

Toward Fair Recommendation in Two-sided Platforms

ARPITA BISWAS, Harvard University, USA

GOURAB K. PATRO and NILOY GANGULY, Indian Institute of Technology Kharagpur, India and
L3S Research Center, Germany

KRISHNA P. GUMMADI, Max Planck Institute for Software Systems, Germany, Germany

ABHIJNAN CHAKRABORTY, Indian Institute of Technology Delhi, India

Many online platforms today (such as Amazon, Netflix, Spotify, LinkedIn, and AirBnB) can be thought of as two-sided markets with producers and customers of goods and services. Traditionally, recommendation services in these platforms have focused on maximizing customer satisfaction by tailoring the results according to the personalized preferences of individual customers. However, our investigation reinforces the fact that such customer-centric design of these services may lead to unfair distribution of exposure to the producers, which may adversely impact their well-being. However, a pure producer-centric design might become unfair to the customers. As more and more people are depending on such platforms to earn a living, it is important to ensure fairness to both producers and customers. In this work, by mapping a fair personalized recommendation problem to a constrained version of the problem of fairly allocating indivisible goods, we propose to provide fairness guarantees for both sides. Formally, our proposed *FairRec* algorithm guarantees Maxi-Min Share of exposure for the producers, and Envy-Free up to One Item fairness for the customers. Extensive evaluations over multiple real-world datasets show the effectiveness of *FairRec* in ensuring two-sided fairness while incurring a marginal loss in overall recommendation quality. Finally, we present a modification of *FairRec* (named as *FairRecPlus*) that at the cost of additional computation time, improves the recommendation performance for the customers, while maintaining the same fairness guarantees.

CCS Concepts: • **Information systems** → **Recommender systems**;

Additional Key Words and Phrases: Fair recommendation, multi-stakeholder recommendation, two-sided markets, fair allocation, maximin share, envy-freeness

Arpita Biswas and Gourab K. Patro equally contributed to the work.

This work was conducted when A. Biswas was a PhD student at the Indian Institute of Science. She gratefully acknowledges the support of a Google PhD Fellowship Award. G. K. Patro acknowledges the support by TCS Research Fellowship. This research was supported in part by an European Research Council (ERC) Advanced Grant for the project “Foundations for Fair Social Computing” (Grant Agreement No. 789373), and an European Research Council (ERC) Marie Skłodowska-Curie grant for the project “NoBIAS—Artificial Intelligence without Bias” (Grant Agreement No. 860630), both funded under the EU’s Horizon 2020.

Authors’ addresses: A. Biswas, Harvard University, Science and Engineering Complex, 150 Western Ave, Boston, MA 02134, USA; email: arpita.biswas@live.in; G. K. Patro and N. Ganguly, Indian Institute of Technology Kharagpur, India and L3S Research Center, Appelstrasse 4, Hannover 30167, Germany; emails: patrogourab@gmail.com, ganguly.niloy@gmail.com; K. P. Gummadi, Max Planck Institute for Software Systems, MPI-SWS, Campus E1 5, D-66123 Saarbrücken, Germany; email: gummadi@mpi-sws.org; A. Chakraborty, Indian Institute of Technology, Hauz Khas, New Delhi - 110016, India; email: abhijnan@iitd.ac.in.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1559-1131/2021/12-ART8 \$15.00

<https://doi.org/10.1145/3503624>

ACM Reference format:

Arpita Biswas, Gourab K. Patro, Niloy Ganguly, Krishna P. Gummadi, and Abhijnan Chakraborty. 2021. Toward Fair Recommendation in Two-sided Platforms. *ACM Trans. Web* 16, 2, Article 8 (December 2021), 34 pages.

<https://doi.org/10.1145/3503624>

1 INTRODUCTION

Popular online platforms such as Netflix, Amazon, Yelp, Spotify, and Google Local provide recommendation services to help their customers browse through the enormous product spaces. These recommendation services often play a huge role in controlling the interaction between the two stakeholders, namely, (i) **producers** of goods and services (e.g., movies on Netflix, products on Amazon, restaurants on Yelp, artists on Spotify) and (ii) **customers** who consume them. Maximization of customer satisfaction has been the traditional focus on these platforms, which is often achieved by tailoring the recommendations according to the personalized preferences of individual customers, largely ignoring the interest of the producers. Several recent studies have shown how such customer-centric designs may undermine the well-being of the producers [1, 20, 27, 36, 38]. As more and more people are depending on two-sided platforms to earn a living, recently platforms have also started showing interest in creating fair marketplaces for all the stakeholders due to multiple reasons: (i) **legal obligation** (e.g., labor bill for the welfare of drivers on Uber and Lyft [50], fair marketplace laws for e-commerce [66]), (ii) **social responsibility or voluntary commitment** (e.g., equality of opportunity to all gender groups in LinkedIn [32], commitment of non-discrimination to hosts and guests by AirBnb [5]), (iii) **business requirement/model** (e.g., minimum business guarantee by AirBnb to attract hosts [6]).

In this article, our focus is on the fairness of personalized recommendation services deployed on the two-sided platforms. Traditionally, platforms employ various state-of-the-art data-driven methods (e.g., neighborhood-based methods [51], latent factorization methods [43, 46], etc.) to estimate the relevance scores of every product-customer pairs, and then recommend k most relevant products to the corresponding customers. While such top- k recommendations achieve high customer utility, our investigation on real-world datasets reinforces the presence of popularity bias [40] that is, they can create a huge disparity in the exposure of the producers (detailed in Section 4.2), which is unfair for the producers, and may also hurt the platforms in the long term.

In these platforms, *exposure often determines the economic opportunities (revenues) for the producers* who depend on it for their livelihood. For instance, high exposure on *Google Maps* can increase the footfall in a local business, thereby increasing their revenue. High exposure on *YouTube*, *Spotify*, or *Last.fm* can increase the traffic to a content producer's channel and hence help them earn better platform royalties or advertisement revenues. However, if only a few producers get most of the exposure, then the other producers would struggle on the platform, which could force them to either quit or switch to other platforms [25, 49, 60]. This, in turn, may limit the choices for the customers, degrading the overall experience on the platform. Thus, it is important to reduce exposure inequalities. However, extremely producer-centric ways of reducing inequality (e.g., recommending the k least exposed producers to the customers) may result in loss and disparity in customer utilities (Section 4.2), making it inefficient as well as unfair to the customers.

To counter such unfairness for both producers and customers, we propose to tackle the challenging task of ensuring two-sided fairness while providing personalized recommendations. Specifically, we propose to ensure a minimum exposure guarantee for every producer such that no producer starves for exposure. Since the exposure guarantee on the producer side could incur losses on the customer side (i.e., reduction in customer utilities), we propose that the loss in utility should be fairly distributed among the customers. Motivated by a vast literature in social choice theory, we

map this problem to the problem of fairly allocating indivisible goods (Section 5). In an allocation problem, there is often a predefined set of items and a set of agents with their valuations (how much an agent values an item), and the task is to allocate the items among the agents. To map our recommendation problem to an allocation problem, we assume the set of customers as the set of agents. Now we can strategically fix the set of items as the one that contains as many copies of each product (or producer) as the chosen exposure guarantee, and then allocate them among the customers. If we have an algorithm that does this task, then the strategic setting of the item-set and their allocation guarantee can, in turn, ensure minimum exposure for the producers. Besides, the algorithm's fairness guarantee for the agents during the allocation would also be able to guarantee customer fairness. Thus, the original recommendation problem becomes an interesting (constrained) extension to the existing fair allocation problem—find an allocation that guarantees minimum exposure (upper bounded by **maximin share of exposure (MMS)**) for the producers, and **envy-free up to one item (EF1)** [19] for the customers.¹ We propose an algorithm *FairRec* (Section 6), which solves this problem and gives guarantees on both producer and customer side (proofs in Section 7). Extensive evaluations over multiple real-world datasets show the effectiveness of FairRec in ensuring two-sided fairness while incurring a marginal loss in recommendation quality.

In summary, we make the following contributions in this article.

- We consider a two-sided fair recommendation problem that not only relates to social or judicial precepts but also to the long-term sustainability of two-sided platforms (Section 4).
- We design an algorithm, *FairRec* (Section 6), exhibiting the desired two-sided fairness by mapping the fair recommendation problem to a fair allocation problem (Section 5). Moreover, it is agnostic to the specifics of the data-driven model (that estimates the product-customer relevance scores), which makes it scalable and easy to adapt.
- In addition to the theoretical guarantees (Section 7), extensive experimentation and evaluation over multiple real-world datasets deliver strong empirical evidence on the effectiveness of our proposal (Section 8).
- Finally, we also present a modified version of FairRec (named as FairRecPlus) that uses an envy-cycle elimination and swapping technique to improve the performance on customer-side metrics, while maintaining the same two-sided fairness guarantees (Section 9).

2 BACKGROUND AND RELATED WORK

We briefly survey related works in two directions: (i) fairness in multi-stakeholder platforms and (ii) fair allocation of goods.

2.1 Fairness in Two-sided Platforms

With the increasing popularity of multi-sided platforms, recently researchers have looked into the issues of unfairness and biases in such platforms. For example, Edelman et al. [27] investigated the possibility of racial bias in guest acceptance by Airbnb hosts, Lambrecht and Tucker [45] studied gender-based discrimination in career ads, Chakraborty et al. [24] proposed to ensure fair representation in crowdsourced recommendations. While these works deal with *group fairness*, Serbos et al. [58] proposed an *envy-free* tour package recommendations on travel booking sites, ensuring *individual fairness for customers*.

On producer fairness, Hannák et al. [38] studied racial and gender bias in freelance marketplaces, and Dash et al. [26] investigated favoritism toward certain producers on e-commerce marketplace.

¹The MMS guarantee ensures that each agent receives a value that is at least their *maximin share* threshold, defined in Equation (1); whereas, EF1 ensures that every agent values her allocation at least as much as any other agent's allocation after (hypothetically) removing the most valuable item from the other agent's allocated bundle.

In a social experiment, Salganik et al. [57] found that the existing popular producers often acquire most of the visibility while new but good ones starve for visibility. Banerjee et al. [9] also found popularity bias in location-based recommendations. Kamishima et al. [41] and Abdollahpouri et al. [2] proposed methods to reduce such popularity bias among producers. Süner et al. [64] proposed to maximize total customer utility while ensuring some exposure for the producers. While these works have proposed to ensure some forms of producer-side fairness, they have not looked into the resulting unfairness on the customer-side, the trade-off between producer and customer fairness, and the cost of achieving one over the other.

Few past works have discussed fairness for both producers and customers. Abdollahpouri and Burke [1] and Burke [20] categorized different types of multi-stakeholder platforms and their desired group fairness properties. Chakraborty et al. [23] and Sühr et al. [63] presented mechanisms for two-sided fairness in matching problems, while Patro et al. [53] addressed fairness issues arising due to frequent updates of platforms. In contrast, our article addresses individual fairness for both producers and customers, which also answers the question of the long-term sustainability of two-sided platforms.

There exists another line of work on fairness in ranking and recommendations, namely, Biega et al. [12], Singh and Joachims [59], which propose to ensure the producers with expected exposures in proportion to their corresponding relevance scores to maintain individual fairness for producers in gig-economy platforms. One of the limitations of these works is that they assume the availability of true relevance scores of producers. However, these relevance scores are often estimated and usually contain noise, which is also highlighted in Raj et al. [56]. The noise is not the only issue here; if the estimated relevance scores themselves exhibit popularity bias, then ensuring exposure in proportion to these relevance scores could cause the same inequalities in producer exposures and can do very little toward producer fairness. Thus, in this work, we try to isolate the considerations of exposure and relevance. Formally, we use the exposure of a producer as the measure of its utility, the relevance of recommended items as the utility of a customer (more details in Section 3), and finally, we define fairness for both sides using the (in)equality in their individual utilities.

2.2 Fair Allocation of Goods

The problem of fair allocation (popularly known as the cake-cutting problem) has been studied extensively in the area of computational social choice theory. The classical notions of fairness for this problem are **envy-freeness (EF)** [31, 65] and **proportional fair share (PFS)** [61]. Recent literature on practical applications of fair allocation [18, 28] has focused on the problem of allocating *indivisible* goods in budgeted course allocation [19], balanced graph partition [16], or allocation of cardinality constrained group of resources [14]. In such instances, no feasible allocation may satisfy EF or PFS fairness guarantees. Thus, the notable work of Budish [19] defined analogous fairness notions that are appropriate for *indivisible* goods, namely, envy-freeness up to one good (EF1) and maximin share guarantee (MMS).

The relevance of EF1 is substantiated by the fact that it is guaranteed to always exist under general monotone valuations and, in fact, such allocations can be obtained in polynomial time [48]. When the valuations are additive, Caragiannis et al. [22] show that a simple greedy round robin algorithm is enough to ensure EF1.

MMS fairness is another solution concept that has been extensively studied in the fair allocation space. In particular, Bouveret et al. [17] showed that an MMS allocation exists when the agents' valuations are additive and binary (valuations are 0 or 1). However, Procaccia and Wang [55] and Kurokawa et al. [44] provided intricate counterexamples to refute the universal existence of MMS allocations, under additive and non-binary valuations. This motivated the study of

approximate maximin share allocations, α -MMS, where each agent obtains a bundle of value at least $\alpha \in (0, 1)$ times her maximin share. The existence of $2/3$ -MMS and accompanying algorithms were developed in a sequence of results [8, 11, 55]. Later, Ghodsi et al. [33] improved the result by providing an efficient algorithm that obtains $3/4$ -MMS allocations.

The vast majority of work in the fair allocation space has solely focused on the unconstrained version of the problem; exceptions include the work of Biswas et al. [14, 15] and Gourvès et al. [34, 35]. Biswas et al. [14] provide algorithms for computing EF1 and $1/3$ -MMS for the allocation problem where items are categorized into groups, and an upper bound restricts the number of items that can be allocated to each agent from each category. This is slightly different from the problem we consider (detailed in Section 5.3). A general version of the category-wise upper bound constraint, namely, laminar matroid constraint, is studied by Biswas et al. [15] and the existence of EF1 is proved for identical valuations. A different problem is considered by Gourvès et al. [34, 35] where the goal is to find MMS fair allocation that *union* of all the allocated goods is an independent set of a given matroid. Although these papers study fair allocation under several combinatorial constraints, they do not directly apply to the problem we consider.

Moreover, all the above-mentioned papers consider fairness among the agents but not among the items. In this work, we consider fairness across the agents as well as the items. We map the problem of fair recommendation to a fair allocation problem, which leads to an interesting extension of previously studied problems owing to the specific constraints pertaining to recommendations (detailed in Section 5.3).

This article is an extended version of our earlier work, titled “FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms” [52]. In Reference [52], we introduced the notions of two-sided fairness in recommendations and proposed the *FairRec* algorithm to ensure fairness for both producers and customers. In this work, we prove that FairRec ensures a stronger theoretical guarantee on the producer-side. Additionally, since we provide a tunable parameter α to the platforms to regulate the minimum exposure guarantee for the producers, we present a detailed analysis on what actually happens when the value of α is changed and its impact on both producer-side and customer-side. In certain scenarios, the platforms might be interested in ensuring different levels of exposure guarantee for different producers; for example, the platform may want to give more exposure guarantee to higher-rated producers than the lower-rated ones. Thus, we also evaluate FairRec in such scenarios by tweaking the FairRec algorithm to entertain such requirements. We add several new baselines, such as MPB19 [3], MSR18 [64], MixedTR- k , MixedTP- k , to compare against FairRec. Finally, we propose a new modification of FairRec (named as FairRecPlus) that utilizes the envy graph to improve the recommendation performance for the customers, but comes at a cost of increased time complexity.

3 PRELIMINARIES

In this section, we define the terminology and notations used throughout the article.

3.1 Products and Producers

In a few two-sided platforms focusing on physical establishments (e.g., Google Maps, Yelp), a producer typically owns one product (e.g., restaurant or shop); whereas in multimedia platforms like Spotify, YouTube, and Netflix, an artist can produce multiple songs or videos; the same is also true for ecommerce platforms like Amazon and Flipkart, where one producer can list many products. To generalize our approach to both types of two-sided platforms, we consider products and producers to be equivalent, and use the terms “product” and “producer” interchangeably. Even for platforms where a producer can have multiple products, ensuring fairness at the product level can

ensure fairness for individual producers—where fairness can be ensured by making the exposure proportional to the producer’s portfolio size.²

3.2 Notations

Let U and P be the sets of customers and producers, respectively, where $|U| = m$, and $|P| = n$. Let k be the number of products to be recommended to every customer. $R_u \subset P$ represents the set of k products recommended to customer u ; $|R_u| = k$.

3.3 Relevance of Products

The *relevance* of a product p to customer u , denoted as $V_u(p)$, represents the likelihood that u would like the product p . Formally, relevance is a function from the set of customers and products to the real numbers $V : U \times P \rightarrow \mathbb{R}$. Usually, the relevance scores are predicted using various data-driven methods (e.g., neighborhood-based methods [51], latent factorization methods [43, 46], etc.), and $V_u(p)$ is a proxy for the utility gained by u if product p is recommended to her.

3.4 Customer Utility

The utility of a recommendation R_u to a customer u is proportional to the sum of relevance scores of products in R_u . Thus, recommending the k most relevant products will give the maximum possible utility. Let R_u^* be the set of top- k relevant products for u . We use a normalized form of customer utility from R_u , defined as: $\phi_u(R_u) = \frac{\sum_{p \in R_u} V_u(p)}{\sum_{p \in R_u^*} V_u(p)}$.

3.5 Producer Exposure

Exposure of a producer/product p is the total amount of attention that p receives from all the customers to whom p has been recommended. In this article, we assume a uniform attention model³ where customers pay similar attention to all k recommended products, and express the exposure of a product p as $E_p = \sum_{u \in U} \mathbb{1}_{R_u}(p)$, where $\mathbb{1}_{R_u}(p)$ is 1 if $p \in R_u$, and 0 otherwise. The sum of exposures of all the products is $\sum_{p \in P} E_p = m \times k$. Note that we assume the relevance of a product does not play any role in producer’s utility (in contrast to Biega et al. [12], Singh and Joachims [59]), and use only the exposure of a producer as her utility.

4 NEED FOR TWO-SIDED FAIRNESS IN PERSONALIZED RECOMMENDATIONS

Traditionally, the goal of personalized recommendation has been to recommend products that would be most relevant to a customer. This task typically requires learning the relevance scoring functions (V), and several state-of-the-art data-driven methods [42, 46, 51, 67, 68] have been developed to estimate the product-customer relevance values. Once these values are obtained, the standard practice, across several recommender systems, is to recommend the top- k (k =size of recommendation) relevant products to corresponding customers. While this approach attempts to maximize the satisfaction of individual customers, it can adversely affect the producers in a two-sided platform, as we explore next.

²In platforms with producers having multiple products, ensuring producer-side fairness other than proportionality remains open for future work.

³There can be more sophisticated attention models considering *position bias* [4], where customers pay more attention to the top ranked products than the lower ranked ones. This being an initial work on two-sided fair recommendation (formulated as a fair-allocation problem), we focused on a basic model setting without position bias.

4.1 Datasets

We consider the impact of customer-centric top- k recommendations on the exposure of the producers using real-world datasets. We use a state-of-the-art relevance scoring model (a very widely used latent factorization method [46]) and also a dataset-specific custom relevance scoring model over the datasets.

4.1.1 Google Local Ratings Dataset (GL). Google Local is a service to find nearby shops, restaurants on Google Maps (as Google Nearby feature) platform. We use the Google Local dataset released by He et al. [39], which contains data about customers, local businesses (producers), and their locations (geographic coordinates), ratings, and so on. We consider the active customers located in New York City and the business entities within 5 miles radius of Manhattan area with at least 10 reviews. The resulting dataset contains 11,172 customers, 855 businesses and 25,686 reviews. We consider the following two relevance scoring functions (V).

- (1) **GL-CUSTOM:** We use a custom relevance scoring function: $V(u, p) = \frac{rating(p)}{distance(u, p)}$, where $rating(p)$ is the average rating of the producer (local business) and $distance(u, p)$ is the distance between customer u and producer p .
- (2) **GL-FACT:** Here, we use the state-of-the-art latent factorization model [43, 46] to predict the relevance scores from the ratings.

From here on, we refer to the above two datasets (relevance score sets) as GL-CUSTOM and GL-FACT, respectively.

4.1.2 Last.fm Dataset (LF). We use the `Last.fm` dataset released by Cantador et al. [21], which contains 1,892 customers, 17,632 artists (producers), and 92,834 records of play counts (the number of times a customer has played songs from an artist). We again use a latent factorization model [43, 46] to find out the relevance scores from the play counts. From here on, we refer to this dataset as **LF**.

4.2 Adverse Impact of Customer-Centric Recommendation

We simulate top- k ($k = 20$) recommendation on all three datasets, and calculate the exposure different producers get. Figures 1(a), 1(b), and 1(c) are the Lorenz curves for producer exposures. In Exposure Lorenz curves, the cumulative fraction of total exposure is plotted against the cumulative fraction of the number of corresponding producers (ranked in increasing order of their exposures). The extent to which the curve goes below a straight diagonal line (or an equality mark) indicates the degree of inequality in the exposure distribution. We observe that the Lorenz curves for top- k recommendations are far below the equal exposure marks, revealing that for conventional top- k recommendation, 50% least exposed producers get only 32%, 5%, and 11% of total available exposure ($m \cdot k$) in *GL-CUSTOM* (using the ratio of rating to distance to estimate relevance), *GL-FACT* (using latent factorization model [46]), and *LF* datasets (using latent factorization model [46]), respectively.

Huge disparity is in nobody's interest: In two-sided platforms, the exposure determines the economic opportunities. Thus, low exposure on a platform often puts many producers at huge losses, forcing them to leave the platform; this may result in fewer available choices for the customers, thereby degrading the overall quality of the platform. Thus, highly skewed exposure distribution of the customer-centric top- k recommendation not only makes it unfair to the producers but also questions the long-term sustainability of the platforms. Thus, there is a need to be fair to the producers while designing recommender systems.

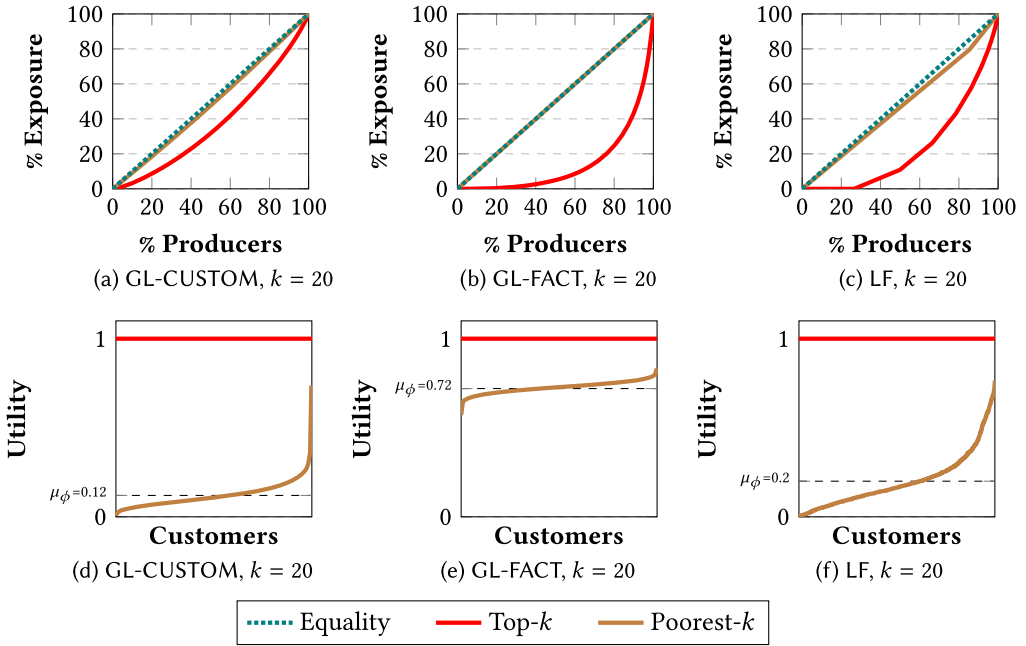


Fig. 1. Lorenz curves (a–c) show high inequality among producer exposures with the top- k recommendation (using the ratio of rating to distance to estimate relevance in GL-CUSTOM, and using latent factorization model [46] in GL-GACT and LF). In these curves, the cumulative fraction of total exposure is plotted against the cumulative fraction of the number of corresponding producers (ranked in increasing order of their exposures). The extent to which the curve goes below a straight diagonal line (or an equality mark) indicates the degree of inequality in the exposure distribution. In panels (d–f), we plot the individual customer utilities sorted in ascending order while μ_ϕ represents the mean customer utility. While poorest- k provides almost equal exposures, it introduces huge loss and disparity in individual customer utilities (d–f).

4.3 Pitfalls of Producer-centric Solution

A naive approach to reduce inequality in producer exposures is to implement a producer-centric recommendation (poorest- k): recommend the least- k exposed products to the customer at any instant. Such producer-centric scheme makes the exposure of all the producers nearly equal, as seen in Figures 1(a), 1(b), and 1(c): Lorenz curves for poorest- k recommendations are closer to the diagonal than those of top- k recommendations. However, such poorest- k recommendation decreases overall customer utilities (as seen in Figures 1(d), 1(e), and 1(f)). Moreover, the poorest- k introduces disparity among individual customer utilities, where some customers may suffer much higher losses than other customers, making it unfair to them.

4.4 Desiderata of Fair Recommendation

To counter the above-mentioned issues, in this work, we propose the following fairness properties to be satisfied by the recommendation to be fair to both producers and customers.

A. Producer Fairness: Mandating a uniform exposure distribution over the producers can be too harsh on the system; it may heavily hamper the quality of the recommendation, and might also kill the existing competition by discouraging the producers from improving the quality of their products or services. Instead, we propose to **ensure a minimum exposure guarantee for**

every producer such that no producer starves for exposure. The proposal is comparable to the fairness of *minimum wage guarantee* (e.g., as required by multiple legislations in the US, starting from Fair Labor Standards Act 1938 to Fair Minimum Wage Act 2007 [30, 37, 54]). Ensuring minimum wage does not itself guarantee equality of income; however, it has been found to decrease income inequality [29, 47]. Similarly, we want to ensure minimum exposure to every producer in the system. Note that we are not mandating a fixed exposure guarantee—the exact value of the guaranteed minimum exposure can be decided by the respective platforms.

B. Customer Fairness: As maintaining producer fairness can cause an overall loss in customer utility, we propose that **the loss in utility should be fairly distributed among the customers**. To ensure this, products need to be recommended in a way such that no customer can gain extra utility by exchanging her set of recommended products with another customer—a property called *envy-freeness*, as detailed in the next section.

5 RE-IMAGINING FAIR RECOMMENDATION AS FAIR ALLOCATION

Given a set of items (say, \mathcal{P}), a set of agents (say, \mathcal{U}), and valuations \mathcal{V} (how much an agent values an item), the **fair allocation problem** aims at distributing the items *fairly* among the agents. In the discrete version of this problem, the items are *discrete* (no item can be broken into pieces) and *non-shareable* (no item can be allocated to multiple agents). If \mathcal{P} contains several copies of the same item, then each copy can be thought of as non-shareable and discrete. The goal is to find a non-shareable and discrete allocation ($\mathcal{A} := \{(A_u)_{u \in \mathcal{U}} : A_u \subseteq \mathcal{P}\}$) while ensuring fairness properties.

5.1 Notions of Fairness in Allocation

The classical fairness notions, such as envy-freeness⁴ (EF) and proportional-fair-share⁵ (PFS), may not be achievable in most instances of the problem. For example, if there are two agents and one item, then the item will be allocated to one of the agents, and the zero allocation to the other agent would violate both EF and PFS. Thus, for *discrete* items, relaxations of EF and PFS have been considered. Two such well-studied notions of fairness in the discrete fair allocation literature are (i) *envy freeness up to one item* (EF1) and (ii) the *maximin share guarantee* (MMS), defined by Budish [19]. Since then, these have been extensively studied in various settings for providing existential and algorithmic guarantees [7, 8, 10, 13–18, 22, 44, 55]. We now formally state these fairness notions:

- An allocation \mathcal{A} is EF1 iff for every pair of agents $u, w \in \mathcal{U}$ there exists an item $p \in A_w$ such that $\mathcal{V}_u(A_u) \geq \mathcal{V}_u(A_w \setminus \{p\})$.
- An allocation is said to satisfy MMS if each agent receives a value greater than or equal to their *maximin share* threshold. This threshold for an agent u is defined as

$$\text{MMS}_u = \max_{\mathcal{A}} \min_{w \in \mathcal{U}} \mathcal{V}_u(A_w). \quad (1)$$

In other words, MMS_u is the maximum value that the agent can guarantee for herself if she were to allocate \mathcal{P} into $|\mathcal{U}|$ bundles and then, from those bundles, receive the minimum valued one. Formally, an allocation \mathcal{A} satisfy MMS fairness iff for all agents $u \in \mathcal{U}$, we have $\mathcal{V}_u(A_u) \geq \text{MMS}_u$.

⁴An allocation is said to satisfy *envy-freeness* if the bundle of items allocated to each agent is as valuable to her as the bundle allocated to any other agent [31, 62, 65].

⁵An allocation is said to satisfy *proportional-fair-share* if each agent receives a bundle of value at least $1/|\mathcal{U}|$ th of her total value for all the items [61].

5.2 Fair Recommendation to Fair Allocation

We propose to see the desired two-sided recommendation problem as a *fair allocation* problem. The set of products P can be thought of as the set of items \mathcal{P} (there can be multiple copies of individual products)⁶; similarly, the set of customers U as the set of agents \mathcal{U} , and the relevance scoring function V as the valuations \mathcal{V} . Now the task of recommending products to customers is the same as allocating items in \mathcal{P} to agents in \mathcal{U} with certain constraints.

- **Setting \mathcal{P} for Producer Fairness:** As the total exposure of the platform is limited ($k \cdot |U|$), the maximum guarantee on minimum possible exposure for the producers is $\lfloor \frac{k \cdot |U|}{|P|} \rfloor$ (this refers to the MMS value for the producers). One way to formally define a lower threshold requirement \bar{E} is by using the notion of *approximate maximin share* (α -MMS). More formally,

Definition 5.1. An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is said to satisfy α -MMS for a fixed value $\alpha \in (0, 1]$ if and only if, for all agents $i \in \{1, \dots, n\}$, the following holds:

$$\mathcal{V}_i(A_i) \geq \alpha \text{MMS}_i, \quad \text{where} \quad \text{MMS}_i = \max_{\mathcal{A}} \operatorname{argmin}_{j \in \{1, \dots, n\}} \mathcal{V}_i(A_j).$$

A platform can decide the exact value of α and provide the minimum exposure guarantee $\bar{E} = \alpha \text{MMS}$ for every producer.

- **Fair Allocation of \mathcal{P} among \mathcal{U} :** Once \mathcal{P} is set according to the desired producer fairness, the entire task of fair recommendation boils down to the allocation of \mathcal{P} among \mathcal{U} while ensuring EF1 for agents/customers (to ensure fairness for customers, as introduced in Section 4.4). However, specific constraints related to the recommendation problem need a novel extension of the traditional fair allocation problem, as explained next.

5.3 Extending the Conventional Fair Allocation Problem

Traditionally, *fair allocation* literature aims at defining and ensuring *fairness* among the agents while allocating all items the set \mathcal{P} exhaustively. However, in the *fair recommendation* problem, along with customer fairness, the challenge is to attain producer (or product) fairness by providing a minimum exposure guarantee (where each product needs to be allocated to at least ℓ different customers). Thus, achieving producer fairness is the same as creating at least ℓ copies of each product and ensuring that all the copies are allocated, along with a feasibility constraint, which enforces that no customer gets more than one copy of the same product. This extension of the problem—where all the items are grouped into disjoint categories and no agent receives more than a pre-specified number of items from the same category—is called cardinality constrained fair allocation problem, proposed in Reference [14]. In this article, we consider a novel extension of the cardinality constrained problem by adding another constraint enforcing that exactly k items are allocated to each customer. This requires tackling hierarchical feasibility constraints—an upper bound cardinality constraint of one on each product and a cardinality constraint of k on the total number of allocated products. Moreover, this additional feasibility constraint makes it difficult to decide how many copies of which product should be made available for a total of ($k \cdot |U|$) allocations, satisfying the feasibility constraints as well as the fairness requirement. Thus, unlike the fair allocation problem, we consider no restriction on the number of copies of each product that are made available. All these contrast points, along with the two-sided fairness guarantees make fair recommendation an interesting extension of the fair allocation problem.

⁶Note that P represents the set of producers or products. However, \mathcal{P} represents the set of items that are to be allocated among the customers; this set of items can be suitably formed by creating and gathering copies of each product based on how much exposure we want to guarantee for that product.

Overall, we aim to recommend a set of k products, denoted as $A_u \subset P$, to each customer $u \in U$ that satisfies the following constraints:

$$\sum_{p \in A_u} V_u(p) \geq \sum_{p \in A_w} V_u(p) - \max_{p \in A_w} V_u(p) \text{ for every pair of customers } u, w \in \mathcal{U}. \quad (2)$$

$$|A_u| = k, \text{ for all customers } u \in U. \quad (3)$$

$$\sum_{p \in P} E_p \geq \left\lfloor \frac{\alpha mk}{n} \right\rfloor, \text{ for each producer } p \in P, \text{ for some } \alpha \in (0, 1]. \quad (4)$$

This formulation leads to a large number of constraints, of the order $|U|^2 + |P|$. Standard Integer Linear Programming techniques for finding a feasible solution, in this large constraint space, do not scale well, especially considering the real-world two-sided platforms. Hence, in the next section, we propose a greedy algorithm to solve the above-mentioned constraint satisfaction problem.

6 FAIRREC: AN ALGORITHM TO ENSURE TWO-SIDED FAIRNESS

In this section, we provide a polynomial-time algorithm *FairRec*, for finding an allocation \mathcal{A} that satisfies the desired two-sided fairness described in Section 4.4 (we prove the theoretical guarantees in Section 7). Note that we consider only the case of $k < |P|$, and leave the trivial case of $k = |P|$ and the infeasible case of $k > |P|$ out of consideration. Also, we consider $|P| \leq k \cdot |U|$, otherwise, at least $(|P| - k \cdot |U|)$ producers can not be allocated to any customer.

FairRec (Algorithm 1) executes in two phases. The first phase ensures EF1 among all the m customers (Lemma 7.2.1) and tries to provide a minimum guarantee on the exposure of the producers (Lemma 7.2.2). However, the first phase may not allocate exactly k products to all the m users, which is then ensured by the second phase while simultaneously maintaining EF1 for customers.

The **first phase** creates $\ell = \lfloor \frac{\alpha mk}{n} \rfloor$ copies of each product. Note that $\lfloor \frac{mk}{n} \rfloor$ is the maximin value of any producer when mk slots are allocated among n producers, and thus ℓ represents the α -MMS value for each producer. The algorithm then initializes each component of the vector S of size $|P|$ to ℓ to ensure that at most ℓ copies from each product are allocated in the first phase. Feasible sets F_u for each customer u are then initialized to ensure that each customer receives at most one copy of the same product. Then, assuming an arbitrary ordering σ of customers, Algorithm 2 is executed and the allocation \mathcal{B} is obtained.

The **second phase** checks if all the customers have received exactly k products (by looking at the number of products allocated to the customer $x + 1$, which is next-in-sequence to the last allocated customer x of the first phase). If the customer $x + 1$ has received k products, then no further allocation is required; if not, then Algorithm 2 is called again with a new ordering obtained by x left-cyclic rotations of σ . The remaining number of items is stored in T , which are to be allocated among the customers. Also, each component of the vector S is updated to $|U|$ to allow allocating any feasible product without any limit on the available number of copies. The second phase retains EF1 fairness among the customers.

Both phases use a modified version of the Greedy-round-robin (Algorithm 2) [14, 22]: it follows the ordering σ in a round-robin fashion (i.e., it selects customers, one after the other, from $\sigma(1)$ to $\sigma(m)$), and iteratively assigns to the selected customer her most desired unallocated product (feasibility maintained by the vector S and sets F_u and ties are broken arbitrarily). This process is repeated over several rounds until one of the two disjoint conditions occur: (i) $T == \emptyset$: a total of T allocations have occurred, or (ii) $p == \emptyset$: no feasible product available (for the current customer $\sigma(i)$, we have $F_{\sigma(i)} \cap \{p : S_p \neq 0\} = \emptyset$). Finally, it returns an allocation B_1, \dots, B_m with each $B_u \subseteq [n]$ for all $u \in [m]$.

ALGORITHM 1: *FairRec* (U, P, k, V)

Input: Set of customers $U = [m]$, set of distinct products $P = [n]$, recommendation set size k (such that $k < n$ and $n \leq k \cdot m$), and the relevance scores $V_u(p)$.

Output: A two-sided fair recommendation.

- 1: Initialize allocation $\mathcal{A}^0 = (A_1^0, \dots, A_m^0)$ with $A_i^0 \leftarrow \emptyset$ for each customer $i \in [m]$.

First Phase:

- 2: Fix an (arbitrary) ordering of the customers $\sigma = (\sigma(1), \sigma(2), \dots, \sigma(m))$.
- 3: Initialize set of feasible products $F_u \leftarrow P$ for each $u \in [m]$.
- 4: Set $\ell \leftarrow \lfloor \frac{\alpha m k}{n} \rfloor$ denoting number of copies of each product.
- 5: Initialize each component of the vector $S = (S_1, \dots, S_n)$ with $S_j \leftarrow \ell, \forall j \in [n]$, this stores the number of available copies of each product.
- 6: Set $T \leftarrow \ell \times n$, total number of items to be allocated.
- 7: $[\mathcal{B}, F, x] \leftarrow \text{Greedy-Round-Robin}(m, n, S, T, V, \sigma, F)$.
- 8: Assign $\mathcal{A} \leftarrow \mathcal{A} \cup \mathcal{B}$.

Second Phase:

- 9: Set $\Lambda = |A_{\sigma((x) \pmod{m} + 1)}|$ denoting the number of items allocated to the customer subsequent to x , according to the ordering σ .
- 10: **if** $\Lambda < k$ **then**
- 11: Update each component of the vector $S = (S_1, \dots, S_n)$ with the value m to allow allocating any product to any customer.
- 12: Set $T \leftarrow 0$.
- 13: **if** $x < m$ **then**
- 14: Set $\sigma'(i) \leftarrow \sigma((i + x - 1) \pmod{m} + 1)$ for all $i \in [m]$.
- 15: $\sigma \leftarrow \sigma'$.
- 16: $T \leftarrow (m - x)$.
- 17: Update $\Lambda \leftarrow \Lambda + 1$.
- 18: **end if**
- 19: $T \leftarrow T + m(k - \Lambda)$ total number of items to be allocated.
- 20: $[C, F, x] \leftarrow \text{Greedy-Round-Robin}(m, n, S, T, V, \sigma, F)$.
- 21: Assign $\mathcal{A} \leftarrow \mathcal{A} \cup C$.
- 22: **end if**
- 23: Return \mathcal{A} .

Note that while we choose minimum exposure guarantee and EF1 as fairness constraints for producers and customers, respectively, it does not necessarily mean that merely satisfying them is enough. Apart from ensuring fairness, we also need to perform well in what any recommender system is originally designed to do, i.e., to provide good (relevant) recommendations to the customers. Thus, in Algorithm 2, which allocates products to customers in greedy-round-robin manner, we try to allocate (in step-7) the best product that is available and feasible in every round. We believe this is the reason why FairRec shows good performance in overall customer utility as well (based on experimental results in Section 8).

7 THEORETICAL GUARANTEES

In this section, we provide a few important properties of Algorithm 2 in Proposition 7.1. Later, we establish the fairness guarantees and time complexity of our proposed algorithm *FairRec* in Theorem 7.2 using Lemmas 7.2.1, 7.2.3, and 7.2.4.

ALGORITHM 2: Greedy-Round-Robin (m, n, S, T, V, σ, F)

Input : Number of customers m , number of producers n , an array with number of available copies of each product S , total number of available products $T > 0$, relevance scores $V_u(p)$ and feasible product set F_u for each customer, and an ordering σ of $[m]$.

Output: An allocation of T products among m customers, the residual feasible set F_u and the last allocated index x .

```

1: Initialize allocation  $\mathcal{B} = (B_1, \dots, B_m)$  with  $B_i \leftarrow \emptyset$  for each customer  $i \in [m]$ .
2: Initiate  $x \leftarrow m$ .
3: Initiate round  $r \leftarrow 0$ .
4: while true do
5:   Set  $r \leftarrow r + 1$ .
6:   for  $i = 1$  to  $m$  do
7:     Set  $p \in \arg \max_{p' \in F_{\sigma(i)}: (S_p \neq 0)} V_{\sigma(i)}(p')$ 
8:     if  $p == \emptyset$  then
9:       Set  $x = i - 1$  only if  $i \neq 1$ .
10:    go to Step 22.
11:   end if
12:   Update  $B_{\sigma(i)} \leftarrow B_{\sigma(i)} \cup p$ .
13:   Update  $F_{\sigma(i)} \leftarrow F_{\sigma(i)} \setminus p$ .
14:   Update  $S_p \leftarrow S_p - 1$ .
15:   Update  $T \leftarrow T - 1$ .
16:   if  $T == 0$  then
17:      $x = i$ .
18:   go to Step 22.
19:   end if
20: end for
21: end while
22: Return  $\mathcal{B} = (B_1, \dots, B_m)$ ,  $F = (F_1, \dots, F_m)$  and index  $x$ .

```

PROPOSITION 7.1. *The allocation obtained by the Greedy-Round-Robin (Algorithm 2) exhibits the following four properties:*

- (P_1) for any two indices x and y , where $x < y$, the customer $\sigma(x)$ (who appears earlier than $\sigma(y)$) according to the ordering σ) does not envy customer $\sigma(y)$, i.e., $V_{\sigma(x)}(B_{\sigma(x)}) \geq V_{\sigma(x)}(B_{\sigma(y)})$.
- (P_2) the allocation \mathcal{B} obtained by Algorithm 2 is EF1.
- (P_3) each customer is allocated at most one item from the same producer, thus ensuring the cardinality constraint is satisfied for each producer (category).
- (P_4) for any two customers, say u and w , the allocation \mathcal{B} obtained by Algorithm 2 satisfies the following: $-1 \leq (|B_u| - |B_w|) \leq 1$.

PROOF. The properties P_1 and P_2 have been observed by Biswas and Barman [14] and Caragianis et al. [22], respectively. For completeness, we repeat the arguments toward these two properties. Let x and y be two indices, such that $1 \leq x < y \leq m$. At each round r , the customer $\sigma(x)$ chooses her most desired product among all the unallocated items before customer $\sigma(y)$. Hence, $V_{\sigma(x)}(p_{\sigma(x)}^r) \geq V_{\sigma(x)}(p_{\sigma(y)}^r)$, where $p_{\sigma(x)}^r$ and $p_{\sigma(y)}^r$ denote the items assigned to customer $\sigma(x)$ and $\sigma(y)$, respectively. Thus, over all the rounds, $\sum_r V_{\sigma(x)}(p_{\sigma(x)}^r) \geq \sum_r V_{\sigma(x)}(p_{\sigma(y)}^r)$, which implies that $V_{\sigma(x)}(B_{\sigma(x)}) \geq V_{\sigma(x)}(B_{\sigma(y)})$ and thus the property P_1 holds.

Property P_2 states that if $\sigma(y)$ envies $\sigma(x)$, it will not violate EF1 property (note: we already saw in P_1 that $\sigma(x)$ does not envy $\sigma(y)$). Now observe that, the value $V_{\sigma(y)}(\cdot)$ of the item allocated to customer $\sigma(y)$ in the r th round is at least that of the item allocated to customer $\sigma(x)$ in the $(r+1)$ th round. Let R denote the total number of rounds, then the following holds:

$$\begin{aligned}
& V_{\sigma(y)}(p_{\sigma(y)}^r) \geq V_{\sigma(y)}(p_{\sigma(x)}^{r+1}) \quad \text{for all } r \in \{1, \dots, R-1\} \\
\Rightarrow & \sum_{r=1}^{R-1} V_{\sigma(y)}(p_{\sigma(y)}^r) \geq \sum_{r=1}^{R-1} V_{\sigma(y)}(p_{\sigma(x)}^{r+1}) \\
\Rightarrow & V_{\sigma(y)}(B_{\sigma(y)}) \geq V_{\sigma(y)}(B_{\sigma(x)}) - V_{\sigma(y)}(p_{\sigma(x)}^1). \tag{5}
\end{aligned}$$

Equation (5) shows that the customer $\sigma(y)$ stops envying $\sigma(x)$ when only one item is (hypothetically) removed from $B_{\sigma(x)}$ (namely, $p_{\sigma(x)}^1$). Thus, the allocation \mathcal{B} is EF1, i.e., P_2 holds.

The property P_3 is satisfied by the use of the feasible sets F_u for each customer u . Each F_u contains the set of producers who have not yet been allocated to the customer u . At any round r , step 7 of Algorithm 2 selects the most relevant producer among the producers who had not been allocated to u in any earlier rounds $r' < r$. Once, a producer p is allocated to a customer u , step 9 of Algorithm 2 removes p from F_u . Thus, each customer is allocated at most one item from the same producer.

The property P_4 states that, for any pair of customers u and w , the number of allocated items $|B(u)|$ and $|B(w)|$, differ by at most 1. It is straightforward to see that, except for the last feasible round, all customers are allocated exactly one item at each round. Thus, all the customers receive the same number of allocations until the second last feasible round. In the last feasible round, some customers may not get any allocation (if there is no available feasible product) and thus may receive one item less than the others. \square

We now state the main theorem (Theorem 7.2) that establishes the fairness guarantees of our proposed algorithm.

THEOREM 7.2. *Given n producers, the proposed polynomial time algorithm, FairRec, returns an EF1 allocation among m customers while allocating exactly k items to each customer, when $k < n \leq mk$. Moreover, it ensures non-zero exposure among all the n producers and α -MMS guarantee among at least $(1 - \frac{1}{m+1} \lfloor \frac{\alpha mk}{n} \rfloor)$ fraction of the producers.*

PROOF. We prove the fairness guarantees of customers and producers in Lemmas 7.2.1 and 7.2.3, respectively. In Lemma 7.2.4, we show that FairRec executes in polynomial time. \square

LEMMA 7.2.1. *Given n producers, m customers, and a positive integer k (such that $k < n \leq mk$), FairRec returns an EF1 allocation among m customers while allocating exactly k items to each customer.*

PROOF. To prove this, we show that both phases of FairRec satisfy EF1. Since Algorithm 2 guarantees EF1 (by property P_2), the allocation \mathcal{A} at step 9 of FairRec is EF1. Thus, for any two customers u and w , there exists an item $j \in B_w$ such that $V_u(B_u) \geq V_u(B_w) - V_w(j)$. Next, the second phase creates $|U|$ copies of each product and calls Algorithm 2 to obtain the allocation \mathcal{C} . Note that the second phase assigns the most valued item to each customer at each round, that is, it allocates top- Λ_u feasible producers to each customer, where $\Lambda_u = k - |B_u|$. Thus, $V_u(C_u) \geq V_u(C_w)$. Thus, $V_u(B_u) + V_u(C_u) \geq V_u(B_w) - V_w(j) + V_u(C_w)$, which implies EF1: $V_u(B_u \cup C_u) \geq V_u(B_w \cup C_w) - V_w(j)$. This completes the proof that FairRec ensures EF1 among all the customers while recommending exactly k products to each customer. \square

We now establish the fairness guarantees of *FairRec* for the producers. For ease of exposition, we first consider $\alpha = 1$ to show exact MMS fairness guarantees are satisfied by at least $(1 - \frac{n}{k})$ fraction of the producers (in Lemma 7.2.2). Subsequently, we provide a stronger guarantee in Lemma 7.2.3 considering any $\alpha \in (0, 1]$.

LEMMA 7.2.2. *Given n producers, m customers, a positive integer k (such that $k < n \leq mk$), and $\alpha = 1$, FairRec ensures non-zero exposure among all the n producers. Moreover, it assures MMS-fairness among at least $n - k$ producers.*

PROOF. We first prove that the first phase guarantees non-zero exposure for producers. The allocation \mathcal{B} obtained by Algorithm 2 in the first phase may have terminated for one of the two conditions

- (1) $T == 0$: this means that all the $\ell = \lfloor \frac{mk}{n} \rfloor$ copies of each producer have been allocated among all the customers. Thus, each producer receives exactly the maximin threshold ℓ . Hence, MMS fairness is achieved by all the n producers.
- (2) $p == \emptyset$: this happens when $T \neq 0$ and $\sum_{p \in F_u} S_p = 0$ for a customer u (at termination). That is, $S_p = 0$ for each producer $p \in F_u$. Thus, all $\ell = \lfloor \frac{mk}{n} \rfloor$ copies of the producers in the set F_u have been allocated, and hence they attain MMS fairness. However, the producers in the set B_u (the set recommended to customer u) is allocated to at least one producer. Thus, a minimum value of 1 is achieved by all the n producers. Also, $|F_u| + |B_u| = n$ and $|B_u| \leq k$, implies that $|F_u| \geq n - k$. Therefore, at least $n - k$ producers attain MMS-fairness.

Since the thresholds are already satisfied in the first phase, adding more allocations in the second phase retains the threshold-based fairness guarantees. This completes the proof that FairRec ensures a non-zero exposure among all the m producers and assures MMS-fairness among at least $n - k$ producers. \square

One consequence of Lemma 7.2.2 is that, when k is much lower than n , a large fraction of producers are guaranteed to attain MMS fairness. We formally state this property of FairRec algorithm in Corollary 7.2.2.1.

COROLLARY 7.2.2.1. *Given n producers, a positive integer k , $\alpha = 1$, and $\beta \in (0, 1)$ such that $k \leq \beta n$, FairRec ensures MMS-fairness among at least $(1 - \beta)n$ producers.*

LEMMA 7.2.3. *Given n producers, m customers, recommendation size k such that $k < n < mk$, and a fixed value $\alpha \in (0, 1]$, FairRec ensures a minimum exposure of 1 for all the producers and an α -MMS guarantee to at least $(1 - \frac{1}{m+1} \lfloor \frac{\alpha mk}{n} \rfloor)$ fraction of the producers.*

PROOF. The first phase of *FairRec* terminates with one of the two (mutually exclusive) options, either $T == 0$ or $p == \emptyset$ (lines 16 and 8 of Algorithm 2). The case $T == 0$ ensures that all the producers achieve α -MMS exposure guarantee. However, in the other case, when $p == \emptyset$, some producers may not achieve α -MMS guarantee. We now provide a lower bound on the fraction of producers who achieve an α -MMS guarantee after the first phase.

Let R be the number of rounds after which the first phase of the algorithm *FairRec* terminates, and Q be the total number of allocations that occurred in the first phase. Therefore,

$$Q \geq mR. \tag{6}$$

Also, let β be the fraction of producers who achieved α -MMS exposure guarantee after the first phase. Thus, the total number of allocation Q can be upper bounded as follows:

$$\begin{aligned} Q &\leq \beta n \left\lfloor \frac{\alpha mk}{n} \right\rfloor + (1 - \beta)n \left(\left\lfloor \frac{\alpha mk}{n} \right\rfloor - 1 \right) \\ &\leq n \left(\left\lfloor \frac{\alpha mk}{n} \right\rfloor + \beta - 1 \right). \end{aligned} \quad (7)$$

Combining Inequalities (6) and (7), we obtain

$$\begin{aligned} mR &\leq n \left(\left\lfloor \frac{\alpha mk}{n} \right\rfloor + \beta - 1 \right), \\ \beta &\geq \frac{mR}{n} - \left\lfloor \frac{\alpha mk}{n} \right\rfloor + 1. \end{aligned} \quad (8)$$

Moreover, using the fact that at least $n - R$ producers achieve α -MMS exposure guarantee, we obtain

$$\begin{aligned} \beta n &\geq n - R \\ \Rightarrow R &\geq (1 - \beta)n. \end{aligned} \quad (9)$$

Using the lower bound on R (Inequality (9)), Inequality (8) can be rewritten as

$$\begin{aligned} \beta &\geq m(1 - \beta) - \left\lfloor \frac{\alpha mk}{n} \right\rfloor + 1 \\ \Rightarrow \beta(1 + m) &\geq m + 1 - \left\lfloor \frac{\alpha mk}{n} \right\rfloor \\ \Rightarrow \beta &\geq 1 - \left\lfloor \frac{\alpha mk}{n} \right\rfloor \left(\frac{1}{m + 1} \right). \end{aligned} \quad (10)$$

Inequality (10) implies that the fraction of producers who achieve α -MMS guarantee is at least $1 - \frac{1}{m+1} \lfloor \frac{\alpha mk}{n} \rfloor$. \square

Finally, in Lemma 7.2.4, we show that FairRec executes in polynomial time.

LEMMA 7.2.4. *The time complexity of FairRec has a worst case bound of $O(mnk)$.*

PROOF. The time complexity of FairRec is the same as the complexity of Algorithm 2. Over the two phases, Algorithm 2 allocates mk items. For each allocation, it finds the maximum possible feasible producer that can be done in at most $O(n)$ time. Thus, the total time complexity of the algorithm is $O(mnk)$. \square

8 EXPERIMENTAL EVALUATION

Experimental Setup and Baselines: We run the proposed *FairRec* algorithm (Section 6) for all three datasets (as listed in Section 4.1) considering different values of the recommendation-size k . For comparison, we use the following methods as baselines.

(1) **Top- k :** This is the traditional way of recommending the top- k relevant products.

(2) **Random- k :** Here, we randomly recommend k products to all customers. Random recommendations can give equal chance to all producers, thus can serve as a baseline that has only producers' interest in mind.

(3) **Poorest- k :** Unlike random- k , this is a deterministic producer-centric method where k least exposed products are recommended to each customer in a round robin manner.

Note that poorest-k and random-k are not real recommendation algorithms, and we consider them as baselines in the article because of their theoretical property of bringing down inequality among producers and serve as some of the most equitable options from the producers' perspective.

(4) **MixedTR-k**: Here, we choose top $\lceil \frac{k}{2} \rceil$ relevant products at first and then the remaining $(k - \lceil \frac{k}{2} \rceil)$ randomly, thus making it a mix of top- $\frac{k}{2}$ and random- $\frac{k}{2}$. In mixedTR-k, the top- $\frac{k}{2}$ could help keep high customer utility while random- $\frac{k}{2}$ could help in improving provider-side performance by giving equal chances to all producers to appear in the second half.

(5) **MixedTP-k**: Here, we choose top $\lceil \frac{k}{2} \rceil$ relevant products at first and then the poorest $(k - \lceil \frac{k}{2} \rceil)$ producers, thus making it a mix of top- $\frac{k}{2}$ and poorest- $\frac{k}{2}$. In mixedTP-k, the top- $\frac{k}{2}$ could help keep high customer utility while poorest- $\frac{k}{2}$ could help in improving the exposure of under-exposed providers by giving them a chance to appear in the second half.

(6) **MPB19**: We use the method proposed by Abdollahpouri et al. [3] as a baseline here. The proposal by Abdollahpouri et al. [3] is to consider the relevance scores as intermediate scores and then add a benefit term for less exposed producers to promote diversity thereby reducing popularity bias. In our problem setting, we consider the relevance score of product p to customer u $V_u(p)$ as intermediate scores and express the modified relevance score as $0.5 \times V_u(p) + 0.5 \times (1 - \frac{E_p}{\sum_{p' \in P} E_{p'}})$; the second part of the modified score promotes less exposed producers.

(7) **MSR18**: Sürer et al. [64] proposed to introduce producer-side constraints similar to exposure guarantees in our article and then optimize overall customer utility. However, as the proposed constrained optimization problem becomes a very hard combinatorial problem, the authors did a Lagrangian relaxation and proposed to use iterative subgradient method to optimize the relaxed problem. Although this methodology is not quite suitable for large scale online platforms as optimizing the hard combinatorial problem could demand huge computing resources and time, we use this as a baseline by limiting the number of iterations to 100.

Experiments: We run three sets of experiments. First, in Section 8.1, we set the exposure guarantee as $\bar{E} = \text{MMS}$ and run FairRec. However, the platforms may not always want to ensure the maximum possible exposure guarantee for the producers as it might cause degradation of customer utility. They might want to set a lower exposure guarantee in such cases. Thus, we set lower exposure guarantees, i.e., by considering $\bar{E} = \alpha \cdot \text{MMS}$ where $0 \leq \alpha \leq 1$ (in Section 8.2). While Sections 8.1 and 8.2 show the efficacies of FairRec, Section 8.3 digs deeper into the functioning of FairRec and discusses how phases 1 and 2 of FairRec work toward better performance on producer and customer sides. Finally, in Section 8.4, we also test FairRec for scenarios where the platforms may want to ensure different levels of exposure guarantee for different categories of producers. For evaluating FairRec and the baselines, we use the following *producer-side* and *customer-side* metrics.

8.0.1 Producer-Side Metrics. The evaluation metrics for capturing the fairness and efficiency among the producers are:

Fraction of Satisfied Producers (H): We call a producer satisfied iff its exposure is more than the minimum exposure guarantee \bar{E} . The fraction of *satisfied* producers can be calculated as follows:

$$H = \frac{1}{|P|} \sum_{p \in P} \mathbb{1}_{E_p \geq \bar{E}}. \quad (11)$$

$\mathbb{1}_{E_p \geq \bar{E}}$ is 1 if $E_p \geq \bar{E}$, otherwise 0. The value of H ranges between 0 and 1. The higher the H , the fairer is the recommender system to producers.

Inequality in Producer Exposures (Z): We earlier observed in Section 4 that conventional top-k recommendation causes huge disparity in individual producer exposures. To capture how unequal

the individual producer exposures are, we employ an entropy-like measure as follows:

$$Z = - \sum_{p \in P} \left(\frac{E_p}{m \times k} \right) \cdot \log_n \left(\frac{E_p}{m \times k} \right). \quad (12)$$

Note that the base of the logarithm above is n , which is the number of producers. Since each of the m customers is given k -sized recommendations, total available exposure is mk . If every producer gets same exposure, i.e., recommended exactly $\frac{mk}{n}$ times, then the above entropy metric will be: $-n \times \frac{mk/n}{mk} \times \log_n \frac{mk/n}{mk} = -n \times \frac{1}{n} \times \log_n \frac{1}{n} = -n \times \frac{1}{n} \times (-1) = 1$. However, if only one producer is allowed to get all the exposure, then the value entropy expression in equation-12 will be $-\frac{mk}{mk} \times \log_n \frac{mk}{mk} = 0$. Thus the range of Z is $[0, 1]$. The lower the Z , the more unequal individual producer exposures are.

Exposure Loss on Producers (L): As *FairRec* tries to ensure minimum exposure guarantee for all the producers, some producers may receive a lower exposure in comparison to what they would have got in top- k recommendations. To capture this, we compute the loss L as the mean amount of impact (loss in exposure) caused by *FairRec*, compared to the top- k recommendations:

$$L^{<method>} = \frac{1}{n} \sum_{p \in P} \max \left(\frac{(E_p^{\text{top-}k} - E_p^{<method>})}{E_p^{\text{top-}k}}, 0 \right). \quad (13)$$

This metric takes the top- k recommendations (which is the regular recommendation based on the estimated relevance scores, but with no additional constraints) as a reference point, and then evaluates how much exposure is lost on average if some other method is used for the recommendations. The lower the exposure loss metric, the lower is the negative impact, and the better is the recommendation algorithm.

8.0.2 Customer-side Metrics. The evaluation metrics for capturing the fairness and efficiency among the customers are:

Mean Average Envy (Y): Although *FairRec* ensures *EF1* guarantee for customers by design, here we capture how effectively this guarantee can reduce overall envy among customers in comparison to the baselines. We define the mean average envy as follows:

$$Y = \frac{1}{n} \sum_{u \in U} \frac{1}{n-1} \sum_{\substack{u' \in U \\ u' \neq u}} \text{envy}(u, u'), \quad (14)$$

where $\text{envy}(u, u') = \max((\phi_u(R_{u'}) - \phi_u(R_u)), 0)$ denoting how much u envies u' , which is the extra utility u would have received if she had received the recommendation that had been given to u' ($R_{u'}$) instead of her own allocated recommendation R_u . The lower the envy (Y), the fairer the recommender system is for the customers.

Mean and Standard Deviation of Customer Utilities (using μ_ϕ, std_ϕ): *FairRec* may not allocate the most relevant products to the customers, which may introduce a loss in customer utilities. This loss can be captured using the expression $\mu_\phi = \frac{1}{m} \sum_{u \in U} \phi_u(R_u)$. Higher the utility (i.e., lower utility loss), the more efficient is the recommender system for the customers. We also calculate the standard deviation of customer utilities, that is, $std_\phi = std_{u \in U}(\phi_u(R_u))$. The lower the std , the lesser is the disparity in individual customer utilities.

8.1 Experiments with MMS Guarantee

Here, we test *FairRec* with exposure guarantee $\bar{E} = \lfloor \frac{mk}{n} \rfloor = \text{MMS}$ (or $\alpha = 1$), recommendation size k in 1 to 20, and discuss the results.

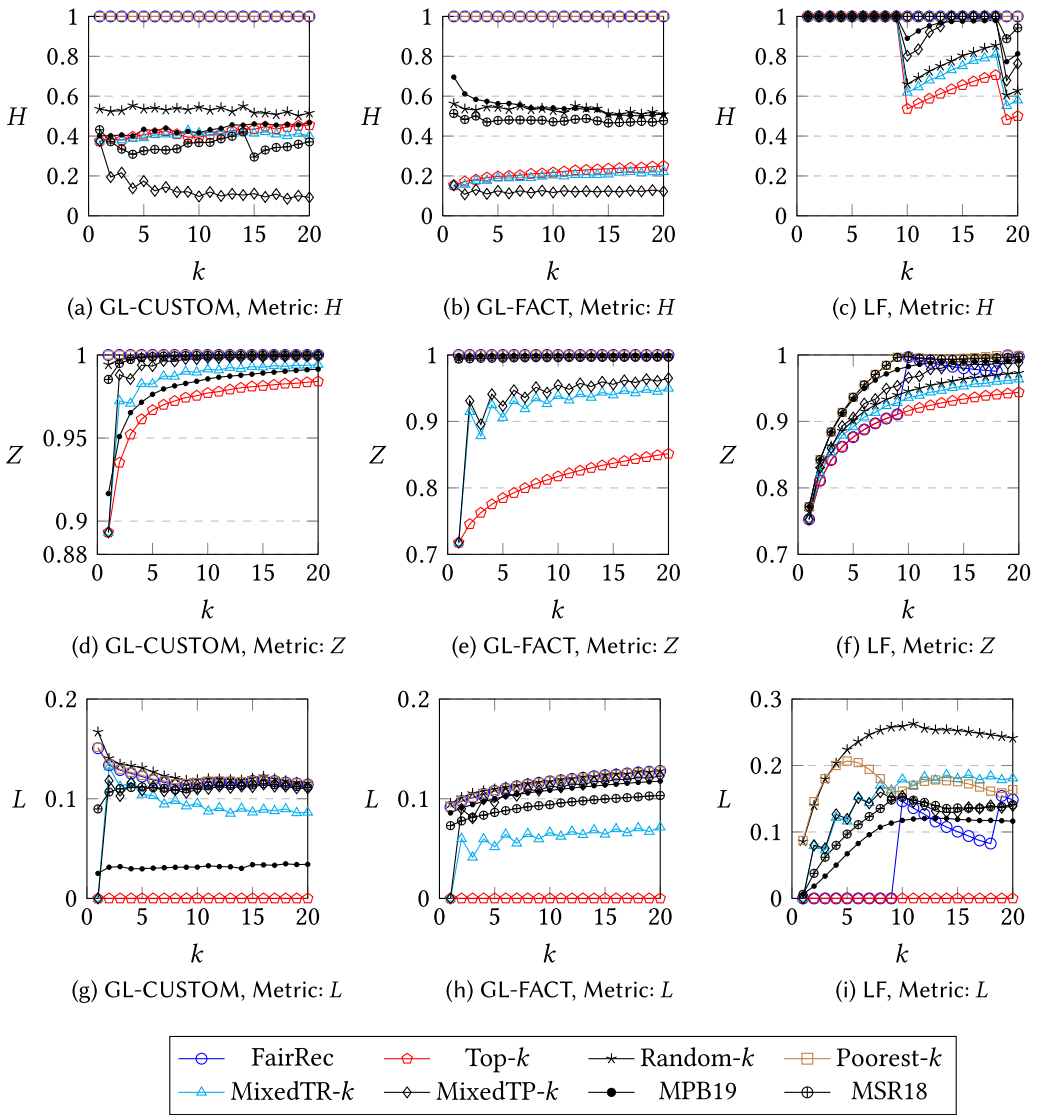


Fig. 2. Producer-side Performances with MMS Guarantee. First row: fraction of satisfied producers (H). Second row: inequality in producer exposures (Z). Third row: exposure loss on producers (L).

8.1.1 Producer-side Results. All producer-side results are plotted in Figure 2.

Producer Satisfaction (H): Figures 2(a), 2(b), and 2(c) show that both *FairRec* and *poorest- k* perform the best while *top- k* performs badly; this is because both *FairRec* and *poorest- k* explicitly try to ensure larger exposure for producers while *top- k* considers only the preferences of the customers. Similar to *mixedTR- k* and *mixedTP- k* , *MPB19* also fails to ensure MMS exposure as it does not explicitly ensure an exposure guarantee. However, *MSR18* has exposure-based constraint for producers similar to the exposure guarantee in *FairRec*. However, *MSR18* fails to exactly satisfy exposure constraint as it follows an approximate approach of constraint optimization while we find that the producer exposures in *MSR18* results are quite close to the MMS exposure even though

not crossing it. Thus, even though MSR18 performs quite bad in H , we find it to be performing almost as good as FairRec in the other two producer-side metrics.

Exposure Inequality (Z): Figures 2(d), 2(e), and 2(f) show that poorest- k has the lowest inequality in exposure while FairRec and random- k perform similar or slightly less than that; however, top- k performs the worst as it is highly customer-centric. Since poorest- k strategy makes the deterministic selection of less exposed producers thereby making smarter choices than random- k , poorest- k ensures less inequality (i.e., high Z) than random- k . Performances of both mixedTR- k and mixedTP- k lie in between the best and the worst lines as they mix customer-centric top half and producer-centric bottom half recommendations. MPB19 shows inconsistent performances across the datasets. However, MSR18 shows low exposure inequality (high Z) very close to that of FairRec and poorest- k , since MSR18 is observed to be ensuring exposures close to MMS for most of the producers. In summary, FairRec, poorest- k , and MSR18 seem to be good at maintaining fairness on the producer-side by keeping exposure inequality low.

Exposure Loss (L): Figures 2(g), 2(h), and 2(i) show that random- k and poorest- k cause the highest amounts of exposure loss in comparison to top- k ; this is because both of them favor equality in producer exposure (random- k gives equal chance to all producers to be recommended while poorest- k tries to increase the exposure of least exposed producer). However, mixedTR- k , mixedTP- k cause smaller losses as only up to half of their recommendations are different from top- k . FairRec causes only up to 0.2 fraction or 20% loss in exposure in comparison to top- k owing to its intelligent selection approach.

It is worth noticing that MMS for LF is low (MMS = 0 for $k < 10$, MMS = 1 for $k \in [10, 18]$, MMS = 2 for $k \in [20, 29], \dots$). MMS is satisfied for all producers until $k = 9$; but at $k=10$, MMS is not guaranteed for all producers, and thus, we see a drop in performance at $k = 10$, which happens again at $k = 19$. Such changes in MMS specific to LF make its plots different from other datasets. In summary, both FairRec, poorest- k perform the best in producer fairness while they cause exposure loss for very popular producers to compensate for the exposure given to less popular producers.

8.1.2 Customer-side Results. All customer-side results are plotted in Figure 3.

Mean Average Envy (Y): Figures 3(a), 3(b), and 3(c) reveal that top- k causes the lowest possible (i.e., 0) mean average envy among the customers; this is because it gives the maximum possible utility of 1 to every customer thereby leaving no chances of envy among customers. Even mixedTR- k , mixedTP- k , MPB19, and MSR18 also show similarly low envy. FairRec generates very low values of envy, which are very comparable to those of top- k here. However, both random- k and poorest- k cause the highest envy as they do not consider customer preferences at all during recommendation.

Mean and Standard Deviation of Customer Utility (μ_ϕ, std_ϕ): From Figures 3(d), 3(e), and 3(f), we see that both random- k and poorest- k cause huge losses in customer utility (i.e., low customer utility) as they neglect customer preferences. The mixedTR- k and mixedTP- k perform moderately. However, FairRec causes very small utility loss and performs almost at par with the customer-centric top- k . The utility losses in MPB19 and MSR18 are higher as they do not guarantee anything on the customer-side. This certifies that FairRec strikes a good balance between customer utility and producer fairness. The standard deviation plots: Figures 3(g), 3(h), and 3(i) reveal that for larger sizes of recommendation, random- k and poorest- k show large disparities in customer utilities while FairRec, mixedTR- k and mixedTP- k show relatively fewer disparities. As top- k is customer-centric and provides the maximum utility of 1 to all the customers, it shows 0 standard

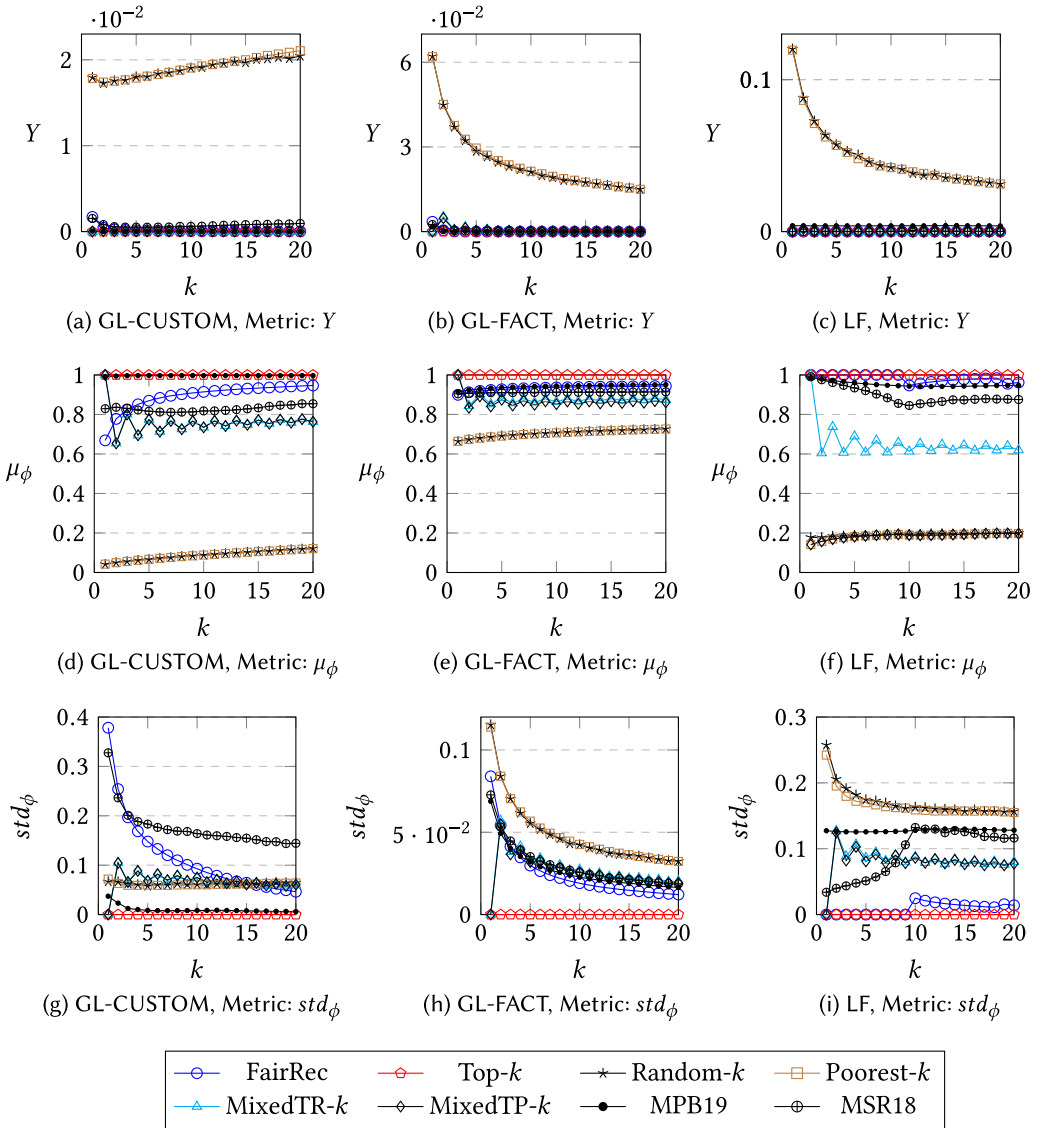


Fig. 3. Customer-side Performances with MMS Guarantee. First row: mean average envy (Y). Second row: mean customer utility (μ_ϕ). Third row: standard deviation of customer utilities (std_ϕ).

deviation. Besides MPB19 and MSR18 show higher disparities on customer-side as they do not specifically guarantee anything on customer-side.

In summary, FairRec strikes a good balance between fairness on both producer-side and customer-side while causing only marginal losses in customer utility.

8.2 Experiments with α -MMS Guarantee

Here, we fix $k = 20$ and test *FairRec* with different values of minimum exposure guarantee, i.e., $\bar{E} = \lfloor \alpha \cdot \frac{mk}{n} \rfloor$ (where $0 \leq \alpha \leq 1$) by varying α in between 0 and 1 (or in other words varying \bar{E} in between 0 and MMS); we plot the results in Figures 4 and 5.

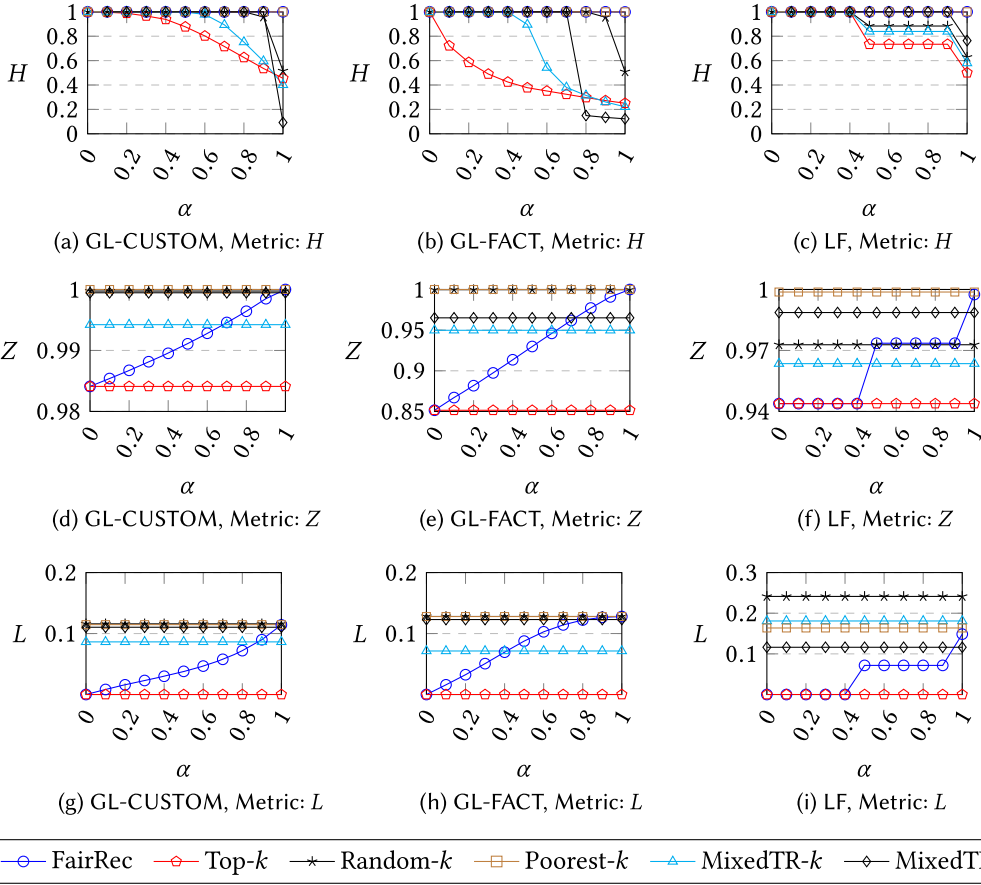


Fig. 4. Producer-side performances with $\bar{E} = \alpha \text{MMS}$ guarantee for $k = 20$. First row: fraction of satisfied producers (H). Second row: inequality in producer exposures (Z). Third row: exposure loss on producers (L).

8.2.1 Producer-side Results. All the relevant producer-side results are plotted in Figure 4.

Producer Satisfaction (H): Figures 4(a), 4(b), and 4(c) reveal that *FairRec* satisfies almost all the producers (as H is close to 1) for all the tested α settings in all the datasets. However, the performances of the baselines in terms of metric H reduces with the increase in α ; this is because the criteria for the producer satisfaction (as $\bar{E} \propto \alpha$) increases with the increase in α while the baseline results do not explicitly change with the change in α .

Exposure Inequality (Z): From Figures 4(d), 4(e), and 4(f), we see that increasing minimum exposure guarantee (i.e., increasing α) results in lower inequality in producer exposures (as high Z signifies lower inequality) for *FairRec* in all the cases. However, as the baselines do not depend on α , their performances remain the same; thus the baseline performances are just horizontal straight lines. At $\alpha = 0$, as the exposure guarantee \bar{E} by *FairRec* becomes 0, the results given by *FairRec* are the same as that of top- k in all the datasets. While *FairRec*'s performance in terms of Z crosses that of mixedTR- k at $\alpha = 0.7$, $\alpha = 0.7$, and $\alpha = 0.5$ in GL-CUSTOM, GL-FACT, and LF, respectively, it becomes close to that of producer-centric poorest- k at $\alpha = 1$ in all the datasets.

Exposure Loss (L): Increasing values of L are observed for increased α in Figures- 4(g), 4(h), 4(i). This suggests that increasing minimum exposure guarantee can cause higher exposure losses for

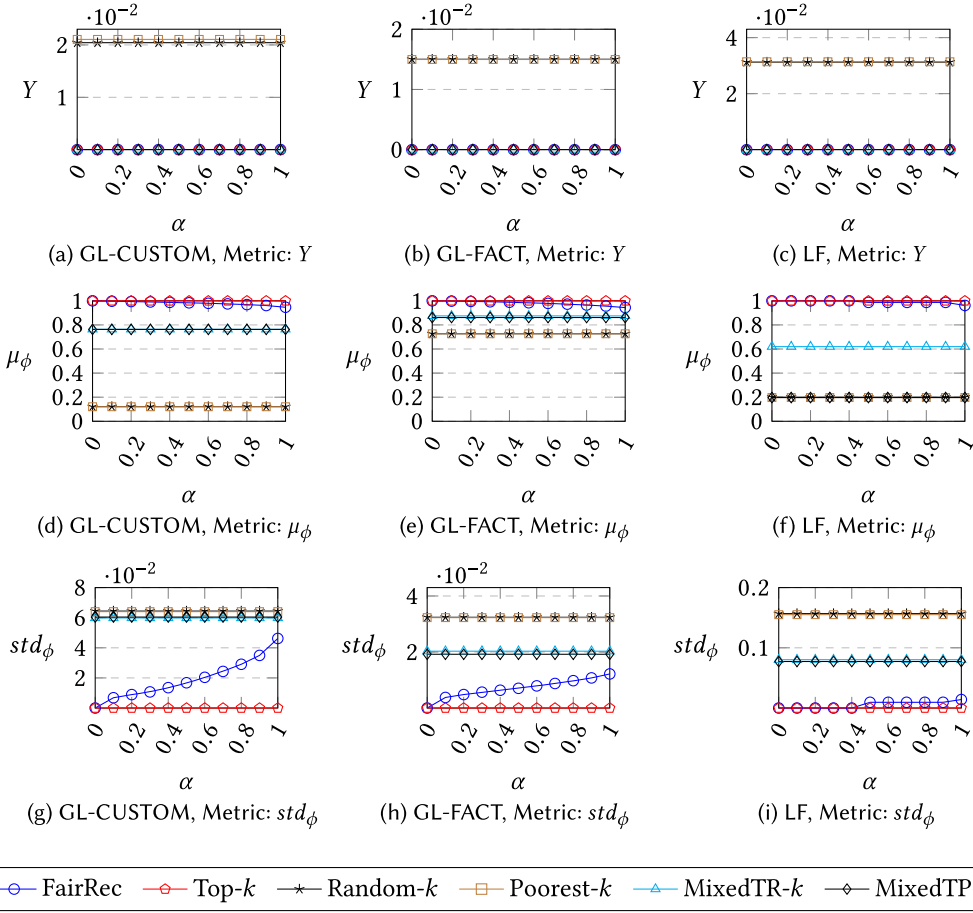


Fig. 5. Customer-side performances with $\bar{E} = \alpha MMS$ guarantee for $k = 20$. First row: mean average envy (Y). Second row: mean customer utility (μ_ϕ). Third row: standard deviation of customer utilities (std_ϕ).

previously popular producers. Just like Z , here also the baseline performances are just horizontal straight lines as they do not depend on α . We see that at $\alpha = 0$, *FairRec* performs same as top- k with no losses. With the increase in α settings, the losses increase and at $\alpha = 1$, the losses in *FairRec* are very close to those of poorest- k (as discussed earlier, poorest- k is the best performing baselines for producer side).

8.2.2 Customer-side Results. All the relevant customer-side results are plotted in Figure 5. As the baselines do not depend on α , the customer-side results of baselines are just horizontal straight lines. We find almost no change in Mean Average Envy (Y) of *FairRec* with the change in α (refer to Figures 5(a), 5(b), and 5(c)). However, with the increase in α (i.e., higher exposure guarantee for producers) there is a small decrease in customer utility (refer to Figures 5(d), 5(e), and 5(f)), and a small increase in the standard deviation of customer utilities.

In summary, although a larger exposure guarantee can help platforms achieve better producer fairness, it might hurt the overall customer satisfaction and also the satisfaction of highly popular

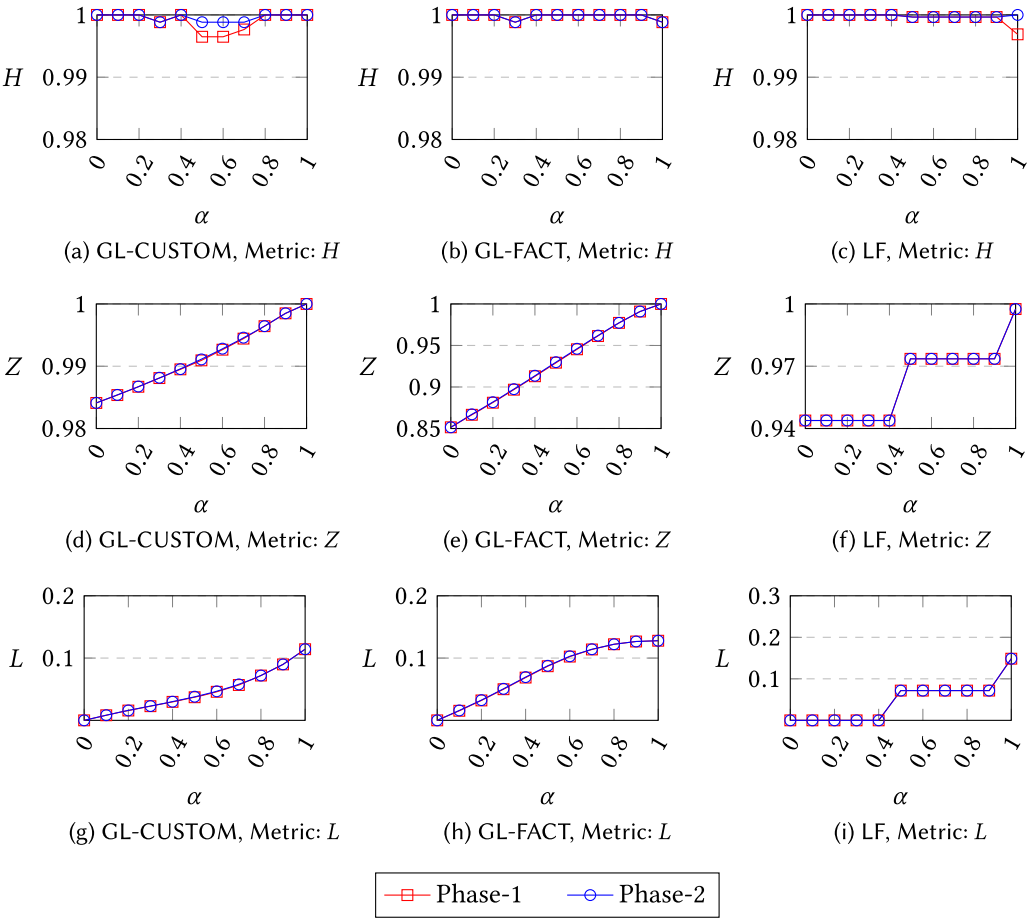


Fig. 6. Phase interplay figures on producer-side for $k = 20$. First row: fraction of satisfied producers (H). Second row: inequality in producer exposures (Z). Third row: exposure loss on producers (L).

producers of the platforms. Thus, the platforms, which are interested in similar minimum exposure guarantees, should not ignore the above trade-offs.

8.3 Interplay Between Phase-1 and Phase-2 of FairRec

To understand how phase-1 and phase-2 of *FairRec* work toward better performance on producer and customer sides, we plot the metrics at the end of both phase-1 and phase-2 (i.e., the end of *FairRec*) in Figures 6 and 7.

8.3.1 Producer-Side Results. All the relevant producer-side results are plotted in Figure 6. We find that the metrics evaluated after phase-1 and phase-2 of *FairRec* have almost the same values in all the datasets; this is because the task of reducing exposure inequality through minimum exposure guarantee to producers happens in phase-1 of *FairRec*, and phase-2 of *FairRec* does not explicitly work toward these goals. Thus, there is very little change in the relevant metrics of producer satisfaction (H in Figures 6(a), 6(b), and 6(c)), exposure inequality (Z in Figures 6(d), 6(e), and 6(f)), and exposure loss (L in Figures 6(g), 6(h), and 6(i)) after phase-2 than those at the end of phase-1. However, at a few settings of α in GL-CUSTOM and LF where phase-1 falls a bit short

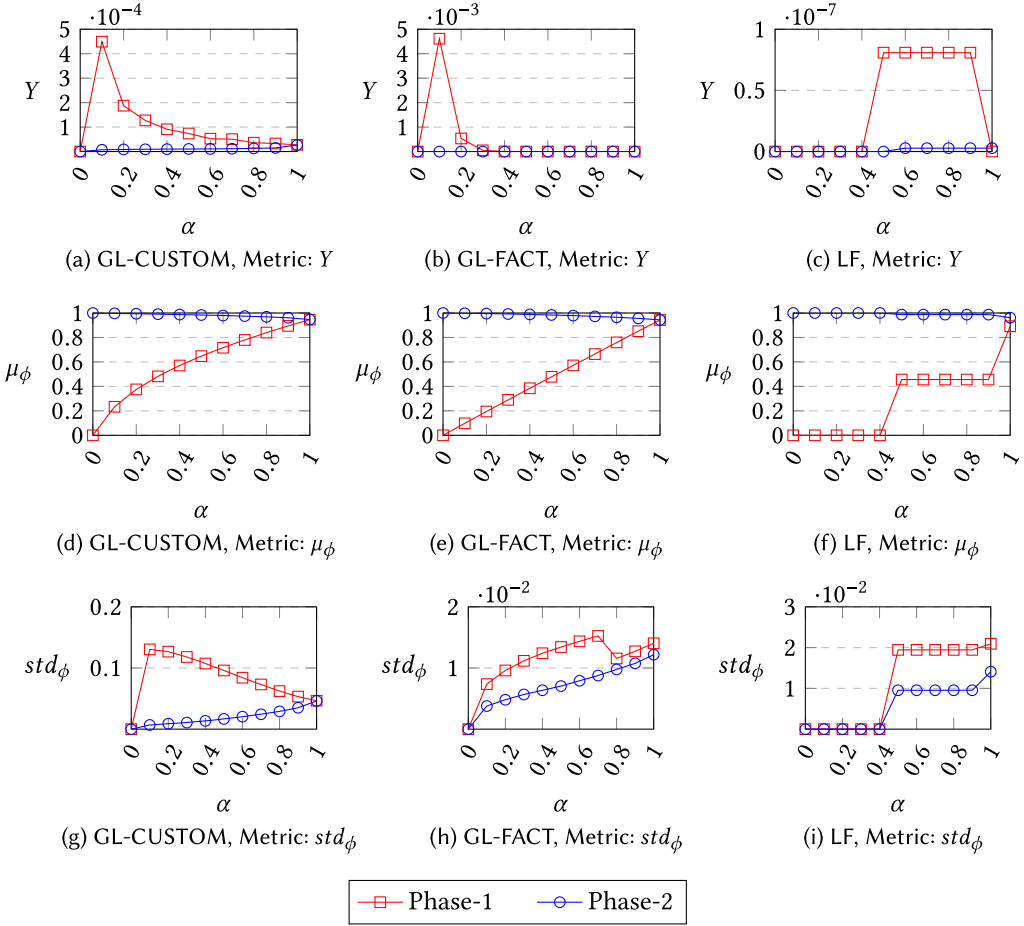


Fig. 7. Phase interplay figures on customer-side for $k = 20$. First row: mean average envy (Y). Second row: mean customer utility (μ_ϕ). Third row: standard deviation of customer utilities (std_ϕ).

in producer satisfaction (H), phase-2 seems to improve it by a small extent (refer to Figures 6(a) and 6(c), respectively).

8.3.2 Customer-side Results. All the relevant customer-side results are plotted in Figure 7. For all the α values that correspond to $\bar{E} = 0$ (like $\alpha = 0$ in GL-CUSTOM and GL-FACT, $\alpha \in \{0.1, 0.2, 0.3, 0.4\}$ in LF), we see that all the metrics Y , μ_ϕ , and std_ϕ correspond to 0 after phase-1 of FairRec as no allocation happens in phase-1; all the allocations happen in phase-2 thus making the final results same as that of top- k ($Y = 0$, $\mu_\phi = 1$, and $std_\phi = 0$ after phase-2) in all the cases. We observe variations in these metrics right from the point where α corresponds to non-zero \bar{E} (i.e., in our experiments $\alpha = 0$ in GL-CUSTOM and GL-FACT, $\alpha = 0.5$ in LF); we describe these observations next.

Mean Average Envy (Y): For smaller non-zero values of α , we see higher envy after phase-1 (Y in Figures 7(a), 7(b), and 7(c)); however, with the increase in α , phase-1 envy decreases; this must be because with the increase in α , the number of rounds of allocation in phase-1 increases, and thereby increasing the chances of canceling out envy to some extent in the Greedy Round

Robin process. However, phase-2 improves the results by reducing Y ; this is because it allows best possible allocations to every customer by setting higher availability of all the products. However, with the increase in α , the amount of improvement happening in phase-2 reduces as more and more allocations happen in phase-1 thereby reducing the number of rounds in phase-2.

Mean and Standard Deviation of Customer Utility (μ_ϕ, std_ϕ): With the increase in α , we observe increase in the mean customer utility of phase-1 allocations (refer to Figures 7(d), 7(e), and 7(f)); this is because the number of allocations in phase-1 increases with the increase in α . Customer utility after phase-2 is higher than that of phase-1 (refer to Figures 7(d), 7(e), and 7(f)); however, the utility improvement in phase-2 reduces with the increase in α as more number of allocations happen in phase-1 and the number of rounds in phase-2 becomes less for higher α . We also observe that phase-2 reduces the disparity in customer utilities than what is observed at the end of phase-1 (refer to Figures 7(g), 7(h), and 7(i)); this is because phase-2 allows best possible allocations to every customer thereby compensating for customer-side inequalities and losses incurred in phase-1.

In summary, phase-1 of *FairRec* tries to achieve better performance from the producer-side, whereas phase-2 mostly improves the performance from the customer-side. In addition to that the customer-side improvement in phase-2 is often more when the α values are smaller as it increases the number of allocations happening in the second phase.

8.4 Experiments with Producer-specific α

The platforms may also need to ensure different levels of exposure guarantee for different producers; for example, the platform may want to give more exposure guarantee to high-rated producers than low-rated ones. Thus, we also test *FairRec* in such a scenario. Here, we fix $\alpha_p = 0.2 \times \lfloor rating(p) \rfloor$; i.e., for $rating(p) = 5$ the exposure guarantee is $\bar{E}_p = 1.0 \times MMS$, and similarly for $rating(p) = 3.2$ the exposure guarantee is $\bar{E}_p = 0.6 \times MMS$. We test *FairRec* with the above settings on GL-CUSTOM and GL-FACT, and plot the results in Figure 8. Figure 8(b) shows the distribution of producers in Google Local dataset with different exposure guarantees (set based on their respective ratings). Figures 8(c) and 8(d) plot H for every group of producers with different exposure guarantees in GL-CUSTOM and GL-FACT, respectively; we find *FairRec* to be satisfying all the producers in both the cases. Metrics Z and L are skipped here, as they are irrelevant in case of different exposure guarantees for different producers. We plot customer-side results in Figures 8(c) and 8(e), GL-CUSTOM and GL-FACT, respectively; we find that *FairRec* is able to achieve high customer utility (μ_ϕ very close to 1) while maintaining customer fairness (small Y and std_ϕ) in both cases.

9 IMPROVING FAIRREC THROUGH ENVY-CYCLE ELIMINATION

While *FairRec* provides two-sided fair recommendations, it can be further tweaked to improve the recommendation performance for the customers. We propose such a modification of *FairRec* in Section 9.1 as *FairRecPlus*, and then evaluate it against *FairRec* in Section 9.2. However, the improvement in *FairRecPlus* comes at a cost of an increased computation time. Thus, in scenarios where platforms can afford more time to compute a recommendation, *FairRecPlus* can be used to improve customer-side performance while still maintaining the same fairness guarantees on both sides.

9.1 FairRecPlus

We now present a modification of *FairRec*: named *FairRecPlus*. *FairRecPlus* executes in two phases, similar to *FairRec*. The **first phase** creates $\ell = \lfloor \frac{mk}{n} \rfloor$ copies of each product and then initializes

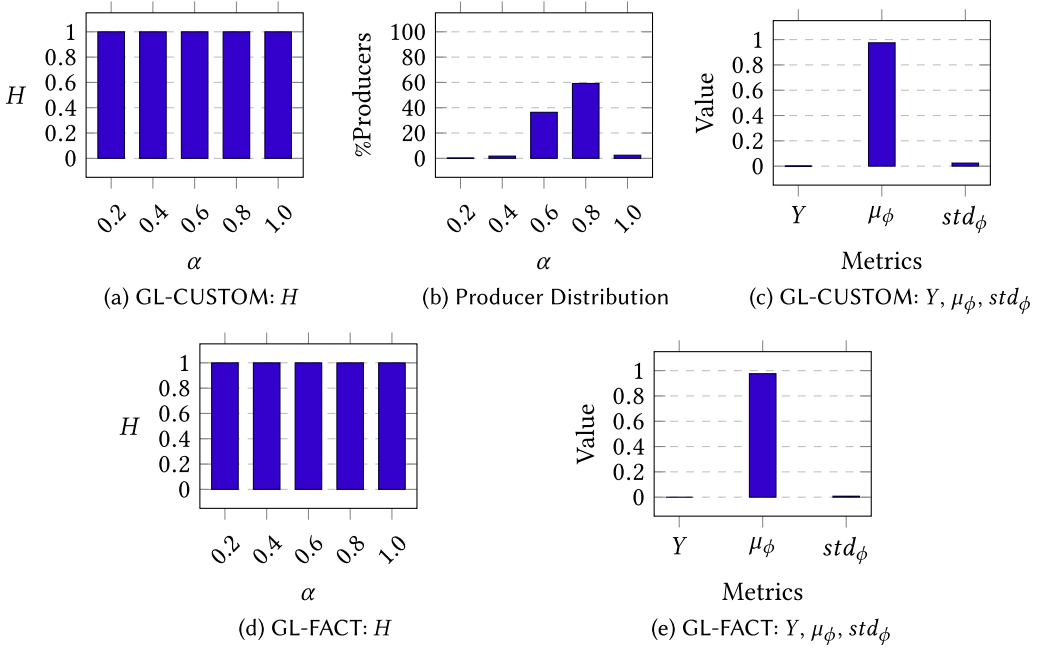


Fig. 8. Results of FairRec with producer-specific α settings in GL-CUSTOM ($k = 20$).

each component of the vector S of size $|P|$ to the value ℓ . Then, assuming an arbitrary ordering σ of customers, the modified greedy algorithm ALG 4 is executed, which is different from our earlier approach. In this modification, at the end of a round of greedy-round-robin allocation, we create an *envy graph* and compute the topological ordering among the agents based on their partial allocation. The high-level idea is to re-order the priorities of the agents during round-robin allocations, aiming to balance the extent of envy between each pair of agents by maintaining an acyclic envy-graph after each round. An envy graph is a directed graph that captures the envy between agents—the nodes in the envy graph represent the agents and it contains a directed edge from i to j if and only if, i envies j , i.e., if and only if $v_i(B_i) < v_i(B_j)$, where B_i and B_j are partial allocations to agents i and j , respectively. It was established in Reference [48] that one can always efficiently update a given partial allocation such that the resulting envy graph is acyclic.

LEMMA 9.0.1 (LIPTON ET AL. [48]). *Given a partial allocation (A_1, \dots, A_m) , we can find another partial allocation $\mathcal{B}=(B_1, \dots, B_m)$ in polynomial time such that*

- (i) *The valuations of the agents for their bundles do not decrease: $v_i(B_i) \geq v_i(A_i)$ for all $i \in [m]$.*
- (ii) *The envy graph $G(\mathcal{B})$ is acyclic.*

PROOF SKETCH. If the envy graph of \mathcal{A} is acyclic, then the claim holds trivially. Otherwise, find a cycle in the graph $G(\mathcal{A})$ (time complexity $O(n + |E|)$, where E is the number of edges in the graph). Let $C = i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_k \rightarrow i_1$ be the cycle. The bundles can be re-allocated as follows: for all agents not in C , i.e., $k \notin \{i_1, i_2, \dots, i_k\}$ set $B_k = A_k$, and for all the agents in the cycle set B_i to be the bundle of their successor in C , i.e., set $B_{i_a} = A_{i_{(a+1)}}$ for $1 \leq a < k$ along with $B_{i_k} = A_{i_1}$. After this reallocation $v_i(B_i) \geq v_i(A_i)$ for all $i \in [n]$. Furthermore, the number of edges in $G(\mathcal{B})$ is strictly less than $G(\mathcal{A})$: the directed edges

ALGORITHM 3: *FairRecPlus* (U, P, k, V)

Input: Set of customers $U = [m]$, set of distinct products $P = [n]$, recommendation set size k (such that $k < m$ and $n \leq k \cdot m$), and the relevance scores $V_u(p)$.

Output: A two-sided fair recommendation.

- 1: Initialize allocation $\mathcal{A}^0 = (A_1^0, \dots, A_m^0)$ with $A_i^0 \leftarrow \emptyset$ for each customer $i \in [m]$.
- 2:
- 3: **First Phase:**
- 4: Fix an (arbitrary) ordering of the customers $\sigma = (\sigma(1), \sigma(2), \dots, \sigma(m))$.
- 5: Initialize set of feasible products $F_u \leftarrow P$ for each $u \in U$.
- 6: Set $\ell \leftarrow \lfloor \frac{\alpha mk}{n} \rfloor$ denoting number of copies of each product.
- 7: Initialize each component of the vector $S = (S_1, \dots, S_n)$ with $S_j \leftarrow \ell, \forall j \in [n]$, this stores the number of available copies of each product.
- 8: Set $T \leftarrow \ell \times n$, total number of items to be allocated.
- 9: $[\mathcal{B}, F] \leftarrow \text{Modified-Greedy-Round-Robin}(m, n, S, T, V, \sigma, F)$.
- 10: Assign $\mathcal{A} \leftarrow \mathcal{A} \cup \mathcal{B}$.
- 11:
- 12: **Second Phase:**
- 13: **for** $i = 1$ to m **do**
- 14: **if** $|A_i| < k$ **then**
- 15: Set H as the top $(k - |A_i|)$ items from F_i (based on $V_i(\cdot)$ scores).
- 16: Update $A_i \leftarrow A_i \cup H$.
- 17: **end if**
- 18: **end for**
- 19: Return \mathcal{A} .

$i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_k \rightarrow i_1$ do not appear in the envy graph of (B_1, \dots, B_n) and if an agent k starts envying an agent in the cycle, say agent i_a , then k must have been envious of i_{a+1} in \mathcal{A} . Edges between agents k and k' , which are not in the cycle C remain unchanged, and edges going out of an agent i in the cycle C can only get removed, since i 's valuation for the bundle assigned to her bundle increases. Therefore, we can repeatedly remove cycles and keep reducing the number of edges in the envy graph to eventually find a partial allocation \mathcal{B} that satisfies the stated claim. \square

The worst-case time complexity of eliminating envy cycles, to obtain a directed acyclic graph at each round, is $O(m^4)$. A topological ordering of the acyclic directed graph $\mathcal{G}(\mathcal{B})$ can be computed for updating the σ in $O(m)$ time. This new ordering is then used for allocating the next round of items in a round-robin manner.

The **second phase** checks if all the customers have received exactly k products. If yes, then no further allocation is required; if not, then allocate each agent i , their most valuable items from F_i until they receive k items.

The time complexity of FairRecPlus is governed by the envy cycle elimination step, which takes $O(m^4 k)$ over k rounds. Moreover, for each of the mk items, finding the maximum valued feasible producer takes $O(n)$ time. Thus, the total time complexity of FairRecPlus is $O(m^4 k + mnk)$. Therefore, for a large number of customers, FairRecPlus would take a huge time to compute a two-sided fair allocation. Although this modification requires more computation, we empirically observe that it improves on the customer-side metrics.

ALGORITHM 4: Modified-Greedy-Round-Robin (m, n, S, T, V, σ, F)

Input : Number of customers m , number of producers n , an array with number of available copies of each product S , total number of available products $T > 0$, relevance scores $V_u(p)$ and feasible product set F_u for each customer, and an ordering σ of $[m]$.

Output: An allocation of T products among m customers and the residual feasible set F_u .

```

1: Initialize allocation  $\mathcal{B} = (B_1, \dots, B_m)$  with  $B_i \leftarrow \emptyset$  for each customer  $i \in [m]$ .
2: Initiate round  $r \leftarrow 0$ .
3: while true do
4:   Set  $r \leftarrow r + 1$ .
5:   for  $i = 1$  to  $m$  do
6:     Set  $p \in \arg \max_{p' \in F_{\sigma(i)}: (S_p' \neq 0)} V_{\sigma(i)}(p')$ 
7:     if  $p == \emptyset$  then
8:       Update  $\mathcal{B} = \{B_1, \dots, B_m\}$  to obtain an acyclic envy graph  $\mathcal{G}(\mathcal{B})$  using Lemma 9.0.1.
9:       Update ordering  $\sigma$  to be the topological ordering of the envy graph  $\mathcal{G}(\mathcal{B})$ .
10:      go to Step 24.
11:     end if
12:     Update  $B_{\sigma(i)} \leftarrow B_{\sigma(i)} \cup p$ .
13:     Update  $F_{\sigma(i)} \leftarrow F_{\sigma(i)} \setminus p$ .
14:     Update  $S_p \leftarrow S_p - 1$ .
15:     Update  $T \leftarrow T - 1$ .
16:     if  $T == 0$  then
17:       Update  $\mathcal{B} = \{B_1, \dots, B_m\}$  to obtain an acyclic envy graph  $\mathcal{G}(\mathcal{B})$  using Lemma 9.0.1.
18:       Update ordering  $\sigma$  to be the topological ordering of the envy graph  $\mathcal{G}(\mathcal{B})$ .
19:       go to Step 24.
20:     end if
21:   end for
22:   Update  $\mathcal{B} = \{B_1, \dots, B_m\}$  to obtain an acyclic envy graph  $\mathcal{G}(\mathcal{B})$  using Lemma 9.0.1.
23: end while
24: Return  $\mathcal{B} = (B_1, \dots, B_m)$  and  $F = (F_1, \dots, F_m)$ .

```

9.2 Results with FairRecPlus

We test the FairRecPlus algorithm on all the datasets with $k = 20$ while varying α from 0 to 1 in separate trials. The results on producer-side and customer-side are plotted in Figures 9 and 10, respectively. While the producer-side plots for FairRec and FairRecPlus (Figure 9) seem to completely overlap, they are marginally different from each other; for example, at $\alpha = 0.5$ in GL-CUSTOM, the H , Z , and L metrics for FairRecPlus are 1, 0.9908, and 0.0376, respectively, while those in FairRec are 0.99, 0.9910, and 0.0380, respectively. The producer-side performances of FairRecPlus is very similar to that of FairRec, since the modification does not change any guarantee on the producer side. However, the modification introduces only the envy-cycle removal rounds, which reduces the envy on customer-side. We observe that FairRecPlus reduces the mean customer envy (check Y in Figures 10(a) to 10(c)). However, the performances of FairRecPlus in other customer-side metrics are similar to those of FairRec (Figures 10(d) to 10(i)); for example, at $\alpha = 0.5$ in GL-CUSTOM, the mean and standard deviation of customer utilities are 0.9841 and 0.0169 in FairRecPlus, against the corresponding values 0.9834 and 0.0167 in FairRec.

10 CONCLUSION

In this work, we propose the notion of two-sided fairness for recommendations in two-sided platforms. For producers, we consider a minimum exposure guarantee while we try to ensure less

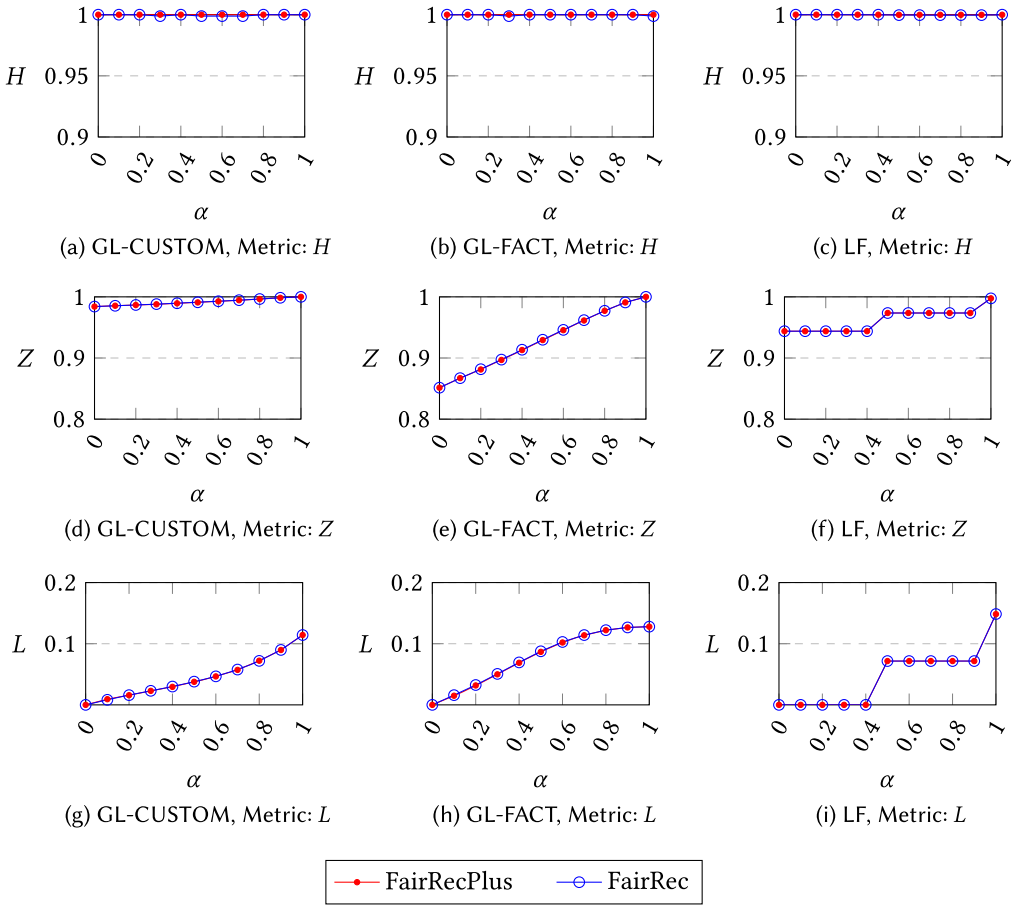


Fig. 9. FairRec modification results on producer-side for $k = 20$. First row: fraction of satisfied producers (H). Second row: inequality in producer exposures (Z). Third row: exposure loss on producers (L).

inequality in customer utilities. Note that we assume the relevance of a product does not play any role in producer's utility (in contrast to Biega et al. [12], Singh and Joachims [59]), and we use only the exposure of a producer as her utility. We provide a scalable and easily adaptable algorithm that exhibits desired two-sided fairness properties while causing a marginal loss in the overall quality of recommendations. We establish theoretical guarantees and provide empirical evidence through extensive evaluations of real-world datasets. Furthermore, we propose a modification of our algorithm and show that it performs better on customer-side metrics while being two-sided fair, but at the cost of additional computation time. Our work can be directly applied to fair recommendation problems in scenarios like mass recommendation/promotion sent through emails and app/web notifications. Though our work considers the offline recommendation scenario where the recommendations are computed for all the registered customers at once, it can also be extended for online recommendation settings by limiting the set of customers to only the active customers at any particular instant. However, developing a more robust realization of the proposed mechanism for a completely online scenario remains future work. Going ahead, we also want to study attention models that can handle position bias [4], where customers pay more attention to the top-ranked products than the lower-ranked ones.

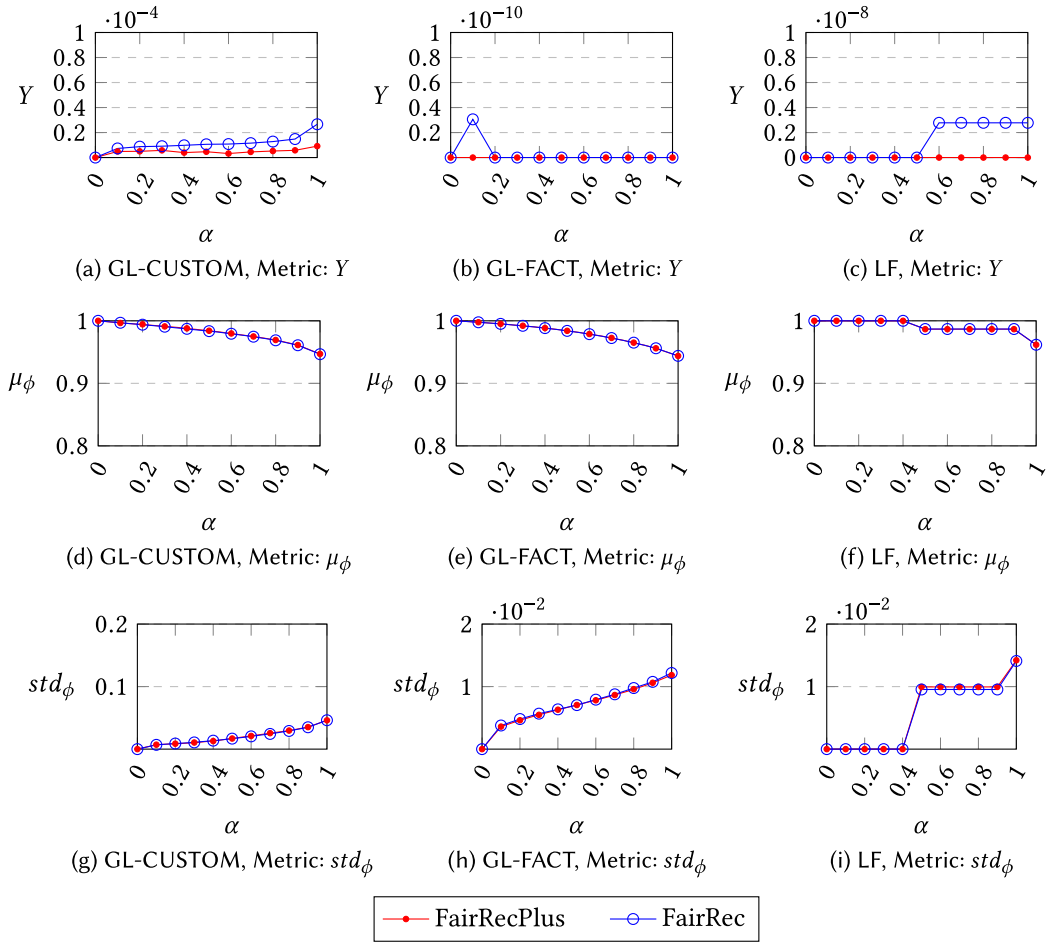


Fig. 10. FairRec modification results on customer-side for $k = 20$. First row: mean average envy (Y). Second row: mean customer utility (μ_ϕ). Third row: standard deviation of customer utilities (std_ϕ).

Reproducibility: Code, dataset, and other details are available at <https://github.com/gourabkumarpatro/FairRec>.

REFERENCES

- [1] Himan Abdollahpouri and Robin Burke. 2019. Multi-stakeholder recommendation and its connection to multi-sided fairness. Retrieved from <https://arXiv:1907.13158>.
- [2] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2017. Controlling popularity bias in learning-to-rank recommendation. In *Proceedings of the 11th ACM Conference on Recommender Systems*. ACM, 42–46.
- [3] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2019. Managing popularity bias in recommender systems with personalized re-ranking. Retrieved from <https://arXiv:1901.07555>.
- [4] Aman Agarwal, Ivan Zaitsev, Xuanhui Wang, Cheng Li, Marc Najork, and Thorsten Joachims. 2019. Estimating position bias without intrusive interventions. In *Proceedings of the 12th ACM International Conference on Web Search and Data Mining*. ACM, 474–482.
- [5] AirBnb. 2019. Retrieved from <https://blog.airbnb.com/the-airbnb-community-commitment/>.
- [6] AirBnb. 2019. Retrieved from <https://www.airbnb.com/help/article/2474/what-is-the-earnings-guarantee-programme-for-new-hosts>.
- [7] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2018. Comparing approximate relaxations of envy-freeness. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'18)*. IJCAI, 42–48.

- [8] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. 2015. Approximation algorithms for computing maximin share allocations. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP'15)*. 39–51.
- [9] Ashmi Banerjee, Gourab K. Patro, Linus W. Dietz, and Abhijnan Chakraborty. 2020. Analyzing “Near Me” Services: Potential for Exposure Bias in Location-based Retrieval. In *Proceedings of the IEEE International Conference on Big Data (Big Data '20)*. IEEE, 3642–3651.
- [10] Siddharth Barman, Arpita Biswas, Sanath Kumar Krishnamurthy, and Y. Narahari. 2018. Groupwise maximin fair allocation of indivisible goods. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 917–924.
- [11] Siddharth Barman and Sanath Kumar Krishnamurthy. 2017. Approximation algorithms for maximin fair division. In *Proceedings of the ACM Conference on Economics and Computation (EC'17)*. ACM, 647–664.
- [12] Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. 2018. Equity of attention: Amortizing individual fairness in rankings. Retrieved from <https://arXiv:1805.01788>.
- [13] Vittorio Bilò, Ioannis Caragiannis, Michele Flammini, Ayumi Igarashi, Gianpiero Monaco, Dominik Peters, Cosimo Vinci, and William S. Zwicker. 2019. Almost envy-free allocations with connected bundles. In *Proceedings of the 11th Innovations in Theoretical Computer Science (ITCS'19)*.
- [14] Arpita Biswas and Siddharth Barman. 2018. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18)*. ijcai.org, 91–97. <https://doi.org/10.24963/ijcai.2018/13>
- [15] Arpita Biswas and Siddharth Barman. 2019. Matroid constrained fair allocation problem. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*. AAAI Press, 9921–9922. <https://doi.org/10.1609/aaai.v33i01.33019921>
- [16] Sylvain Bouveret, Katarína Cechlárová, Edith Elkind, Ayumi Igarashi, and Dominik Peters. 2017. Fair division of a graph. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'17)*. 135–141.
- [17] Sylvain Bouveret and Michel Lemaître. 2014. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems*. 1321–1328.
- [18] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. 2016. *Handbook of Computational Social Choice*. Cambridge University Press. <https://doi.org/10.1017/CBO9781107446984>
- [19] Eric Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. Politic. Econ.* 119, 6 (2011), 1061–1103.
- [20] Robin Burke. 2017. Multisided fairness for recommendation. Retrieved from <https://arXiv:1707.00093>.
- [21] Iván Cantador, Peter Brusilovsky, and Tsvi Kuflik. 2011. 2nd workshop on information heterogeneity and fusion in recommender systems (HetRec 2011). In *Proceedings of the 5th ACM conference on Recommender systems (RecSys'11)*. ACM, New York, NY.
- [22] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2016. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the ACM Conference on Economics and Computation (EC'16)*. ACM, 305–322.
- [23] Abhijnan Chakraborty, Aniko Hannak, Asia J. Biega, and Krishna P. Gummadi. 2017. Fair sharing for sharing economy platforms.
- [24] Abhijnan Chakraborty, Gourab K. Patro, Niloy Ganguly, Krishna P. Gummadi, and Patrick Loiseau. 2019. Equality of voice: Towards fair representation in crowdsourced top-k recommendations.
- [25] CNBC.com. 2017. Retrieved from <https://www.cnbc.com/2017/04/20/only-4-percent-of-uber-drivers-remain-after-a-year-says-report.html>.
- [26] Abhisek Dash, Abhijnan Chakraborty, Saptarshi Ghosh, Animesh Mukherjee, and Krishna P. Gummadi. 2021. When the umpire is also a player: Bias in private label product recommendations on e-commerce marketplaces. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*. 873–884.
- [27] Benjamin Edelman, Michael Luca, and Dan Svirsky. 2017. Racial discrimination in the sharing economy: Evidence from a field experiment. *Amer. Econ. J.: Appl. Econ.* 9, 2 (2017).
- [28] Ulle Endriss. 2017. *Trends in Computational Social Choice*. Lulu.com.
- [29] Niklas Engbom and Christian Moser. 2018. Earnings inequality and the minimum wage: Evidence from brazil. *Federal Reserve Bank of Minneapolis-Opportunity and Inclusive Growth Institute Working Paper* 7 (2018), 18–50.
- [30] Armin Falk, Ernst Fehr, and Christian Zehnder. 2006. Fairness perceptions and reservation wages—the behavioral effects of minimum wage laws. *Quart. J. Econ.* 121, 4 (2006), 1347–1381.
- [31] D. C. Foley. 1967. Resource allocation in the public sector. *Yale Econ. Essays* 7 (1967), 73–76.
- [32] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD'19)*.
- [33] Mohammad Ghodsi, MohammadTaghi HajiAghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2018. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the ACM Conference on Economics and Computation*. ACM, 539–556.

- [34] Laurent Gourvès and Jérôme Monnot. 2017. Approximate maximin share allocations in matroids. In *Proceedings of the 10th International Conference on Algorithms and Complexity (CIAC'17)*. 310–321. https://doi.org/10.1007/978-3-319-57586-5_26
- [35] Laurent Gourvès, Jérôme Monnot, and Lydia Tlilane. 2014. Near fairness in matroids. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI'14)*. IOS Press, 393–398. <https://doi.org/10.3233/978-1-61499-419-0-393>
- [36] Mark Graham, Isis Hjorth, and Vili Lehdonvirta. 2017. Digital labour and development: Impacts of global digital labour platforms and the gig economy on worker livelihoods. *Transfer: Eur. Rev. Labour Res.* 23, 2 (2017), 135–162.
- [37] David A. Green and Kathryn Harrison. 2010. *Minimum Wage Setting and Standards of Fairness*. Technical Report. IFS working papers.
- [38] Anikó Hannák, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson. 2017. Bias in online freelance marketplaces: Evidence from taskrabbit and fiverr. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW'17)*.
- [39] Ruining He, Wang-Cheng Kang, and Julian McAuley. 2017. Translation-based recommendation. In *Proceedings of the 11th ACM Conference on Recommender Systems*. ACM, 161–169.
- [40] Dietmar Jannach, Lukas Lerche, Iman Kamehkhosh, and Michael Jugovac. 2015. What recommenders recommend: An analysis of recommendation biases and possible countermeasures. *User Model. User-Adapt. Interact.* 25, 5 (2015), 427–491.
- [41] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2014. Correcting popularity bias by enhancing recommendation neutrality. In *Poster Proceedings of the 8th ACM Conference on Recommender Systems, RecSys 2014, Foster City, Silicon Valley, CA, USA*, Li Chen and Jalal Mahmud (Eds.). CEUR-WS.org. <http://ceur-ws.org/Vol-1247>.
- [42] Donghyun Kim, Chanyoung Park, Jinoh Oh, Sungyoung Lee, and Hwanjo Yu. 2016. Convolutional matrix factorization for document context-aware recommendation. In *Proceedings of the 10th ACM Conference on Recommender Systems*. 233–240.
- [43] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 8 (2009), 30–37.
- [44] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2016. When can the maximin share guarantee be guaranteed? In *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI, 523–529.
- [45] Anja Lambrecht and Catherine Tucker. 2019. Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Manage. Sci.* 65, 7 (2019), 2966–2981. DOI : [10.1287/mnsc.2018.3093](https://doi.org/10.1287/mnsc.2018.3093)
- [46] Dawen Liang, Jaan Altosaar, Laurent Charlin, and David M. Blei. 2016. Factorization meets the item embedding: Regularizing matrix factorization with item co-occurrence. In *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, 59–66.
- [47] Carl Lin and Myeong-Su Yun. 2016. The effects of the minimum wage on earnings inequality: Evidence from China. In *Income Inequality Around the World*. Emerald Group Publishing Limited, 179–212.
- [48] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC'04)*. ACM, 125–131. <https://doi.org/10.1145/988772.988792>
- [49] Medium.com. 2019. Retrieved from <https://medium.com/public-market/in-their-own-words-why-sellers-are-fed-up-with-amazon-e97da44f7f18>.
- [50] NewYorkTimes. 2019. Retrieved from <https://www.nytimes.com/2019/09/09/business/economy/uber-lyft-california.html>.
- [51] Xia Ning, Christian Desrosiers, and George Karypis. 2015. A comprehensive survey of neighborhood-based recommendation methods. In *Recommender Systems Handbook*. Springer, 37–76.
- [52] Gourab K. Patro, Arpita Biswas, Niloy Ganguly, Krishna P. Gummadi, and Abhijnan Chakraborty. 2020. FairRec: Two-sided fairness for personalized recommendations in two-sided platforms. In *Proceedings of the Web Conference*. 1194–1204.
- [53] Gourab K. Patro, Abhijnan Chakraborty, Niloy Ganguly, and Krishna P. Gummadi. 2020. Incremental fairness in two-sided market platforms: On smoothly updating recommendations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 181–188.
- [54] Robert Pollin, Mark Brenner, Stephanie Luce, and Jeannette Wicks-Lim. 2008. *A Measure of Fairness: The Economics of Living Wages and Minimum Wages in the United States*. Cornell University Press.
- [55] Ariel D. Procaccia and Junxing Wang. 2014. Fair enough: Guaranteeing approximate maximin shares. In *ACM Conference on Economics and Computation, EC*. 675–692.
- [56] Amifa Raj, Connor Wood, Ananda Montoly, and Michael D. Ekstrand. 2020. Comparing fair ranking metrics. Retrieved from <https://arXiv:2009.01311>.
- [57] Matthew J. Salganik, Peter Sheridan Dodds, and Duncan J. Watts. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311, 5762 (2006), 854–856. <https://doi.org/10.1126/science.1121066>

- [58] Dimitris Serbos, Shuyao Qi, Nikos Mamoulis, Evaggelia Pitoura, and Panayiotis Tsaparas. 2017. Fairness in package-to-group recommendations. In *Proceedings of the World Wide Web Conference (WWW'17)*.
- [59] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2219–2228.
- [60] Slate.com. 2019. Retrieved from <https://slate.com/technology/2019/06/billion-dollar-bully-documentary-yelp.html>.
- [61] Hugo Steinhaus. 1948. The problem of fair division. *Econometrica* 16 (1948), 101–104.
- [62] Walter Stromquist. 1980. How to cut a cake fairly. *Amer. Mathe. Monthly* 87, 8 (1980), 640–644.
- [63] Tom Sühr, Asia J. Biega, Meike Zehlike, Krishna P. Gummadi, and Abhijnan Chakraborty. 2019. Two-sided fairness for repeated matchings in two-sided markets: A case study of a ride-hailing platform. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD'19)*.
- [64] Özge Sürer, Robin Burke, and Edward C. Malthouse. 2018. Multistakeholder recommendation with provider constraints. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 54–62.
- [65] Hal R. Varian. 1974. Equity, envy, and efficiency. *J. Econ. Theory* 9, 1 (1974), 63–91.
- [66] Wired.com. 2019. Retrieved from <https://www.wired.com/story/europes-new-rules-aim-make-online-marketplaces-more-fair/>.
- [67] Hong-Jian Xue, Xinyu Dai, Jianbing Zhang, Shujian Huang, and Jiajun Chen. 2017. Deep matrix factorization models for recommender systems. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'17)*, Vol. 17. 3203–3209.
- [68] Ke Zhou, Shuang-Hong Yang, and Hongyuan Zha. 2011. Functional matrix factorizations for cold-start recommendation. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 315–324.

Received November 2020; revised October 2021; accepted October 2021