

---

# Two-sided fairness in rankings via Lorenz dominance

---

Virginie Do<sup>1,2</sup>, Sam Corbett-Davies<sup>1</sup>, Jamal Atif<sup>2</sup>, Nicolas Usunier<sup>1</sup>

<sup>1</sup>Facebook AI

<sup>2</sup>LAMSADE, Université PSL, Université Paris Dauphine, CNRS, France

virginiedo@fb.com, scd@fb.com, jamal.atif@dauphine.psl.eu, usunier@fb.com

## Abstract

We consider the problem of generating rankings that are fair towards both users and item producers in recommender systems. We address both usual recommendation (e.g., of music or movies) and reciprocal recommendation (e.g., dating). Following concepts of distributive justice in welfare economics, our notion of fairness aims at increasing the utility of the worse-off individuals, which we formalize using the criterion of *Lorenz efficiency*. It guarantees that rankings are Pareto efficient, and that they maximally redistribute utility from better-off to worse-off, at a given level of overall utility. We propose to generate rankings by maximizing concave welfare functions, and develop an efficient inference procedure based on the Frank-Wolfe algorithm. We prove that unlike existing approaches based on fairness constraints, our approach always produces fair rankings. Our experiments also show that it increases the utility of the worse-off at lower costs in terms of overall utility.

## 1 Introduction

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, house, or people to connect with. While the objective of recommender systems is usually defined as maximizing the quality of recommendations from the user’s perspective, the recommendations also have an 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. *Two-sided fairness in rankings* is the problem of generating personalized recommendations by fairly mediating between the interests of users and items. It involves a complex multidimensional trade-off. Fairness towards item producers requires boosting the exposure of small producers (e.g., to avoid winner-take-all effects and popularity biases [1]) at the expense of average user utility. Fairness towards users aims at increasing the utility of the least served users (e.g., so that least served users do not support the cost of item-side fairness), once again at the expense of average user utility. The goal of this paper is to provide an algorithmic framework to generate rankings that achieve a variety of these trade-offs, leaving the choice of a specific trade-off to the practitioner.

The leading approach to fairness in rankings is to maximize user utility under constraints of equal item exposure (or equal quality-weighted exposure) [54, 7] or equal user satisfaction [6]. When these constraints imply an unacceptable decrease in average user utility, so-called “trade-offs between utility and fairness” [65, 41] are obtained by relaxing the fairness constraints, leading to the optimization of a trade-off between average user utility and a measure of users’ or items’ inequality.

Thinking about fairness in terms of optimal utility/inequality trade-offs has, however, two fundamental limitations. First, the optimization of a utility/inequality trade-off is not necessarily Pareto-efficient from the point of view of users and items: it sometimes chooses solutions that decrease the utility of some individuals without making anybody else better off. We argue that reducing inequalities by decreasing the utility of the better-off is not desirable if it does not benefit anyone. The second

limitation is that focusing on a single measure of inequality does not address the question of how inequality is reduced, and in particular, which fraction of the population benefits or bears the cost of reducing inequalities.

In this paper, we propose a new framework for two-sided fairness in rankings grounded in the analysis of generalized Lorenz curves of user and item utilities. Widely used to study efficiency and equity in cardinal welfare economics [53], these curves plot the cumulative utility obtained by fractions of the population ordered from the worst-off to the best-off. A curve that is always above another means that all fractions of the populations are better off. We define fair rankings as those with non-dominated generalized Lorenz curves for users and items. First, this definition guarantees that fair rankings are Pareto efficient. Second, examining the entirety of the generalized Lorenz curves provides a better understanding of which fractions of the population benefit from an intervention, and which ones have to pay for it. We present our general framework based on Lorenz dominance in usual recommendation settings (e.g., music or movie recommendation), and also show how extend it to *reciprocal recommendation* tasks such as dating applications or friends recommendation, where users are recommended to other users.

We present a new method for generating rankings based on the maximization of concave welfare functions of users’ and items’ utilities. The parameters of the welfare function control the relative weight of users and items, and how much focus is given to the worse-off fractions of users and items. We show that rankings generated by maximizing our welfare functions are fair for every value of the parameters. Our framework does not aim at defining what parameters are suitable in general — rather, the choice of a specific trade-off depends on the application.

From an algorithmic perspective, two-sided fairness is challenging because items’ utilities depend on the rankings of all users, requiring global inference. Previous work on item-side fairness 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, addressing both our welfare maximization approach and existing item-side fairness penalties.

We demonstrate that our welfare function approach enjoys stronger theoretical guarantees than existing methods. While it always generates rankings with non-dominated generalized Lorenz curves, many other approaches do not. We show that one of the main criteria of the literature, called equity of attention by Biega et al. [7], can lead to decrease user utility, while *increasing* inequalities of exposure between items. Moreover, equal user satisfaction criteria in reciprocal recommendation can lead to decrease the utility of *every user*, even the worse-off. Our notion of fairness prevents these undesirable behaviors. We report experimental results on music and friend recommendation tasks, where we analyze the trade-offs obtained by different methods by looking at different points of their Lorenz curves. Our welfare approach generates a wide variety of trade-offs, and is, in particular, more effective at improving the utility of worse-off users than the baselines.

We present our formal framework in Section 2. We discuss the theoretical properties of previous approaches in Section 3, and present our ranking algorithm in Section 4. Our experiments are described in Section 5, and the related work is discussed in Section 6.

## 2 Two-sided fairness via Lorenz dominance

### 2.1 Formal framework

**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 users 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  $n \in \mathbb{N}$ , we denote by  $\llbracket n \rrbracket = \{1, \dots, n\}$ . The set of users  $\mathcal{N}$  is identified with  $\{1, \dots, |\mathcal{N}|\}$  and the set of items  $\mathcal{I}$  is identified with  $\{|\mathcal{N}| + 1, \dots, n\}$  where  $n = |\mathcal{N}| + |\mathcal{I}|$ . For  $(i, j) \in \mathcal{N} \times \mathcal{I}$ , we denote by  $\mu_{ij}$  the value of item  $j$  to user  $i$ .

A (deterministic) ranking  $\sigma : \mathcal{I} \rightarrow \llbracket |\mathcal{I}| \rrbracket$  is a one-to-one mapping from items  $j$  to their rank  $\sigma(j)$ . Following [54], we use *stochastic rankings* because they allow us to perform inference using convex optimization (see Section 4). 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  $k$ . We denote by  $\mathcal{P}$  the set of ranking tensors.




Figure 1: Generalized Lorenz curves for usual (left) and reciprocal (right) recommendation.

Utilities of users and items are defined through a position-based model, as in previous work [54, 7, 63]. Let  $v \in \mathbb{R}^{|\mathcal{I}|}$ , where  $v_k$  is the exposure weight at rank  $k$ . We assume that lower ranks receive more exposure, so that  $\forall k \in [|\mathcal{I}| - 1], v_k \geq v_{k+1} \geq 0$ .<sup>1</sup> Given a user  $i$  and a ranking  $\sigma_i$ , the *user-side utility* of  $i$  is the sum of the  $\mu_{ij}$ s weighted by the exposure weight of their rank  $\sigma_i(j)$ :  $u_i(\sigma_i) = \sum_{j \in \mathcal{I}} v_{\sigma_i(j)} \mu_{ij}$ . Given an item  $j$ , the *item-side utility* of  $j$  is the sum over users  $i$  of the exposure of  $j$  to  $i$ . These definitions extend to stochastic rankings by taking the expectation over rankings, written in matrix form:<sup>2</sup>

$$\text{user-side utility: } u_i(P) = \sum_{j \in \mathcal{I}} \mu_{ij} P_{ij} v \quad \text{item-side utility (exposure): } u_j(P) = \sum_{i \in \mathcal{N}} P_{ij} v$$

We denote by  $\mathbf{u}(P) = (u_i(P))_{i=1}^n$  the utility profile for  $P$ , and by  $\mathcal{U} = \{\mathbf{u}(P) : P \in \mathcal{P}\}$  the set of feasible profiles. For  $\mathbf{u} \in \mathcal{U}$ ,  $\mathbf{u}_{\mathcal{N}} = (u_i)_{i \in \mathcal{N}}$  and  $\mathbf{u}_{\mathcal{I}} = (u_i)_{i \in \mathcal{I}}$  denote the utility profiles of users and items respectively.

**Two-sided fairness in rankings.** In practice, values of  $\mu_{ij}$  are not known to the recommender system. Ranking algorithms use an estimate  $\hat{\mu}$  of  $\mu$  based on historical data. We address here the problem of *inference*: the task is to compute the ranking tensor given  $\hat{\mu}$ , with the goal of making fair trade-offs between (true) user and item utilities. Notice that the user-side utility depends only on the ranking of the user, but for every item, the exposure depends on the rankings of *all* users. Thus, accounting for both users' and items' utilities in the recommendations is a global inference problem.

**More general item utilities** We consider exposure as the item-side utility to follow prior work and for simplicity. Our framework and algorithm readily applies in a more general case of *two-sided preferences*, where items also have preferences over users (for instance, in hiring, job seekers have preferences over which recruiters they are recommended to). Denoting  $\mu_{ji}$  the value of user  $i$  to item  $j$ , the item side-utility is then  $u_j(P) = \sum_{i \in \mathcal{N}} \mu_{ji} P_{ij} v$ .

## 2.2 Lorenz efficiency and the welfare function approach

Our notion of fairness aims at improving the utility of the worse-off users and items. Since this does not prescribe exactly which fraction of the worse-off users/items should be prioritized, the assessment of trade-offs requires looking at all fractions of the population. This is captured by the generalized Lorenz curve used in cardinal welfare economics [53]. Formally, given a utility profile  $\mathbf{u}$ , let  $(u_{(i)})_{i=1}^n$  be the sorted values in  $\mathbf{u}$  from smallest to largest, i.e.,  $u_{(1)} \leq \dots \leq u_{(n)}$ , then the generalized Lorenz curve plots  $(U_i)_{i=1}^n$  where  $U_i = u_{(1)} + \dots + u_{(i)}$ . To assess the fairness of trade-offs, we rely on the following dominance relations on utility profiles:

**Pareto-dominance**  $\succ_P$ .  $\mathbf{u} \succ_P \mathbf{u}' \iff (\forall i \in [n], u_i \geq u'_i \text{ and } \exists i \in [n], u_i > u'_i)$ .

**Lorenz-dominance**  $\succ_L$ . Then  $\mathbf{u} \succ_L \mathbf{u}' \iff \mathbf{U} \succ_P \mathbf{U}'$ .

We write  $\succeq_L$  for non-strict Lorenz dominance (i.e.,  $\forall i, U_i \geq U'_i$ ). Notice that Pareto-dominance implies Lorenz-dominance. Our notion of fairness, which we call *Lorenz efficiency*, states that a ranking is fair if the utility profiles for users and for items are not jointly Lorenz-dominated:

<sup>1</sup>We use a user-independent  $v$  for simplicity. Considering user-dependent weights is straightforward.

<sup>2</sup>We consider  $P_{ij}$  as a row vector in the formula, so that  $P_{ij}v = \sum_{k=1}^{|\mathcal{I}|} P_{ijk}v_k$ .

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

We consider that Lorenz-dominated profiles are undesirable (and unfair) because the utility of worse-off fractions of the population could have been increased at no cost for total utility. Examples of Lorenz-curves of users and items are given in Fig. 1. The blue solid, green dotted and orange dashed curves are all non-dominated (the blue solid ranking has higher user utility but high item inequality, the green dotted and orange dashed curves have similar item exposure profiles, but user curves that intersect). On the other hand, the red dot/dashed curve is an unfair ranking: compared to the green dotted and orange dashed curve, all fractions of the worse off users have lower utility, together with less exposure for worse-off items.

A fundamental result from cardinal welfare economics is that concave welfare functions of utility profiles order profiles according to Lorenz dominance [3, 53]. The choice of the welfare function specifies which (fair) trade-off is desirable in a specific context. This result holds when all utilities are comparable. In our case where there are users and items, we propose the following welfare function parameterized by  $\theta = (\lambda, \alpha_1, \alpha_2)$ :<sup>3</sup>

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

Inference is carried out by maximizing  $W_{\theta}$  (an efficient algorithm is proposed in Section 4):

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

In  $W_{\theta}$ ,  $\lambda \in [0, 1]$  controls the relative weight of users and items. The motivation for the specific choice of  $\psi$  is that it appears in scale invariant welfare functions [43], but other families can be used as long as the functions are *increasing* and *concave*. Monotonicity implies that maxima of  $W_{\theta}$  are Pareto-efficient. For  $\alpha_1 < 1$  and  $\alpha_2 < 1$ ,  $W_{\theta}$  is strictly concave. Then,  $W_{\theta}$  exhibits *diminishing returns*, which is the key to Lorenz efficiency: an increment in utility for a worse-off user/item increases welfare more than the same increment for a better-off user/item. The effect of the parameters is shown in Fig. 1 (left): For *item fairness* we obtain more item equality by using  $\alpha_1 < 1$  (here,  $\alpha_1 = 0.5$ ) and increasing  $\lambda$  (see blue solid vs orange dashed curve). The parameter  $\alpha_2$  controls *user fairness*: smaller values yield more user utility for the worse-off users at the expense of total utility, with similar item exposure curve (green dotted vs orange dashed curves). Let  $\Theta = \{(\lambda, \alpha_1, \alpha_2) \in (0, 1) \times (-\infty, 1)^2\}$ . For every  $\theta \in \Theta$ ,  $W_{\theta}$  is strictly concave, and users and items have non-zero weight. We then have (the result is a straightforward consequence of diminishing returns, see Appendix C):

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

**Relationship to inequality measures** A well-known measure of inequality is the Gini index, defined as  $1 - 2 \times \text{AULC}$ , where AULC is the area under the Lorenz curve. The difference between Lorenz and generalized Lorenz curves is that the former is normalized by the cumulative utility. This difference is fundamental: we can decrease inequalities while dragging everyone's utility to 0. However, this would lead to dominated *generalized* Lorenz curves. Interestingly, for *item-side* fairness, the cumulative exposure is a constant and thus trade-offs between user utility and item exposure inequality are not really problematic. However, for user-side fairness, the total utility is not constant and reducing inequalities might require dragging the utility of some users down for the benefit of no one.

**Additional theoretical results** In App. C.2, we show that as  $\alpha_1, \alpha_2 \rightarrow -\infty$ , utility profiles tend to leximin-optimal solutions [43]. Leximin optimality corresponds to increasing the utility of the worst-off users/items one a time, similarly to a lexical order. In App. C.3, we present an excess risk bound, which provides theoretical guarantees on the *true* welfare when computing rankings based on *estimated* preferences, depending on the quality of the estimates.

<sup>3</sup> $W_{\theta}(\mathbf{u}) = -\infty$  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.

## 2.3 Extension to reciprocal recommendation

In reciprocal recommendation problems such as dating, the users are also items. The notion of fairness simplifies to increasing the utility of the worse-off users, which can in practice be done by boosting the exposure of worse-off users. Our framework above applies readily by taking  $\mathcal{N} = \mathcal{I}$  and  $n = |\mathcal{N}|$ . The critical step however is to redefine the utility of a user to account for the fact that (1) the user utility comes from both the recommendation they receive and who they are recommended to, and (2) users have preferences over who they are recommended to.

To define this *two-sided utility*, let us denote by  $\mu_{ij}$  the mutual preference value between  $i$  and  $j$ , and our examples follow the common assumption that  $\mu_{ij} = \mu_{ji}$  (see e.g., [45]). For instance, when recommending CVs to recruiters,  $\mu_{ij}$  can be the probability of interview, while in dating, it can be that of a “match”. The two-sided utility is then the sum of the user-side utility and item-sided utility of the user:

$$\begin{array}{ccc} \text{user-side utility} & \text{item-side utility} & \text{(two-sided) utility} \\ \text{(} j \text{ recommended to } i \text{)} & \text{(} i \text{ recommended to } j \text{)} & \\ \overbrace{\bar{u}_i(P) = \sum_{j \in \mathcal{I}} \mu_{ij} P_{ij} v} & \overbrace{\bar{v}_i(P) = \sum_{j \in \mathcal{N}} \mu_{ij} P_{ji} v} & \overbrace{u_i(P) = \bar{u}_i(P) + \bar{v}_i(P)} \end{array}$$

With this definition of two-sided utility, our previous framework can be readily applied using  $\mathcal{N} = \mathcal{I}$ . A (two-sided) utility profile  $\mathbf{u} \in \mathcal{U}$  is *Lorenz-efficient* if there is no  $\mathbf{u}' \in \mathcal{U}$  such that  $\mathbf{u}' \succ_L \mathbf{u}$ . The welfare function simplifies to  $W_\theta(\mathbf{u}) = \sum_{i=1}^n \psi(u_i, \alpha)$ , and Proposition 1 also holds true in this setting: maximizing the welfare function always yields Lorenz-efficient rankings.

Fig. 1 (right) illustrates how decreasing  $\alpha$  increases utilities for the worse-off users at the expense of total utility. It also shows a Lorenz-dominated (unfair) profile, in which all fractions from the worst-off to the better-off users have lower utility.

From now on, we refer to *one-sided* recommendation for non-reciprocal recommendation.

## 3 Comparison to utility/inequality trade-off approaches

As stated in the introduction, leading approaches to fairness in ranking are based on utility/inequality trade-offs. We describe here the representative approaches we consider as baselines in our experiments. We then present theoretical results illustrating the undesirable behavior of some of them.

### 3.1 Objective functions

**One-sided recommendation** In one-sided recommendation, the leading approach is to define exposure-based criteria for item fairness [54, 7]. The first criterion, *equality of exposure*, aims at equalizing exposure across items. The second one, *quality-weighted exposure*<sup>4</sup>, which is advocated by many authors, defines the *quality* of an item as the sum of user values  $q_j = \sum_{i \in \mathcal{N}} \mu_{ij}$  and aims for item exposure proportional to quality. The motivation of quality-weighted exposure is to take user utilities into account in the extreme case where the constraint is strictly enforced. Interestingly, as we show later, this approach has bad properties in terms of trading off user and item utilities.

In our experiments, we use the standard deviation as a measure of inequality. Denoting by  $E = |\mathcal{N}| \|v\|_1$  the total exposure and by  $Q = \sum_{j \in \mathcal{I}} q_j$  the total quality:

$$\text{quality-weighted exposure} \quad F_\beta^{\text{qua}}(\mathbf{u}) = \sum_{i \in \mathcal{N}} u_i - \beta \sqrt{D^{\text{qua}}(\mathbf{u})} \quad \text{with} \quad D^{\text{qua}}(\mathbf{u}) = \frac{1}{n} \sum_{j \in \mathcal{I}} \left( u_j - \frac{q_j E}{Q} \right)^2.$$

$$\text{equality of exposure} \quad F_\beta(\mathbf{u}) = \sum_{i \in \mathcal{N}} u_i - \beta \sqrt{D(\mathbf{u})} \quad \text{with} \quad D(\mathbf{u}) = \sum_{j \in \mathcal{I}} \frac{1}{n} \left( u_j - \frac{1}{|\mathcal{I}|} \sum_{j' \in \mathcal{I}} u_{j'} \right)^2.$$

Some authors use  $D'(\mathbf{u}) = \sum_{(j,j') \in \mathcal{I}^2} \left| \frac{u_j}{q_j} - \frac{u_{j'}}{q_{j'}} \right|$  instead of  $\sqrt{D^{\text{qua}}}$  [55, 42, 6].  $D^{\text{qua}}$  and  $D'$  have qualitatively the same behavior. We propose  $D^{\text{qua}}(\mathbf{u})$  as a computationally efficient alternative to  $D'$ , since it involves only a linear number of terms and  $\sqrt{D^{\text{qua}}}$  is convex and differentiable except on 0.

<sup>4</sup>We use here the terminology of [63]. This criterion has also been called “disparate treatment” [54], “merit-based fairness” [55] and “equity of attention” [7].

**Reciprocal recommendation** For reciprocal recommendation, we consider as competing approach a trade-off between total (two-sided) utility and inequality of utilities, as measured by the standard deviation:

$$\text{equality of utility} \quad F_\beta(\mathbf{u}) = \sum_{i \in \mathcal{N}} u_i - \beta \sqrt{D(\mathbf{u})} \quad \text{with } D(\mathbf{u}) = \sum_{j \in \mathcal{I}} \frac{1}{n} \left( u_j - \frac{1}{|\mathcal{I}|} \sum_{j' \in \mathcal{I}} u_{j'} \right)^2.$$

### 3.2 Inequity and inefficiency of some of the previous approaches

We point out here to two deficiencies of previous approaches.

First, for one-sided recommendation, we show that in some cases, compared to the welfare approach with any choice of the parameter  $\theta \in \Theta$ , quality-weighted exposure leads to the undesirable behavior of *decreasing user utility while increasing inequalities of exposure* between items. This is formalized by the proposition below, which uses the following notation: for  $\theta \in \Theta$ , let  $\mathbf{u}^\theta = \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}} W_\theta(\mathbf{u})$ , and for  $\beta > 0$ , let  $\mathcal{U}_\beta^{\text{qua}} = \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}} F_\beta^{\text{qua}}(\mathbf{u})$ .

**Proposition 2.** *The following claims hold irrespective of the choice of  $\mathbf{u}^{\text{qua}, \beta} \in \mathcal{U}_\beta^{\text{qua}}$ .*

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

$$(\exists \beta > 0, \mathbf{u}_\mathcal{N}^\theta \succ_L \mathbf{u}_\mathcal{N}^{\text{qua}, \beta} \text{ and } \mathbf{u}_\mathcal{I}^\theta \succ_L \mathbf{u}_\mathcal{I}^{\text{qua}, \beta}) \quad \text{and} \quad \lim_{\beta \rightarrow \infty} \frac{\sum_{i \in \mathcal{N}} u_i^{\text{qua}, \beta}}{\sum_{i \in \mathcal{N}} u_i^\theta} \xrightarrow{d \rightarrow \infty} \frac{5}{6}.$$

Second, in reciprocal recommendation, striving for pure equality can even lead to 0 utility for *every user*, even that of the worst-off user. More precisely, we show that in some cases, compared to the welfare approach with any choice of parameter  $\theta \in \Theta$ , there exists  $\beta > 0$  such that equality of utility has lower utility for every user, eventually leading to 0 utility for everyone in the limit  $\beta \rightarrow \infty$ .

**Proposition 3.** *For  $\beta > 0$ , let  $\mathcal{U}_\beta^{\text{eq}} = \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}} F_\beta(\mathbf{u})$ . The claim below holds irrespective of the choice of  $\mathbf{u}^{\text{eq}, \beta} \in \mathcal{U}_\beta^{\text{eq}}$ . Let  $n \geq 5$ . There is a reciprocal recommendation task with  $n$  users such that:*

$$\forall \theta \in \Theta, \mathbf{u}^\theta, \exists \beta > 0 : \forall i \in \llbracket n \rrbracket, u_i^\theta > u_i^{\text{eq}, \beta} \quad \text{and} \quad \lim_{\beta \rightarrow \infty} \sum_{i \in \mathcal{N}} u_i^{\text{eq}, \beta} = 0.$$

**Proofs and additional results** All proofs are deferred to App. D, where we provide several additional results regarding the use of quality-weighted exposure and equality of exposure in reciprocal recommendation: We show in Prop. 8 that there are cases where both approaches lead to user utility profiles with Lorenz-dominated curves, and significantly lower total user utility than the welfare approach for any choice of the parameters.


## 4 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:

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

where for every  $i$ ,  $\Phi_i : \mathbb{R}_+ \rightarrow \mathbb{R}$  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 [54, 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 [42, 7].

Our approach is based on the Frank-Wolfe algorithm [18], which was previously used in machine learning in e.g., structured output prediction or low-rank matrix completion [30], but to the best of



(a) Examples of generalized Lorenz curves achieved by *welf*. (b) Summary of trade-offs (c) Dominated curve

Figure 2: Summary of results on Lastfm-2k, focusing on the user utility/item inequality trade-off.

our knowledge not for ranking. Denoting  $\langle X | Y \rangle = \sum_{ijk} X_{ijk} Y_{ijk}$  the dot product between tensors, the algorithm iterates  $P^{(t)}$  by first computing  $\tilde{P} = \operatorname{argmax}_{P \in \mathcal{P}} \langle P | \nabla W(P^{(t)}) \rangle$  and then updating  $P^{(t)} = (1 - \gamma^{(t)})P^{(t-1)} + \gamma^{(t)}\tilde{P}$  with  $\gamma^{(t)} = \frac{2}{t+2}$  [13]. Starting from an initial solution<sup>5</sup>, the algorithm always stays in the feasible region without any additional projection step. Our main contribution of this section is to show that  $\operatorname{argmax}_{P \in \mathcal{P}} \langle P | \nabla W(P^{(t)}) \rangle$  can be computed efficiently, requiring only one sort operation per user after computing the utilities. In the result 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.<sup>6</sup>

**Theorem 1.** Let  $\tilde{\mu}_{ij} = \Phi'_i(u_i(P^{(t)}))\mu_{ij} + \Phi'_j(u_j(P^{(t)}))\mu_{ji}$ . Let  $\tilde{P}$  such that:

$$\forall i \in \mathcal{N}, \forall \tilde{\sigma}_i \in \mathfrak{S}(\tilde{P}_i): \tilde{\sigma}_i(j) < \tilde{\sigma}_i(j') \implies \tilde{\mu}_{ij} \geq \tilde{\mu}_{ij'}.$$

$P \in \mathcal{P}$

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 [13], assuming each  $\Phi'_i$  is bounded, we have:

**Proposition 4.** Let  $B = \max_{i \in \llbracket n \rrbracket} \|\Phi''_i\|_\infty$  and  $U = \max_{\mathbf{u} \in \mathcal{U}} \|\mathbf{u}\|_2^2$ . Let  $K$  be the maximum index of a nonzero value in  $v$  (or  $|\mathcal{I}|$ ). Then  $\forall t \geq 1, W(P^{(t)}) \geq \max_{P \in \mathcal{P}} W(P) - O(\frac{BU}{t})$ . Moreover, for each user, an iteration costs  $O(|\mathcal{I}| \ln K)$  operations and requires  $O(K)$  additional bytes of storage.

## 5 Experiments

### 5.1 One-sided recommendation

We first present experiments on movie recommendation task. We report here our experiments with the Lastfm-2k dataset [9, 47], which contains the music listening histories of 1.9k users. We present in App. F.2 experiments on a larger portion of the Last.fm dataset, and in App. F.3 results using the MovieLens-20m dataset [24]. Our results are qualitatively similar across the three datasets.

We select the top 2500 items most listened to, and estimate preferences with a matrix factorization algorithm using a random sample of 80% of the data. All experiments are carried out with three repetitions for this subsample. The details of the experimental protocol are in App. F.1. Since the goal is to analyze the behavior of the ranking algorithms rather than the quality of the preference estimates, we consider the estimated preferences as ground truth when computing user utilities and comparing methods, following previous work. We compare our welfare approach (*welf*) to three baselines. The first one is the algorithm of [47] (referred to as *Patro et al.* in the figures), who consider envy-freeness for user-side fairness and, for item-side fairness, a constraint that the minimum exposure of an item is  $\beta \frac{E}{|\mathcal{I}|}$  where  $\beta$  is the trade-off parameter. The other baselines are quality-weighted exposure (*qua.-weighted*) and equality of exposure (*eq. exposure*) as described in Sec. 3.

<sup>5</sup>In our experiments, we initialize with the utilitarian ranking (Proposition 6).

<sup>6</sup>Formally,  $\mathfrak{S}(P_i) = \{\sigma : \mathcal{I} \rightarrow \llbracket |\mathcal{I}| \rrbracket \mid \sigma \text{ is one-to-one, and } \forall j \in \mathcal{I}, P_{i_j \sigma(j)} > 0\}$ .

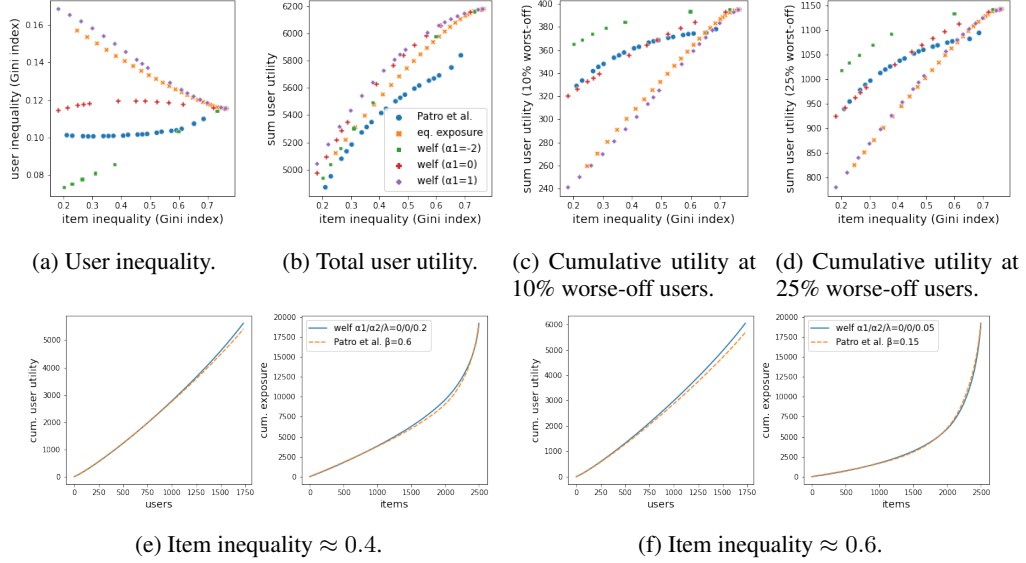


Figure 3: Summary of results on Lastfm-2k for two-sided fairness: effect of varying  $\alpha_1$ .

**Item-side fairness** We first study in isolation item-side fairness, defined as improving the exposure of the worse-off item (producers). To summarize the trade-offs, we show the trade-offs by looking at exposure inequalities as measured by the Gini index (see Sec. 2.2). The results are given in Fig. 2:

- *Generating user utility/item inequality trade-offs* is performed with our approach by keeping  $\alpha_1 = \alpha_2 = 0$  and varying the relative weight of items  $\lambda$ . Fig. 2a plots some trade-offs achieved by our approach. As expected, the user utility curve degrades as we increase the weight of items, while at the same time the curve of item exposure moves towards the straight line, which corresponds to strict equality of exposure. Fig. 6 in the appendix provides analogous curves for all methods, obtained by varying the weight  $\beta$  of the inequality measure.
- *qua.-weighted yields unfair trade-offs* Fig. 2c shows a *welf* ranking that dominates a *qua.-weighted* ranking on both user and item curves. This is in line with the discussion of Section 3, *qua.-weighted* can lead to unfair rankings on utility/item inequality trade-offs.
- *welf dominates the user utility/item inequality (Gini) trade-offs* as seen on Fig. 2b: while all methods have the same total user utility when accepting high item inequality, *welf* dominates *Patro et al.*, *eq. exposure* and *qua.-weighted* as soon as  $\text{Gini} \leq 0.5$ . Note, however, that the Gini index is only one measure of inequality. When measuring item inequalities with the standard deviation, *eq. exposure* becomes optimal since our implementation optimizes a trade-off with this measure (see Fig. 8 in App. F.1). Overall, *welf* and *eq. exposure* yield different fair trade-offs.

**Two-sided fairness** Fig. 3 shows the effect of the user curvature  $\alpha_1 \in \{-2, 0, 1\}$ , keeping  $\alpha_2 = 0$ . Fig. 8 in App. F.1 shows similar plots when the item inequality is measured by the standard deviation rather than the Gini index.

- *Smaller  $\alpha_1$  reduce user inequalities at the expense of total user utility, at various levels of item inequality.* This is observed by comparing the results for  $\alpha_1 \in \{-2, 0, 1\}$  in Fig. 3a and Fig. 3b.
- *welf  $\alpha_1 = 0$  is better than Patro et al.*, which can be seen by jointly looking at Fig. 3c, 3d and Fig. 3b which give the cumulative utility at different points of the Lorenz curve (10%, 25% and 100% of the users respectively). We observe that *welf  $\alpha_1 = 0$*  is similar to *Patro et al.* at the 10% and 25% levels, but has higher total utility. Example curves are given in Fig. 3e and 3f which plot *welf  $\alpha_1 = 0$*  and *Patro et al.* at two levels of item inequality. *welf  $\alpha_1 = 0$*  obtains similar curves to *Patro et al.*, except that it performs better at the end of the curve. A similar comparison can be made with *welf  $\alpha_1 = 1$*  and *eq. exposure*.
- *More user inequalities is not necessarily unfair* as seen in Fig. 3a comparing *welf  $\alpha_1 = 0$*  and *Patro et al.*. We observe that *welf  $\alpha_1 = 0$*  has slightly higher Gini index, but this is not unfair: as seen in Fig. 3e and 3f, this is due to the higher utility at the end of the generalized Lorenz curve of *welf*, but the worse-off users have similar utilities with *welf* and *Patro et al.*.



## 5.2 Reciprocal recommendation

We now present results on a reciprocal recommendation task, where fairness refers to increasing the utility of the worse-off users (this can be done by boosting their exposure at the expense of total utility). Since there is no standard benchmark for reciprocal recommendation, we generate an artificial task based on the Higgs Twitter dataset [15], which contains follower links, and address the task of finding mutual followers (i.e., “matches”). We keep users having at least 20 mutual links, resulting in a subset of 13k users. We build estimated match probabilities using matrix factorization. The experimental protocol is detailed in App. F.4. We also present in App. F.5 additional experiments using the Epinions dataset [49]. The results are qualitatively similar.

Our main baseline is equal utility (*eq. utility*) defined in Section 3. We also compare to quality-weighted exposure, and equality of exposure as baselines that ignore the reciprocal nature of the task. The results are summarized in Fig. 4:

- *Example of trade-offs obtained by varying  $\alpha$*  are plotted in Fig. 4a. As  $\alpha$  decreases, the utility increases for the worse-off users at the expense of better-off users. We note that increasing the utility of worse-off users has a massive cost on total user utility: looking at the exact numbers we observe that  $\alpha = -5$  has more than doubled the cumulative utility of the 10% worse off users compared to  $\alpha = 1$  (120 vs 280), but at the cost of more than 60% of the total utility (17k vs 6.4k). Fig. 6 in Appendix F.4 contains plots of the trade-offs achieved by the other methods.
- *qua.-weighted and eq. exposure are dominated by welf on a large range of hyperparameters.* An example is given in Fig. 4b, where *welf*  $\alpha = 0.5$  already dominates some of their models, even though in this region of  $\alpha$  there is little focus on worse-off users. More generally, all values of  $\beta \geq 0.1$  for *qua.-weighted* and *eq. exposure* lead to rankings with dominated curves. This is expected since they ignore the reciprocal nature of the task.
- *eq. utility is dominated by welf near strict equality* as illustrated in Fig. 4c: for large values of  $\beta$ , it is not possible to increase the utility of the worse off users, and *eq. utility* only drags utility of better-off users down.
- *welf is more effective at increasing utility of the worse-off users* as can be seen in Fig. 4e-g, which plots the total utility as a function of the cumulative utility at different points of the Lorenz curve (10%, 20%, 50% worse-off users respectively). For total utilities larger than 50% of the maximum achievable, *welf* significantly dominates *eq. utility* in terms of utility of worse-off users (10% and 25%) at a given level of total utility. *welf* also dominates *eq. utility* on the 50% worse-off users (Fig. 4h) in the interesting region where the total utility is within 20% of the maximum.
- *More inequality is not necessarily unfair* As shown in Fig. 4d, we see that for the same utility for the 10% worse-off users, *welf* models have higher inequalities than *eq. utility*. As seen before, this higher inequality is due to a higher total utility (and higher total utilities for the 25% worse-off users). The analysis of these Lorenz curves allow us to conclude that these larger inequalities are not due to unfairness. They arise because *welf* optimizes the utility of the worse-off users at lower cost in terms of average utility than *eq. utility*.

## 6 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 [57, 33, 21, 40, 35]. Our work follows the subsequent work on ranking algorithms that promote fairness of exposure for individual or sensitive groups of items [10, 8, 7, 54, 42, 65]. The goal is often to prevent winner-take-all effects, combat popularity bias [1] or promote smaller producers [39, 41]. Section 3 is devoted to the comparison with this type of approaches. Most of these works use a notion of fairness oriented towards items only. Towards two-sided fairness, Wang and Joachims [60] promote user-side fairness using concave functions of user utilities, similarly to us. Other works use equality constraints to define user-side fairness [6, 63]. These three approaches rely on the definitions of item-side fairness discussed in Section 3. Patro et al. [47] 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 one-sided and reciprocal recommendation within the same framework. There is less existing work studying the fairness of rankings in the reciprocal setting. Xia et al. [64] aim at equalizing user utility between groups, which suffers from the problems discussed in

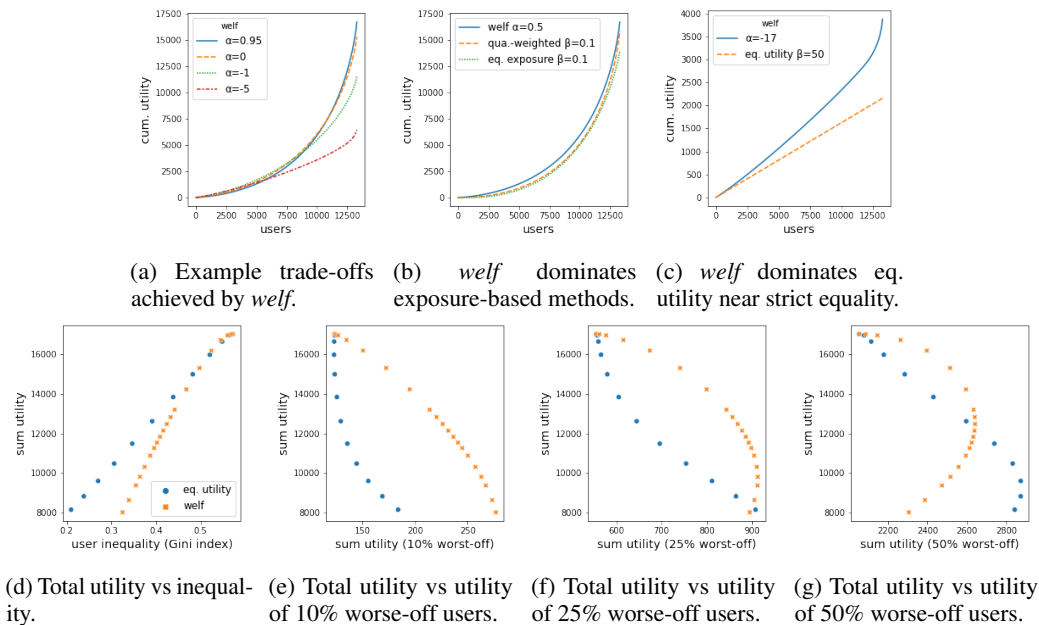


Figure 4: Results on the twitter dataset.

Section 3. Jia et al. [31] generate rankings using a welfare function approach, but optimizing only the utility of users *being recommended*. Paraschakis and Nilsson [46] postprocess rankings to correct for inconsistencies between estimated and declared preferences of users. In contrast, 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. Fairness is also studied in the context of ridesharing applications [62, 37, 44], but they address matching rather than ranking problems.

There is growing interest in making the relationship between fairness in machine learning and social choice theory [25, 59, 4, 20, 27, 12, 16, 17], and welfare economics in particular [56, 28, 34, 36, 67]. In line with Hu and Chen [28], who focused on classification and parity penalties, we argue that Pareto-efficiency should be part of fairness assessments. We are the first to propose concave welfare functions and Lorenz dominance to address two-sided fairness in recommendation.

## 7 Conclusion

We view fairness in rankings as optimizing the distribution of user and item utilities, giving priority to the worse-off. Following this view, we defined fair rankings as having non-dominated generalized Lorenz curves of user and item utilities, and develop a new conceptual and algorithmic framework for fair ranking. The generality of the approach is showcased on several recommendation tasks, including reciprocal recommendation.

The expected positive societal impact of this work is to provide more principled approaches to mediating between several parties on a recommendation platform. Yet, we did not address several questions that are critical for the deployment of our approach. In particular, true user preferences are often not directly available, and we only observe proxies to them, such as clicks or likes. Second, interpersonal comparisons of utilities are critical in this work. It is thus necessary to make sure that the proxies we choose lead to meaningful comparisons of utilities between users. Third, estimating preferences or their proxies is itself not trivial in recommendation because of partial observability. The true fairness of our approach is bound to a careful analysis of (at least) these additional steps.

## Acknowledgments and Disclosure of Funding

We thank David Lopez-Paz, Jérôme Lang, as well as the anonymous NeurIPS reviewers for their constructive feedback on early versions of the paper.

## 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] O. Celma. *Music Recommendation and Discovery in the Long Tail*. Springer, 2010.
- [12] A. Chakraborty, G. K. Patro, N. Ganguly, K. P. Gummadi, and P. Loiseau. Equality of voice: Towards fair representation in crowdsourced top-k recommendations. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 129–138. ACM, 2019.
- [13] K. L. Clarkson. Coresets, sparse greedy approximation, and the frank-wolfe algorithm. *ACM Transactions on Algorithms (TALG)*, 6(4):1–30, 2010.
- [14] D. Cossock and T. Zhang. Statistical analysis of bayes optimal subset ranking. *IEEE Transactions on Information Theory*, 54(11):5140–5154, 2008.
- [15] M. De Domenico, A. Lima, P. Mougél, and M. Musolesi. The anatomy of a scientific rumor. *Scientific reports*, 3(1):1–9, 2013.
- [16] V. Do, S. Corbett-Davies, J. Atif, and N. Usunier. Online certification of preference-based fairness for personalized recommender systems. *arXiv preprint arXiv:2104.14527*, 2021.
- [17] 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. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 489–503, 2021.
- [18] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- [19] C. Gini. Measurement of inequality of incomes. *The economic journal*, 31(121):124–126, 1921.
- [20] P. Gözl, A. Kahng, and A. D. Procaccia. Paradoxes in fair machine learning. *Advances in Neural Information Processing Systems (NeurIPS)*. Forthcoming, 2019.
- [21] 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.
- [22] M. Hardt, E. Price, N. Srebro, et al. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, pages 3315–3323, 2016.
- [23] G. H. Hardy, J. E. Littlewood, and G. Pólya. Inequalities. 2nd ed. Cambridge, Engl.: At the University Press. XII, 324 p. (1952)., 1952.

- [24] 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.
- [25] 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.
- [26] H. Heidari, M. Loi, K. P. Gummadi, and A. Krause. A moral framework for understanding fair ml through economic models of equality of opportunity. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 181–190, 2019.
- [27] 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.
- [28] 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.
- [29] 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.
- [30] M. Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *International Conference on Machine Learning*, pages 427–435. PMLR, 2013.
- [31] 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.
- [32] C. C. Johnson. Logistic matrix factorization for implicit feedback data. *Advances in Neural Information Processing Systems*, 27(78):1–9, 2014.
- [33] 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.
- [34] J. Kleinberg, J. Ludwig, S. Mullainathan, and A. Rambachan. Algorithmic fairness. In *Aea papers and proceedings*, volume 108, pages 22–27, 2018.
- [35] 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.
- [36] M. S. A. Lee, L. Floridi, and J. Singh. Formalising trade-offs beyond algorithmic fairness: lessons from ethical philosophy and welfare economics. *AI and Ethics*, pages 1–16, 2021.
- [37] 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.
- [38] D. Lim, J. McAuley, and G. Lanckriet. Top-n recommendation with missing implicit feedback. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 309–312, 2015.
- [39] 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.
- [40] 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.
- [41] 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.
- [42] M. Morik, A. Singh, J. Hong, and T. Joachims. Controlling fairness and bias in dynamic learning-to-rank. *arXiv preprint arXiv:2005.14713*, 2020.
- [43] H. Moulin. *Fair division and collective welfare*. MIT press, 2003.
- [44] 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.

- [45] 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.
- [46] 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.
- [47] 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.
- [48] 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.
- [49] M. Richardson, R. Agrawal, and P. Domingos. Trust management for the semantic web. In *International semantic Web conference*, pages 351–368. Springer, 2003.
- [50] J. E. Roemer and A. Trannoy. Equality of opportunity: Theory and measurement. *Journal of Economic Literature*, 54(4):1288–1332, 2016.
- [51] A. Sen. Equality of what? *The Tanner lecture on human values*, 1, 1979.
- [52] A. F. Shorrocks. The class of additively decomposable inequality measures. *Econometrica: Journal of the Econometric Society*, pages 613–625, 1980.
- [53] A. F. Shorrocks. Ranking income distributions. *Economica*, 50(197):3–17, 1983.
- [54] 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.
- [55] A. Singh and T. Joachims. Policy learning for fairness in ranking. In *Advances in Neural Information Processing Systems*, pages 5427–5437, 2019.
- [56] 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.
- [57] L. Sweeney. Discrimination in online ad delivery. *Queue*, 11(3):10, 2013.
- [58] P. D. Thistle. Ranking distributions with generalized lorenz curves. *Southern Economic Journal*, pages 1–12, 1989.
- [59] 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.
- [60] L. Wang and T. Joachims. Fairness and diversity for rankings in two-sided markets. *arXiv preprint arXiv:2010.01470*, 2020.
- [61] M. Wang, M. Gong, X. Zheng, and K. Zhang. Modeling dynamic missingness of implicit feedback for recommendation. *Advances in neural information processing systems*, 31:6669, 2018.
- [62] 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.
- [63] 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.
- [64] 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.
- [65] M. Zehlike and C. Castillo. Reducing disparate exposure in ranking: A learning to rank approach. In *Proceedings of The Web Conference 2020*, pages 2849–2855, 2020.
- [66] 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.
- [67] M. Zimmer, C. Glanois, U. Siddique, and P. Weng. Learning fair policies in decentralized cooperative multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 12967–12978. PMLR, 2021.

## A Outline of the appendices

These appendices are structured as follows:

- In Appendix B, we present how our fairness framework can be applied to sensitive groups of users or categories of items.
- In Appendix C, we present a deeper analysis of the trade-offs achieved by the welfare approach. We also provide a theoretical guarantee relating the true welfare obtained by maximizing the welfare using estimated preferences, depending on the quality of the estimates.
- In Appendix D, we present the proofs for the theoretical results comparing our results and previous criteria of fairness in rankings. In addition, in Appendix D.3, we describe how to extend the criteria of equality of exposure and quality-weighted exposure in a reciprocal recommendation setting. This is the extension used in our experiments on reciprocal recommendation. In Proposition 8, we present an additional result regarding the inefficiency of these criteria in reciprocal recommendation.
- In Appendix E, we present the more general version of the Frank-Wolfe algorithm, which we use both to optimize the welfare function over stochastic rankings, as well as the penalty-based baselines. This appendix also contains the proofs of the results in Section 4. In addition, this appendix contains fundamental lemmas that are used in other appendices.
- Appendix F gives the details of the experiments presented in Section 5, as well as many additional experiments (two additional, larger scale datasets on one-sided recommendation, and an additional dataset for reciprocal recommendation)
- Appendix G briefly discusses the difference between the penalty we use in our implementation of the baseline approaches and an alternative penalty used by some authors.
- Finally, Appendix H discusses the difference between applying item-side fairness criteria for every ranking, compared to what we do in the paper, which defines item-side utility as an aggregate over the rankings of all users.

## B Fairness towards sensitive groups rather than individuals

In all the paper we focus on fairness towards individual users and items rather than groups of users or items. Prior work [54, 42, 55] considered the utility of a group as the sum or the average utility of its members. Using this definition of group utility, our framework directly extends to groups rather than individuals. In this section we describe the case of one-sided recommendation with groups of users and item categories. The case of reciprocal recommendation (with user groups only) is similar but simpler.

Let  $\mathcal{S} = (s_p)_{p=1}^{|\mathcal{S}|}$  be (possibly overlapping) user groups, i.e.,  $\forall p \in [|\mathcal{S}|], s_p \subseteq \mathcal{N}$  and  $\cup_{p \in [|\mathcal{S}|]} s_p = \mathcal{N}$ . Similarly, let  $\mathcal{C} = (c_q)_{q=1}^{|\mathcal{C}|}$  be (possibly overlapping) item categories, i.e.,  $\forall q \in [|\mathcal{C}|], c_q \subseteq \mathcal{I}$  and  $\cup_{q \in [|\mathcal{C}|]} c_q = \mathcal{I}$ . On the user side, such groups would typically correspond to demographic groups considered sensitive for the application at hand [57]. On the item side, groups can represent a single producer for the case where we want to be fair to producers based on the aggregate utility they obtain from their products [41], or demographic groups as well [33].

In all cases, we redefine the user-side utility for groups and the item-side utility for categories:

$$u_{s_p}^{\text{gr}}(P) = \sum_{i \in s_p} u_i(P) \qquad u_{c_q}^{\text{cat}}(P) = \sum_{j \in c_q} u_j(P)$$

Let  $\mathbf{u}^{\text{gr}}(P) = (u_{s_p}^{\text{gr}}(P))_{p=1}^{|\mathcal{S}|}$  and  $\mathbf{u}^{\text{cat}}(P) = (u_{c_q}^{\text{cat}}(P))_{q=1}^{|\mathcal{C}|}$  be the utility profiles of user groups and item categories associated to  $P$  respectively. The two-sided Lorenz efficiency for groups and categories is defined as:

**Definition 2.** *Let  $\mathcal{S}$  be a set of user groups and  $\mathcal{C}$  a set of item categories. Let  $P \in \mathcal{P}$ .  $P$  is  $(\mathcal{S}, \mathcal{C})$ -Lorenz efficient if there is no  $P' \in \mathcal{P}$  such that either condition holds:*

1.  $\mathbf{u}^{\text{gr}}(P') \succeq_{\text{L}} \mathbf{u}^{\text{gr}}(P)$  and  $\mathbf{u}^{\text{cat}}(P') \succ_{\text{L}} \mathbf{u}^{\text{cat}}(P)$ , or

2.  $\mathbf{u}^{\text{cat}}(P') \succeq_{\mathcal{L}} \mathbf{u}^{\text{cat}}(P)$  and  $\mathbf{u}^{\text{gr}}(P') \succ_{\mathcal{L}} \mathbf{u}^{\text{gr}}(P)$ .

The welfare function associated to  $(\mathcal{S}, \mathcal{C})$ , still parametrized by  $\theta = (\lambda, \alpha_1, \alpha_2) \in \Theta$ , is defined as

$$W_{\theta}^{\text{gr}}(P) = (1 - \lambda) \sum_{s \in \mathcal{S}} \psi(u_s^{\text{gr}}(P), \alpha_1) + \lambda \sum_{c \in \mathcal{C}} \psi(u_c^{\text{cat}}(P), \alpha_2)$$

The welfare function follows the general form of objective function used for the algorithm in Appendix E, so the optimization of  $W_{\theta}^{\text{gr}}$  requires similar computational complexity as  $W_{\theta}$ .

Finally, the extension of Proposition 1 is straightforward. Its proof is similar to the proof presented in Appendix C.

**Proposition 5.**  $\forall \theta \in \Theta, \forall P^* \in \operatorname{argmax}_{P \in \mathcal{P}} W_{\theta}^{\text{gr}}(P)$ ,  $P^*$  is  $(\mathcal{S}, \mathcal{C})$ -Lorenz efficient.

Note that this way of treating groups is not necessarily optimal. In particular, it does not account for within-group fairness. The separate consideration of within-group and between-group fairness has been studied extensively in the literature on equality of opportunity [50], which has inspired several works on algorithmic fairness [22, 26]. Yet, how to apply these principles to two-sided fairness in recommendation is still open, and is left as future work.

## C More on welfare functions

This appendix provides an in-depth analysis of the trade-offs that are achievable by the welfare approach. We first prove the proposition of Section 2.2, and analyze the utilitarian rankings (obtained with  $\alpha_1 = \alpha_2 = 1$ ). We then analyze how to obtain leximin optimal solutions on the side of the items in Appendix C.2, as mentioned in Section 2.2. Finally, we prove Theorem 2 in Appendix C.3, which provides a regret bound relating the true welfare achieved when maximizing welfare on estimated preferences. Some results in this section use Lemma 3 of Appendix 4, which is proved in Appendix 4.

Throughout the appendices, we use the more general version of item utilities (two-sided preferences), described at the end of Section 2.1. Moreover, to clarify the notation, we remind that a *ranking tensor* is a three-way tensor  $P$  where  $P_{ijk}$  is the probability that item  $j$  is recommended to user  $i$  at rank  $k$ . We consider  $P$  as an  $n \times n \times |\mathcal{I}|$  tensor, where irrelevant entries are set to 0. With this notation, the utility for both users and items can be written with the same formula:

$$\forall i \in \llbracket n \rrbracket, u_i(P) = \sum_{j=1}^n \mu_{ij}(P_{ij} + P_{ji})v.$$

Note that this formula also corresponds to the two-sided utility in reciprocal recommendation. In general, the results in this appendix can be extended to reciprocal recommendation with minimal changes to their proofs, using  $\mathcal{N} = \mathcal{I} = \llbracket n \rrbracket$  and the formula above for the utility.

### C.1 Lorenz efficiency and utilitarian ranking

We first prove Proposition 1:

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

*Proof.* It is well known that if  $\Phi$  is increasing and strictly concave, then  $F(\mathbf{u}) = \sum_{i=1}^n \Phi(u_i)$  is monotonic with respect to Lorenz dominance [53, 58]:  $\mathbf{u} \succ_{\mathcal{L}} \mathbf{u}' \implies F(\mathbf{u}) > F(\mathbf{u}')$ .

In the case of  $W_{\theta}$ , for every  $\theta = (\lambda, \alpha_1, \alpha_2) \in \Theta$ , both  $\psi(\cdot, \alpha_1)$  and  $\psi(\cdot, \alpha_2)$  are strictly concave by the definition of  $\Theta$  (recall that in  $\Theta$ , we have  $\alpha_1, \alpha_2 < 1$ ).

The partial function<sup>7</sup>  $\mathbf{u}'_{\mathcal{N}} \mapsto W_{\theta}((\mathbf{u}_{\mathcal{I}}, \mathbf{u}'_{\mathcal{N}}))$  is, up to a constant, of the form of  $F$  and likewise for the partial function  $\mathbf{u}'_{\mathcal{I}} \mapsto W_{\theta}((\mathbf{u}'_{\mathcal{I}}, \mathbf{u}_{\mathcal{N}}))$ . We now prove the result by contradiction. Assume that  $\mathbf{u} \in \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}} W_{\theta}(\mathbf{u})$  is not Lorenz efficient. Then there is  $\mathbf{u}' \in \mathcal{U}$  such that  $(\mathbf{u}'_{\mathcal{N}} \succeq_{\mathcal{L}} \mathbf{u}_{\mathcal{N}}$  and

<sup>7</sup>We denote by  $(\mathbf{u}_{\mathcal{I}}, \mathbf{u}'_{\mathcal{N}})$  the vector  $\mathbb{R}^d$  such that  $(\mathbf{u}_{\mathcal{I}}, \mathbf{u}'_{\mathcal{N}})_i = u_i$  if  $i \in \mathcal{I}$  and  $(\mathbf{u}_{\mathcal{I}}, \mathbf{u}'_{\mathcal{N}})_i = u'_i$  if  $i \in \mathcal{N}$ .

$\mathbf{u}'_{\mathcal{I}} \succ_L \mathbf{u}_{\mathcal{I}}$ ) or ( $\mathbf{u}'_{\mathcal{N}} \succ_L \mathbf{u}_{\mathcal{N}}$  and  $\mathbf{u}'_{\mathcal{I}} \succeq_L \mathbf{u}_{\mathcal{I}}$ ). Let us assume ( $\mathbf{u}'_{\mathcal{N}} \succeq_L \mathbf{u}_{\mathcal{N}}$  and  $\mathbf{u}'_{\mathcal{I}} \succ_L \mathbf{u}_{\mathcal{I}}$ ), the other case is dealt with similarly. We then have:

$$\begin{aligned} W_{\theta}(\mathbf{u}') &\geq W_{\theta}((\mathbf{u}'_{\mathcal{I}}, \mathbf{u}_{\mathcal{N}})) && \text{(because } \mathbf{u}'_{\mathcal{N}} \succeq_L \mathbf{u}_{\mathcal{N}}) \\ &> W_{\theta}((\mathbf{u}_{\mathcal{I}}, \mathbf{u}_{\mathcal{N}})) && \text{(because } \mathbf{u}'_{\mathcal{I}} \succ_L \mathbf{u}_{\mathcal{I}}) \end{aligned}$$

which contradicts the maximality of  $\mathbf{u}$ .  $\square$

The analogous for Proposition 1 for reciprocal recommendation is a direct consequence of standard results that concave welfare functions are monotonic with respect to Lorenz dominance [53, 58].

**Utilitarian ranking** Proposition 6 below generalizes to two-sided utilities the well-known result that maximizing user-side utility is achieved by sorting  $j \in \mathcal{I}$  by decreasing  $\mu_{ij}$  (see e.g., [14]). For a ranking tensor  $P$  and a user  $i$ , we denote by  $\mathfrak{S}(P_i)$  the support of  $P_i$  in ranking space.<sup>8</sup> We remind that  $\sigma(j)$  is the rank of item  $j$ , and that lower ranks are better. For a user  $i$  and item  $j$ , we use  $\mu_{ji} = 1$ .

**Proposition 6** (Utilitarian ranking). *Assume  $\forall k \in \llbracket n-1 \rrbracket, v_k > v_{k+1} \geq 0$  and let*

$$P^* \in \operatorname{argmax}_{P \in \mathcal{P}} W_{\frac{1}{2}, 1, 1}(P) = \operatorname{argmax}_{P \in \mathcal{P}} \sum_{i \in \llbracket n \rrbracket} u_i(P).$$

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  $\forall (i, j) \in \llbracket n \rrbracket^2, \mu_{ij} = \mu_{ji}$ , then  $\tilde{\mu}_{ij} \geq \tilde{\mu}_{ij'} \iff \mu_{ij} \geq \mu_{ij'}$ .

When mutual preferences are symmetric (i.e.,  $\mu_{ij} = \mu_{ji}$ ), the utilitarian ranking is the same as the usual sort by decreasing  $\mu_{ij}$ . This also obviously holds when we consider exposure as item utility ( $\mu_{ji} = 1$ ). 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 [45].

For the proof of Proposition 6, the main part is the following lemma:

**Lemma 1.** *Let  $F(\mathbf{u}(P)) = \sum_{i=1}^n u_i(P)$  and  $\tilde{\mu}_{ij} = \mu_{ij} + \mu_{ji}$ . Assume  $\forall k \in \llbracket n-1 \rrbracket, v_k \geq v_{k+1} \geq 0$ .*

*If  $P^* \in \mathcal{P}$  is such that  $\forall \sigma \in \mathfrak{S}(P_i^*), \forall j, j', \sigma(j) < \sigma(j') \implies \tilde{\mu}_{ij} \geq \tilde{\mu}_{ij'}$  then  $P^* \in \operatorname{argmax}_{P \in \mathcal{P}} \mathbf{u}(P)$ .*

*Moreover, if  $\forall k \in \llbracket n-1 \rrbracket, v_k > v_{k+1} \geq 0$ , then the reciprocal is true.*

*Proof.* Notice that, thanks to the completion of  $P$  with zeros on irrelevant entries and formula C,  $F(\mathbf{u}(P))$  can be rewritten as:

$$F(\mathbf{u}(P)) = \sum_{i=1}^n u_i(P) = \sum_{i=1}^n \sum_{j=1}^n \mu_{ij}(P_{ij} + P_{ji})v = \sum_{i=1}^n \sum_{j=1}^n (\mu_{ij} + \mu_{ji})P_{ij}v$$

where the last equality is obtained by swapping  $i$  and  $j$  in the second sum, which is possible since  $i$  and  $j$  span the same range.

The result is then a direct consequence of Lemma 3 in Appendix E, using  $A_{ij} = \mu_{ij} + \mu_{ji}$ .  $\square$

The first of statement of Proposition 6 assumes that the exposure weights  $v$  are non-negative and strictly decreasing as per the second point of Lemma 1. Lemma 1 above gives the statement for the more general case of non-increasing  $v$ .

*Proof of Proposition 6.* The first statement is the consequence of Lemma 1 above, noticing that  $F(\mathbf{u}(P))$  in Lemma 1 always has the same argmax. The second statement is obvious from the assumptions.  $\square$

<sup>8</sup>Formally,  $\mathfrak{S}(P_i) = \{\sigma : \mathcal{I} \rightarrow \llbracket |\mathcal{I}| \rrbracket \mid \sigma \text{ is one-to-one, and } \forall j \in \mathcal{I}, P_{i_j \sigma(j)} > 0\}$ .



## C.2 Item-side leximin optimality

The most egalitarian trade-off achievable by our method is described by the leximin order [51]. Given two utility profiles  $\mathbf{u}$  and  $\mathbf{u}'$ ,  $\mathbf{u} \geq_{\text{lex}} \mathbf{u}'$  if  $\mathbf{U}$  is greater than  $\mathbf{U}'$  according to the lexicographic order.<sup>9</sup> 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.

The proposition below formalizes how leximin optimal solutions on the side of items are found. It shows that item-side leximin solutions are obtained by having  $\alpha_2 \rightarrow -\infty$  and  $\lambda \rightarrow 1$  at the same time. The proposition gives a formal statement of the rate at which  $\lambda$  should converge to 1 relative to  $\alpha$ .

In the statement of the proposition, given two functions  $F$  and  $G$ , we use  $F(\alpha) \underset{\alpha \rightarrow -\infty}{\geq} G(\alpha)$  as a shorthand for  $F(\alpha) \geq G(\alpha)$  for  $\alpha$  sufficiently small.<sup>10</sup>

**Proposition 7.** Let  $\mathcal{U}_{\text{lex}}^{\text{item}} = \{\mathbf{u} \in \mathcal{U} : \forall \mathbf{u}' \in \mathcal{U}, \mathbf{u}_{\mathcal{I}} \geq_{\text{lex}} \mathbf{u}'_{\mathcal{I}}\}$  and let  $\mathbf{u}^* = \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}_{\text{lex}}^{\text{item}}} \sum_{i \in \mathcal{N}} \psi(u_i, \alpha_1)$ .

$$\forall \eta > \max(1, \|\mathbf{u}^*_{\mathcal{I}}\|_{\infty}), \forall \mathbf{u} \in \mathcal{U} : W_{1-\eta^\alpha, \alpha_1, \alpha}(\mathbf{u}^*) \underset{\alpha \rightarrow -\infty}{\geq} W_{1-\eta^\alpha, \alpha_1, \alpha}(\mathbf{u}).$$

This means that among the leximin-optimal item-side utility profiles,  $\alpha_1$  still controls the redistribution profile on the user side, since it is possible that  $|\mathcal{U}_{\text{lex}}^{\text{item}}| > 1$  in one-sided recommendation. A similar result holds for user-side item leximin.

*Proof.* Let  $\mathbf{u}^* = \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}_{\text{lex}}^{\text{item}}} \sum_{i \in \mathcal{N}} \psi(u_i, \alpha_1)$  and  $\mathbf{u} \in \mathcal{U}$ . Let  $\theta = (\lambda, \alpha_1, \alpha)$  and take  $\alpha < \min(0, \alpha_1)$ .

Let  $(j_1, j_2, \dots, j_{|\mathcal{I}|})$  be the ranking of  $\mathbf{u}^*_{\mathcal{I}}$  in increasing order:  $u^*_{j_1} \leq \dots \leq u^*_{j_{|\mathcal{I}|}}$ . Likewise, let  $(j'_1, j'_2, \dots, j'_{|\mathcal{I}|})$  be the ranking of  $\mathbf{u}_{\mathcal{I}}$  in increasing order:  $u_{j'_1} \leq \dots \leq u_{j'_{|\mathcal{I}|}}$ .

Let  $m = \max\{k \in [|\mathcal{I}|] \cup \{0\} : \forall \ell \leq k, u^*_{j_\ell} = u_{j'_\ell}\} + 1$ , be the last index (+1) such that the smallest values of  $\mathbf{u}^*$  and  $\mathbf{u}$  are equal ( $m = 1$  if the smallest values are different).

Let  $C(\alpha) = W_{1-\eta^\alpha, \alpha_1, \alpha}(\mathbf{u}^*) - W_{1-\eta^\alpha, \alpha_1, \alpha}(\mathbf{u})$ .

Let  $K = \sum_{i \in \mathcal{N}} (\psi(u_i^*, \alpha_1) - \psi(u_i, \alpha_1))$ .

**case 1:**  $m = |\mathcal{I}| + 1$ . Then  $C(\alpha) = (1 - \eta^\alpha)K \geq 0$  since  $\mathbf{u}^*_{\mathcal{I}} = \mathbf{u}_{\mathcal{I}}$  and  $\mathbf{u}^*$  maximizes the user-side welfare.

**case 2:**  $m < |\mathcal{I}|$ . Then, we have  $u_{j'_m} < u^*_{j_m}$  by the leximin optimality of  $\mathbf{u}^*_{\mathcal{I}}$ . We then have:

$$\begin{aligned} C(\alpha) &= (1 - \eta^\alpha)K + \eta^\alpha \sum_{j \in \mathcal{I}} -(u_j^*)^\alpha + (u_j)^\alpha \\ &= -(1 - \eta^\alpha)(u_{j_m}^*)^\alpha \underbrace{\left( \frac{K}{1 - \eta^\alpha} \left( \frac{\eta}{u_{j_m}^*} \right)^\alpha + 1 \right)}_{\xrightarrow{\alpha \rightarrow -\infty} 0} + \sum_{k > m} \underbrace{\left( \frac{u_{j_k}^*}{u_{j_m}^*} \right)^\alpha}_{\xrightarrow{\alpha \rightarrow -\infty} 0} - \underbrace{\left( \frac{u_{j'_m}}{u_{j_m}^*} \right)^\alpha}_{\xrightarrow{\alpha \rightarrow -\infty} +\infty} - \sum_{k > m} \underbrace{\left( \frac{u_{j'_k}}{u_{j_m}^*} \right)^\alpha}_{\geq 0} \end{aligned}$$

which implies  $\lim_{\alpha \rightarrow -\infty} C(\alpha) = +\infty$  and thus the desired result.  $\square$

## C.3 Guarantees when performing inference with estimated preferences

In practice, inference is carried out on an estimate  $\hat{\mu}$  of  $\mu$ , meaning that, denoting  $\hat{\mathbf{u}}$  the resulting estimated utility<sup>11</sup> the system output  $\hat{P} = \operatorname{argmax}_{P \in \mathcal{P}} W_\theta(\hat{\mathbf{u}}(P))$ . The following result extends

<sup>9</sup>Formally,  $\mathbf{u} >_{\text{lex}} \mathbf{u}'$  if  $(\exists k \in [d] \text{ s.t. } \forall i < k, U_i = U'_i \text{ and } U_k > U'_k)$ .  $\mathbf{u} \geq_{\text{lex}} \mathbf{u}' \iff \neg(\mathbf{u}' >_{\text{lex}} \mathbf{u})$ .

<sup>10</sup>Formally,  $F(\alpha) \underset{\alpha \rightarrow -\infty}{\geq} G(\alpha) \iff \exists \alpha_0 \in \mathbb{R}, \forall \alpha \leq \alpha_0, F(\alpha) \geq G(\alpha)$ .

<sup>11</sup>We have  $\hat{\mathbf{u}}_i(P) = \sum_{j \in \mathcal{I}} \hat{\mu}_{ij} P_{ij} v$  for  $i \in \mathcal{N}$ .

surrogate regret bounds that exist in classification [5, 66] and learning to rank [14, 48, 2] 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  $\mathbf{u}(\hat{P})$  (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.<sup>12</sup>

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

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

$$W_\theta(\mathbf{u}(P^*)) - W_\theta(\mathbf{u}(\hat{P})) \leq 4B(\hat{\mu}) \sqrt{n \|v\|_2^2} \sqrt{\sum_{(i,j) \in \mathcal{N} \times \mathcal{I}} (\hat{\mu}_{ij} - \mu_{ij})^2}.$$

The existing results closest to our Theorem 2 are Theorem 2 of [14]. Here the result is substantially more difficult to prove because of the concave function and the fact that utilities are two-sided, calling for considering the rankings of multiple users at once.

*Proof.* We have:

$$\begin{aligned} W_\theta(\mathbf{u}(P^*)) - W_\theta(\mathbf{u}(\hat{P})) &= W_\theta(\mathbf{u}(P^*)) - \underbrace{W_\theta(\hat{\mathbf{u}}(\hat{P}))}_{\geq W_\theta(\hat{\mathbf{u}}(P^*))} + W_\theta(\hat{\mathbf{u}}(\hat{P})) - W_\theta(\mathbf{u}(\hat{P})) \\ &\leq \underbrace{W_\theta(\mathbf{u}(P^*)) - W_\theta(\hat{\mathbf{u}}(P^*))}_{=C_1} + \underbrace{W_\theta(\hat{\mathbf{u}}(\hat{P})) - W_\theta(\mathbf{u}(\hat{P}))}_{=C_2} \end{aligned}$$

Let  $B_1(\hat{\mu}) = \max_{i \in [n]} \psi'(\hat{\mathbf{u}}_i(P^*), \alpha)$ .

We first prove:

$$C_1 \leq 2B_1(\hat{\mu}) \sqrt{n \|v\|_2^2} \sqrt{\sum_{(i,j) \in [n]^2} (\hat{\mu}_{ij} - \mu_{ij})^2}. \quad (2)$$

To prove (2), we start by using the concavity of  $\psi(\cdot, \alpha)$  for  $\alpha \leq 1$ . Let  $\Phi(\cdot) = \frac{1}{2}\psi(\cdot, \alpha)$ . We have:

$$\begin{aligned} C_1 &= \sum_{i=1}^n (\Phi(\mathbf{u}(P^*)) - \Phi(\hat{\mathbf{u}}(P^*))) \leq \sum_{i=1}^n \Phi'(\hat{\mathbf{u}}_i(P^*)) (\mathbf{u}(P^*) - \hat{\mathbf{u}}(P^*)) \\ \text{thus } C_1 &\leq \sum_{i=1}^n \sum_{j=1}^n \Phi'(\hat{\mathbf{u}}_i(P^*)) (\mu_{ij} - \hat{\mu}_{ij}) (P_{ij}^* + P_{ji}^*) v \\ &= \sum_{i=1}^n \sum_{j=1}^n \underbrace{(\Phi'(\hat{\mathbf{u}}_i(P^*)) (\mu_{ij} - \hat{\mu}_{ij}) + \Phi'(\hat{\mathbf{u}}_j(P^*)) (\mu_{ji} - \hat{\mu}_{ji}))}_{=A_{ij}} P_{ij}^* v \end{aligned}$$

where, similarly to the proof of Lemma 1, we swapped the indexed  $(i, j)$  in the  $\Phi'(\hat{\mathbf{u}}_i(P^*)) \mu_{ij} P_{ji}^* v$ , which is possible because  $i$  and  $j$  span the same range in the sum.

Notice that the terms  $A_{ij} P_{ij}^* v$  are all zero except if  $i \in \mathcal{N}$  and  $j \in \mathcal{I}$  (because  $P_{ijk}^* = 0$  otherwise). For  $i \in \mathcal{N}$ , let  $\sigma_i$  be a ranking which ranks  $(A_{ij})_{j \in \mathcal{I}}$  in decreasing order, i.e.,  $\sigma_i(j) < \sigma_i(j') \implies A_{ij} \geq A_{ij'}$ . Using Lemma 3 in Appendix E, we have:

$$C_1 \leq \max_{P \in \mathcal{P}} \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{I}} A_{ij} P_{ij} v = \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{I}} A_{ij} v_{\sigma_i(j)}$$

Now let  $V = [v_{\sigma_i(j)}]_{\substack{i \in \mathcal{N} \\ j \in \mathcal{I}}}$ . By Cauchy-Shwarz inequality and denoting  $\|X\|_F = \sqrt{\sum_{ij} X_{ij}^2}$

the Frobenius norm of matrix  $X$ , we have  $\|V\|_F = \sqrt{n \|v\|_2^2}$  and  $\|A\|_F \leq B_1(\hat{\mu})(\|\mu - \hat{\mu}\|_F +$

<sup>12</sup>The dependency on  $\hat{\mu}$  in  $B(\hat{\mu})$  is because  $\psi'(\cdot, \alpha)$  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  $B < \psi'(\eta, \alpha)$ .

$\|\mu^\top - \hat{\mu}^\top\|_F$ ), leading to:

$$C_1 \leq \sqrt{n \|v\|_2^2} \|A\|_F \leq 2B_1(\hat{\mu}) \sqrt{n \|v\|_2^2} \|\mu - \hat{\mu}\|_F$$

which proves (2).

Similarly, using  $B_2(\hat{\mu}) = \max_{i \in \llbracket n \rrbracket} \psi'(u_i(\hat{P}), \alpha)$  and the same arguments as above, we obtain:

$$C_2 \leq \sqrt{n \|v\|_2^2} \|A\|_F \leq 2B_2(\hat{\mu}) \sqrt{n \|v\|_2^2} \|\mu - \hat{\mu}\|_F$$

which yields the desired result.  $\square$

## D Comparison to utility/inequality trade-offs

In this appendix, we provide the proofs of Section 3, and describe more precisely how we applied quality-weighted exposure and equality of exposure in reciprocal recommendation.

### D.1 One-sided recommendation: quality-weighted exposure

We prove here Proposition 2 of Section 3. The result shows 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 5 gives an example.

**Proposition 2.** *The following claims hold irrespective of the choice of  $\mathbf{u}^{\text{qua},\beta} \in \mathcal{U}_\beta^{\text{qua}}$ .*

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

$$(\exists \beta > 0, \mathbf{u}_N^\theta \succ_L \mathbf{u}_N^{\text{qua},\beta} \text{ and } \mathbf{u}_T^\theta \succ_L \mathbf{u}_T^{\text{qua},\beta}) \quad \text{and} \quad \lim_{\beta \rightarrow \infty} \frac{\sum_{i \in \mathcal{N}} u_i^{\text{qua},\beta}}{\sum_{i \in \mathcal{N}} u_i^\theta} \xrightarrow{d \rightarrow \infty} \frac{5}{6}.$$

*Proof.* We prove it for  $N = 1$ , the more general case is just obtained by repeating the pattern with  $d + 1$  items and  $d + 1$  users.

Let  $i_1, \dots, i_{d+1}$  be the indexes of the users and  $j_1, \dots, j_{d+1}$  the indexes of the items. The preferences have the following pattern:

$$\forall k \in \llbracket d + 1 \rrbracket, \mu_{i_k j_k} = 1 \quad \forall k \in \llbracket d \rrbracket, \mu_{i_k j_{d+1}} = \frac{1}{2}$$

all other  $\mu_{ij}$  (for user  $i$  and item  $j$ ) are set to 0 (note that we are in a problem with one-sided preferences, which means  $\mu_{ji} = 1$  for every item  $j$  and user  $i$ ).

We consider a task with a single recommendation slot ( $v_1 = 1, v_2 = \dots = v_{|I|} = 0$ ). On that problem, the optimal ranking for every  $\theta \in \Theta$  is to show item  $j_k$  to user  $i_k$ , which leads to perfect equality in terms of item exposure, and maximizes every user utility. It is thus leximin optimal for both users and items for every  $\theta \in \Theta$ .

Then, the qualities are equal to:

$$\forall k \in \llbracket d \rrbracket, q_{j_k} = 1 \quad q_{j_{d+1}} = \frac{1}{2}d + 1$$

the target exposure is thus  $t_{j_k} = \frac{d+1}{\frac{3}{2}d+1}$  for  $k \in \llbracket d \rrbracket$  and  $t_{j_{d+1}} = (d + 1) \frac{\frac{1}{2}d+1}{\frac{3}{2}d+1}$ .

Since the problem is symmetric in the users  $i_1, \dots, i_d$ , by the concavity of  $F_\beta^{\text{qua}}(\mathbf{u}(P))$  with respect to  $P$ , there is an optimal ranking described by a single probability  $p$  as:

$$\forall k \in \llbracket d \rrbracket, P_{i_k j_k} = 1 - p \quad P_{i_k j_{d+1}} = p \quad P_{i_{d+1} j_{d+1}} = 1$$

Note that for such a  $P$ ,  $\forall k \in \llbracket d \rrbracket, u_{i_k}^{\text{qua},\beta}(P) = 1 - \frac{1}{2}p$ , and it is clear that there is  $\beta > 0$  such that  $p > 0$ , which then implies  $\mathbf{u}^\theta \succ_L \mathbf{u}_N^{\text{qua},\beta}$  and  $\mathbf{u}^\theta \succ_L \mathbf{u}_T^{\text{qua},\beta}$ .

Now, as  $\beta \rightarrow \infty$ ,  $p$  is such that exposure equals its target, which leads to the following equation:

$$dp + 1 = (d + 1) \frac{\frac{1}{2}d + 1}{\frac{3}{2}d + 1}.$$

We thus get  $p = \frac{d+1}{d} \frac{d+2}{3d+2} - \frac{1}{d} \xrightarrow{d \rightarrow \infty} \frac{1}{3}$ , which gives the result  $u_{i_k}^{\text{qua},\beta}(P) = 1 - \frac{1}{2}p \xrightarrow{p \rightarrow \frac{1}{3}} \frac{5}{6}$ .

Notice that similarly to Proposition 3, the result does not depend on the choice of  $\mathbf{u}^{\text{qua},\beta}$  because the sum of user utilities converges.  $\square$

## D.2 Reciprocal recommendation: equality of exposure

We now prove Proposition 3.

**Proposition 3.** For  $\beta > 0$ , let  $\mathcal{U}_\beta^{\text{eq}} = \text{argmax}_{\mathbf{u} \in \mathcal{U}} F_\beta(\mathbf{u})$ . The claim below holds irrespective of the choice of  $\mathbf{u}^{\text{eq},\beta} \in \mathcal{U}_\beta^{\text{eq}}$ . Let  $n \geq 5$ . There is a reciprocal recommendation task with  $n$  users such that:

$$\forall \theta \in \Theta, \mathbf{u}^\theta, \exists \beta > 0 : \forall i \in \llbracket n \rrbracket, u_i^\theta > u_i^{\text{eq},\beta} \quad \text{and} \quad \lim_{\beta \rightarrow \infty} \sum_{i \in \mathcal{N}} u_i^{\text{eq},\beta} = 0.$$

*Proof.* The example is given in Figure 5. We still consider a recommendation task with a single recommendation slot.

Let us rename the users by  $i_1, i_2, \dots, i_5$ . The preference patterns are  $\mu_{i_1 i_2} = \mu_{i_1 i_3} = 1$  and  $\mu_{i_4 i_5} = 1$ . Apart from  $\mu_{ij} = \mu_{ji}$ , other  $\mu_{ij}$ s are 0. In this proof, we show that  $u_{i_1}^{\text{eq},\beta} = 2u_{i_2}^{\text{eq},\beta}$  for every  $\beta$ , which implies that  $u_{i_1}^{\text{eq},\beta} \xrightarrow{\beta \rightarrow \infty} 0$  because 0 utility for every user is feasible. On this task, the leximin ranking also maximizes the sum of users utilities (as shown in Figure 5), so the optimal ranking is the same for every  $\theta \in \Theta$ , and every user has a two-sided utility of at least 1.5.

Since  $F_\beta(\mathbf{u})$  is strictly Schur-concave for  $\beta > 0$ ,  $i_2$  and  $i_3$  always have the same utility in an optimal utility profile (because they play a symmetric role).  $i_4$  and  $i_5$  also have the same utility. Note that the interest of  $i_4$  and  $i_5$  in that problem is to make it possible to recommend them to  $i_1$ , which has 0 value.

Similarly to the problem in one-sided recommendation, the only way to decrease the penalty is to reduce the utility of  $i_1, i_4, i_5$ . However, reducing the utility of  $i_1$  can only be done by either recommending  $i_4$  or  $i_5$  to  $i_1$ , or recommending  $i_4/i_5$  to  $i_2/i_3$ . In all cases, decreasing  $i_1$ 's utility decreases  $i_2/i_3$ 's utilities.

More precisely, because of the symmetries and the concavity of  $F_\beta(\mathbf{u}(P))$  with respect to  $P$ , for every  $\beta > 0$ , there is an optimal ranking tensor described by three probabilities  $p, q, q'$  such that:<sup>13</sup>

$$\begin{aligned} P_{i_1 i_2} = P_{i_1 i_3} &= \frac{1}{2}p & P_{i_2 i_1} = P_{i_3 i_1} &= q & P_{i_4 i_5} = P_{i_5 i_4} &= q' \\ P_{i_1 i_4} = P_{i_1 i_5} &= \frac{1}{2}(1-p) & P_{i_2 i_3} = P_{i_2 i_4} = P_{i_2 i_5} &= \frac{1}{3}(1-q) & P_{i_4 i_1} = P_{i_4 i_2} = P_{i_4 i_3} &= \frac{1}{3}(1-q') \\ & & P_{i_3 i_2} = P_{i_3 i_4} = P_{i_3 i_5} &= \frac{1}{3}(1-q) & P_{i_5 i_1} = P_{i_5 i_2} = P_{i_5 i_3} &= \frac{1}{3}(1-q') \end{aligned}$$

In all cases, the two-sided utility are

$$u_{i_1}(P) = \underbrace{p}_{P_{i_1 i_2} \mu_{i_1 i_2} + P_{i_1 i_3} \mu_{i_1 i_3}}_{\text{user-side utility}} + \underbrace{2q}_{P_{i_2 i_1} \mu_{i_2 i_1} + P_{i_3 i_1} \mu_{i_3 i_1}}_{\text{item-side utility}} \quad \text{and} \quad u_{i_2}(P) = q + \frac{1}{2}p$$

Thus, in an optimal ranking for  $F_\beta(\mathbf{u})$ , we must have  $u_{i_1}(P) = 2u_{i_2}(P)$ . Equality, which is achieved at  $\beta \rightarrow \infty$  can then only be at 0 utility for every user (since 0 is feasible).

The task used in the proof contains only 5 users. Any number of users can be added to the group  $\{i_4, i_5\}$ , with a ‘‘complete’’ preference profile ( $\mu_{ij} = 1$  for all pair  $i, j$  in that group).  $\square$

<sup>13</sup>Since there is a single recommendation slot, we identify  $P_{i_j 1}$  with  $P_{i_j}$

The Lorenz efficiency of our welfare approach guarantees that it cannot exhibit the undesirable behaviors of equality or quality-weighted exposure penalties described in Propositions 2 and 8.

### D.3 Equality of exposure and quality-weighted exposure in reciprocal recommendation

In one-sided recommendation with one-sided preferences, equality of exposure is the same as equality of utility. More generally, let  $e_j(P) = \sum_{i \in \mathcal{N}} P_{ij} v$  the total exposure of item  $j$ . Equality of exposure is defined by:

$$F_{\beta}^{\text{expo}}(P) = \sum_{i \in \mathcal{N}} \bar{u}_i(P) - \beta \sqrt{\sum_{j \in \mathcal{I}} \left( e_j(P) - \frac{|\mathcal{N}|}{|\mathcal{I}|} \|v\|_1 \right)^2}$$

In one-sided recommendation, parity of exposure is relatively well behaved because the exposure target  $\frac{|\mathcal{N}|}{|\mathcal{I}|} \|v\|_1$  is constant. Driving towards equality can thus not lead to a decrease of the total exposure budget, which was the problem with equality of utility in settings with two-sided preferences (driving towards equality of utility leads to a decrease of total utility), as we described in Section 3.

The formula allows us to extend parity of exposure in the next section and in our experiments, since it is also valid in reciprocal recommendation. Likewise, the formula of quality-weighted exposure that is also valid in reciprocal recommendation is given by:

$$F_{\beta}^{\text{qua}}(P) = \sum_{i \in \mathcal{N}} \bar{u}_i(P) - \beta \sqrt{\sum_{j \in \mathcal{I}} \left( e_j(P) - \frac{q_j E}{Q} \right)^2}.$$

The result below shows that equality of exposure and quality-weighted exposure lead to inefficiencies in reciprocal recommendation settings:

**Proposition 8.** *For every  $n \in \mathbb{N}_*$ , there is a reciprocal recommendation task with  $n$  users such that:*

$$\forall \theta \in \Theta, \exists \beta > 0 : \quad \mathbf{u}^{\theta} \succ_{\text{L}} \mathbf{u}^{\text{expo}, \beta} \quad \text{and} \quad \mathbf{u}^{\theta} \succ_{\text{L}} \mathbf{u}^{\text{qua}, \beta}.$$

$$\text{Moreover, } \lim_{\beta \rightarrow \infty} \sum_{i \in \mathcal{N}} u_i^{\text{expo}, \beta} = \frac{2}{n} \sum_{i \in \mathcal{N}} u_i^{\theta} \text{ and } \lim_{\beta \rightarrow \infty} \sum_{i \in \mathcal{N}} u_i^{\text{qua}, \beta} = \frac{2+n}{2n} \sum_{i \in \mathcal{N}} u_i^{\text{sum}}.$$

*Proof.* An example of extreme case is with  $n$  users when there is a “leader” who is the only possible match with other users. We consider a single recommendation slot. The preferences are:

$$\forall j \in \{2, \dots, n\}, \mu_{1j} = \mu_{j1} = 1 \quad \forall (i, j) \in \{2, \dots, n\}^2, \mu_{ij} = 0.$$

On this task, the for every  $\theta \in \Theta$ , the optimal ranking is given by:

$$\forall j \in \{2, \dots, n\}, P_{1j} = \frac{1}{n-1} \quad \forall i \in \{2, \dots, n\} P_{i1} = 1.$$

The reason it is the only possible optimal ranking is because it is leximin optimal and has the maximum achievable sum of utilities. The utilities are then  $u_1(P) = n$  and  $u_i(P) = 1 + \frac{1}{n-1}$ , which leads to  $\sum_{i=1}^n u_i = 2n$ .

**Equality of exposure** Driving towards equality of exposure requires to reduce the exposure of user 1, which in turn reduces the utility of user 1 and the utilities of those who user 1 is less exposed to. Thus, there is  $\beta > 0$  such that  $\mathbf{u}^{\theta} \succ_{\text{L}} \mathbf{u}^{\text{expo}, \beta}$  because of the loss of efficiency. Finally, by the concavity of the objective with respect to  $P$ , and by the symmetry of the problem with respect to  $i_2, \dots, i_n$ , we can conclude that an optimal way to achieve perfect equality of exposure is to recommend, to every user  $i$ , every user  $j \neq i$  with probability  $\frac{1}{n-1}$ . The utility is then  $u_1(P) = 1 + (n-1)\frac{1}{n-1}$  and  $u_i(P) = \frac{2}{n-1}$  for  $i \geq 2$ , leading to  $\sum_{i=1}^n u_i = 4$ , which gives the result.

**Quality-weighted exposure** On the same example, the qualities are  $q_1 = n-1$  and  $q_i = 1$  for  $i \geq 2$ . The total exposure targets are then  $t_1 = \frac{1}{2}n$  and  $t_i = \frac{n}{2(n-1)}$ . These exposure targets mean less exposure for 1 than in the leximin ranking. Thus  $\beta$  sufficiently large has the effect of reducing 1’s exposure<sup>14</sup>, which reduces the utility of 1 and the users to whom 1 is less recommended. Thus

<sup>14</sup>Direct calculations of the derivatives show that when  $\beta > 0$  is too small the penalty has no effect.




Figure 5: **Left:** Example of a reciprocal recommendation task where equality of utility leads to 0 utility (see the proof of Prop. 3 in App. D). 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  $\theta$ . **Right:** Example where quality-weighted exposure reduces user utility while increasing inequalities between items.

$\mathbf{u}^\theta \succ_L \mathbf{u}^{\text{qua},\beta}$ . By the symmetry of the problem, as  $\beta \rightarrow \infty$ , quality weighted exposure is achieved by setting:

$$\begin{aligned} \forall j \in \{2, \dots, n\} : P_{1j} &= \frac{1}{n-1} & P_{j1} &= \frac{n}{2(n-1)} \\ \forall j' \in \{2, \dots, n\}, j' \neq j, P_{jj'} &= \frac{1 - \frac{n}{2(n-1)}}{n-2} = \frac{1}{2(n-1)} \end{aligned}$$

The utilities are then  $u_1(P) = 1 + (n-1)\frac{n}{2(n-1)} = 1 + \frac{n}{2}$  and  $u_i(P) = \frac{n}{2(n-1)} + \frac{1}{n-1} = \frac{n+2}{2(n-1)}$ . The total utility is thus  $2 + n$ , which gives the result.  $\square$

The Lorenz efficiency of our welfare approach guarantees that it cannot exhibit the undesirable behaviors of parity or quality-weighted exposure penalties described in Propositions 2 and 8.

## E A generic Frank-Wolfe algorithm for ranking

In this section, we present a general form of our algorithm presented in Section 4, as well as the proofs of the claims.

Let  $F : \mathbb{R}^n \rightarrow \mathbb{R}$ , concave, and we want to find

$$P^* \in \operatorname{argmax}_{P \in \mathcal{P}} F(\mathbf{u}(P)). \quad (3)$$

Let  $\langle X | Y \rangle = \sum_{ijk} X_{ijk} Y_{ijk}$  be the dot product between three-way tensors, and let  $\nabla(F \circ \mathbf{u})(P)$  be the gradient of  $P \mapsto F(\mathbf{u}(P))$  taken at  $P$ , i.e.,  $(\nabla(F \circ \mathbf{u}))_{ijk} = \frac{\partial F \circ \mathbf{u}}{\partial P_{ijk}}$

Starting from  $P^{(0)} \in \mathcal{P}$  (in our experiments we always use a utilitarian ranking  $P^{(0)} \in \operatorname{argmax}_{P \in \mathcal{P}} \sum_{i=1}^n u_i(P)$ ), the Frank-Wolfe algorithm alternates two steps for  $t \geq 1$ :

1. let  $\tilde{P} \in \operatorname{argmax}_{P \in \mathcal{P}} \langle P | \nabla(F \circ \mathbf{u})(P^{t-1}) \rangle$
2.  $P^{(t)} = (1 - \gamma^{(t)})P^{(t-1)} + \gamma^{(t)}\tilde{P}$  with  $\gamma^{(t)} = \frac{2}{t+2}$

The stepsize  $\frac{2}{t+2}$  is from Clarkson [13, Section 3], which avoids a line search and in our experiments seemed to yield acceptable results. Irrespective of the step size, the fundamental results which allows to use Frank-Wolfe in the setting of (3) are the two following lemmas:

**Lemma 2.** Recall that  $u_i(P) = \sum_{j=1}^n \mu_{ij}(P_{ij} + P_{ji})v$ . Let  $\frac{\partial F}{\partial u_i}$  denote the derivative of  $F$  with respect to its  $i$ -th argument and  $\frac{\partial F}{\partial u_i}(\mathbf{u}(P))$  the value of this derivative at  $\mathbf{u}(P)$ .

Then,  $\forall i \in \mathcal{N}, \forall j \in \mathcal{I}, \forall k \in \llbracket \mathcal{I} \rrbracket$ , we have:

$$\frac{\partial F \circ \mathbf{u}}{\partial P_{ijk}}(P) = \left( \mu_{ij} \frac{\partial F}{\partial u_i}(\mathbf{u}(P)) + \mu_{ji} \frac{\partial F}{\partial u_j}(\mathbf{u}(P)) \right) v_k.$$

*Proof.* The result is a consequence of the chain rule:

$$\frac{\partial F \circ \mathbf{u}}{\partial P_{ijk}}(P) = \sum_{p=1}^n \frac{\partial F}{\partial u_p}(\mathbf{u}(P)) \frac{\partial u_p(P)}{\partial P_{ijk}}$$

With

$$u_p(P) = \sum_{q=1}^n \mu_{pq} \sum_{r=1}^{|\mathcal{I}|} (P_{pqr} + P_{qpr}) v_k.$$

Thus  $\frac{\partial u_p(P)}{\partial P_{ijk}} = (\mu_{ij} \mathbb{1}_{\{p=i\}} + \mu_{ji} \mathbb{1}_{\{p=j\}}) v_k$ , which gives the desired result.  $\square$

**Lemma 3.** *Let  $A$  be an  $n \times n$  matrix with  $A_{ij} \in \mathbb{R}$  (not necessarily non-negative). Let  $v \in \mathbb{R}^{|\mathcal{I}|}$  with non-negative and non-increasing entries, i.e.,  $\forall k \in [|\mathcal{I}| - 1], v_k \geq v_{k+1} \geq 0$ . Let  $K$  be the last index such that  $v_K > 0$  (or  $K = |\mathcal{I}|$  if there is no such index).*

Let  $P \in \mathcal{P}$  such that:

$$\forall i, \forall \sigma_i \in \mathfrak{S}(P_i), \forall (j, j') \in \mathcal{I}^2 : \left( \sigma_i(j) \leq K \text{ and } \sigma_i(j) < \sigma_i(j') \implies A_{ij} \geq A_{ij'} \right).$$

And let  $X$  be the  $n \times n \times |\mathcal{I}|$  tensor defined as  $X_{ijk} = A_{ij} v_k$ .

Then  $P \in \operatorname{argmax}_{P \in \mathcal{P}} \langle P | X \rangle$ .

Moreover, if  $\forall k \in [|\mathcal{I}| - 1], v_k > v_{k+1} \geq 0$ , then for every  $P \in \operatorname{argmax}_{P \in \mathcal{P}} \langle P | X \rangle$ , we have:

$$\forall i, \forall \sigma_i \in \mathfrak{S}(P_i), \forall (j, j') \in \mathcal{I}^2 : \left( \sigma_i(j) < \sigma_i(j') \implies A_{ij} \geq A_{ij'} \right).$$

*Proof.* The result stems from the rearrangement inequality (also known as the Hardy-Littlewood inequality [23]), which states that for two vectors  $a \in \mathbb{R}_+^n$ , and  $b \in \mathbb{R}^n$ ,  $\operatorname{argmax}_{\nu} \sum_{j=1}^n a_{\nu(j)} b_j$ , where  $\nu$  spans the permutations of  $[n]$ , is the set of permutations such that  $b$  is ordered similarly to  $(a_{\nu(i)})_{i=1}^n$ . If the  $a_k$ s are non-increasing, then every permutation that sorts  $b$  in decreasing order is in the  $\operatorname{argmax}$ . We need the reciprocal statement for the second part of our Lemma: if the  $a_i$ s are strictly decreasing, then only the permutations that sort  $b$  in decreasing order are in  $\operatorname{argmax}_{\nu} \sum_{j=1}^n a_{\nu(j)} b_j$ . Note that these arguments are well-known in learning to rank [see, e.g., 14].

In our case, notice that

$$\langle P | X \rangle = \sum_{i \in \mathcal{N}} \left( \sum_{j \in \mathcal{I}} A_{ij} P_{ijk} v_k \right)$$

The maximization over  $P$  can then be performed over each user  $i$  (and thus each bistochastic matrix  $P_i$  separately). Now, if  $P_i$  is such that every  $\sigma_i \in \mathfrak{S}(P_i)$  orders  $A_{ij}$  in decreasing order, then by the rearrangement inequality  $\sigma_i \in \operatorname{argmax}_{\nu} \sum_{j \in \mathcal{I}} A_{ij} v_{\nu(j)}$ . Notice that if only the  $K$  first elements of  $v$  are non-zero, we only need a top- $K$  ranking. This gives us the first part of the theorem.

The second part of the theorem follows from the reciprocal of the rearrangement inequality, since for  $P_i$  to be an optimal stochastic ranking for  $\sum_{j \in \mathcal{I}} A_{ij} P_{ijk} v_k$ , every permutation  $\sigma_i$  in its support must be in  $\operatorname{argmax}_{\nu} \sum_{j \in \mathcal{I}} A_{ij} v_{\nu(j)}$ .  $\square$

## E.1 Proof of Theorem 1

Lemma 2 and 3 together are sufficient to give algorithms for the inference of stochastic rankings using our welfare function (1) and using the penalties of Section 3, by computing the partial derivatives  $\frac{\partial F}{\partial u_i}$ . The main result of Section 4, which we prove now, instantiates this principle for the welfare function approach:

**Theorem 1.** *Let  $\tilde{\mu}_{ij} = \Phi'_i(u_i(P^{(t)})) \mu_{ij} + \Phi'_j(u_j(P^{(t)})) \mu_{ji}$ . Let  $\tilde{P}$  such that:*

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

*Proof.* Notice that with  $W(P) = F(\mathbf{u}(P)) = \sum_{i=1}^n \Phi_i(u_i(P))$ , then  $\frac{\partial F}{\partial u_i}(\mathbf{u}(P)) = \Phi'_i(u_i(P))$ . By Lemma 2, we have that  $\langle P | \nabla F(P^{(t)}) \rangle$  is of the form  $\langle P | X \rangle$  with  $X_{ijk} = A_{ij} v_k$  with  $A_{ij} = \tilde{\mu}_{ij}$ , so the result is implied by Lemma 3.  $\square$

## E.2 Proof of Proposition 4

**Proposition 4.** Let  $B = \max_{i \in [n]} \|\Phi_i'\|_\infty$  and  $U = \max_{\mathbf{u} \in \mathcal{U}} \|\mathbf{u}\|_2^2$ . Let  $K$  be the maximum index of a nonzero value in  $v$  (or  $|\mathcal{I}|$ ). Then  $\forall t \geq 1$ ,  $W(P^{(t)}) \geq \max_{P \in \mathcal{P}} W(P) - O(\frac{BU}{t})$ . Moreover, for each user, an iteration costs  $O(|\mathcal{I}| \ln K)$  operations and requires  $O(K)$  additional bytes of storage.

*Proof.* Note that  $\mathcal{P}$  is a simplex over ranking tensors containing one deterministic ranking for each user. Using [13, Section 3], the Frank-Wolfe algorithm with our step-size converges in  $O(\frac{C_W}{t})$ , where, using [13, Equation 11] and denoting by  $\nabla^2 W$  the Hessian of  $W$ , we have

$$C_W \leq \sup_{\substack{\mathbf{u}, \mathbf{u}' \in \mathcal{U} \\ \tilde{\mathbf{u}} \in \mathcal{U}}} -\frac{1}{2}(\mathbf{u} - \mathbf{u}')^\top \nabla^2 W(\tilde{\mathbf{u}})(\mathbf{u} - \mathbf{u}') \leq \frac{B}{2} \sup_{\mathbf{u}, \mathbf{u}' \in \mathcal{U}} \|\mathbf{u} - \mathbf{u}'\|_2^2 \leq 2BU.$$

where we used  $\|\mathbf{u} - \mathbf{u}'\|_2^2 \leq 2\|\mathbf{u}\|_2^2 + 2\|\mathbf{u}'\|_2^2$ .

For the computation cost, we use Lemma 3, which is more precise than Theorem 1, to see that finding the argmax only requires a top- $K$  ranking. While technically any  $P \in \mathcal{P}$  should contain a whole bistochastic matrix, it is not necessary to store a completion of the top- $K$  rankings because they have no impact on the utility. As such, storing each  $\tilde{P}$  only costs  $O(K)$  bytes per user, which contain the indices of the top- $K$  items in the ranking found by Theorem 2.

Computing the two-sided utilities costs  $O(|\mathcal{N}||\mathcal{I}|)$ , and thus  $O(|\mathcal{I}|)$  per user. Moreover, computing the top- $K$  ranking costs  $O(|\mathcal{I}| \ln K)$  in the worst case, with a streaming method that maintains a min-heap of the top- $K$  elements seen so far, and finish with sorting the top- $K$  elements.  $\square$

Notice that for faster average performance, the top- $K$  sort can be performed using a fast selection algorithm (such as quickselect), to obtain the top- $K$  elements with  $O(|\mathcal{I}|)$  expected time complexity, and then sorting, yielding  $O(|\mathcal{I}| + K \ln K)$  expected time complexity per user at each iteration.



## F Additional experimental results

Our experiments are fully implemented in Python 3.9 using PyTorch<sup>15</sup>. We provide the code as supplementary material. We compare our welfare maximization approach with the fairness penalties presented in Section 3.

We also compare ourselves to the algorithm FairRec from Patro et al. [47] (referred to as *Patro et al.* in the figures and description), who consider envy-freeness as user-side fairness criterion, and max-min share of exposure as item-side fairness criterion. Envy-freeness states that every user should prefer their recommendation list to that of any other user. The max-min exposure criterion on the item side means that each user should receive an exposure of at least  $\beta \frac{E}{|I|}$ , where  $\beta$  is a parameter allowing to control how much exposure is guaranteed to items. We vary this parameter in our experiments to show the trade-offs achieved by *Patro et al.*. Since *Patro et al.* does not produce rankings, we took the recommendation list with the given order as a ranked list.

### F.1 One-sided recommendation: Lastfm-2k dataset

We describe in this section the details of the experiments presented in Section 5.1. We use a dataset from the online music service Last.fm<sup>16</sup>. In the main paper, we presented results on **Lastfm-2k** from Cantador et al. [9] which contains real play counts of  $2k$  users for  $19k$  artists, and was used by Patro et al. [47] who also study two-sided fairness in recommendation. We filter the top 2,500 items most listened to. Following [32], we pre-process the raw counts with log-transformation. We split the dataset into train/validation/test sets, each including 70%/10%/20% of the user-item play counts. We create three different splits using three random seeds. One-sided preferences are estimated using the standard matrix factorization algorithm<sup>17</sup> of Hu et al. [29] trained on the train set, with hyperparameters selected on the validation set by grid search. The number of latent factors is chosen in [16, 32, 64, 128], the regularization in [0.1, 1., 10., 20., 50.], and the confidence weighting parameter in [0.1, 1., 10., 100.]. The estimated preferences we use are the positive part of the resulting estimates.

Rankings are inferred from these estimated preferences. The exposure weights we use in the computation of utilities are the standard weights of the *discounted cumulative gain* (DCG) (also used in e.g., [54, 7, 42]):  $\forall k \in [|I|], v_k = \frac{1}{\log_2(1+k)}$ . For each ranking approach, the Frank-Wolfe algorithm is run with 5000 iterations to make sure we are close to convergence, and the number of recommendation slots is set to 40.

We evaluate rankings on estimated preferences, considered as ground truth, following many works on fair recommendation [54, 47, 60, 63]. This is because the goal is to evaluate the fairness of ranking algorithms themselves, rather than biases in preference estimates. All results are averaged over three random seeds. To obtain various trade-offs, for *welf* we vary  $\lambda$  in [0.001, 0.01, 0.05, 0.075, 0.1, 0.125, 0.15, 0.2, 0.3, 0.325, 0.35] and [0.4, 0.45, 0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8, 0.9, 0.95, 0.99, 0.999]. For *Patro et al.* we vary  $\beta$  in [0.01, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4] and [0.45, 0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95, 1], and for other methods we vary  $\beta$  in [0.001, 0.005, 0.01, 0.015, 0.0175, 0.02, 0.025, 0.03, 0.035, 0.04, 0.045, 0.05, 0.055, 0.06] and [0.065, 0.07, 0.075, 0.08, 0.085, 0.09, 0.095, 0.1, 0.105, 0.11, 0.2, 0.5, 1, 2, 5, 10, 20, 30, 40, 50, 70, 100].

**Item-side fairness** Figure 6 presents the various trade-offs achieved by each method in one-sided recommendation, as discussed in Section 5.1. We observe that only *qua.-weighted* is unable to reach equal exposure because of its quality-weighted exposure target: perfectly equal exposure is only permitted when all items have the same quality.

**Two-sided fairness** Figure 7 shows the effect of varying  $\alpha_1$  and  $\lambda$  on user fairness as in Figure 3 of the main paper, but with results repeated over three random seeds. We observe the same trade-offs and conclude again that *welf* is better than *Patro et al.* and *eq. exposure*, in terms of its impact on worse-off users.

<sup>15</sup><http://pytorch.org>

<sup>16</sup><https://www.last.fm/>

<sup>17</sup>Using the Python library Implicit: <https://github.com/benfred/implicit> (MIT License).




Figure 6: representative trade-offs achieved by the various compared methods on Lastfm-2k. The trade-offs achieved by the different methods look alike, except that *qua.-weighted* does not aim at reaching equality of exposure for extreme values of  $\beta$ . See Section 5.1 for the discussions on the differences between the trade-offs achieved by the different approaches.

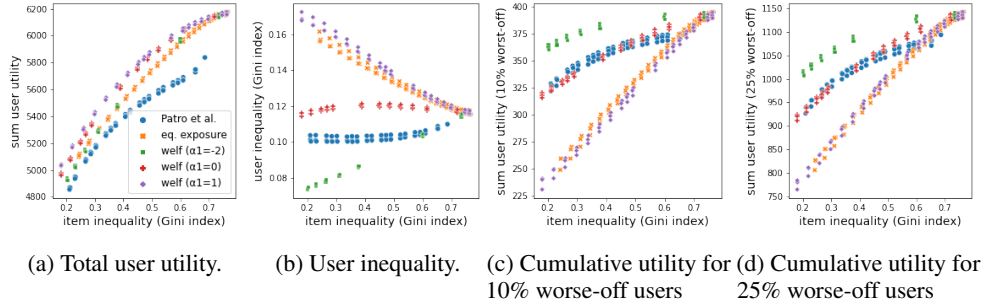


Figure 7: Focus on user fairness on Lastfm-2k: effect of varying  $\alpha_1$  (user-side curvature of the welfare function) keeping  $\alpha_2 = 0$ . The figure shows all the results obtained with a repetition of three seeds. Overlapping points correspond to the same model parameter across different seeds. We can see that the variance is negligible compared to the observed differences.

**The importance of considering the whole Lorenz curve** In Fig. 8 we show the results of the same models as before, but changing the way we measure the item inequality: using the standard deviation of exposure rather than the Gini index. Now, *eq. exposure* dominates the total utility/item inequality plot, since the plot corresponds exactly to the objective function of the algorithm. Comparing *eq. exposure* with *welf*  $\alpha_1 = 1$ , we now see that the trade-offs are different, with *eq. exposure* performing better on the worse-off users. Comparing *welf*  $\alpha_1 = 0$  and *Patro et al.*, we see that they still exhibit similar behaviors, with *welf*  $\alpha_1 = 0$  being better for better off users. Finally, *welf*  $\alpha_1 = -2$  still dominates the other methods in terms of performance on the worse-off users.

## F.2 One-sided recommendation: Lastfm-15k dataset

We replicate the experiments on a larger dataset to verify our conclusions at a larger scale. We consider another Lastfm dataset from Celma [11], which includes 360k users and 180k items (artists). We select the top 15,000 users and items having the most interactions, so we refer to this dataset as Lastfm-15k. We apply exactly the same experimental protocol as for Lastfm-2k, with the same range of hyperparameters for the different methods.

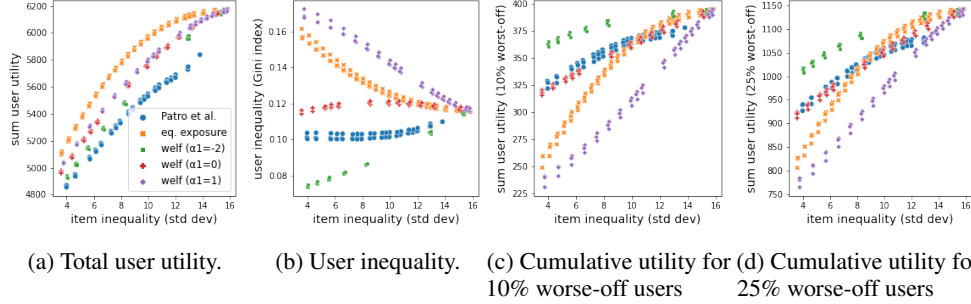


Figure 8: Focus on user fairness on Lastfm-2k, measuring item inequality with standard deviation rather than Gini index. We observe a similar relative behavior between *welf* and *Patro et al.*, but now equality of exposure is optimal on the total utility/item inequality trade-off since it corresponds exactly to the objective of the algorithm. Nonetheless, *welf*  $\alpha_1 = -2$  still obtains higher performance on 10%-25% worse-off users, showing that *welf* offers a larger range of trade-offs than *eq. exposure*.

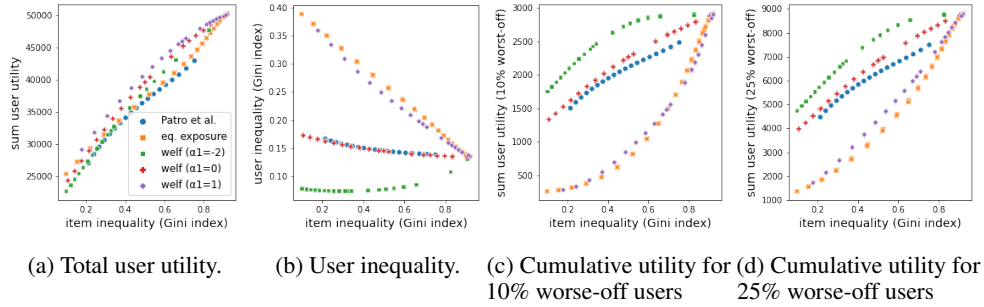


Figure 9: Results on Lastfm-15k when measuring the inequality between items with the Gini coefficient.

**Results** Fig. 11 and 10 show the results obtained by *welf*, *Patro et al.* and *eq. exposure*. The conclusions are similar to those on Lastfm-2k, with the results of *welf*  $\alpha = 0$  being more uniformly better than those of *Patro et al.*, even though overall similar. *welf*  $\alpha_1 = -2$  dominates in terms of user utility on worse-off users. *welf* and *eq. exposure* still find different trade-offs, with *welf* dominating *eq. exposure* when inequality between items is measured by the Gini index, and *eq. exposure* dominating *welf* when inequality is measured by the standard deviation.

### E.3 One-sided recommendation: Movielens dataset

We provide additional results on the **MovieLens-20m** dataset [24], which contains ratings on a 5-star scale of movies by real users. To simulate a collaborative filtering task with implicit feedback similar




Figure 10: Results on Lastfm-15k when measuring inequalities between items with the standard deviation.




Figure 11: Results on Movielens when measuring the inequality between items with the Gini coefficient.

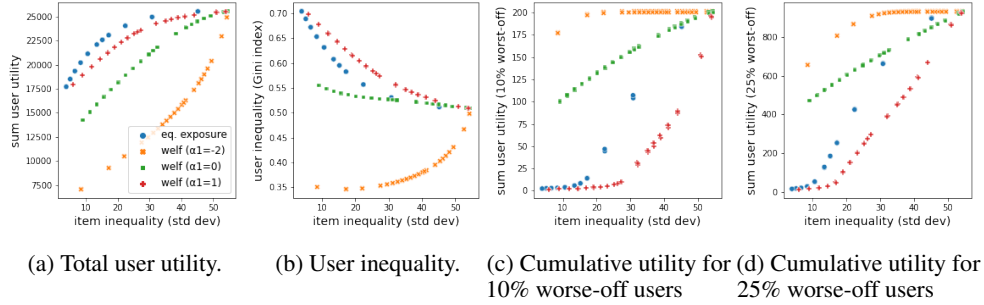


Figure 12: Results on Movielens when measuring inequalities between items with the standard deviation.

to Last.fm, we consider missing ratings as negative feedback and the task is to predict positive values. Since ratings  $< 3$  are usually considered as negative [38, 61], we set ratings  $< 3$  to zero, resulting in a dataset with preference values among  $\{0, 3, 3.5, 4, 4.5, 5\}$ . As for Lastfm-15k, we select the top 15,000 users and items with the most interactions. For the inference and evaluation of rankings, we follow the same protocols as for Last.fm.

The experimental protocol is the same as for Lastfm-2k and Lastfm-15k except that we do not run the algorithm by [47] because its runtime was prohibitive.

**results** The results are qualitatively similar to those on Lastfm-2k and Lastfm-15k except that the trends are magnified.  $welf \alpha = 1$  and  $eq. exposure$  seem more similar, with  $welf \alpha = 0$  dominating the trade-off total utility/item inequality when item inequality is measured with the Gini index, and  $eq. exposure$  dominating this trade-off when item inequality is measured with standard deviation.  $welf \alpha = -2$  has great performance on worse-off users compared to  $eq. exposure$  or  $welf$  with larger  $\alpha$ , but also comes at a significant cost in terms of total user utility, which is very rapidly driven down.

#### F.4 Reciprocal recommendation: Twitter-13k dataset

We now provide the full details of the experiments on Twitter presented in Section 5.2 of the main body. Given the lack of common benchmark for reciprocal recommendation [45], we generate a reciprocal recommendation task for people-to-people recommendation problems based on the social network Twitter. We use the Higgs Twitter-13k dataset which includes (directed) follower relationships between users.<sup>18</sup> We keep users having at least 20 mutual follows, resulting in a subset of 13k users. We use the directed links to estimate the probability  $\phi_{ij}$  that  $i$  follows  $j$ , and the (symmetric) probability of a mutual follow, which is  $\mu_{ij} = \phi_{ij} \times \phi_{ji}$ . As in the experiments for one-sided recommendation, we split the dataset into train/validation/test sets, each including 70%/10%/20% of the *directed* follower links. We create three random uniform splits, corresponding to three different seeds.

<sup>18</sup>It was collected following the discovery of the Higgs boson in July, 2012.

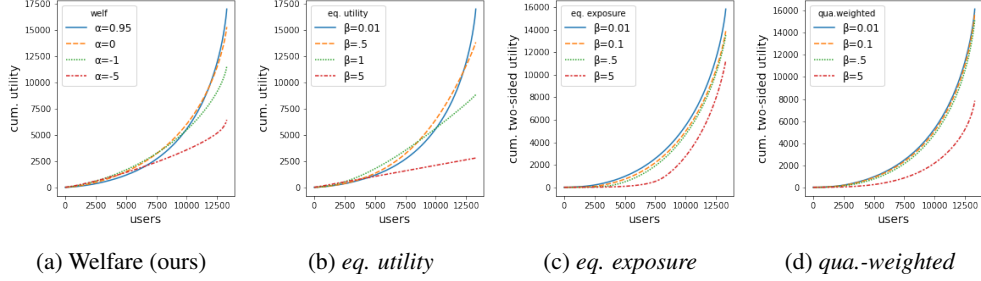


Figure 13: representative trade-offs achieved by the various compared methods on Twitter-13k. Exposure-based approaches (*qua.-weighted* and *eq. exposure*) do not yield interesting trade-offs as they are unable to increase the utility of worse-off users. The trade-offs achieved by the *welf* and *eq. utility* are different. Equal utility rapidly generates near-flat curves without really focusing on the very first users, while *welf* increases the utility of the worst-off users while keeping the total utility relatively high.

Estimates  $\hat{\phi}_{ij}$  are built with logistic matrix factorization<sup>19</sup> [32] trained on the train set with hyperparameter selection on the validation set. The number of latent factors is chosen in [16, 32, 64, 128], the regularization in [0.1, 1., 10., 20., 50.], and the confidence weighting parameter in [0.1, 1., 10., 100.]. Rankings are inferred from all estimated mutual preferences  $\hat{\mu}_{ij} = \max(\hat{\phi}_{ij}\hat{\phi}_{ji}, 0)$ . For each ranking method, the Frank-Wolfe algorithm is run with 5000 iterations, and the number of recommendation slots is set to 40. As for one-sided recommendation, rankings are estimated on estimated mutual preferences taken as ground truth.

We generate different trade-offs with *welf* by varying  $\alpha$  in [0.99, 0.9, 0.75, 0.5, 0.25, 0, -0.25, -0.5, -0.6, -0.7, -0.8, -0.9, -1.0], [-1.1, -1.25, -1.5, -1.75, -2.0, -2.5, -3, -5, -10, -15, -16, -17, -18]. For all other methods, we vary  $\beta$  in [0.01, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1, 1.1, 1.25, 1.5, 2, 5, 10, 50, 100].

All presented results are obtained by averaging performance over the three seeds.

**Results** Figure 13 presents the trade-offs achieved by the different methods on Twitter-13k. As expected, *qua.-weighted* and *eq. exposure* do not exhibit a good behavior: stronger penalties lead to more dominated curves where the utility of every user is decreased. This is because constraining item exposure is not meaningful in reciprocal recommendation, where the relevant utility is the two-sided utility. The trade-offs achieved by the *welf* and *eq. utility* are different. Equal utility rapidly generates near-flat curves without really focusing on the very first users, while *welf* increases the utility of the worst-off users while keeping the total utility relatively high.

## F.5 Reciprocal recommendation: Epinions dataset

We present additional experiments on reciprocal recommendation with the Epinions dataset [49]. Epinions.com is a consumer review site with a who-trust-whom network, and the dataset gathers (directed) trust relationships between members of the platform. Here, we consider the task of finding mutual trust links. We keep users having at least 20 mutual trust links, resulting in a subset of 800 entities. For the inference and evaluation of rankings, we use the same protocols as for the Twitter experiments described in the previous subsection. The experimental parameters are the same as for the Twitter-13k dataset.

**Results** Figure 14 presents the trade-offs achieved by the different methods on Epinions. As expected, *qua.-weighted* and *eq. exposure* do not exhibit a good behavior: stronger penalties lead to more dominated curves where the utility of every user is decreased. In Figure 15 plots the equivalent of Fig. 4. The results are similar: all of *qua.-weighted*, *eq. exposure* and *eq. utility* have dominated curves. We also observe that in the more interesting region where we are closer to the maximum achievable utility, *welf* optimizes better the utility of worse-off users. Yet, in that region, there is no strict dominance of *welf* over *eq. utility*.

<sup>19</sup>Using the Python library Implicit (MIT License).

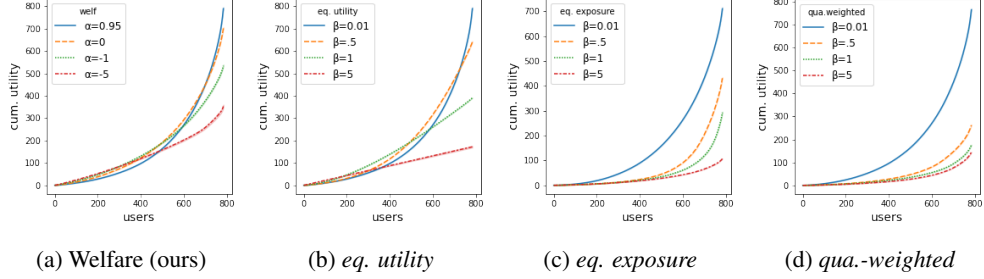


Figure 14: representative trade-offs achieved by the various compared methods on Epinions. The results are qualitatively similar to those on Twitter-13k.




Figure 15: Results on the epinions dataset.

## G Pairwise vs pointwise penalties

Our penalty-based approach uses the penalty  $\sqrt{D(\mathbf{u})}$  with:

$$D(\mathbf{u}) = \sum_{j \in \mathcal{I}} \left( u_j - \frac{1}{|\mathcal{I}|} \sum_{j' \in \mathcal{I}} u_{j'} \right)^2.$$

Some authors use  $D'(\mathbf{u}) = \sum_{(j,j') \in \mathcal{I}^2} |u_j - u_{j'}|$  instead of  $\sqrt{D(\mathbf{u})}$  [55, 42, 6], but it is less computationally efficient than our penalty because it involves a quadratic number of terms.

The penalties are similar in that they are related to well-known measures of inequalities:

- $\frac{D'(\mathbf{u})}{2|\mathcal{I}| \sum_{j \in \mathcal{I}} u_j}$  is the Gini index of  $\mathbf{u}_{\mathcal{I}}$  [19], which, up to an affine transform is the area under the Lorenz curve.
- $D(\mathbf{u})$ , which is (up to a constant) the variance of  $\mathbf{u}_{\mathcal{I}}$  is part of the family of additively decomposable inequality measures [52]. We use  $\sqrt{D(\mathbf{u})}$  to scale the penalty with the sum of users' utilities.

Note that  $\sqrt{D(\mathbf{u})}$  and  $D'(\mathbf{u})$  have the same dependency on the overall scale of the utilities (i.e., multiplying all utilities by a constant factor has the effect of multiplying both penalties by the same

factor). Since both penalties drive towards equality, it is straightforward to show that the results of Section 3 as  $\beta \rightarrow \infty$  also apply to  $D'(\mathbf{u})$ .

## H Exposure constraints at the level of every ranking

The notions of fairness of exposure are sometimes defined with item-side constraints defined at the level of *every ranking* [54, 6]. We give here the examples of constraints for equality of exposure and quality-weighted exposure:

$$\begin{array}{ll} \text{equality of} & P^{\text{expo}} \in \operatorname{argmax}_{P \in \mathcal{P}} \sum_{i \in \mathcal{N}} u_i(P) \quad \text{u.c. } \forall (i, j) \in \mathcal{N} \times \mathcal{I}, P_{ij} v = \frac{\|v\|_1}{|\mathcal{I}|} \\ \text{exposure} & \\ \text{quality-weighted} & P^{\text{qua}} \in \operatorname{argmax}_{P \in \mathcal{P}} \sum_{i \in \mathcal{N}} u_i(P) \quad \text{u.c. } \forall (i, j) \in \mathcal{N} \times \mathcal{I}, P_{ij} v = \frac{\mu_{ij} \|v\|_1}{\sum_{j' \in \mathcal{I}} \mu_{ij'}} \\ \text{exposure} & \end{array}$$

The advantage of this formulation is that it leads to optimization problems that can be solved locally for every user, since there is no dependency between user rankings through item utility anymore.

However, applying the exposure criterion at the level of every ranking effectively applies a different notion of fairness. In our setting, this corresponds to defining a separate recommendation task for every user, i.e., taking  $|\mathcal{N}| = 1$ . The welfare function then mediates, within a single ranking, between the user utility and the utility of the different items.

When evaluated on exposures aggregated over all users, as we do in the paper, applying the fairness constraints at the level of individual rankings can lead to drastic reductions of user utility for no benefit in terms of total item exposure. This is summarized in the following result, which shows that there exists problems for which the optimal rankings for every  $\theta \in \Theta$  satisfy the constraints of equality of exposure and quality-weighted exposure as we define them in Section 3, but when applying the constraints at the level of every ranking, it has the effect of reducing user utility. In the proposition, we use the notation of the objective function for parity of exposure  $F_\beta$  and  $F_\beta^{\text{qua}}$  of Section 3.

**Proposition 9.** *For every  $d \in \mathbb{N}_*$  and every  $N \in \mathbb{N}_*$ , there is a one-sided recommendation task with  $d + 1$  items and  $N(d + 1)$  users such that,  $\forall \theta \in \Theta$ :*

$$\forall \mathbf{u}^\theta \in \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}} W_\theta(\mathbf{u}), \forall \beta > 0 \text{ we have: } \mathbf{u}^\theta \in \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}} F_\beta(\mathbf{u}) \text{ and } \mathbf{u}^\theta \in \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}} F_\beta^{\text{qua}}(\mathbf{u}),$$

and

$$\sum_{i \in \mathcal{N}} u_i(P^{\text{expo}}) = \frac{2}{d+1} \sum_{i \in \mathcal{N}} u_i^\theta \quad \text{and} \quad \sum_{i \in \mathcal{N}} u_i(P^{\text{qua}}) = \left(\frac{1}{2} + \frac{1}{d}\right) \sum_{i \in \mathcal{N}} u_i^\theta.$$

In other words, applying the constraints at the level of every ranking might lead to a drastic decrease of user utilities, even in tasks where satisfying the constraints on average over users (as we do in this paper) does not conflict with the optimal ranking.

*Proof.* We describe the problem with  $N = 1$ , the general case is obtained by repeating the preference pattern. Let us consider a task with  $d + 1$  users,  $d + 1$  items and a single recommendation slot. Let  $i_1, \dots, i_{d+1}$  be the user indexes, and  $j_1, \dots, j_{d+1}$  the item indexes. The preferences are defined as:

$$\forall k \in \llbracket d + 1 \rrbracket, \mu_{i_k j_k} = 1 \quad \forall j \neq j_k, \mu_{i_k j} = \frac{1}{d}.$$

All items have the same quality. For every  $\theta \in \Theta$ ,  $\mathbf{u}^\theta$  is given by the utilitarian ranking, which gives probability 1 to item  $j_k$  for user  $i_k$ , which leads to optimal user utility  $u_i^\theta = 1$  and equal exposure to every item  $u_j^\theta = 1$ . Since the quality is the same for all items (equal to  $1 + d \frac{1}{d}$ ), the ranking for  $\mathbf{u}^\theta$  satisfies both equality of exposure and quality-weighted exposure constraints. Thus, for every  $\beta > 0$ ,  $\mathbf{u}^\theta \in \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}} F_\beta(\mathbf{u})$  and  $\mathbf{u}^\theta \in \operatorname{argmax}_{\mathbf{u} \in \mathcal{U}} F_\beta^{\text{qua}}(\mathbf{u})$ .

On the other hand, satisfying equality of exposure at the level of every ranking requires  $P_{ij}^{\text{expo}} = \frac{1}{d+1}$  for every user  $i$  and item  $j$ , which leads to  $u_i(P^{\text{expo}}) = \frac{1}{d+1} + d \times \frac{1}{d} \times \frac{1}{d+1} = \frac{2}{d+1}$  for every user.

For quality-weighted exposure for every ranking, it leads to:

$$\forall k \in \llbracket d + 1 \rrbracket, P_{i_k j_k}^{\text{qua}} = \frac{1}{2} \qquad \forall j \neq j_k, P_{i_k j}^{\text{qua}} = \frac{1}{d}$$

and thus a user utility  $u_i(P^{\text{qua}}) = \frac{1}{2} + d \times \frac{1}{d} \times \frac{1}{d} = \frac{1}{2} + \frac{1}{d}$ . □

Notice that in the examples of the proof, the global exposure of items is constant in  $P^{\text{expo}}$  and  $P^{\text{qua}}$ , as well as in the ranking given by optimal welfare. So from the point of view of our definitions of utility, applying the constraints at the level of every ranking only decreased user utility for the benefit of no items. Yet, we re-iterate that applying item-side fairness at the level of every ranking might be meaningful in some contexts. The goal of this section is to highlight the difference between using global and local definitions of item utilities.