Two-sided fairness in rankings via Lorenz dominance

Anonymous Author(s)
Affiliation
Address
email

Abstract

We propose a new approach to infer ranked lists of recommendations that are fair towards both consumers of the ranking ("users") and the "items" being ranked. We consider recommendations fair when they satisfy *efficiency* (it is not possible to make someone better off without making someone else worse off) and *equity* (they are maximally redistributive at a given level of overall utility). These two conditions are satisfied when a recommendation produces utility profiles for users and items with non-dominated generalized Lorenz curves. We achieve this by finding stochastic rankings that maximize a concave welfare function of utility profiles. The curvature of the welfare function and the relative weight of users and items control the degree of redistribution between users' and items' utilities. Compared to existing algorithms which maximize a trade-off between user utility and inequality of item utility, we provide theoretical and experimental support that our approach achieves superior trade-offs on problems with two-sided preferences. While standard approaches can destroy utility instead of redistributing it, our method efficiently redistributes utility between users and items.

1 Introduction

Online recommender systems have a growing impact on the information we see and on our life opportunities, as they help us browse news articles, find a new job, a house, or people to connect with. While the objective of recommender systems is traditionally defined as maximizing the quality of recommendations from the user's perspective, the recommendations also have impact on the recommended "items". News outlets rely on exposure to generate revenue, finding a job depends on which recruiter gets to see our resume, and the effectiveness of a dating application also depends on who we are recommended to—and if we are being recommended, then someone else is not. We define *two-sided fairness in rankings* as the problem of generating personalized recommendations by fairly mediating between the interests of all parties: users and items.

The leading approach to fairness in rankings is to define a state of "perfect fairness" through target amounts of exposure an item should receive—for instance, guaranteeing equal exposure, or exposure proportional to some estimate of overall quality [45, 36, 7]. Then, recommendations aim to maximize user utility under constraints. Intermediate trade-offs are obtained by relaxing the constraints, leading to user utility/item-side "fairness" trade-offs, but the fairness of the trade-off itself is undefined.

In this paper we propose an approach grounded in a branch of cardinal welfare economics, where the recommender system is seen as a social planner which decides on personalized rankings by trading-off utilities of users and producers. Cardinal welfare economics defines two criteria that a desirable fairness trade-off must satisfy: *efficiency* and *equity*. Efficiency means that a fairness trade-off is desirable only if making someone better off would make someone else worse off. Equity is a concept of optimal redistribution: keeping the sum of utility constant, it is always desirable to transfer utility from the better-off to the worse-off. Together, these concepts are captured by the notion of Lorenz dominance, which we introduce in Section 3.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

We propose to generate rankings by maximizing an increasing concave welfare function of users' and items' utilities. Monotonicity ensures efficiency and concavity ensures preference for equity. The parameters of the welfare function control the relative weight of users and items, and how much redistribution is desirable within users and within items. We show that our approach always yields rankings that are efficient and equitable, while existing approaches do not. First, one the main criterion of the literature, called equity of attention by Biega et al. [7], can lead to decrease the utility of every user, while *increasing* inequalities of exposure between items. Second, in reciprocal recommendation (such as dating applications), the existing criterion can lead to decrease the utility of every user, even the worse-off. Our notion of fairness prevents these undesirable behaviors.

We describe our approach to two-sided fairness in a framework that captures both usual recommendation tasks, such as music recommendation, as well as reciprocal recommendation problems such as dating where users are also items (Section 2). In Section 3, we formalize our notion of fairness and describe the abstract inference procedure that leads to fair trade-offs. We also provide performance guarantees when inference is carried out from estimated user/item preferences. In Section 4, we discuss existing methods from the perspective of efficiency and preference for equity. In Section 5, we present an efficient algorithm to generate fair rankings. Two-sided fairness is challenging from an algorithmic perspective because items' utilities depend on the rankings of *all users*, requiring global inference. Previous work addressed this issue with heuristic methods without guarantees or control on the achievable trade-offs. We show how the Frank-Wolfe algorithm can be leveraged to make inference tractable. Finally, in Section 6, we compare our approach and previous ones, showing that we achieve superior trade-offs on reciprocal and non-reciprocal recommendation tasks.

Related work The question of fairness in rankings originated from independent audits on recommender systems or search engines, which showed that results could exhibit bias against relevant social groups [48, 27, 17, 34, 30]. Our work follows the subsequent work on ranking algorithms that promote fairness of exposure for individual or sensitive groups of items [32, 10, 7, 45, 46, 36, 8, 50, 6, 41, 50, 53]. The goal is often to prevent winner-take-all effects, combat popularity bias [1] or promote smaller producers [33, 35]. Section 4 is devoted to the comparison with this type of approaches. Most of these works use a notion of fairness oriented towards items only. Wang and Joachims [50] propose to promote user-side fairness using concave functions of user utilities, similarly to us. Other works use equality constraints to define user-side fairness [6, 53]. These three approaches rely on the definitions of item-side fairness discussed in Section 4. Patro et al. [41] generate rankings that are envy-free on the user side, and guarantees the fair min-share for items. This approach is not amenable to controllable trade-offs between user and item utilities.

We are the first to address reciprocal and non-reciprocal recommendation within the same framework. There is less existing work studying the fairness of rankings in the reciprocal setting. Xia et al. [54] aim at equalizing user utility between groups, which suffers from the problems discussed in Section 4. Jia et al. [25] generate rankings using a welfare function approach, but optimizing only the utility of users *being recommended*. Paraschakis and Nilsson [40] postprocess rankings to correct for inconsistencies between estimated and declared preferences of users. We do not aim at correcting biases in preference estimates through post-processing. Rather, we aim at fair trade-offs between user and item utilities, under the assumption that biases in the preference estimates have been addressed earlier in the recommendation pipeline. Algorithmic fairness is also studied in the context of ridesharing applications [52, 31, 38], but they address matching rather than ranking problems.

There is growing interest in bridging the gap between fairness in machine learning and social choice theory [49, 4, 16, 21], and welfare economics in particular [47, 22, 28]. In line with Hu and Chen [22], who focused on classification and parity penalties, we argue that Pareto-efficiency should be part of the definition of algorithmic fairness. To the best of our knowledge, we are the first to propose Lorenz dominance as an equity criterion in fair machine learning.

2 A common framework for one-sided and reciprocal recommendation

Terminology and notation. We identify an item with its producer, so that "item utility" means "item producer's utility". The main paper focuses on fairness towards individual user and items. We describe in Appendix B the extension of our approach to sensitive groups of users or items. $|\mathcal{X}|$ denotes the cardinal of the set $\mathcal{X}$. Given $n \in \mathbb{N}$, $[\![n]\!] = \{1, \dots, n\}$. For an $n \times m \times k$ tensor P, P_i is the i-th horizontal slice, an $m \times k$ matrix, and P_{ij} is (i, j)-th frontal fiber.

In a recommendation task, there is a set of n entities (users or items), identified with $[\![n]\!]$. We denote by $\mathcal{N} \subseteq [\![n]\!]$ the set of users and by $\mathcal{I} \subseteq [\![n]\!]$ the set of items. In **one-sided recommendation**, users and items are different entities, i.e., $\mathcal{N} \cap \mathcal{I} = \emptyset$, while in **reciprocal recommendation**, users are also items so $\mathcal{N} = \mathcal{I}$. We denote by μ the $n \times n$ matrix where $\mu_{ij} \geq 0$ is is the value to entity i of the recommendation involving j. μ_{ij} is either the value to user i if they are recommended item j, or the value to i when it is recommended to user j. An important special case is **one-sided preferences**, where only users specify preferences. This is the usual setting of one-sided recommendation of music or movies. In that case, $\mu_{ij} \geq 0$ is the value to user i of item j, and we use $\mu_{ji} = 1$ because items only value the quantity of exposure they receive. In contrast, with **two-sided preferences**, our examples use $\mu_{ij} = \mu_{ji}$, which is the common practice because the value of the recommendation depends on the preferences of both parties. For instance, when recommending resumes to recruiters, $\mu_{ij} = \mu_{ji}$ can be a probability of interview, while in dating, it can be the probability of a "match".

A (deterministic) ranking $\sigma : \mathcal{I} \to [\![|\mathcal{I}|]\!]$ is a one-to-one mapping from items j to their rank $\sigma(j)$. Following [45], we use *stochastic rankings* because they allow us to perform inference using convex optimization (see Section 5). A stochastic ranking is represented by an $|\mathcal{I}| \times |\mathcal{I}|$ bistochastic matrix where each entry (j, k) is the probability that j is at rank k. The recommender system produces one stochastic ranking per user, represented by a 3-way *ranking tensor* P where P_{ijk} is the probability that j is recommended to i at rank k. To keep the notation coherent, P is an $n \times n \times |\mathcal{I}|$ tensor where irrelevant entries are set to 0 (using an $|\mathcal{N}| \times |\mathcal{I}| \times |\mathcal{I}|$ tensor would break the indexing of users and items in one-sided recommendation). We denote by $\mathcal{P}$ the set of ranking tensors.[1]

Utilities of users and items are defined through a position-based model, as in previous work [45, 7, 53]. Let $v \in \mathbb{R}^{|\mathcal{I}|}$, where v_k is the exposure weight at rank k. Since lower ranks receive more exposure $\forall k \in [\![|\mathcal{I}| - 1]\!], v_k \geq v_{k+1} \geq 0$.[2] Given a user i and a ranking σ_i, the *user-side utility* of i is the sum of the μ_{ij}s weighted by the exposure weight of their rank $\sigma_i(j)$: $\overline{u}_i(\sigma_i) = \sum_{j \in \mathcal{I}} v_{\sigma_i(j)} \mu_{ij}$. The *item-side utility* of i, denoted by $\overline{v}_i(P)$, is the sum over users j of the exposure of i to j weighted by μ_{ij}. The two-sided utility, or simply *utility* sums both. These definitions extend to stochastic rankings by taking the expectation over rankings, written in matrix form:[3]

$$\overbrace{\overline{u}_i(P) = \sum_{j \in \mathcal{I}} \mu_{ij} P_{ij} v}^{\substack{\text{user-side utility} \\ (j \text{ recommended to } i)}} \qquad \overbrace{\overline{v}_i(P) = \sum_{j \in \mathcal{N}} \mu_{ij} P_{ji} v}^{\substack{\text{item-side utility} \\ (i \text{ recommended to } j)}} \qquad \overbrace{u_i(P) = \overline{u}_i(P) + \overline{v}_i(P)}^{\text{(two-sided) utility}}$$

$\overline{v}_i(P)$ depends on the rankings of all users, while $\overline{u}_i(P)$ depends on P_i only. Thus, in general, finding rankings while taking into account users' and items' utilities leads to a global inference problem.

In one-sided recommendation, the user-side utility of an item is 0 and the item-side utility of a user is 0, so that the utility of every entity is the relevant one-sided utility. In reciprocal recommendation, the utility aggregates the values of the recommendations received and the value of being recommended to other users. This two-sided utility, which we introduce in this paper, is the key to addressing one-sided and reciprocal recommendation within a single framework.

Two-sided fairness in rankings. In practice, values of μ_{ij} are not known to the recommender system. Ranking algorithms use an estimate $\hat{\mu}$ of μ based on historical data. We address here the problem of *inference*: the goal is to compute the ranking tensor given $\hat{\mu}$, with the goal of making fair trade-offs between (true) user and item utilities.

3 Fairness via Lorenz dominance

3.1 Two-sided fairness via Lorenz-dominance

Our formal definition of fairness criterion focuses on *utility profiles* $\boldsymbol{u}(P) = (u_i(P))_{i=1}^n$. We denote by $\mathcal{U} = \{\boldsymbol{u}(P) : P \in \mathcal{P}\}$ the set of feasible profiles. For $\boldsymbol{u} \in \mathcal{U}$, $\boldsymbol{u}_{\mathcal{N}} = (u_i)_{i \in \mathcal{N}}$ and $\boldsymbol{u}_{\mathcal{I}} = (u_i)_{i \in \mathcal{I}}$ are the utility profiles of users and items respectively.

[1] Formally, $\mathcal{P}$ is the set of $n \times n \times |\mathcal{I}|$, $[0, 1]$-valued tensors P such that $P_{ijk} = 0$ if $i \notin \mathcal{N}$ or $j \notin \mathcal{I}$, and such that $\forall i \in \mathcal{N}$, we have: $\forall k \in [\![|\mathcal{I}|]\!] : \sum_{j \in \mathcal{I}} P_{ijk} = 1$ and, $\forall j \in \mathcal{I}, \sum_{k=1}^{n_i} P_{ijk} = 1$.

[2] We use a user-independent v for simplicity. Considering user-dependent weights is straightforward.

[3] We consider P_{ij} as a row vector in the formula, so that $P_{ij} v = \sum_{k=1}^{|\mathcal{I}|} P_{ijk} v_k$.

Given two utility profiles $\boldsymbol{u}$ and $\boldsymbol{u'}$, preference for efficiency means that it is socially desirable to improve the situation of someone if it is not at the expense of someone else. Equity states that if $\boldsymbol{u}$ is obtained from $\boldsymbol{u'}$ by redistributing utility from a better-off to a worse-off while keeping the sum of utilities constant, then $\boldsymbol{u}$ is socially preferred. These preference relations over utility profiles are formalized by Pareto- and Lorenz-dominance respectively. Let $\boldsymbol{u}$ and $\boldsymbol{u'}$ be two utility profiles:

Pareto-dominance $\succ_{\mathrm{P}}$. $\boldsymbol{u} \succ_{\mathrm{P}} \boldsymbol{u'} \iff (\forall i \in [\![n]\!], u_i \geq u'_i \text{ and } \exists i \in [\![n]\!], u_i > u'_i)$.

Lorenz-dominance $\succ_{\mathrm{L}}$. Let $(u_{(i)})_{i=1}^n$ be the sorted values in $\boldsymbol{u}$ from smallest to largest, i.e., $u_{(1)} \leq \ldots \leq u_{(n)}$, and let $U_i = u_{(1)} + \ldots + u_{(i)}$. Then $\boldsymbol{u} \succ_{\mathrm{L}} \boldsymbol{u'} \iff \boldsymbol{U} \succ_{\mathrm{P}} \boldsymbol{U'}$.

Non-strict Lorenz-dominance is denoted by $\boldsymbol{u} \succeq_{\mathrm{L}} \boldsymbol{u'}$.[4] Lorenz-dominance captures *efficiency* since $\boldsymbol{u} \succ_{\mathrm{P}} \boldsymbol{u'} \implies \boldsymbol{u} \succ_{\mathrm{L}} \boldsymbol{u'}$. It also captures *equity* through $\boldsymbol{U}$, which represents the cumulative utility of fractions of the populations from the "poorest" to the "richest". A transfer from "rich" to "poor" does not change the cumulative utility owned up to the "rich" but increases that of the "poor", leading to a Lorenz-dominant profile. A utility profile with perfect equality cannot be Lorenz dominated.

Lorenz-dominance relies on cardinal utilities, which implies that interpersonal comparisons are meaningful. We take the point of view that utilities are *a priori* comparable between users and between items, but defining how to trade-off user and item utilities is context dependent. We thus define our notion of two-sided fairness, Lorenz efficiency, by requiring that the utility profile of the users and the utility profile of the items are not jointly dominated. Lorenz-efficient rankings are Pareto-efficient rankings that perform an optimal redistribution within users and within items.

Definition 1 (Lorenz efficiency)**.** *A utility profile $\boldsymbol{u} \in \mathcal{U}$ is* Lorenz-efficient *if there is no $\boldsymbol{u'} \in \mathcal{U}$ such that either ($\boldsymbol{u}'_{\mathcal{I}} \succeq_{\mathrm{L}} \boldsymbol{u}_{\mathcal{I}}$ and $\boldsymbol{u}'_{\mathcal{N}} \succ_{\mathrm{L}} \boldsymbol{u}_{\mathcal{N}}$) or ($\boldsymbol{u}'_{\mathcal{N}} \succeq_{\mathrm{L}} \boldsymbol{u}_{\mathcal{N}}$ and $\boldsymbol{u}'_{\mathcal{I}} \succ_{\mathrm{L}} \boldsymbol{u}_{\mathcal{I}}$).*

A ranking tensor $P \in \mathcal{P}$ is Lorenz-efficient if its utility profile $\boldsymbol{u}(P)$ is Lorenz-efficient.

Additive concave welfare functions of utility profiles order profiles according to Lorenz dominance [3]. This result holds when all utilities are comparable. In our case where there are users and items, we propose the following welfare function parametrized by $\theta = (\lambda, \alpha_1, \alpha_2)$:[5]

$$\forall \boldsymbol{u} \in \mathbb{R}_+^n : \quad W_\theta(\boldsymbol{u}) = (1-\lambda) \sum_{i \in \mathcal{N}} \psi(u_i, \alpha_1) + \lambda \sum_{j \in \mathcal{I}} \psi(u_j, \alpha_2) \quad \text{with } \psi(x, \alpha) = \begin{cases} x^\alpha & \text{if } \alpha > 0 \\ \log(x) & \text{if } \alpha = 0 \\ -x^\alpha & \text{if } \alpha < 0 \end{cases}.$$

In reciprocal recommendation, we choose $\alpha_1 = \alpha_2$ so that $W_\theta(\boldsymbol{u}) = \sum_{i=1}^n \psi(u_i, \alpha)$. In all cases, inference is carried out by maximizing W_θ (an efficient algorithm is proposed in Section 5):

$$\textit{(ranking procedure)} \qquad P^* \in \operatorname*{argmax}_{P \in \mathcal{P}} W_\theta(\boldsymbol{u}(P)) \tag{1}$$

In W_θ, $\lambda \in [0, 1]$ controls the relative weight of users and items. The motivation for the specific choice of ψ is that it appears in welfare economics in scale invariant welfare functions [37], but other families can be used as long as the functions are *increasing* and *concave*. Monotonicity implies that maxima of W_θ are Pareto-efficient. For $\alpha_1 < 1$ and $\alpha_2 < 1$, W_θ is strictly concave. Then, W_θ exhibits diminishing returns with respect to individual utilities. Constant-sum transfers from "rich" to "poor" within users and items then increase W_θ, which yields Lorenz-efficiency. Note that transfers between users and items need not be constant sum. How W_θ distributes utility between users and items is controlled by λ and the difference between α_1 and α_2. Let $\Theta = \{(\lambda, \alpha_1, \alpha_2) \in (0, 1) \times (-\infty, 1)^2\}$. For every $\theta \in \Theta$, W_θ is strictly concave, and users and items have non-zero weight. We then have:

Proposition 1. *$\forall \theta \in \Theta, \forall P^* \in \operatorname*{argmax}_{P \in \mathcal{P}} W_\theta(P)$, P^* is Lorenz-efficient.*

The proof is straightforward from the diminishing return property (see Appendix C).

The most egalitarian trade-off achievable by our method is described by the leximin order [44]. Given two utility profiles $\boldsymbol{u}$ and $\boldsymbol{u'}$, $\boldsymbol{u} \geq_{\text{lex}} \boldsymbol{u'}$ if $\boldsymbol{U}$ is greater than $\boldsymbol{U'}$ according to the lexicographic order.[6] The leximin optimal profile is egalitarian in the sense that it maximizes the utility of individuals in sequence, from the worse-off to the better-off. Depending on the set of feasible profiles, this may not lead to equal utility for everyone, but any further reduction of inequality can only be achieved by making people worse off for the benefit of no other, in violation of Pareto-dominance.

[4] I.e., $\boldsymbol{u} \succeq_{\mathrm{L}} \boldsymbol{u'} \iff \forall i \in [\![n]\!], U_i \geq U'_i$.

[5] $W_\theta(\boldsymbol{u}) = -\infty$ if $\alpha \leq 0$ and $\exists i, u_i = 0$. In practice, we use $\psi(x + \eta, \alpha)$ for $\eta > 0$ to avoid this case.

[6] Formally, $\boldsymbol{u} >_{\text{lex}} \boldsymbol{u'}$ if $(\exists k \in [\![d]\!]$ s.t. $\forall i < k, U_i = U'_i$ and $U_k > U'_k)$. $\boldsymbol{u} \geq_{\text{lex}} \boldsymbol{u'} \iff \neg(\boldsymbol{u'} \geq_{\text{lex}} \boldsymbol{u})$.

In Appendix C.2, we show that item-side leximin-optimal rankings are found by keeping α_1 constant and having $\alpha_2 \to -\infty$, in cases where λ converges to 1 faster than α_2 tends to $-\infty$. On the other end of achievable trade-offs lies the utilitarian ranking ($\alpha_1 = \alpha_2 = 1$), which we analyze below. For a ranking tensor P and a user i, we denote by $\mathfrak{S}(P_i)$ the support of P_i in ranking space.[7] We remind that $\sigma(j)$ is the rank of item j, and that lower ranks are better.

Proposition 2 (Utilitarian ranking)**.** *Let* $P^* \in \underset{P\in\mathcal{P}}{\operatorname{argmax}} W_{\frac{1}{2},1,1}(P) = \underset{P\in\mathcal{P}}{\operatorname{argmax}} \sum_{i\in\mathcal{N}} u_i(P) + \sum_{j\in\mathcal{I}} u_j(P)$.

1. $\forall i \in \mathcal{N}, \forall \sigma \in \mathfrak{S}(P_i^*) : \sigma(j) < \sigma(j') \implies \tilde{\mu}_{ij} \geq \tilde{\mu}_{ij'}$ *with* $\tilde{\mu}_{ij} = \mu_{ij} + \mu_{ji}$.
2. *If* $\forall (i,j) \in \mathcal{N} \times \mathcal{I}, \mu_{ji} = 1$ *or* $\forall (i,j) \in [\![n]\!]^2, \mu_{ij} = \mu_{ji}$, *then* $\tilde{\mu}_{ij} \geq \tilde{\mu}_{ij'} \iff \mu_{ij} \geq \mu_{ij'}$.
3. *Under the condition of 2., if* $\mathcal{N} \cap \mathcal{I} = \emptyset$ *then* $\forall \boldsymbol{u'} \in \mathcal{U}, \boldsymbol{u}_{\mathcal{N}}(P^*) \geq_{\text{lex}} \boldsymbol{u}'_{\mathcal{N}}$.

Proposition 2 generalizes to two-sided utilities the well-known result that maximizing user-side utility $\overline{u}_i(P)$ is achieved by sorting $j \in \mathcal{I}$ by decreasing μ_{ij} (see e.g., [12]). The proof is in Appendix C. The second part applies to practical cases where either only item exposure is important to items, or when the value is symmetric between users and items. There, two remarks apply. First, the utilitarian ranking is the same as the usual sort by decreasing μ_{ij}. This means that without considerations of two-sided fairness ($\alpha_1, \alpha_2 < 1$), the optimal ranking for two-sided utilities is the same as the usual ranking. This might explain why the two-sided utility has never been studied before, even in reciprocal recommendation [39]. Second, in one-sided recommendation, the utilitarian ranking is also the leximin optimal on the user-side. Any attempt to reduce inequality in user utilities can only decrease the utility of some users for the benefit of no other. This shows that fundamentally, the concept of user-side fairness (defined as driving towards more equity in user utility) is of practical significance only in a context of two-sided fairness. In that case, preference for equity on the user side means that transfers between users and items mostly comes at the expense the most satisfied users.

3.2 Guarantees when performing inference with estimated preferences

In practice, inference is carried out on an estimate $\hat{\mu}$ of μ, meaning that, denoting $\hat{\boldsymbol{u}}$ the resulting estimated utility[8] the system output $\hat{P} = \operatorname{argmax}_{P\in\mathcal{P}} W_\theta(\hat{\boldsymbol{u}}(P))$. The following result extends surrogate regret bounds that exist in classification [5, 55] and learning to rank [12, 42, 2, 29] to the case of welfare functions and global stochastic rankings. It makes the link between the quality of the estimate $\hat{\mu}$ and an optimality guarantee for $\boldsymbol{u}(\hat{P})$ (i.e., the true welfare of the ranking inferred on the estimated values). We prove the result for $\theta = (\frac{1}{2}, \alpha, \alpha)$ for $\alpha \leq 1$ to simplify notation.[9]

Theorem 1. *Let* $\alpha \leq 1$ *and* $\theta = (\frac{1}{2}, \alpha, \alpha) \in \Theta$. *Let* $\hat{\mu} \in \mathbb{R}_+^{n\times n}$, $\hat{P} = \operatorname{argmax}_{P\in\mathcal{P}} W_\theta(\hat{\boldsymbol{u}}(P))$, *and* $P^* = \operatorname{argmax}_{P\in\mathcal{P}} W_\theta(\boldsymbol{u}(P))$.

Let furthermore $B(\hat{\mu}) = \max\left(\max_{i\in[\![n]\!]} \psi'(u_i(\hat{P}), \alpha), \max_{i\in[\![n]\!]} \psi'(\hat{u}_i(P^*), \alpha)\right)$. *We have:*

$$W_\theta(\boldsymbol{u}(P^*)) - W_\theta(\boldsymbol{u}(\hat{P})) \leq 4B(\hat{\mu})\sqrt{n\,\|v\|_2}\sqrt{\sum_{(i,j)\in[\![n]\!]^2} (\hat{\mu}_{ij} - \mu_{ij})^2}.$$

4 The unfairness of trading off user utility and inequality in item utilities

Previous work on fair rankings mostly studied item-side fairness defined as equal utility or quality-weighted exposure, in one-sided recommendation. We show that these approaches, as well as their natural extension to reciprocal recommendation that is permitted by our definition of (two-sided) utility, do not produce trade-offs that satisfy the desirable properties of efficiency and equity.

4.1 The inefficiency of equal item utility

The main approach to item-side fairness is what we call *equality of item utility*[10] [45, 36, 6, 51, 6, 53], which stipulates $u_j = u_{j'}$ for every two items j and j', or, equivalently, that every u_j should be equal

[7] Formally, $\mathfrak{S}(P_i) = \left\{\sigma : \mathcal{I} \to [\![|\mathcal{I}|]\!] \,\middle|\, \sigma \text{ is one-to-one, and } \forall j \in \mathcal{I}, P_{ij\sigma(j)} > 0\right\}$.

[8] We have $\hat{u}_i(P) = \sum_{j=1}^n \hat{\mu}_{ij}(P_{ij} + P_{ji})v$.

[9] The dependency on $\hat{\mu}$ in $B(\hat{\mu})$ is because $\psi'(., \alpha)$ is not bounded in general. In practice, we use $\psi(x+\eta, \alpha)$ for a small $\eta > 0$ to avoid the singular point at 0, in which case $B < \psi'(\eta, \alpha)$.

[10] also called "disparate treatment" constraints in the literature [45, 6]

to the average item utility. To generate trade-offs between user utility and equality, we relax the constraint using the following objective, where $\beta > 0$ controls the trade-off:

$$\textit{equality of item utility} \qquad F_\beta(\boldsymbol{u}) = \sum_{i\in\mathcal{N}} u_i - \beta\sqrt{D(\boldsymbol{u})} \qquad \text{with } D(\boldsymbol{u}) = \sum_{j\in\mathcal{I}} \Big(u_j - \frac{1}{|\mathcal{I}|}\sum_{j'\in\mathcal{I}} u_{j'}\Big)^2.$$

The square root makes the penalty on the same scale as user utilities. Some authors use $D'(\boldsymbol{u}) = \sum_{(j,j')\in\mathcal{I}^2} |u_j - u'_j|$, instead of $D(\boldsymbol{u})$ [46, 36, 6]. For the purpose of this section, D and D' have qualitativaly the same behavior: they both drive towards pure equality rather than leximin optimality, they may thus lead to inefficient solutions. We propose D as a computationally efficient alternative to D', since it involves only a linear number of terms and $\sqrt{D}$ is convex and differentiable except on 0.

We first prove the existence of special cases where, for any choice of parameter $\theta \in \Theta$ of our approach, we obtain solutions with utility profiles for both users and items that Lorenz dominate those found by equality of item utility. We also show that in reciprocal recommendation, striving for pure equality can even lead to 0 utility for every user, leading to a decrease of every user's utility, even that of the worse off. Figure 1 gives an example. The proposition describes extreme cases, our experiments (Secion 6) study how the approaches behave on real data. All proofs of this section are in Appendix D. For $\theta \in \Theta$, let $\boldsymbol{u}^\theta = \operatorname{argmax}_{\boldsymbol{u}\in\mathcal{U}} W_\theta(\boldsymbol{u})$. For $\beta > 0$, let $\mathcal{U}^{\text{eq}}_\beta = \operatorname{argmax}_{\boldsymbol{u}\in\mathcal{U}} F_\beta(\boldsymbol{u})$.

Proposition 3. *The following claims hold irrespectively of the choice of $\boldsymbol{u}^{\text{eq},\beta} \in \mathcal{U}^{\text{eq}}_\beta$.*

- *For every $N \geq 2$ and $M \geq 2$, there is a one-sided recommendation task with N users and M items, with $\forall i, j, \mu_{ij} = \mu_{ji} \in [0, 1]$ and $\forall j \in \mathcal{I}, \max_{i\in\mathcal{N}} \mu_{ij} = 1$, such that, $\forall\theta \in \Theta$, we have:*
$$(\forall\beta > 0,\ \ \boldsymbol{u}^\theta_{\mathcal{N}} \succ_{\text{L}} \boldsymbol{u}^{\text{eq},\beta}_{\mathcal{N}} \ \textit{ and }\ \boldsymbol{u}^\theta_{\mathcal{I}} \succ_{\text{L}} \boldsymbol{u}^{\text{eq},\beta}_{\mathcal{I}}) \quad \textit{and} \quad \lim_{\beta\to\infty} \sum_{i\in\mathcal{N}} u_i^{\text{eq},\beta} = \frac{1}{M-1}\sum_{i\in\mathcal{N}} u_i^\theta.$$
- *For every $n \geq 5$, there is a reciprocal recommendation task with n users such that:*
$$\forall\theta \in \Theta, \boldsymbol{u}^\theta, \forall\beta > 0 :\ \ \forall i \in [\![n]\!], u_i^\theta > u_i^{\text{eq},\beta} \qquad \textit{and} \qquad \lim_{\beta\to\infty} \sum_{i\in\mathcal{N}} u_i^{\text{eq},\beta} = 0.$$

4.2 One-sided preferences: the inequity of quality-weighted exposure

In one-sided recommendation with one sided preferences, item utility is equal to the item's exposure $u_j(P) = \sum_{i\in\mathcal{N}} P_{ij} v$, and we refer to the criterion above as *equality of exposure*. Several authors advocated for a variant of it which accounts for the relevance of the items in the fairness constraints [45, 46, 7, 53]. The new criterion, called *quality-weighted exposure*[11] aims for an exposure proportional to the quality of the item, defined as $q_j = \sum_{i\in\mathcal{N}} \mu_{ij}$. Denoting by $Q = \sum_{j\in\mathcal{I}} q_j$ the total quality and noticing that the total exposure is $E = |\mathcal{N}|\,\|v\|_1$, we obtain the following objective:

$$\textit{quality-weighted exposure} \quad F^{\text{qua}}_\beta(\boldsymbol{u}) = \sum_{i\in\mathcal{N}} u_i(P) - \beta\sqrt{D^{qua}(\boldsymbol{u})} \ \text{ with } \ D^{qua}(\boldsymbol{u}) = \sum_{j\in\mathcal{I}} \Big(u_j(P) - \frac{q_j E}{Q}\Big)^2.$$

We show here that in some cases, compared to any choice of the parameter $\theta \in \Theta$ of the welfare approach, quality-weighted exposure leads to the undesirable behavior of *decreasing user utility* while *increasing inequalities of exposure* between items. Figure 1 gives an example.

Proposition 4. *For $\beta > 0$, let $\mathcal{U}^{\text{qua}}_\beta = \operatorname{argmax}_{\boldsymbol{u}\in\mathcal{U}} F^{\text{qua}}_\beta(\boldsymbol{u})$. The following claim holds irrespective of the choice of $\boldsymbol{u}^{\text{qua},\beta} \in \mathcal{U}^{\text{qua}}_\beta$.*

For every $d \in \mathbb{N}_*$ and every $N \in \mathbb{N}_*$, there is a one-sided recommendation problem with one-sided preferences, with $d + 1$ items and $N(d + 1)$ users, such that $\forall\theta \in \Theta$, we have:

$$\big(\forall\beta > 0, \boldsymbol{u}^\theta \succ_{\text{L}} \boldsymbol{u}^{\text{qua},\beta}_{\mathcal{N}} \textit{ and } \boldsymbol{u}^\theta \succ_{\text{L}} \boldsymbol{u}^{\text{qua},\beta}_{\mathcal{I}}\big) \qquad \textit{and} \qquad \lim_{\beta\to\infty} \sum_{i\in\mathcal{N}} u_i^{\text{qua},\beta} \xrightarrow[d\to\infty]{} \frac{5}{6}\sum_{i\in\mathcal{N}} \boldsymbol{u}^\theta.$$

Interestingly, equality of exposure penalties do not suffer from the flaws described in this section because the total budget of exposure E is constant.[12] In Appendix D, we discuss additional results concerning the extension or parity or quality-weighted criteria to reciprocal recommendation.

[11] We use here the terminology of [53]. This criterion has also been called "disparate treatment" [45], "merit-based fairness" Singh and Joachims [46] and "equity of attention" [7].

[12] Constant total exposure is a consequence of the position-based user model chosen to define utility. This is the standard choice [45, 36, 6, 34**?** , 53], even if it does not model the case where the relevance of recommendations affect the user's browsing behavior. Taking such feedback loops into account is out of the scope of this work.

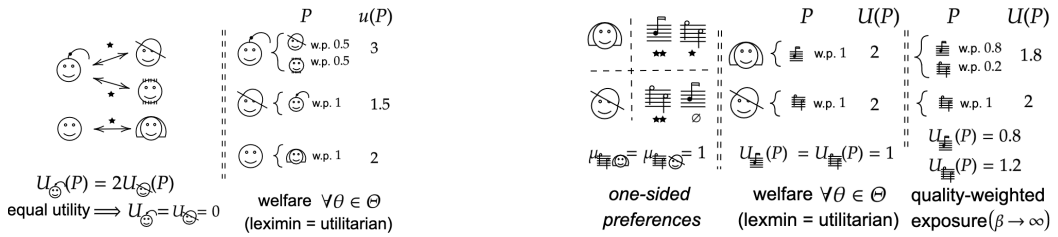

Figure 1: **Left:** Example of a reciprocal recommendation task where equality of utility leads to 0 utility. There is one recommendation slot per user. We give the recommendation probabilities and utilities for the utilitarian ranking and three users, the other ones are obtained by the symmetry of the problem. The utilitarian ranking is also leximin optimal, so our approach yields the same recommendations for all θ. The two top-left users are equivalent, so the utility of the top left user is twice the utility of the top right user, making equality is possible only at 0 utility. **Right:** Example where quality-weighted exposure reduces user utility while increasing inequalities between items.

5 Efficient inference of fair rankings with the Frank-Wolfe algorithm

We now present our inference algorithm for (1). Appendix E contains the proofs of this section and describes a similar approach for the objective functions of the previous section. From an abstract perspective, the goal is to find a maximum P^* such that:

$$P^* \in \operatorname*{argmax}_{P \in \mathcal{P}} W(P) \quad \text{with} \quad W(P) = \sum_{i=1}^{n} \Phi_i\Big(\sum_{j=1}^{n} \mu_{ij}(P_{ij} + P_{ji})v\Big)$$

where for every i, $\Phi_i : \mathbb{R}_+ \to \mathbb{R}$ is concave increasing, $\mu_{ij} \geq 0$ and v is a vector of non-negative non-increasing values. Since W is concave and $\mathcal{P}$ is defined by equality constraints, the problem above is a convex optimization problem. However, this is a global optimization problem over the rankings of all users, so a naive approach would require $|\mathcal{N}||\mathcal{I}|^2$ parameters and $2|\mathcal{N}||\mathcal{I}|$ linear constraints. The same problem arises with the penalties of previous work. In the literature, authors either considered applying the item-fairness constraints to each ranking individually [45, 6], which leads to inefficiencies with our definition of utility (see Appendix H), or resort to heuristics to compute the rankings one by one without guarantees on the trade-offs that are achieved [36, 7].

Our approach is based on the Frank-Wolfe algorithm [15], which was previously used in machine learning in e.g., structured output prediction or low-rank matrix completion [24], but to the best of our knowledge not for ranking. Denoting $\langle X \mid Y \rangle = \sum_{ijk} X_{ijk} Y_{ijk}$ the dot product between tensors, the algorithm creates iterates $P^{(t)}$ by first computing $\tilde{P} = \operatorname{argmax}_{P \in \mathcal{P}} \langle P \mid \nabla W(P^{(t)}) \rangle$ and then updating $P^{(t)} = (1 - \gamma^{(t)})P^{(t-1)} + \gamma^{(t)}\tilde{P}$ with $\gamma^{(t)} = \frac{2}{t+2}$ [11]. Starting from an initial solution[13], the algorithm always stays in the feasible region without any additional projection step. Our main contribution of this section is to show that $\operatorname{argmax}_{P \in \mathcal{P}} \langle P \mid \nabla W(P^{(t)}) \rangle$ can be computed efficiently, requiring only one sort operation per user after computing the utilities.

Theorem 2. *Let* $\tilde{\mu}_{ij} = \Phi_i'\big(u_i(P^{(t)})\big)\mu_{ij} + \Phi_j'\big(u_j(P^{(t)})\big)\mu_{ji}$. *Let* $\tilde{P}$ *such that:*

$$\forall i \in \mathcal{N}, \forall \tilde{\sigma}_i \in \mathfrak{S}(\tilde{P}): \quad \tilde{\sigma}_i(j) < \tilde{\sigma}_i(j') \implies \tilde{\mu}_{ij} \geq \tilde{\mu}_{ij'}. \text{ Then } \tilde{P} \in \operatorname*{argmax}_{P \in \mathcal{P}} \langle P \mid \nabla W(P^{(t)}) \rangle.$$

Moreover, it produces a compact representation of the stochastic ranking as a weighted sum of permutation matrices. The number of iterations of the algorithm allows to control the trade-off between memory requirements and accuracy of the solution. Using previous convergence results for the Frank-Wolfe algorithm [11], assuming each Φ_i'' is bounded, we have:

Proposition 5. *Let K be the maximum index of a nonzero value in v (or $|\mathcal{I}|$). The iterates $P^{(t)}$ satisfy $W(P^{(t)}) \geq \max_{P \in \mathcal{P}} W(P) - O(n/t)$. Moreover, each iteration costs $O(|\mathcal{I}| \ln K)$ operations per user and only require $O(K)$ additional bytes of storage per user.*

[13] In our experiments, we initialize with the utilitarian ranking (Proposition 2).

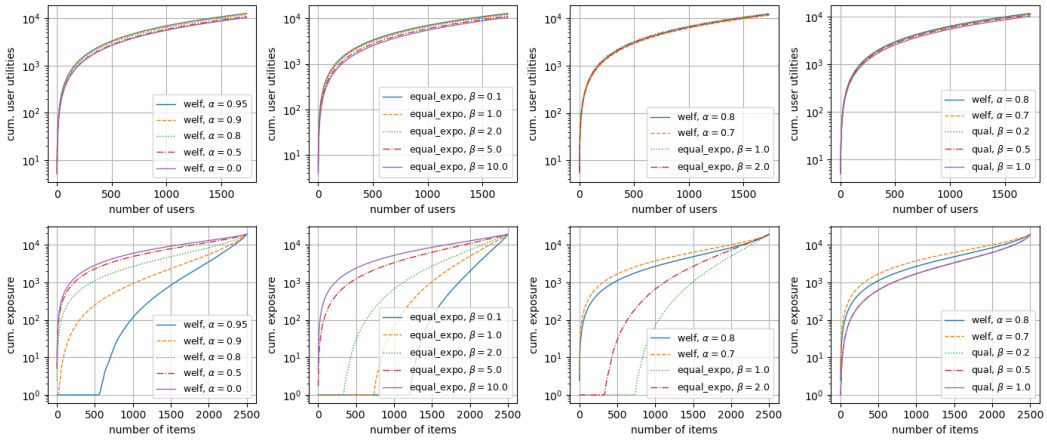

Figure 2: Trade-offs obtained on Last.fm. (left, middle left): generalized Lorenz curves of our method, equality of exposure and quality-weighted exposure on estimated preferences. (middle right) zoom on trade-offs where equity of exposure is greatly increased without dramatic effect on user utility. (right) quality weighted exposure rapidly converges to a fixed exposure pattern with clearly dominated user and item curves.

6 Experiments

6.1 One-sided recommendation with one-sided preferences: Last.fm and MovieLens

We first present experiments on one-sided recommendations with one-sided preferences, carried out on the Last.fm dataset [9] (also used by [41]) which contains the music listening histories of $1.9k$ users. We select the top 2500 items most listened to, and estimate $\hat{\mu}_{ij}$ using a standard matrix factorization algorithm [23] trained on 80% of the available interactions, including hyperparameter selection using cross-validation. All results are averaged over three seeds which correspond to three train/test splits. We provide the full details of the experiments, as well as a deeper analysis of the effect of the parameters of our method, in Appendix F.1. Additional results with qualitatively similar results on movie recommendation using the MovieLens20m dataset [19] are presented in Appendix F.2.

Since the different ranking algorithms operate on estimated preferences only, we focus here on comparing the trade-offs they obtain assuming the estimation is the ground truth. The appendix contains additional evaluations on the test set of the remaining 20% interactions. The algorithm presented in Section 5 is used to perform inference for all methods.

We compare the different trade-offs in terms of their generalized Lorenz curves, which plots the fraction of utility owned by the percentiles of the individuals with less utility. We plot the y-axis in log-scale for better visibility at the beginning of the curve. Since users and items do not have comparable utilities, we focus on joint dominance of curves: a parameter setting dominates another if its curve for one of the groups (users or items) is strictly above the other method, while being at least as high for the other group.

We compare the trade-offs obtained by our approach (*welf*) to the penalties of Section 4. The item-side utility is then the amount of exposure an item receives. Equality of utility is then called *equality of exposure* (*equal_expo*). We also compare to quality-weighted exposure (*qual*), which is the criterion used in the experiments of previous work [7, 36, 53]. For our method, we fix λ to balance the user and item weights and use $\alpha_1 = \alpha_2 = \alpha$. For the penalty-based methods, we vary β.

The results are shown in Figure 2. We see on the first two left columns the different trade-offs achieved by *welf* and *equal_expo*, which both span the entire range of exposure. The third column zooms in on trade-offs achievable while keeping the total user utility near the maximum. We observe that both *welf* and *equal_expo* significantly increase equality in exposure with little degradation of user utility, which is an interesting operating point in practice. In the last column, we see that quality of exposure gradually converges to an exposure profile that is dominated by *welf*, both in terms of

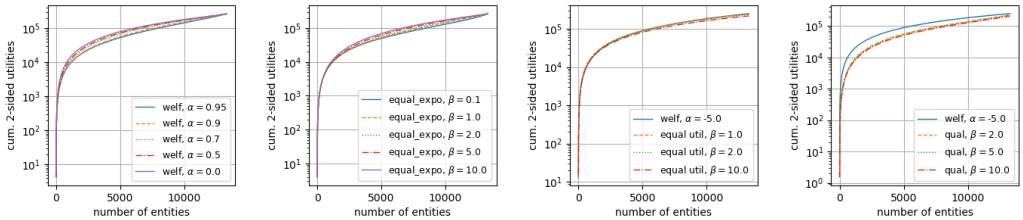

Figure 3: Trade-offs obtained in reciprocal recommendation on Twitter. (Left) Lorenz curves computed on test. (Right) Lorenz curves computed on estimated preferences.

user utility and item exposure. Overall, our results show that our method successfully makes fair trade-offs, and suggest that the welfare approach and equality of exposure penalties lead to better trade-offs than quality-weighted exposure.

6.2 Reciprocal recommendation: Twitter and Epinions

We perform a similar analysis on reciprocal recommendation. Since there is no widely accepted benchmark for reciprocal recommendation, we generate a reciprocal recommendation task from people-to-people recommendation problems. We use the Higgs **Twitter** dataset [13], which contains follower links, and address the task of finding mutual followers. For each dataset, we keep users having at least 20 mutual links, resulting in a subset of 13k users. Estimates of follow probabilities $\hat{\phi}_{ij} \in [0, 1]$ are learnt with logistic matrix factorization [26] trained on a 80% sample of directed links. The estimated symmetric preferences are $\hat{\mu}_{ij} = \hat{\phi}_{ij}\hat{\phi}_{ji}$. The hyperparameters of the logistic matrix factorization are chosen by cross-validation on this 80% sample. the remaining 20% of links are used to compute the test utility.

We compare our method (*welf*) to equal utility (*equal_util*) and the extension of equality of exposure and quality-weighted exposure. Here, there are only users and the two-sided utility accounts for the user-side and item-side utility, for both the welfare approach and equal utility. Appendix F.3 contains the full details of the experiments and additional results, while Appendix F.4 provides additional experiments on the Epinions dataset [43].

Figure 3 depicts generalized Lorenz curves for the two-sided utility for one value of our method compared to various values of equal utility and quality of exposure. The values have been chosen to represent relatively egalitarian profiles. The two left plots show different trade-offs on estimated preferences for *welf* and parity of exposure. The two other plots compare *welf* and equality of utility and quality-weighted exposure. On this task, we observe a slight dominance of *welf* against equality of utility, and a clear dominance on quality-weighted exposure.

Overall, the experiments conclude that quality-weighted exposure is suboptimal for mediating between user and item utilities. For equality of exposure and equality of utility, despite the drive towards equality that might lead to inefficient solution, they may still lead to reasonable trade-offs in practice.

7 Conclusion

We presented a new approach to fairness in rankings rooted in cardinal welfare economics. Our notion of fairness based on preference for efficiency and equity. Based on a theoretical and experimental analysis, we argue that our approach to fairness in rankings produces trade-offs that are preferable to those that can result from the equality constraints that are popular in the literature.

The expected positive societal impact of this work is to provide more principled approaches to mediating between several parties on a recommendation platform. Yet, there are several limitations that should be addressed in future work. First, the notion of fairness relies on carefully designed utilities. The careful design of these utilities is critical for the practical fairness of the approach [44]. Second, our approach is not meant to be a substitute for addressing fairness at all levels of the pipeline, such as in the estimation of preferences.

References

[1] H. Abdollahpouri, M. Mansoury, R. Burke, and B. Mobasher. The unfairness of popularity bias in recommendation. *arXiv preprint arXiv:1907.13286*, 2019.

[2] S. Agarwal. Surrogate regret bounds for bipartite ranking via strongly proper losses. *The Journal of Machine Learning Research*, 15(1):1653–1674, 2014.

[3] A. B. Atkinson. On the measurement of inequality. *Journal of economic theory*, 2(3):244–263, 1970.

[4] M.-F. Balcan, T. Dick, R. Noothigattu, and A. D. Procaccia. Envy-free classificatoion. *arXiv preprint arXiv:1809.08700*, 2018.

[5] P. L. Bartlett, M. I. Jordan, and J. D. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.

[6] K. Basu, C. DiCiccio, H. Logan, and N. E. Karoui. A framework for fairness in two-sided marketplaces. *arXiv preprint arXiv:2006.12756*, 2020.

[7] A. J. Biega, K. P. Gummadi, and G. Weikum. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international acm sigir conference on research & development in information retrieval*, pages 405–414, 2018.

[8] R. Burke. Multisided fairness for recommendation. *arXiv preprint arXiv:1707.00093*, 2017.

[9] I. Cantador, P. Brusilovsky, and T. Kuflik. 2nd workshop on information heterogeneity and fusion in recommender systems (hetrec 2011). In *Proceedings of the 5th ACM conference on Recommender systems*, RecSys 2011, New York, NY, USA, 2011. ACM.

[10] L. E. Celis, D. Straszak, and N. K. Vishnoi. Ranking with fairness constraints. *arXiv preprint arXiv:1704.06840*, 2017.

[11] K. L. Clarkson. Coresets, sparse greedy approximation, and the frank-wolfe algorithm. *ACM Transactions on Algorithms (TALG)*, 6(4):1–30, 2010.

[12] D. Cossock and T. Zhang. Statistical analysis of bayes optimal subset ranking. *IEEE Transactions on Information Theory*, 54(11):5140–5154, 2008.

[13] M. De Domenico, A. Lima, P. Mougel, and M. Musolesi. The anatomy of a scientific rumor. *Scientific reports*, 3(1):1–9, 2013.

[14] J. Finocchiaro, R. Maio, F. Monachou, G. K. Patro, M. Raghavan, A.-A. Stoica, and S. Tsirtsis. Bridging machine learning and mechanism design towards algorithmic fairness. *arXiv preprint arXiv:2010.05434*, 2020.

[15] M. Frank, P. Wolfe, et al. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.

[16] P. Gölz, A. Kahng, and A. D. Procaccia. Paradoxes in fair machine learning. *Advances in Neural Information Processing Systems (NeurIPS). Forthcoming*, 2019.

[17] A. Hannak, G. Soeller, D. Lazer, A. Mislove, and C. Wilson. Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 conference on internet measurement conference*, pages 305–318, 2014.

[18] G. H. Hardy, J. E. Littlewood, and G. Pólya. Inequalities. 2nd ed. Cambridge, Engl.: At the University Press. XII, 324 p. (1952)., 1952.

[19] F. M. Harper and J. A. Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.

[20] H. Heidari, C. Ferrari, K. Gummadi, and A. Krause. Fairness behind a veil of ignorance: A welfare analysis for automated decision making. In *Advances in Neural Information Processing Systems*, pages 1265–1276, 2018.

[21] S. Hossain, A. Mladenovic, and N. Shah. Designing fairly fair classifiers via economic fairness notions. In *Proceedings of The Web Conference 2020*, pages 1559–1569, 2020.

[22] L. Hu and Y. Chen. Fair classification and social welfare. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 535–545, 2020.

[23] Y. Hu, Y. Koren, and C. Volinsky. Collaborative filtering for implicit feedback datasets. In *2008 Eighth IEEE International Conference on Data Mining*, pages 263–272. Ieee, 2008.

[24] M. Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *International Conference on Machine Learning*, pages 427–435. PMLR, 2013.

[25] Y. Jia, X. Liu, and W. Xu. When online dating meets nash social welfare: Achieving efficiency and fairness. In *Proceedings of the 2018 World Wide Web Conference*, pages 429–438, 2018.

[26] C. C. Johnson. Logistic matrix factorization for implicit feedback data. *Advances in Neural Information Processing Systems*, 27(78):1–9, 2014.

[27] M. Kay, C. Matuszek, and S. A. Munson. Unequal representation and gender stereotypes in image search results for occupations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 3819–3828. ACM, 2015.

[28] J. Kleinberg, J. Ludwig, S. Mullainathan, and A. Rambachan. Algorithmic fairness. In *Aea papers and proceedings*, volume 108, pages 22–27, 2018.

[29] W. Kotlowski, K. Dembczynski, and E. Huellermeier. Bipartite ranking through minimization of univariate loss. In *ICML*, 2011.

[30] A. Lambrecht and C. Tucker. Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads. *Management Science*, 65(7):2966–2981, 2019.

[31] N. S. Lesmana, X. Zhang, and X. Bei. Balancing efficiency and fairness in on-demand ridesourcing. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[32] L. Li, R. Munos, and C. Szepesvári. Toward minimax off-policy value estimation. 2015.

[33] W. Liu, J. Guo, N. Sonboli, R. Burke, and S. Zhang. Personalized fairness-aware re-ranking for microlending. In *Proceedings of the 13th ACM Conference on Recommender Systems*, pages 467–471, 2019.

[34] R. Mehrotra, A. Anderson, F. Diaz, A. Sharma, H. Wallach, and E. Yilmaz. Auditing search engines for differential satisfaction across demographics. In *Proceedings of the 26th international conference on World Wide Web companion*, pages 626–633, 2017.

[35] R. Mehrotra, J. McInerney, H. Bouchard, M. Lalmas, and F. Diaz. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In *Proceedings of the 27th acm international conference on information and knowledge management*, pages 2243–2251, 2018.

[36] M. Morik, A. Singh, J. Hong, and T. Joachims. Controlling fairness and bias in dynamic learning-to-rank. *arXiv preprint arXiv:2005.14713*, 2020.

[37] H. Moulin. *Fair division and collective welfare*. MIT press, 2003.

[38] V. Nanda, P. Xu, K. A. Sankararaman, J. Dickerson, and A. Srinivasan. Balancing the tradeoff between profit and fairness in rideshare platforms during high-demand hours. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2210–2217, 2020.

[39] I. Palomares, C. Porcel, L. Pizzato, I. Guy, and E. Herrera-Viedma. Reciprocal recommender systems: Analysis of state-of-art literature, challenges and opportunities towards social recommendation. *Information Fusion*, 69:103–127, 2021.

[40] D. Paraschakis and B. J. Nilsson. Matchmaking under fairness constraints: a speed dating case study. In *International Workshop on Algorithmic Bias in Search and Recommendation*, pages 43–57. Springer, 2020.

[41] G. K. Patro, A. Biswas, N. Ganguly, K. P. Gummadi, and A. Chakraborty. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. *arXiv preprint arXiv:2002.10764*, 2020.

[42] P. Ravikumar, A. Tewari, and E. Yang. On ndcg consistency of listwise ranking methods. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 618–626. JMLR Workshop and Conference Proceedings, 2011.

[43] M. Richardson, R. Agrawal, and P. Domingos. Trust management for the semantic web. In *International semantic Web conference*, pages 351–368. Springer, 2003.

[44] A. Sen. Equality of what? *The Tanner lecture on human values*, 1, 1979.

[45] A. Singh and T. Joachims. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2219–2228. ACM, 2018.

[46] A. Singh and T. Joachims. Policy learning for fairness in ranking. In *Advances in Neural Information Processing Systems*, pages 5427–5437, 2019.

[47] T. Speicher, H. Heidari, N. Grgic-Hlaca, K. P. Gummadi, A. Singla, A. Weller, and M. B. Zafar. A unified approach to quantifying algorithmic unfairness: Measuring individual &group unfairness via inequality indices. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2239–2248. ACM, 2018.

[48] L. Sweeney. Discrimination in online ad delivery. *Queue*, 11(3):10, 2013.

[49] B. Ustun, Y. Liu, and D. Parkes. Fairness without harm: Decoupled classifiers with preference guarantees. In *International Conference on Machine Learning*, pages 6373–6382, 2019.

[50] L. Wang and T. Joachims. Fairness and diversity for rankings in two-sided markets. *arXiv preprint arXiv:2010.01470*, 2020.

[51] Y. Wang, J.-Y. Audibert, and R. Munos. Algorithms for infinitely many-armed bandits. In *Advances in Neural Information Processing Systems*, pages 1729–1736, 2009.

[52] O. Wolfson and J. Lin. Fairness versus optimality in ridesharing. In *2017 18th IEEE International Conference on Mobile Data Management (MDM)*, pages 118–123. IEEE, 2017.

[53] Y. Wu, J. Cao, G. Xu, and Y. Tan. Tfrom: A two-sided fairness-aware recommendation model for both customers and providers. *arXiv preprint arXiv:2104.09024*, 2021.

[54] B. Xia, J. Yin, J. Xu, and Y. Li. We-rec: A fairness-aware reciprocal recommendation based on walrasian equilibrium. *Knowledge-Based Systems*, 182:104857, 2019.

[55] T. Zhang et al. Statistical behavior and consistency of classification methods based on convex risk minimization. *The Annals of Statistics*, 32(1):56–85, 2004.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 7
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] We discussed some possible avenues for negative impact in Section 7
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A] The total compute is on the order of several hundreds of CPU.days
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Code is in supplement
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A] The datasets we use are widely accepted as benchmarks by th research community.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] no personally identifiable information nor offensive content
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]