

FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms

Gourab K Patro*
IIT Kharagpur, India

Arpita Biswas*
Indian Institute of Science, India

Niloy Ganguly
IIT Kharagpur, India

Krishna P. Gummadi
MPI-SWS, Germany

Abhijnan Chakraborty
MPI-SWS, Germany

ABSTRACT

We investigate the problem of fair recommendation in the context of two-sided online platforms, comprising customers on one side and producers on the other. 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 reveals that such customer-centric design may lead to unfair distribution of exposure among the producers, which may adversely impact their well-being. On the other hand, a producer-centric design might become unfair to the customers. Thus, we consider fairness issues that span both customers and producers. Our approach involves a novel mapping of the fair recommendation problem to a constrained version of the problem of fairly allocating indivisible goods. Our proposed *FairRec* algorithm guarantees at least Maximin Share (MMS) of exposure for most of the producers and Envy-Free up to One Good (EF1) fairness for every customer. Extensive evaluations over multiple real-world datasets show the effectiveness of *FairRec* in ensuring two-sided fairness while incurring a marginal loss in the overall recommendation quality.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Fair Recommendation, Two-Sided Markets, Fair Allocation, Maximin Share, Envy-Freeness

ACM Reference Format:

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 2020 (WWW '20)*, April 20–24, 2020, Taipei, Taiwan. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3366423.3380196>

1 INTRODUCTION

Popular online platforms such as Netflix, Amazon, Yelp, Spotify, Google Local provide recommendation services to help their

*Authors contributed equally to this work.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '20, April 20–24, 2020, Taipei, Taiwan

© 2020 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-7023-3/20/04.

<https://doi.org/10.1145/3366423.3380196>

customers browse through the enormous product spaces. By providing these services, the platforms control 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. These platforms have traditionally focused on maximizing customer satisfaction by tailoring the results 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, 17, 22, 28, 30]. As more and more people are depending on two-sided platforms to earn a living, recently platforms have 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 [39], fair marketplace laws for e-commerce [51]), (ii) **social responsibility** or **voluntary commitment** (e.g., equality of opportunity to all gender groups in LinkedIn [27], commitment of non-discrimination to hosts and guests by AirBnb [4]), (iii) **business requirement/model** (e.g., minimum business guarantee by AirBnb to attract hosts [5]).

In this paper, 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 [40], latent factorization methods [33, 36], 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 reveals that they can create a huge disparity in the exposure of the producers (detailed in §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. On the other hand, if only a few producers get most of the exposure, then the other producers would struggle on the platform, which will force them to either quit or switch to other platforms [21, 38, 46]. 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. Naive ways of reducing inequality (e.g., producer-centric poorest- k : selecting k least exposed producers) may result in loss and disparity in customer utilities (§4.2), making

it inefficient as well as unfair to the customers. We postulate that while being fair to the producers, the platforms should also attempt to fairly distribute the loss in utility among all the customers.

Considering this, we tackle the challenging task of ensuring two-sided fairness while giving personalized recommendations. Motivated by a vast literature in social choice theory, we map this problem to the problem of fairly allocating indivisible goods (§5). However, due to various constraints pertaining to the requirements of recommender systems, the problem becomes an interesting extension to the existing fair allocation problem—find an allocation that guarantees minimum exposure (upper bounded by maximin share of exposure or MMS) for the producers, and envy-free up to one item (EF1) [16] for the customers (proofs in §7). The MMS guarantee ensures that each agent receives a value which is at least their *maximin share* threshold, defined in §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.

Contributions: (i) we consider two-sided notions of fairness which not only relate to social or judicial precepts but also to the long-term sustainability of two-sided platforms (§4); (ii) we design an algorithm, *FairRec* (§6), exhibiting the desired two-sided fairness by mapping the fair recommendation problem to a fair allocation problem (§5); moreover, it is agnostic to the specifics of the data-driven model (that estimates the product-customer relevance scores) which makes it more scalable and easy to adapt; (iii) in addition to the theoretical guarantees (§7), extensive experimentation and evaluation over multiple real-world datasets deliver strong empirical evidences on the effectiveness of our proposal (§8).

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.

Fairness in Two-Sided Platforms: With the increasing popularity of multi-sided platforms, recently multiple researchers have looked into the issues of unfairness and biases in such platforms. For example, Edelman et al. [22] investigated the possibility of racial bias in guest acceptance by Airbnb hosts, Lambrecht and Tucker [35] studied gender-based discrimination in career ads. While these works deal with *group fairness*, Serbos et al. [45] proposed an *envy free* tour package recommendations on travel booking sites, ensuring *individual fairness for customers*.

On producer fairness, Hannák et al. [30] studied racial and gender bias in freelance marketplaces. In a social experiment, Salganik et al. [44] found that the existing popular producers often acquire most of the visibility while new but good ones starve for visibility. Biega et al. [9] considered individual producer fairness in ranking in gig-economy platforms. Kamishima et al. [32] and Abdollahpouri et al. [2] reduced popularity bias among producers while Patro et al. [41] addressed fairness issues arising due to frequent updates of platforms. However, these papers did not study the trade-off between producer and customer fairness, and the cost of achieving one over the other.

A few papers have considered group-fairness among the producers and customers. Abdollahpouri et al. [1] and Burke [17] categorized different types of multi-stakeholder platforms and their

desired group fairness properties, Chakraborty et al. [20] and Sühr et al. [49] presented mechanisms for two-sided fairness in matching problems. In contrast, our paper addresses individual fairness for both producers and customers, which also answers the question of the long-term sustainability of two-sided platforms.

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) [26, 48, 50] and proportional fair share (PFS) [47]. Recent literature on practical applications of fair allocation [15, 23] have focused on the problem of allocating *indivisible* goods in budgeted course allocation [16], balanced graph partition [13], or allocation of cardinality constrained group of resources [11]. In such instances, no feasible allocation may satisfy EF or PFS fairness guarantees. Thus, the notable work of Budish [16] defined analogous fairness notions which are appropriate for *indivisible* goods—namely, envy-freeness up to one good (EF1) and maximin share guarantee (MMS). A rich literature has focused on providing existence and algorithmic guarantees for computing fair allocations [6–8, 10, 12, 14, 19, 34, 43]. In this work, 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 §5).

3 PRELIMINARIES

Next, we define the terminology and notations used in the paper.

3.1 Producers and Products

In platforms like Google Maps and Yelp, a producer typically owns one product (restaurant or shop); whereas in multimedia platforms like Spotify, YouTube, and Netflix, the songs/videos generated by the artists are the products, and one producer can list many products; same is true for ecommerce platforms such as Amazon and Flipkart. To generalize all such 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 would also ensure fairness for individual producers; here, fairness can be ensured by making the exposure proportional to the producer’s portfolio size.

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, 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 paper, we assume a uniform attention model¹ 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 $\sum_{p \in P} E_p = m \times k$.

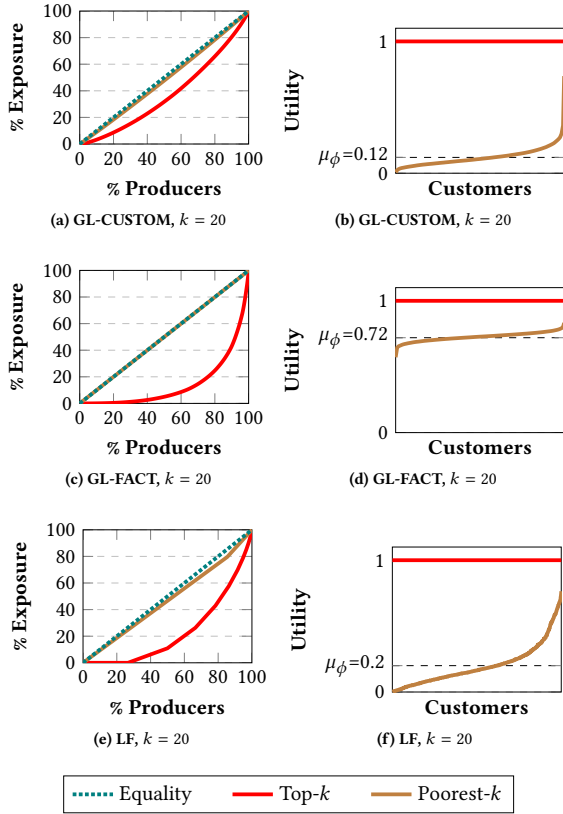


Figure 1: Lorenz curves show high inequality among producer exposures with top- k recommendation. While poorest- k provides almost equal exposures, it introduces huge loss and disparity in individual customer utilities.

¹This being the first work on two-sided-fair-recommendation posed as a fair-allocation problem, we focused on a basic setting without *position bias* [3], where customers pay more attention to the top ranked products than the lower ranked ones.

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 [36, 40] 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 the above approach is followed to maximize the satisfaction of individual customers, it can adversely affect the producers in a two-sided platform, as we explore next.

4.1 Datasets

We consider the impact of customer-centric top- k recommendations on producer exposures using real-world datasets. We use a state-of-the-art relevance scoring model (a very widely used latent factorization method [36]) 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 places on Google Maps (as Google Nearby feature) platform. We use the Google Local dataset released by He et al. [31], which contains data about customers, local businesses (producers), and their locations (geographic coordinates), ratings, etc. 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 11172 customers, 855 businesses and 25686 reviews. We consider the following two relevance scoring functions (V).

A. 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 .

B. GL-FACT: Here we use the state-of-the-art latent factorization model [33, 36] to predict the relevance scores from the ratings.

4.1.2 Last.fm Dataset (LF).

We use the `Last.fm` dataset released by Cantador et al. [18], which contains 1892 customers, 17632 artists (producers), and 92834 records of play counts (the number of times a customer has played songs from an artist). We again use a latent factorization model [33, 36] to find out the relevance scores from the play counts.

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-1a,1c,1e are the Lorenz curves for producer exposures. In Exposure Lorenz curves, 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*, *GL-FACT*, and *LF* datasets, 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.

4.3 Pitfalls of Naive 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-1a,1c,1e: 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-1b,1d,1f). Moreover, poorest- k introduces disparity among individual customer utilities, where some customers may suffer much higher loss than other customers, making it unfair to them.

4.4 Desiderata of Fair Recommendation

Intuitively, we need 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 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) [25, 29, 42]). Ensuring minimum wage does not itself guarantee equality of income; however it has been found to decrease income inequality [24, 37]. The exact value of the minimum exposure guarantee (\bar{E}) can be decided by the respective platforms.

B. Customer Fairness: As maintaining producer fairness can cause an overall loss in customer utility; that loss 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 next).

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, 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-freeness² (EF) and proportional-fair-share³ (PFS), may not be achievable in most instances of the problem. For example, if there is one item and two agents, 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 [16]. Since then, these have been extensively studied in various settings for providing existential and algorithmic guarantees [6–8, 10–15, 19, 34, 43]. 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 $\mathcal{V}_u(A_u) \geq \text{MMS}_u$.

5.2 Fair Recommendation to Fair Allocation

We propose to see the problem of fair recommendation 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); 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 remains 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). Thus, based on desired exposure guarantee \bar{E} ($\bar{E} \leq \text{MMS}$), we can make \bar{E} copies of each product (for producer fairness as in §4.4) to construct \mathcal{P} .
- **Fair Allocation of \mathcal{P} among \mathcal{U} :** Once \mathcal{P} is set according to desired producer fairness, the entire task of fair recommendation boils down to the allocation of \mathcal{P} among \mathcal{U} while

²An 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 [26, 48, 50].

³An 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 [47].

ensuring EF1 for agents/customers (for customer fairness, §4.4).

5.3 Extending the Conventional Fair Allocation Problem

Traditionally, *fair allocation* literature aims at defining and ensuring *fairness* among the agents while allocating all the items belonging to the set \mathcal{P} . 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 (say, each product should 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 [11]. In this paper, 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 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 two-sided fairness guarantees make fair recommendation an interesting extension of the fair allocation 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} which satisfies the desired two-sided fairness described in §4.4 (we prove the theoretical guarantees in §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.

The algorithm 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 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 $\ell = \lfloor \frac{mk}{n} \rfloor$ copies of each product (note that $\lfloor \frac{mk}{n} \rfloor$ is the maximin value of any producer while mk slots are allocated among n producers). It 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

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 $\ell \leftarrow \lfloor \frac{mk}{n} \rfloor$ 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 in order 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: $[\mathcal{C}, F, x] \leftarrow \text{Greedy-Round-Robin}(m, n, S, T, V, \sigma, F)$.
 - 21: Assign $\mathcal{A} \leftarrow \mathcal{A} \cup \mathcal{C}$.
 - 22: **end if**
 - 23: **Return** \mathcal{A} .
-

receives at most one copy of the same product. Then, assuming an arbitrary ordering σ of customers, ALG 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 ALG 2 is called again with a new ordering obtained by x left-cyclic rotations of σ . The remaining number of items are 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.

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  $x$ .

```

Both phases use a modified version of the Greedy-round-robin algorithm (ALG 2) [11, 19]: 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) $T == 0$: 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 $u \in [m]$.

7 THEORETICAL GUARANTEES

In this section, we provide a few important properties of ALG 2 in Proposition 7.1. Later, we establish the fairness guarantees and time complexity of our proposed algorithm *FairRec* in Theorem 7.2 using Lemma 7.2.1, 7.2.2 and 7.3.1. Note that, for all the proofs, we have fixed $\alpha = 1$ or $\bar{E} = \text{MMS}$.

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

- (P₁) 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., $V_{\sigma(x)}(B_{\sigma(x)}) \geq V_{\sigma(x)}(B_{\sigma(y)})$.

(P₂) the allocation \mathcal{B} obtained by ALG 2 is EF1.

(P₃) each customer is allocated at most one item from the same producer, thus ensuring the cardinality constraint is satisfied for each producer (category).

(P₄) for any two customers, say u and w , the allocation \mathcal{B} obtained by ALG 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 [11] and Caragiannis et al. [19], respectively. For completeness, we repeat the arguments towards 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:

$$\begin{aligned}
V_{\sigma(y)}(p_{\sigma(y)}^r) &\geq V_{\sigma(y)}(p_{\sigma(x)}^{r+1}) \quad \text{for all } r \in \{1, \dots, R-1\} \\
\Rightarrow \sum_{r=1}^{R-1} V_{\sigma(y)}(p_{\sigma(y)}^r) &\geq \sum_{r=1}^{R-1} V_{\sigma(y)}(p_{\sigma(x)}^{r+1}) \\
\Rightarrow V_{\sigma(y)}(B_{\sigma(y)}) &\geq V_{\sigma(y)}(B_{\sigma(x)}) - V_{\sigma(y)}(p_{\sigma(x)}^1) \quad (2)
\end{aligned}$$

Equation 2 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 ALG 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 ALG 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 till 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 MMS guarantee among at least $n - k$ producers.*

PROOF. We prove the fairness guarantees of customers and producers in Lemma 7.2.1 and 7.2.2, respectively. In Lemma 7.3.1, 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 ALG 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 ALG 2 to obtain the allocation 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

LEMMA 7.2.2. *Given n producers, m customers, and a positive integer k (such that $k < n \leq mk$), 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 ALG 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 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. On the other hand, the producers in the set B_u (the set recommended to customer u) is allocated to at least one producer. Thus, 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 7.3.

COROLLARY 7.3. *Given n producers, a positive integer k , and $\beta \in (0, 1)$ such that $k \leq \beta n$, FairRec ensures MMS-fairness among at least $(1 - \beta)n$ producers.*

Finally, in Lemma 7.3.1, we show that FairRec executes in polynomial time.

LEMMA 7.3.1. *The time complexity of FairRec has a worst case bound of $O(mnk)$.*

PROOF. The time complexity of FairRec is same as the complexity of ALG 2. Over the two phases, ALG 2 allocates mk items. For each allocation, it finds the maximum possible feasible producer which 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 (§6) for all the datasets (as listed in §4.1) considering different values of the recommendation-size k . For comparison, we use the following methods as baselines:

- (1) **Top- k** : recommending the top- k relevant products,
- (2) **Random- k** : randomly recommending k products,
- (3) **Mixed- k** : choosing top $\lfloor \frac{k}{2} \rfloor$ relevant products at first and then the remaining $(k - \lfloor \frac{k}{2} \rfloor)$ randomly.
- (4) **Poorest- k (PR- k)**: this is a producer-centric method where k least exposed products are recommended to each customer in a round robin manner.

We run two sets of experiments. First, we set the exposure guarantee $\bar{E} = \text{MMS}$ (in §8.1). Next, we set lower exposure guarantees i.e., by considering $\bar{E} = \alpha \cdot \text{MMS}$ where $0 \leq \alpha \leq 1$ (in §8.2). 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 below.

$$H = \frac{1}{|P|} \sum_{p \in P} \mathbb{1}_{E_p \geq \bar{E}} \quad (3)$$

$\mathbb{1}_{E_p \geq \bar{E}}$ 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 §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 below.

$$Z = - \sum_{p \in P} \left(\frac{E_p}{m \times k} \right) \cdot \log_n \left(\frac{E_p}{m \times k} \right) \quad (4)$$

The range of Z is $[0, 1]$; when all the producers get equal exposure, then Z becomes 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

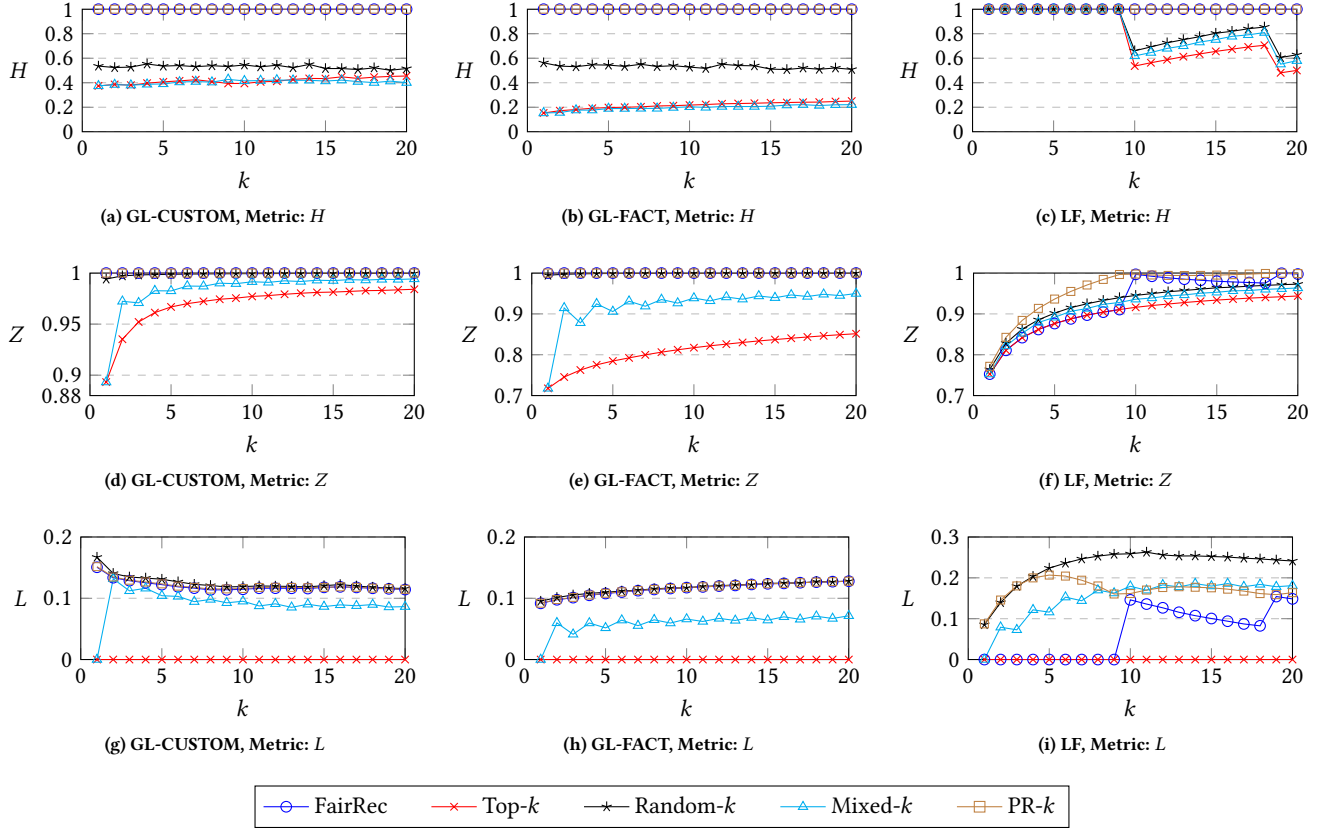


Figure 2: Producer-Side Performances with MMS Guarantee. First row: fraction of satisfied producers (H). Second row: inequality in producer exposures (Z). Third row: exposure loss on producers (L).

the loss L as the mean amount of impact (loss in exposure) caused by *FairRec*, compared to the top- k recommendations.

$$L^{\langle \text{method} \rangle} = \frac{1}{n} \sum_{p \in P} \max \left(\frac{(E_p^{\text{top-}k} - E_p^{\langle \text{method} \rangle})}{E_p^{\text{top-}k}}, 0 \right) \quad (5)$$

The lower the negative impact, 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 below.

$$Y = \frac{1}{n} \sum_{u \in U} \frac{1}{n-1} \sum_{\substack{u' \in U \\ u' \neq u}} \text{envy}(u, u') \quad (6)$$

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.

Loss and Disparity in 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 $\mu_\phi = \frac{1}{m} \sum_{u \in U} \phi_u(R_u)$. The lower is the 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 , 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.

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

Producer Satisfaction (H): Figures 2a, 2b, and 2c show that both *FairRec* and PR- k perform the best while top- k , mixed- k perform the worst; this is because both *FairRec* and PR- k try to ensure larger exposure for producers while top- k , mixed- k consider only the preferences of the customers;

Exposure Inequality (Z): Figures 2d, 2e, and 2f show that PR- k

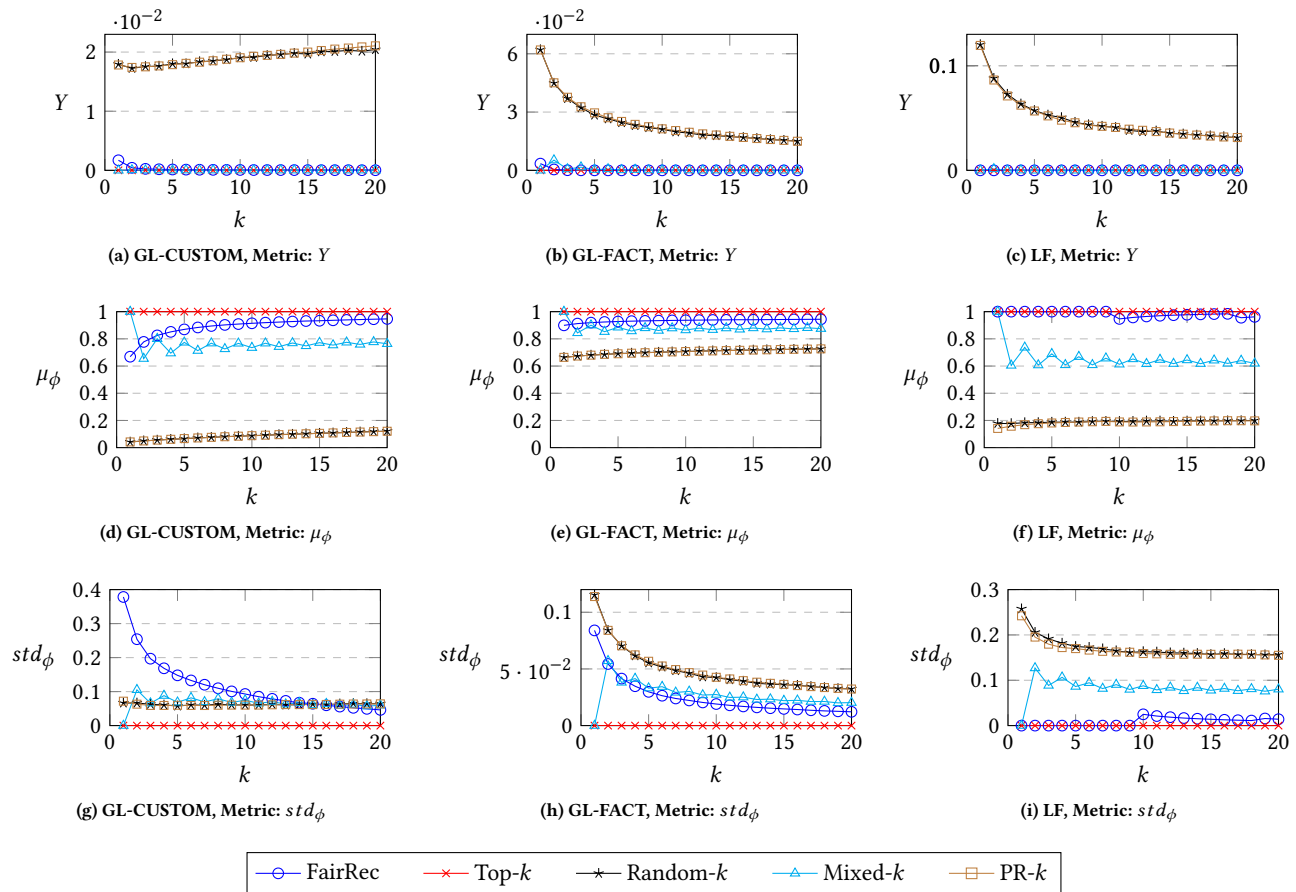


Figure 3: Customer-Side Performances with MMS Guarantee. First row: mean average envy (Y). Second row: mean customer utility (μ_ϕ). Third row: standard deviation of customer utilities (σ_ϕ).

has lowest inequality in exposure while *FairRec* and random- k perform similar or slightly less than that; on the other hand top- k and mixed- k perform the worst as they are highly customer-centric.

Exposure Loss (L): Figures 2g, 2h, and 2i show that random- k and PR- 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 PR- k tries to increase the exposure of least exposed producer). On the other hand, mixed- k causes smaller losses as only up to half of its recommendations are random. *FairRec* causes only up to 0.2 fraction or 20% loss in exposure in comparison to top- k owing to the intelligent selection approach of *FairRec*.

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 $k \in [20, 29]$,...). MMS is satisfied for all producers till $k = 9$; 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*, PR- 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 3a, 3b, and 3c reveal that top- k causes lowest possible mean average envy among the customers; this is because it gives maximum possible utility of 1 to every customer thereby leaving no chances of envy among customers. Other customer-centric mixed- k also shows similarly low envy. *FairRec* generates very low values of envy which are very comparable to those of top- k here. On the other hand both random- k and PR- k cause the highest envy as they do not consider customer preferences at all during recommendation.

Loss and Disparity in Customer Utility (μ_ϕ, std_ϕ): From Figures 3d, 3e, and 3f, we see that both random- k and PR- k cause huge loss in customer utility as they neglect customer preferences. On the other hand, while mixed- k performs moderately, *FairRec* causes very less utility loss and performs almost at par with the customer-centric top- k . This certifies that *FairRec* strikes a good balance between customer utility and producer fairness. The standard deviation plots: Figures-3g, 3h, and 3i reveal that for larger sizes of recommendation, random- k and PR- k show large disparities in customer utilities while *FairRec* and mixed- k show relatively less

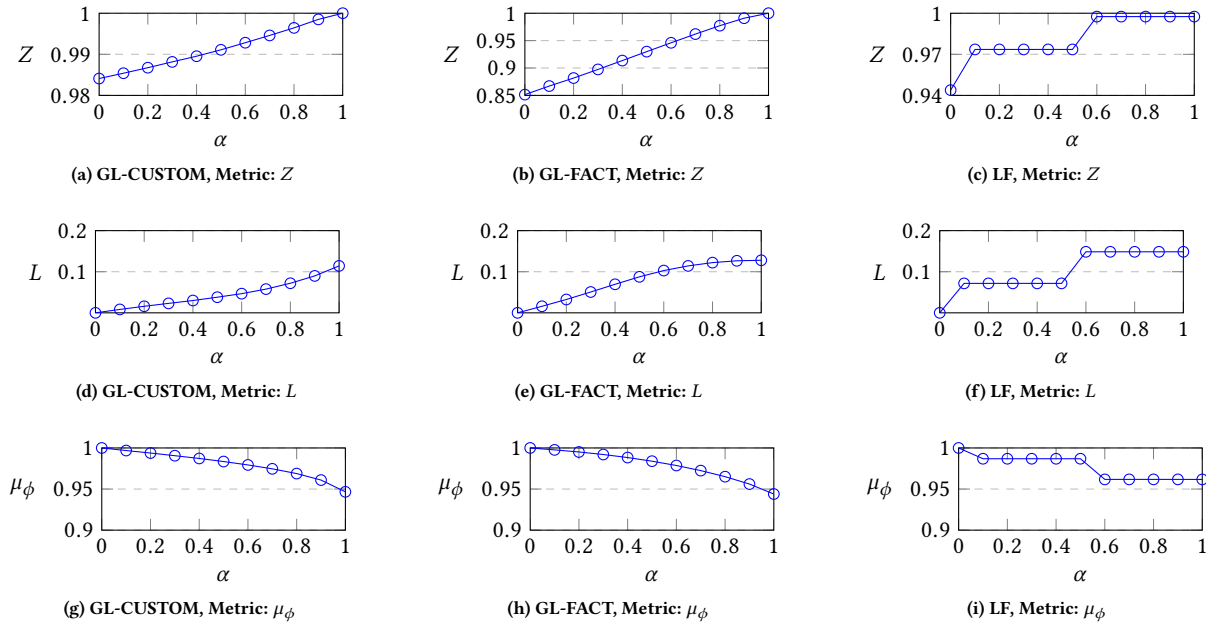


Figure 4: Performances with $\bar{E} = \alpha$ MMS guarantee for $k = 20$. Plots show that higher exposure guarantee can achieve better producer fairness but can cause exposure loss for very popular producers and utility loss for the customers.

disparities. As top- k is customer-centric and provides the maximum utility of 1 to all the customers, it shows standard deviation of 0.

8.2 Experiments with α -MMS Guarantee

Here we fix $k = 20$, and test *FairRec* with different values of minimum exposure guarantee i.e., $\bar{E} = \left\lfloor \alpha \cdot \frac{mk}{n} \right\rfloor$ (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); although we do not see much change in H , Y , std_ϕ , we do find some interesting insights from the variations observed in other metrics (detailed next).

(i) Increasing minimum exposure guarantee results in lower inequality in producer exposures. We observe the direct correlation of Z with the α value (refer Figures-4a,4b,4c). High Z signifies lower inequality.

(ii) Increasing minimum exposure guarantee can cause higher exposure losses for previously popular producers. Increasing values of L or mean exposure loss are observed (refer Figures-4d, 4e, 4f) for higher α .

(iii) Higher exposure guarantee for producers can negatively impact overall customer utility. Figures-4g, 4h, 4i show that mean customer utility decreases with higher α or exposure guarantee for the producers.

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

9 CONCLUSION

In this work, 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 evidences through extensive evaluations on real-world datasets. Our work can be directly applied to fair recommendation problems in scenarios like mass recommendation/promotion sent through emails, 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 [3], where customers pay more attention to the top ranked products than the lower ranked ones.

ACKNOWLEDGMENTS

G. K Patro is supported by TCS Research Fellowship. A. Biswas gratefully acknowledges the support of a Google Ph.D. Fellowship Award. This research was supported in part by a European Research Council (ERC) Advanced Grant for the project “Foundations for Fair Social Computing”, funded under the EU Horizon 2020 Framework Programme (grant agreement no. 789373).

Reproducibility: All the codes and datasets details are available at https://github.com/gourabkumarpatro/FairRec_www_2020.

REFERENCES

- [1] Himan Abdollahpour and Robin Burke. 2019. Multi-stakeholder Recommendation and its Connection to Multi-sided Fairness. *arXiv preprint arXiv:1907.13158* (2019).
- [2] Himan Abdollahpour, Robin Burke, and Bamshad Mobasher. 2017. Controlling popularity bias in learning-to-rank recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. ACM, 42–46.
- [3] Aman Agarwal, Ivan Zaitsev, Xuanhui Wang, Cheng Li, Marc Najork, and Thorsten Joachims. 2019. Estimating position bias without intrusive interventions. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. ACM, 474–482.
- [4] AirBnb. 2019. <https://blog.airbnb.com/the-airbnb-community-commitment/>. (2019).
- [5] AirBnb. 2019. <https://www.airbnb.com/help/article/2474/what-is-the-earnings-guarantee-programme-for-new-hosts>. (2019).
- [6] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2018. Comparing Approximate Relaxations of Envy-Freeness. In *International Joint Conference on Artificial Intelligence*. IJCAI, 42–48.
- [7] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. 2015. Approximation Algorithms for Computing Maximin Share Allocations. In *International Colloquium on Automata, Languages, and Programming, ICALP*. 39–51.
- [8] Siddharth Barman, Arpita Biswas, Sanath Kumar Krishnamurthy, and Y. Narahari. 2018. Groupwise Maximin Fair Allocation of Indivisible Goods. In *AAAI Conference on Artificial Intelligence*. 917–924.
- [9] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of Attention: Amortizing Individual Fairness in Rankings. *arXiv preprint arXiv:1805.01788* (2018).
- [10] 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 *11th Innovations in Theoretical Computer Science (ITCS)*.
- [11] Arpita Biswas and Siddharth Barman. 2018. Fair Division Under Cardinality Constraints. In *International Joint Conference on Artificial Intelligence*. IJCAI, 91–97.
- [12] Arpita Biswas and Siddharth Barman. 2019. Matroid Constrained Fair Allocation of Goods. In *AAAI Conference on Artificial Intelligence*. AAAI, 9921–9922.
- [13] Sylvain Bouveret, Katarina Cechlarová, Edith Elkind, Ayumi Igarashi, and Dominik Peters. 2017. Fair Division of a Graph. In *International Joint Conference on Artificial Intelligence*. IJCAI, 135–141.
- [14] Sylvain Bouveret and Michel Lemaître. 2014. Characterizing Conflicts in Fair Division of Indivisible Goods using a Scale of Criteria. In *International Conference on Autonomous Agents and Multi-Agent Systems*. 1321–1328.
- [15] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. 2016. *Handbook of Computational Social Choice, Chapter 12*. Cambridge University Press. <https://doi.org/10.1017/CBO9781107446984>
- [16] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- [17] Robin Burke. 2017. Multisided fairness for recommendation. *arXiv preprint arXiv:1707.00093* (2017).
- [18] 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* (Chicago, IL, USA) (*RecSys 2011*). ACM, New York, NY, USA.
- [19] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *ACM Conference on Economics and Computation, EC*. ACM, 305–322.
- [20] Abhijnan Chakraborty, Aniko Hannak, Asia J Biega, and Krishna P Gummadi. 2017. Fair Sharing for Sharing Economy Platforms. (2017).
- [21] CNBC.com. 2017. <https://www.cnbc.com/2017/04/20/only-4-percent-of-uber-drivers-remain-after-a-year-says-report.html>. (2017).
- [22] Benjamin Edelman, Michael Luca, and Dan Svirsky. 2017. Racial discrimination in the sharing economy: Evidence from a field experiment. *American Economic Journal: Applied Economics* 9, 2 (2017).
- [23] Ulle Endriss. 2017. *Trends in Computational Social Choice*. Lulu. com.
- [24] 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.
- [25] Armin Falk, Ernst Fehr, and Christian Zehnder. 2006. Fairness perceptions and reservation wages—the behavioral effects of minimum wage laws. *The Quarterly Journal of Economics* 121, 4 (2006), 1347–1381.
- [26] DC Foley. 1967. Resource Allocation in the Public Sector. *Yale Economic Essays* 7 (1967), 73–76.
- [27] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. 2019. Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search. In *ACM KDD*.
- [28] 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: European Review of Labour and Research* 23, 2 (2017), 135–162.
- [29] David A Green and Kathryn Harrison. 2010. *Minimum wage setting and standards of fairness*. Technical Report. IFS working papers.
- [30] Anikó Hannák, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson. 2017. Bias in online freelance marketplaces: Evidence from taskrabit and fiverr. In *ACM CSCW*.
- [31] Ruining He, Wang-Cheng Kang, and Julian McAuley. 2017. Translation-based recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. ACM, 161–169.
- [32] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2014. Correcting Popularity Bias by Enhancing Recommendation Neutrality.. In *RecSys Posters*.
- [33] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 8 (2009), 30–37.
- [34] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2016. When Can the Maximin Share Guarantee Be Guaranteed?. In *AAAI Conference on Artificial Intelligence*. AAAI, 523–529.
- [35] Anja Lambrecht and Catherine Tucker. 2019. Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads. *Management Science* (2019).
- [36] Dawen Liang, Jaan Altonaar, 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.
- [37] 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.
- [38] Medium.com. 2019. <https://medium.com/public-market/in-their-own-words-why-sellers-are-fed-up-with-amazon-e97da44f7f18>. (2019).
- [39] NewYorkTimes. 2019. <https://www.nytimes.com/2019/09/09/business/economy/uber-lyft-california.html>. (2019).
- [40] Xia Ning, Christian Desrosiers, and George Karypis. 2015. A comprehensive survey of neighborhood-based recommendation methods. In *Recommender systems handbook*. Springer, 37–76.
- [41] Gourab K Patro, Abhijnan Chakraborty, Niloy Ganguly, and Krishna P Gummadi. 2020. Incremental Fairness in Two-Sided Market Platforms: On Smoothly Updating Recommendations. *AAAI, Feb* (2020).
- [42] 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.
- [43] Ariel D Procaccia and Junxing Wang. 2014. Fair enough: Guaranteeing approximate maximin shares. In *ACM Conference on Economics and Computation, EC*. 675–692.
- [44] 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>
- [45] Dimitris Serbos, Shuyao Qi, Nikos Mamoulis, Evaggelia Pitoura, and Panayiotis Tsaparas. 2017. Fairness in package-to-group recommendations. In *WWW*.
- [46] Slate.com. 2019. <https://slate.com/technology/2019/06/billion-dollar-bully-documentary-yelp.html>. (2019).
- [47] Hugo Steinhaus. 1948. The Problem of Fair Division. *Econometrica* 16 (1948), 101–104.
- [48] Walter Stromquist. 1980. How to Cut a Cake Fairly. *The American Mathematical Monthly* 87, 8 (1980), 640–644.
- [49] 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 *ACM KDD*.
- [50] Hal R Varian. 1974. Equity, Envy, and Efficiency. *Journal of Economic Theory* 9, 1 (1974), 63–91.
- [51] Wired.com. 2019. <https://www.wired.com/story/europes-new-rules-aim-make-online-marketplaces-more-fair/>. (2019).