\}$ is the indicator function.

²The sum of each row and column is 1.

Table 1: A Illustrative Ranking Problem showing the Exposure Allocations and Impact Fairness of Different Ranking Policies

| (a) True Relevance Table | | | (b) Max: $U(\pi_{\max}) = 1.30$ | | | (c) Fair: $U(\pi_{\text{expo-fair}}) = 1.23$ | | | (d) Uniform: $U(\pi_{\text{unif}}) = 1.00$ | | |
|--------------------------|-------|-------|--|-------|-------|--|-------|-------|---|-------|-------|
| | i_1 | i_2 | | i_1 | i_2 | | i_1 | i_2 | | i_1 | i_2 |
| $r(u_1, i)$ | 0.8 | 0.3 | $\text{Exp}_i(X^{\pi_{\max}} u_1)$ | 1.0 | 0.0 | $\text{Exp}_i(X^{\pi_{\text{expo-fair}}} u_1)$ | 1.0 | 0.0 | $\text{Exp}_i(X^{\pi_{\text{unif}}} u_1)$ | 0.5 | 0.5 |
| $r(u_2, i)$ | 0.5 | 0.4 | $\text{Exp}_i(X^{\pi_{\max}} u_2)$ | 1.0 | 0.0 | $\text{Exp}_i(X^{\pi_{\text{expo-fair}}} u_2)$ | 0.3 | 0.7 | $\text{Exp}_i(X^{\pi_{\text{unif}}} u_2)$ | 0.5 | 0.5 |
| Merit_i | 1.3 | 0.7 | $\text{Exp}_i(X^{\pi_{\max}})$ | 2.0 | 0.0 | $\text{Exp}_i(X^{\pi_{\text{expo-fair}}})$ | 1.3 | 0.7 | $\text{Exp}_i(X^{\pi_{\text{unif}}})$ | 1.0 | 1.0 |
| | | | $\text{Imp}_i(X^{\pi_{\max}}_{*,i,*})$ | 1.3 | 0.0 | $\text{Imp}_i(X^{\pi_{\text{expo-fair}}}_{*,i,*})$ | 0.95 | 0.28 | $\text{Imp}_i(X^{\pi_{\text{unif}}}_{*,i,*})$ | 0.65 | 0.35 |

user utility is equal to the sum of impacts under the specialization

$$U(\pi) = \sum_{i \in I} \text{Imp}_i(X^{\pi}_{*,i,*}),$$

making both equally measurable and comparable.

3.3 Fairness Axioms

We are now in a position to state axioms that ensure individual fairness with respect to the impact on each item as quantified by $\text{Imp}_i(X^{\pi}_{*,i,*})$. Making connections to well-established principles of fair division [8, 18, 22, 31], we treat the rank positions as the limited resource to allocate, and impact $\text{Imp}_i(X^{\pi}_{*,i,*})$ as the natural valuation of the items for each resource allocation. We start with the axiom of *envy-freeness*.

Definition 3.1. (Envy-Freeness) A ranking policy π is said to be *envy-free* if no item prefers another item's position allocation of π , i.e., $\text{Imp}_i(X^{\pi}_{*,i,*}) \geq \text{Imp}_i(X^{\pi}_{*,j,*})$, $\forall i, j \in I$.

An envy-free ranking policy ensures that no item would prefer the positions allocated to another item over their own allocation. If an allocation is envy-free, no item will want to switch allocations with another item.

The second axiom is *dominance over uniform ranking*, which requires that a fair ranking policy does not provide worse impact $\text{Imp}_i(X^{\pi}_{*,i,*})$ to any item than π_{unif} , which samples every possible permutation uniformly at random (uniform ranking policy).

Definition 3.2. (Dominance over Uniform Ranking) A ranking policy π is said to *dominate the uniform random policy* π_{unif} if π provides better or equal impact on every item compared to π_{unif} , i.e., $\text{Imp}_i(X^{\pi}_{*,i,*}) \geq \text{Imp}_i(X^{\pi_{\text{unif}}}_{*,i,*})$, $\forall i \in I$. Moreover, at least one item should gain impact strictly better than that under π_{unif} .

We regard π_{unif} as a baseline policy, because it provides an impact distribution that one can achieve without implementing any optimization or policy learning procedure. If a policy makes some items worse off than under π_{unif} , it inflicts active harm on those items and they are likely to abandon the platform.

In addition to the aforementioned fairness desiderata, we also want the ranking policy to be *optimal* in the following sense.

Definition 3.3. (Pareto Optimality) A ranking policy π is said to be *Pareto optimal* if no alternative policy can make some items strictly better off without making any others strictly worse off.

Pareto optimality ensures that the user utility, and equivalently the sum of item impacts, cannot easily be improved. This codifies that the policy should not needlessly sacrifice utility to the users

or aggregate impact to the items, and that there is no obviously avoidable harm to any of the items.

3.4 Fairness of Exposure Violates Axioms

To motivate the need for new algorithms to achieve fairness of impact, we now show in detail that the conventional fairness-of-exposure framework can cause envy among the items and create an unfair impact distribution in light of our axioms. Exposure-based fairness imposes a constraint to allocate exposure proportional to the amortized merit [3]. Here, we describe a more general constraint [32], where merit is quantified through an application-dependent link function $f(\cdot) > 0$,

$$\frac{\text{Exp}_i(X^{\pi})}{f(\text{Merit}_i)} = \frac{\text{Exp}_j(X^{\pi})}{f(\text{Merit}_j)}, \quad \forall i, j \in I, \quad (3)$$

where $\text{Merit}_i := \sum_{u \in \mathcal{U}} r(u, i)$ is the *amortized merit*, i.e., the relevance of item i amortized over all users. In contrast, $\text{Exp}_i(X^{\pi}) := \sum_u \text{Exp}_i(X^{\pi}|u) = \sum_u \sum_k e(k) r(u, i) X^{\pi}_{u,i,k}$ is the total amount of exposure allocated to item i under policy π , i.e., the *amortized exposure*. The link function $f(\cdot)$ maps the amortized merit to a positive merit value. A typical choice is $f(x) = x$, which leads to a linear program to optimize the rankings [3, 27, 28].

Algorithms for computing exposure-based fair rankings find a doubly stochastic matrix that maximizes the utility to the users (or equivalently the sum of impacts), subject to Eq. (3) as follows.

$$\begin{aligned}
& \pi_{\text{expo-fair}} \\
& = \arg \max_{\{X^{\pi}_{*,i,*}\}_{i \in I}} \sum_{u \in \mathcal{U}} \sum_{i \in I} \sum_{k=1}^n e(k) r(u, i) X^{\pi}_{u,i,k} \quad \left(= \sum_{i \in I} \text{Imp}_i(X^{\pi}_{*,i,*}) \right), \\
& \text{s.t.} \quad \frac{\text{Exp}_i(X^{\pi})}{f(\text{Merit}_i)} = \frac{\text{Exp}_j(X^{\pi})}{f(\text{Merit}_j)}, \quad \forall i, j \in I, \\
& \quad \sum_{k=1}^n X^{\pi}_{u,i,k} = 1, \quad \forall (u, i) \\
& \quad \sum_{i \in I} X^{\pi}_{u,i,k} = 1, \quad \forall (u, k) \\
& \quad 0 \leq X^{\pi}_{u,i,k} \leq 1, \quad \forall (u, i, k)
\end{aligned}$$

We use $\pi_{\text{expo-fair}}$ to denote a policy that solves the above optimization and call it the exposure-based fair ranking policy.

While widely used as the desideratum for fair ranking [3, 14, 20, 23, 27, 28, 33, 35, 36], the following provides a counterexample which illustrates that a fair allocation of exposure cannot avoid a *disparate impact* among individual items. Our counterexample consists of two users ($\mathcal{U} = \{u_1, u_2\}$) and two items ($I = \{i_1, i_2\}$).

We know the true relevance labels of every user-item pair, which is given in Table 1a. For simplicity, we also assume that only the top-ranked item is exposed to the users ($e(1) = 1, e(2) = 0$). We consider three ranking policies, (i) the utility-maximizing policy π_{max} , (ii) the exposure-based fair ranking policy $\pi_{expo-fair}$, and, (iii) the uniform random policy π_{unif} .

Tables 1b to 1d provide the exposure allocation of the three policies and how much impact each item receives. π_{max} allocates all exposure on the item having highest relevance (i_1). As a result, it achieves the highest user utility ($U(\pi_{max}) = 1.3$) while leading to the unfair situation where i_2 obtains no exposure despite its substantial relevance.

$\pi_{expo-fair}$ deals with this unfair allocation by imposing the exposure constraint. Specifically, it allocates some amount of exposure to the less relevant item (i_2) and ensures that the exposure is allocated proportional to merit as shown in Table 1c. However, our example indicates that $\pi_{expo-fair}$ violates impact-based fairness. First, $\pi_{expo-fair}$ is *not* envy-free. Indeed, i_2 envies the allocation of i_1 , because impact on i_2 can be improved by swapping its allocation with that of i_1 .³ Second, $\pi_{expo-fair}$ does *not* dominate π_{unif} for all items. It is easy to see that impact on i_2 under $\pi_{expo-fair}$ (0.28) is smaller than under π_{unif} (0.35). This implies that $\pi_{expo-fair}$ improves impact on i_1 at the cost of impact on i_2 , creating a substantial disparity.

Our counterexample indicates that exposure-based fairness allows a policy to produce envy and an unfair distribution of impact. The following theorem more formally illustrates the possible disparate impact of the existing exposure-based framework.

THEOREM 3.4. *There exist two-sided markets with $|\mathcal{U}| = |\mathcal{I}| = n$ such that, even under the exposure-based fair ranking policy $\pi_{expo-fair}$ with $f(x) = x$, there exists item i whose maximum envy grows while impact compared to π_{unif} diminishes with the market size n , i.e.,*

$$\frac{\max_{j \in \mathcal{I}} \text{Imp}_i(X_{*,j,*}^{\pi_{expo-fair}})}{\text{Imp}_i(X_{*,i,*}^{\pi_{expo-fair}})} = \Omega(n), \quad \frac{\text{Imp}_i(X_{*,i,*}^{\pi_{expo-fair}})}{\text{Imp}_i(X_{*,i,*}^{\pi_{unif}})} = O(n^{-1}).$$

for any $e(\cdot)$ such that the exposure-fair constraint in Eq. (3) is feasible.

Theorem 3.4 indicates that exposure-based fair rankings can produce large amounts of envy for some items. It can also actively harm some items beyond a reasonable baseline. Singh and Joachims [27] describe the *disparate impact* constraint aiming at allocating impact proportional to merit in the context of *group* fairness. However, this constraint merely leads to the uniform ranking in the case of *individual* fairness [27]. The disparate impact caused by the exposure-based framework and the inapplicability of the disparate impact constraint validate that our framework of *impact-based* individual fairness is fundamentally new and different from conventional exposure-based fairness.

4 COMPUTING RANKINGS WITH FAIRNESS OF IMPACT GUARANTEES

While the previous section introduced the axiomatization of fairness-of-impact for ranking and showed that the conventional fairness-of-exposure approach does not satisfy these axioms, it is not yet clear

³We know that $\text{Imp}_{i_2}(X_{*,i_2,*}^{\pi_{expo-fair}}) = 0.28$ from Table 1c. If we replaced $X_{*,i_2,*}^{\pi_{expo-fair}}$ with $X_{*,i_1,*}^{\pi_{expo-fair}}$, impact on i_2 would increase, i.e., $\text{Imp}_{i_2}(X_{*,i_1,*}^{\pi_{expo-fair}}) = 0.42$.

whether an efficient algorithm exists to compute fair rankings – and whether fair rankings even exist. To approach this algorithmic question, we formulate the following optimization problem which extends the concept of *Nash Social Welfare* (NSW) to fair rankings.

$$\begin{aligned} \pi_{NSW} = \arg \max_{\{X_{*,i,*}^{\pi}\}_{i \in \mathcal{I}}} & \prod_{i \in \mathcal{I}} \text{Imp}_i(X_{*,i,*}^{\pi}), \\ \text{s.t.} & \sum_{k=1}^n X_{u,i,k}^{\pi} = 1, \quad \forall (u, i) \\ & \sum_{i \in \mathcal{I}} X_{u,i,k}^{\pi} = 1, \quad \forall (u, k) \\ & 0 \leq X_{u,i,k}^{\pi} \leq 1, \quad \forall (u, i, k) \end{aligned}$$

The objective maximizes the *product* of impacts on the items, and the constraints express that each item needs to have probability 1 to be placed in some position, and each position needs to have probability 1 of receiving an item. This constraint structure makes this optimization problem different from the standard NSW in fair division [5, 17, 22]. However, we retain that the optimization problem is convex and thus efficiently solvable, which is easy to see by equivalently replacing the product of impacts in objective with the sum of their logarithms, $\sum_{i \in \mathcal{I}} \log \text{Imp}_i(X_{*,i,*}^{\pi})$.⁴ We call the policy that solves the optimization problem the NSW policy, which we denote by π_{NSW} .

A useful intuition is that the NSW does not allow any item to have zero impact, thus achieving a more equitable impact distribution. In particular, the NSW becomes zero if there is even a single item with zero impact. This contrasts with the conventional objective of maximizing user utility (which equals the sum of impacts), which may be maximized even if there are some items receiving no impact. The following shows more formally that this intuition is correct, and that π_{NSW} guarantees (approximate) dominance over π_{unif} . Furthermore, the following shows that π_{NSW} is (approximately) envy-free and Pareto optimal.

THEOREM 4.1. *If $K = 1$ and only the top-ranked item contributes to the utility and impact (e.g., DCG@1), π_{NSW} is exactly Pareto optimal, is envy-free, and dominates π_{unif} . For the more general case with $K > 1$, π_{NSW} is still Pareto optimal. Moreover, let us call a pair of items i and j “ ϵ -twins” if $\max_u |r(u, i) - r(u, j)| \leq \epsilon$. Then, if every item has at least $K + 1$ ϵ -twins and $n/m = O(1)$, we have*

$$\begin{aligned} \max_{j \in \mathcal{I}} \text{Imp}_i(X_{*,j,*}^{\pi_{NSW}}) - \text{Imp}_i(X_{*,i,*}^{\pi_{NSW}}) &= O(\epsilon), \\ \text{Imp}_i(X_{*,i,*}^{\pi_{unif}}) - \text{Imp}_i(X_{*,i,*}^{\pi_{NSW}}) &= O(\epsilon), \end{aligned}$$

for all item $i \in \mathcal{I}$.

In the case of $K = 1$, the solution of our problem coincides with the allocation computed via the *competitive equilibrium from equal incomes* (CEEI), a classic model of market equilibrium [9, 31], which is known to be Pareto optimal, to be envy-free, and to dominate a uniform random allocation. When $K > 1$, the classical result from CEEI does not necessarily hold due to the additional constraints in our problem [22]. However, Theorem 4.1 characterizes that, even for the general case, π_{NSW} can approximately satisfy our fairness axioms if there are items that are similar to each other, which

⁴This type of program is known as the Eisenberg-Gale (EG) program [17].

Table 2: Allocation of the NSW Policy: $U(\pi_{NSW}) = 1.20$

| | i_1 | i_2 |
|--------------------------------|-------|-------|
| $Exp_i(X^{\pi_{NSW}} u_1)$ | 1.0 | 0.0 |
| $Exp_i(X^{\pi_{NSW}} u_2)$ | 0.0 | 1.0 |
| $Exp_i(X^{\pi_{NSW}})$ | 1.0 | 1.0 |
| $Imp_i(X^{\pi_{NSW}}_{*,i,*})$ | 0.8 | 0.4 |

Table 3: Exposure Allocations and Impact Fairness of the α -NSW Policies with Different Hyper-parameter Values

| (a) 1-NSW: $U(\pi_{1-NSW}) = 1.21$ | | | (b) 2-NSW: $U(\pi_{2-NSW}) = 1.24$ | | |
|------------------------------------|-------|-------|------------------------------------|-------|-------|
| | i_1 | i_2 | | i_1 | i_2 |
| $Exp_i(X^{\pi_{1-NSW}} u_1)$ | 1.0 | 0.0 | $Exp_i(X^{\pi_{2-NSW}} u_1)$ | 1.0 | 0.0 |
| $Exp_i(X^{\pi_{1-NSW}} u_2)$ | 0.1 | 0.9 | $Exp_i(X^{\pi_{2-NSW}} u_2)$ | 0.4 | 0.6 |
| $Exp_i(X^{\pi_{1-NSW}})$ | 1.1 | 0.9 | $Exp_i(X^{\pi_{2-NSW}})$ | 1.4 | 0.6 |
| $Imp_i(X^{\pi_{1-NSW}}_{*,i,*})$ | 0.85 | 0.36 | $Imp_i(X^{\pi_{2-NSW}}_{*,i,*})$ | 1.00 | 0.24 |

is reasonable for large markets.⁵ Note that Theorem 4.1 assumes the availability of the true relevance labels, but in Section 5, we empirically show that π_{NSW} robustly produces fair distributions of impact even with predicted relevance values.

Table 2 illustrates these guarantees of π_{NSW} on the same example used earlier. First, we can see that π_{NSW} produces an *envy-free* allocation of positions, i.e., neither item prefers the other item's allocation. In addition, π_{NSW} dominates π_{unif} , since every item gains higher impact than under π_{unif} .

4.1 Controlling the Fairness/Utility Trade-Off

In some situations, we may want to exert explicit control over how much user utility to sacrifice for stronger fairness guarantees to the items. For example, in Tables 2 and 1, the user utility of π_{NSW} (1.20) is lower than that of π_{max} (1.30) and $\pi_{expo-fair}$ (1.23), and we may want to put more emphasis on user utility. To enable explicit control, we extend the NSW to what we call the α -NSW as follows.

$$\begin{aligned}
 \pi_{\alpha-NSW} = \arg \max_{\{X^{\pi}_{*,i,*}\}_{i \in I}} & \prod_{i \in I} Imp_i(X^{\pi}_{*,i,*})^{Merit_i^{\alpha}}, \\
 \text{s.t.} & \sum_{k=1}^n X^{\pi}_{u,i,k} = 1, \forall (u, i) \\
 & \sum_{i \in I} X^{\pi}_{u,i,k} = 1, \forall (u, k) \\
 & 0 \leq X^{\pi}_{u,i,k} \leq 1, \forall (u, i, k)
 \end{aligned}$$

$\alpha \geq 0$ is a hyper-parameter that controls the balance between maximizing user utility (large α) and guaranteeing impact-based fairness (small α). When $\alpha = 0$, the α -NSW objective is reduced to the standard NSW. Note that we can again take the logarithm to rewrite the objective to $\sum_{i \in I} Merit_i^{\alpha} \log Imp_i(X^{\pi}_{*,i,*})$ and solve the program efficiently. The α -NSW policy ensures only a weaker version of envy-freeness and dominance over π_{unif} as detailed in the

⁵Note that it may be reasonable to assume that ϵ decreases with the size of the market, since it should be easier to find a sufficient number of ϵ -twins in a larger market. This means we could expect π_{NSW} to more accurately satisfy our axioms with an increasing number of items.

appendix, which means it can achieve higher user utility compared to the standard NSW.

Table 3 illustrates the allocations of the α -NSW policy $\pi_{\alpha-NSW}$ for $\alpha = 1$ and $\alpha = 2$ on the example problem. We see that a larger value of α leads to a better user utility, i.e., $U(\pi_{NSW}) < U(\pi_{1-NSW}) < U(\pi_{expo-fair}) < U(\pi_{2-NSW})$, while a smaller value leads to a more balanced impact distribution. What is particularly notable here is that π_{1-NSW} still dominates π_{unif} and remains envy-free, while achieving a better utility compared to π_{NSW} . This suggests that an appropriate choice of α may allow us to improve user utility over π_{NSW} without sacrificing envy-freeness and dominance over π_{unif} on particular problem instances, even if we no longer have a-priori guarantees that hold over all problem instances. In the next section, we empirically explore how α controls this trade-off.

5 EXPERIMENTS

We first present experiments on synthetic data where we can control the popularity pattern in the market and the accuracy of the relevance prediction. In addition, we use real-world extreme classification datasets to evaluate how our method works with realistic relevance predictions. Our experiment implementation is available at <https://github.com/usaito/kdd2022-fair-ranking-nsw>.

5.1 Synthetic Data

To generate synthetic data, we first define the ground-truth relevance between user u and item i as

$$r_{true}(u, i) = (1 - \lambda) \cdot r_{unif}(u, i) + \lambda \cdot r_{pop}(u, i), \quad (4)$$

where $r_{unif}(u, i)$ is an independent and uniform draw within range $[0, 1]$. The second term infuses *popularity bias*, which we define as

$$r_{pop}(u, i) := \begin{cases} \frac{n-i+1}{n} \cdot \frac{m-u+1}{m} & (\text{randomly sampled } 70\% \text{ of the items}) \\ \frac{n-i+1}{n} \cdot \frac{u}{m} & (\text{rest of the items}) \end{cases}$$

$\lambda \in [0, 1]$ controls the popularity pattern, where a larger value of λ leads to more severe popularity bias. We also simulate that we typically only have access to (inaccurately) predicted relevances, which we model as

$$r_{pred}(u, i) := \text{clip}(r_{true}(u, i) + \eta, 0, 1), \quad \eta \sim \text{Unif}(-c, c),$$

where η is independent and uniform noise, and $c \geq 0$ controls the accuracy of the predicted relevance values. We use the predicted relevance $r_{pred}(u, i)$ to optimize the rankings, and evaluate the item fairness and user utility of the policies using the ground-truth relevance $r_{true}(u, i)$. We use the inverse examination function $e(k) := \mathbb{I}\{k \leq K\}/k$ in the definition of utility and impact, but we also present experiments with a different function in the appendix.

We evaluate the degree of individual item fairness of a ranking policy π using the following measures:

- (1) Mean Max Envy (smaller is better):

$$\frac{1}{|I|} \sum_{i \in I} \left\{ \max_{j \in I} Imp_i(X^{\pi}_{*,j,*}) - Imp_i(X^{\pi}_{*,i,*}) \right\}$$

- (2) Proportion (%) of items for which π improves impact by over 10% compared to π_{unif} (larger is better):

$$\frac{100}{|I|} \sum_{i \in I} \mathbb{I} \left\{ Imp_i(X^{\pi}_{*,i,*}) / Imp_i(X^{\pi_{unif}}_{*,i,*}) \geq 1.1 \right\}$$




Figure 1: Fairness and user utility of the NSW policy compared to baseline policies on synthetic data for varying levels of popularity bias.




Figure 2: Fairness and user utility relative to the exposure-based fair policy $\pi_{\text{expo-fair}}$ for various α -NSW policies. Note that we omit $\lambda = 0.0$ from (a) and $\lambda = 0.0, 0.2, 0.4$ from (c), because the measures of $\pi_{\text{expo-fair}}$ are (almost) zero for these cases.

- (3) Proportion (%) of items for which π decreases impact by over 10% compared to π_{unif} (smaller is better):

$$\frac{100}{|I|} \sum_{i \in I} \mathbb{I} \left\{ \text{Imp}_i(X_{*,i,*}^{\pi}) / \text{Imp}_i(X_{*,i,*}^{\pi_{\text{unif}}}) \leq 0.9 \right\}$$

The following evaluates π_{NSW} in comparison to π_{max} , $\pi_{\text{expo-fair}}$, and π_{unif} . We also evaluate how the hyper-parameter α of $\pi_{\alpha\text{-NSW}}$ controls the trade-off of item fairness and user utility. Note that we use $|\mathcal{U}| = 100$, $|\mathcal{I}| = 50$, $c = 0.05$, $\lambda = 0.5$, $K = 5$ as defaults, and vary each experiment configuration to see how each affects the behavior of the ranking policies.⁶ Results are averaged over 10 simulation runs performed with different random seeds.

5.1.1 How do the ranking policies perform with different popularity patterns? Here, we test how the fairness measures and user utility of the ranking policies change with different popularity patterns. For this purpose, we vary the experiment parameter λ in Eq. (4) in the range of $\{0.0, 0.2, \dots, 1.0\}$. First, Figure 1 (a) shows the amount of envy produced under different levels of popularity bias. The figure suggests that π_{NSW} is always almost envy-free, while π_{max} and $\pi_{\text{expo-fair}}$ produce larger amounts of envy for markets with greater popularity bias. Next, Figure 1 (b) shows how many items gain substantially better impact compared to π_{unif} under different levels of popularity bias. We observe that π_{NSW} substantially

improves impact on almost all items in all cases, while π_{max} and $\pi_{\text{expo-fair}}$ fail to achieve this desideratum. In addition, Figure 1 (c) indicates that π_{max} and $\pi_{\text{expo-fair}}$ decrease impacts on some items over 10% from those under π_{unif} . In particular, when $\lambda = 1.0$, 80% and 40% of the items experience substantial loss of impact under π_{max} and $\pi_{\text{expo-fair}}$, respectively. Instead, π_{NSW} does not substantially decrease any item's impact below π_{unif} , even if there exists severe popularity bias. These results demonstrate that π_{NSW} avoids producing envy and provides fair improvements in impact for almost all items. On the other hand, both π_{max} and $\pi_{\text{expo-fair}}$ can lead to an unfair impact distribution and produce substantial amounts of envy, especially in the presence of severe popularity bias.

Although π_{NSW} is the most desirable in terms of item fairness, Figure 1 (d) suggests the expected trade-off between maximizing the user utility and satisfying our impact-based axioms. Figure 1 (d) shows how much user utility each policy gains under different popularity patterns. We observe that π_{max} achieves the largest user utility followed by $\pi_{\text{expo-fair}}$ and π_{NSW} . This trend does not change across different levels of popularity bias, however, the gap in user utility among those policies varies. When $\lambda = 0$ and there is no popularity bias in the market, all policies except for π_{unif} achieve almost the same user utility. However, for larger values of λ , the gap in user utility between π_{max} , $\pi_{\text{expo-fair}}$, and π_{NSW} becomes greater. In particular, π_{max} and $\pi_{\text{expo-fair}}$ improve user utility, while that of π_{NSW} decreases. An explanation is that an increase in popularity bias makes it easier to maximize the user utility by allocating large

⁶The results with varying number of items (n) and lengths of ranking (K) are reported in Appendix C

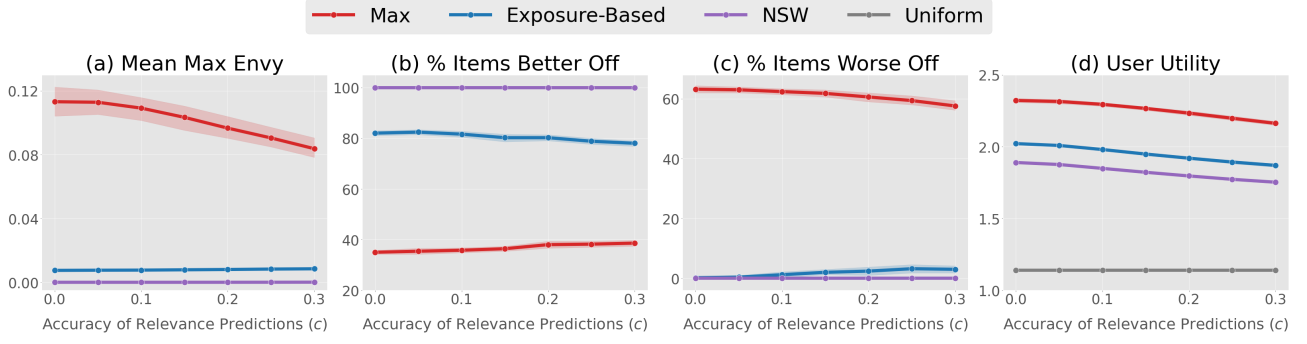


Figure 3: Fairness and user utility of the NSW policy compared to baseline policies on synthetic data for varying levels of relevance-prediction accuracy.




Figure 4: Fairness and user utility on the Delicious dataset for varying lengths of the ranking.

amounts of exposure to popular items. Conversely, π_{NSW} shows the opposite trend since it needs to ever more strongly counteract the popularity bias to maintain impact fairness.

5.1.2 How does the hyper-parameter of $\pi_{\alpha-NSW}$ balance item fairness and user utility? This experiment investigates how effectively one can control the balance between item fairness and user utility through the hyper-parameter α of $\pi_{\alpha-NSW}$. To this end, we evaluate $\pi_{\alpha-NSW}$ with $\alpha \in \{0.0, 0.5, 1.0, 2.0\}$ (when $\alpha = 0$, $\pi_{\alpha-NSW}$ is identical to π_{NSW}) in comparison to $\pi_{expo-fair}$. Figure 2 reports the fairness measures and user utility of $\pi_{\alpha-NSW}$ relative to those of $\pi_{expo-fair}$. Overall, the result demonstrates that $\pi_{\alpha-NSW}$ is able to choose a range of trade-offs through the hyper-parameter. The most interesting trade-off is achieved with $\alpha = 1.0$. Figure 2 (a) indicates that π_{1-NSW} reduces the amount of envy over 40% compared to $\pi_{expo-fair}$ in all cases. Moreover, Figure 2 (b) and (c) demonstrate that π_{1-NSW} provides a far more equitable distribution of impact compared to $\pi_{expo-fair}$. Finally, Figure 2 (d) shows that π_{1-NSW} achieves almost the same user utility as $\pi_{expo-fair}$ (about 4% drop even when $\lambda = 1.0$). This result suggests that $\pi_{\alpha-NSW}$ has the potential to achieve a substantially fairer impact distribution than $\pi_{expo-fair}$ while achieving a comparable level of user utility.

5.1.3 How do the ranking policies perform under different relevance-prediction accuracies? We also empirically investigate how robust the ranking policies are to inaccurate relevance labels, given that one would typically use predicted relevances in real-world applications. Figure 3 illustrates fairness and user utility for varying values

of the accuracy parameter c . The plots show that user utility slightly decreases as the relevance labels become more inaccurate, while fairness remain almost constant for all policies. Notably, π_{NSW} is almost envy-free and achieves a fairer impact distribution compared to the other policies in all situations, indicating that the NSW approach is robust to inaccurate relevance predictions.

5.2 Real-World Data

To further evaluate how our method performs with predicted relevances on real-world data, we adapt two multilabel extreme classification datasets, namely Delicious and Wiki10-31K from the Extreme Classification Repository.⁷ We regard each data as a user and each label as an item. If a data belongs to a label, then they are considered relevant. As preprocessing, we randomly sample 100 labels,⁸ and then split the data into 90% training and 10% test sets. We train a Logistic Regression model on the training set, and predict the probabilities of each data belonging to the labels, which correspond to the relevance predictions. Based on the predictions, the ranking policies optimize the ordering of the labels for each data. We finally evaluate item fairness and user utility of the ranking policies on the test set using the true class labels as the ground-truth relevance. The fairness measures are the same as in the previous section. We use the inverse examination function and vary $K \in \{1, 3, 5, 10, 20\}$.

⁷<http://manikvarma.org/downloads/XC/XMLRepository.html>

⁸We sample label i with probability $\mu \cdot \mathbb{I}\{i \text{ is one of top 100 frequent labels}\} + (1 - \mu) \cdot (100/|I|)$ (where we set $\mu = 0.5$) to introduce some popularity bias in the data.

We iterate the simulation 10 times with different seeds and report the aggregated results.

Figure 4 compares fairness and user utility of π_{max} , π_{NSW} , π_{1-NSW} , $\pi_{expo-fair}$ and π_{unif} on the Delicious dataset. First, π_{NSW} achieves almost envy-free and fair impact distribution with increasing K , suggesting that our desiderata can be achieved even with realistically inaccurate relevance predictions. In addition, π_{1-NSW} succeeds in finding a utility-fairness trade-off that is more desirable than that of $\pi_{expo-fair}$. Specifically, π_{1-NSW} obtains user utility slightly better than $\pi_{expo-fair}$ while leading to a fairer impact distribution. Figure 4 (a) suggests that π_{1-NSW} produces a larger amount of envy compared to $\pi_{expo-fair}$, but the amount is controllable by tuning α . Appendix C reports qualitatively similar results on Wiki10-31K.

6 CONCLUSION

We provide a new conceptualization of fairness in ranking as a resource allocation problem and propose impact-based axioms for individual item fairness. This allows us to build upon well-established principles of fairness from economics, removing the need to choose a difficult-to-justify link function as required for exposure-based fairness. We also contribute an efficient algorithm for computing impact-fair rankings, adapting the Nash Social Welfare to ranking problems. Furthermore, we develop a practical extension of the NSW, which enables us to control the trade-off between user utility and item fairness. This work opens a wide range of new research directions, ranging from the development of end-to-end policy learning methods for impact-fair ranking, to the use of partial-information data and extensions to sequential ranking problems.

ACKNOWLEDGMENTS

This research was supported in part by NSF Awards IIS-1901168 and IIS-2008139. Yuta Saito was supported by the Funai Overseas Scholarship. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

REFERENCES

- [1] Abolfazl Asudeh, HV Jagadeish, Julia Stoyanovich, and Gautam Das. 2019. Designing Fair Ranking Schemes. In *Proceedings of the 2019 International Conference on Management of Data*. 1259–1276.
- [2] Maria-Florina F Balcan, Travis Dick, Ritesh Noothigattu, and Ariel D Procaccia. 2019. Envy-Free Classification. In *Advances in Neural Information Processing Systems*, Vol. 32. 1240–1250.
- [3] Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 405–414.
- [4] Garrett Birkhoff. 1940. *Lattice Theory*. Vol. 25. American Mathematical Soc.
- [5] Benjamin Birnbaum, Nikhil R Devanur, and Lin Xiao. 2011. Distributed Algorithms via Gradient Descent for Fisher Markets. In *Proceedings of the 12th ACM conference on Electronic commerce*. 127–136.
- [6] Arpita Biswas and Siddharth Barman. 2018. Fair Division under Cardinality Constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. 91–97.
- [7] Amanda Bower, Hamid Eftekhari, Mikhail Yurochkin, and Yuekai Sun. 2020. Individually Fair Rankings. In *International Conference on Learning Representations*.
- [8] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- [9] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- [10] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation (TEAC)* 7, 3 (2019), 1–32.
- [11] L Elisa Celis, Anay Mehrotra, and Nisheeth K Vishnoi. 2020. Interventions for Ranking in The Presence of Implicit Bias. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 369–380.
- [12] L Elisa Celis, Damian Straszak, and Nisheeth K Vishnoi. 2017. Ranking with Fairness Constraints. *arXiv preprint arXiv:1704.06840* (2017).
- [13] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An Experimental Comparison of Click Position-bias Models. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*. 87–94.
- [14] Fernando Diaz, Bhaskar Mitra, Michael D Ekstrand, Asia J Biega, and Ben Carterette. 2020. Evaluating Stochastic Rankings with Expected Exposure. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 275–284.
- [15] Virginie Do, Sam Corbett-Davies, Jamal Atif, and Nicolas Usunier. 2021. Online Certification of Preference-based Fairness for Personalized Recommender Systems. *arXiv preprint arXiv:2104.14527* (2021).
- [16] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. 214–226.
- [17] Edmund Eisenberg and David Gale. 1959. Consensus of Subjective Probabilities: The Aari-Mutuel Method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168.
- [18] Duncan Karl Foley. 1966. *Resource Allocation and the Public Sector*. Yale University.
- [19] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated Gain-based Evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)* 20, 4 (2002), 422–446.
- [20] Olivier Jeunen and Bart Goethals. 2021. Top-k Contextual Bandits with Equity of Exposure. In *Fifteenth ACM Conference on Recommender Systems*. 310–320.
- [21] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. 2017. Unbiased Learning-to-Rank with Biased Feedback. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. 781–789.
- [22] Christian Kroer and Alexander Peysakhovich. 2019. Scalable Fair Division for ‘At Most One’ Preferences. *arXiv preprint arXiv:1909.10925* (2019).
- [23] Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. 2020. Controlling Fairness and Bias in Dynamic Learning-to-Rank. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 429–438.
- [24] Harrie Oosterhuis. 2021. Computationally Efficient Optimization of Plackett-Luce Ranking Models for Relevance and Fairness. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [25] Gourab K Patro, Arpita Biswas, Niloy Ganguly, Krishna P Gummadi, and Abhinjan Chakraborty. 2020. Fairrec: Two-sided Fairness for Personalized Recommendations in Two-Sided Platforms. In *Proceedings of The Web Conference 2020*. 1194–1204.
- [26] Stephen E Robertson. 1977. The Probability Ranking Principle in IR. *Journal of documentation* (1977).
- [27] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2219–2228.
- [28] Ashudeep Singh and Thorsten Joachims. 2019. Policy Learning for Fairness in Ranking. In *Advances in Neural Information Processing Systems*, Vol. 32.
- [29] Ashudeep Singh, David Kempe, and Thorsten Joachims. 2021. Fairness in Ranking Under Uncertainty. In *Advances in Neural Information Processing Systems*, Vol. 34.
- [30] Yi Su, Magd Bayoumi, and Thorsten Joachims. 2022. Optimizing Rankings for Recommendation in Matching Markets. In *Proceedings of the ACM Web Conference 2022*. 328–338.
- [31] Hal R Varian et al. 1974. Equity, Envy, and Efficiency. *Journal of Economic Theory* 9, 1 (1974), 63–91.
- [32] Lequn Wang, Yiwei Bai, Wen Sun, and Thorsten Joachims. 2021. Fairness of Exposure in Stochastic Bandits. In *Proceedings of the 38th International Conference on Machine Learning*, Vol. 139. PMLR, 10686–10696.
- [33] Himank Yadav, Zhengxiao Du, and Thorsten Joachims. 2021. Policy-Gradient Training of Fair and Unbiased Ranking Functions. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1044–1053.
- [34] Ke Yang and Julia Stoyanovich. 2017. Measuring Fairness in Ranked Outputs. In *Proceedings of the 29th International Conference on Scientific and Statistical Database Management*. 1–6.
- [35] Tao Yang and Qingyao Ai. 2021. Maximizing Marginal Fairness for Dynamic Learning to Rank. In *Proceedings of the ACM Web Conference 2021*. 137–145.
- [36] Meike Zehlke and Carlos Castillo. 2020. Reducing Disparate Exposure in Ranking: A Learning to Rank Approach. In *Proceedings of The Web Conference 2020*. 2849–2855.

A CONNECTING EXPOSURE-BASED AND IMPACT-BASED NOTIONS OF FAIRNESS

The existing exposure-based fairness and our impact-based axioms have an interesting connection. Here, we show that the equality of attention of Biega et al. [3] is a special case of our notion of envy-freeness.

THEOREM A.1. *If the impact function is defined as $v_i(u, k) := e(k)$, then envy-freeness coincides with the equality of attention.*

PROOF. Consider a pair of items $i, j \in \mathcal{I} (i \neq j)$. Given that $v_i(u, k) := e(k)$, envy-freeness mandates that

$$\text{Imp}_i(X_{*,i,*}^\pi) \geq \text{Imp}_j(X_{*,j,*}^\pi) \implies \text{Exp}_i(X^\pi) \geq \text{Exp}_j(X^\pi). \quad (5)$$

Analogously, for the opposite direction, we should have that

$$\text{Imp}_j(X_{*,j,*}^\pi) \geq \text{Imp}_i(X_{*,i,*}^\pi) \implies \text{Exp}_j(X^\pi) \geq \text{Exp}_i(X^\pi), \quad (6)$$

Thus, from Eqs (5) and (6), envy-freeness coincides with the equality of attention, which requires that the amortized exposure should be equal for all items, i.e., $\text{Exp}_i(X^\pi) = \text{Exp}_j(X^\pi), \forall i, j \in \mathcal{I}$. \square

B PROOF SKETCHES

We do not have space enough to describe the complete proofs, so here we provide only their sketches. The complete proofs are included in the appendix of the arXiv version of the paper.

Proof Sketch of Theorem 3.4. Consider the true relevance function $r(u, i) = (n - u + 1)(n - i + 1)/n^2$, which creates severe popularity bias. Then, we have $\text{Imp}_n(X_{*,n,*}^{\pi_{\text{expo-fair}}}) \leq 2TE/n(n+1)$, $\text{Imp}_n(X_{*,n,*}^{\pi_{\text{unif}}}) = (n+1)TE/2n^2$ where $TE := \sum_k e(k)$. We can also show that $\text{Imp}_n(X_{*,1,*}^{\pi_{\text{expo-fair}}})/\text{Imp}_n(X_{*,n,*}^{\pi_{\text{expo-fair}}}) \geq (n+1)/2$.

Proof Sketch of Theorem 4.1. If $K = 1$, our problem is reduced to the classical fair division of divisible goods with additive utility. Thus, the NSW-maximizing solution is known to satisfy Pareto optimality, envy-freeness, and dominance over π_{unif} (which corresponds to PFS) [31]. For a general case with $K > 1$, our problem is reduced to fair division with *constrained groups* described in Appendix A of Kroer and Peysakhovich [22]. Thus, our result follows from a version of their Theorem 1.

Note that $\pi_{\alpha\text{-NSW}}$ satisfies the following relaxed version of envy-freeness and dominance over π_{unif} in the case of $K = 1$.

$$\text{Weighted Envy-Freeness: } \frac{\text{Imp}_i(X_{*,i,*}^\pi)}{\text{Merit}_i^\alpha} \geq \frac{\text{Imp}_j(X_{*,j,*}^\pi)}{\text{Merit}_j^\alpha}, \forall j \in \mathcal{I},$$

$$\text{Weighted Dominance: } \text{Imp}_i(X_{*,i,*}^\pi) \geq \frac{n \text{Merit}_i^\alpha}{\sum_j \text{Merit}_j^\alpha} \text{Imp}_i(X_{*,i,*}^{\pi_{\text{unif}}}),$$

for all $i \in \mathcal{I}$.

C EXPERIMENT DETAILS AND RESULTS

C.1 Synthetic Data

In this section, we employ the same default setting as in Section 5.1 and vary some additional configurations.

Table 4: Statistics of the real-world datasets after preprocessing and the relevance prediction accuracy on the test set

| | Delicious | Wiki10-31k |
|--------------------------------------|-----------|------------|
| Training Data Size | 9,123 | 7,639 |
| Test Data Size ($ \mathcal{U} $) | 1,014 | 849 |
| Number of Labels ($ \mathcal{I} $) | 100 | 100 |
| AUC (on test set) | 0.81 | 0.87 |
| LogLoss (on test set) | 0.31 | 0.18 |

C.1.1 How do the ranking policies perform with different lengths of ranking? Here, we validate how the policies perform as the length of ranking (K) varies. Figure 5 reports the results with $K \in \{1, 3, 5, 10, 20\}$. First, we can see that all policies improve user utility with increasing K in Figure 5 (d), as more items contribute to the user utility. As for the fairness measures, π_{NSW} remains envy-free and improves impact on all items over 10% from π_{unif} , while $\pi_{\text{expo-fair}}$ produces a larger amount of envy and a less equitable impact distribution with increasing K .

C.1.2 How do the ranking policies perform with different market sizes? Next, we vary the size of the market. Figure 6 shows how the fairness measures and user utility of the ranking policies change with varying numbers of items and a fixed number of users. We observe that π_{max} produces a larger amount of envy and a less equitable impact distribution with the growing market size. In contrast, $\pi_{\text{expo-fair}}$ slightly improves both the fairness measures and user utility, but π_{NSW} is much fairer for all cases, always being almost envy-free and dominating π_{unif} in terms of impact distribution. $\pi_{\text{expo-fair}}$ does not achieve these desiderata for all market sizes.

C.1.3 How does the examination function influence the relative performance? Finally, we evaluate whether a different examination function results in a different conclusion. We use the inverse function $e(k) := \mathbb{I}\{k \leq K\}/k$ in the main text, but there are some other types of examination functions we can consider [30]. Here we use the exponential examination function defined as $e(k) := \mathbb{I}\{k \leq K\}/\exp(k-1)$, which has a steeper drop off in probability than the inverse function, and as a result, assumes that users are likely to only see the top few items. Figure 7 shows the item fairness measures and user utility with varying popularity patterns ($\lambda \in \{0.0, 0.2, \dots, 1.0\}$) and with the exponential examination function. In general, we observe trends similar to Figure 1 (which is obtained with the inverse examination function) for all metrics. Because we use a steeper examination function, the user utility decreases for all policies compared to Figure 1, but the relative fairness and utility remain similar.

C.2 Real-World Data

Table 4 summarizes some statistics of the datasets after applying the preprocessing described in the main text. Figure 8 shows the fairness measures and user utility of π_{max} , π_{NSW} , $\pi_{1\text{-NSW}}$, $\pi_{\text{expo-fair}}$, and π_{unif} with varying K on Wiki10-31K. We observe that π_{NSW} is almost envy-free and achieves an equitable impact distribution for all values of K . In addition, $\pi_{1\text{-NSW}}$ achieves similar levels of user utility as $\pi_{\text{expo-fair}}$, while leading to a fairer impact distribution.




Figure 5: Fairness and user utility on synthetic data with varying lengths of ranking.

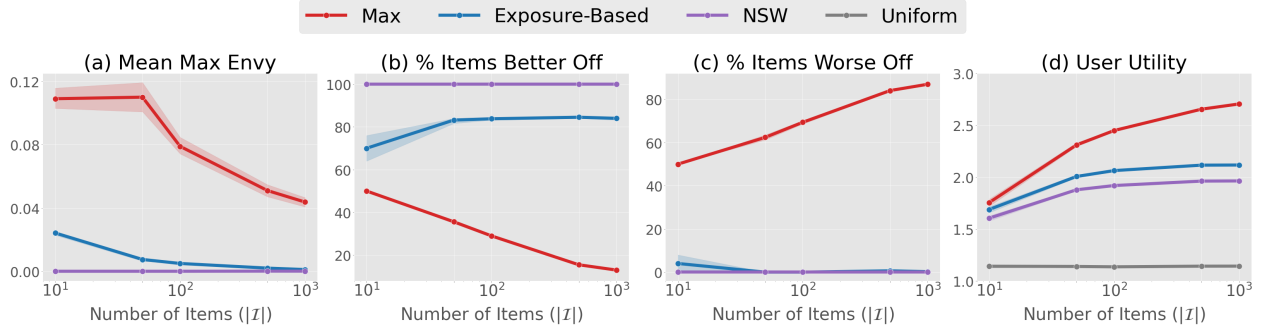


Figure 6: Fairness and user utility on synthetic data with varying numbers of items.




Figure 7: Fairness and user utility with different popularity patterns and with exponential examination function.




Figure 8: Fairness and user utility on the Wiki10-31K dataset with varying lengths of ranking.