

Incremental Fairness in Two-Sided Market Platforms: On Smoothly Updating Recommendations

Gourab K Patro

Max Planck Inst. for Software Systems, Germany
Indian Institute of Technology Kharagpur, India

Abhijnan Chakraborty

Max Planck Inst. for Software Systems, Germany

Niloy Ganguly

Indian Institute of Technology Kharagpur, India

Krishna P.Gummadi

Max Planck Inst. for Software Systems, Germany

Abstract

Major online platforms today can be thought of as two-sided markets with producers and customers of goods and services. There have been concerns that over-emphasis on customer satisfaction by the platforms may affect the well-being of the producers. To counter such issues, few recent works have attempted to incorporate fairness for the producers. However, these studies have overlooked an important issue in such platforms – to supposedly improve customer utility, the underlying algorithms are frequently updated, causing abrupt changes in the exposure of producers. In this work, we focus on the fairness issues arising out of such frequent updates, and argue for incremental updates of the platform algorithms so that the producers have enough time to adjust (both logistically and mentally) to the change. However, naive incremental updates may become unfair to the customers. Thus focusing on recommendations deployed on two-sided platforms, we formulate an ILP based online optimization to deploy changes incrementally in η steps, where we can ensure smooth transition of the exposure of items while guaranteeing a minimum utility for every customer. Evaluations over multiple real world datasets show that our proposed mechanism for platform updates can be efficient and fair to both the producers and the customers in two-sided platforms.

1 Introduction

Many popular online platforms today can be thought of as *two-sided markets*, such as, sharing economy platforms like Uber, Lyft or Airbnb, e-commerce sites like Amazon, news aggregation services like Google News, location-based review and recommendation services like Yelp, Google Local, employment sites like LinkedIn, Indeed, or hotel aggregators like Booking.com. There are three stakeholders in these markets: (i) *producers* of goods and services (e.g., sellers in Amazon, hosts in Airbnb), (ii) *customers* who pay for them, and (iii) the *platform* at the center of the ecosystem. Services on these platforms have traditionally been designed to maximize customer satisfaction, since they are the ones directly contributing to the platform revenue, largely ignoring the interest of the other key stakeholder – *producers*.

Several recent studies have shown how sole focus on the customers may adversely affect the well-being of the producers, as more and more people are depending on two-sided platforms to earn a living (Edelman, Luca, and Svirsky 2017; Hannák et al. 2017; Abdollahpouri and Burke 2019; Chakraborty et al. 2017; Burke 2017; Graham, Hjorth, and Lehdonvirta 2017). Subsequently, few research works have attempted to reduce unfairness in these platforms (Sühr et al. 2019; Süreer, Burke, and Malthouse 2018; Geyik, Ambler, and Kenthapadi 2019). However, existing works have overlooked an important issue. They assume that the underlying platform algorithms remain unchanged; whereas, to offer supposedly higher customer utility, the platform algorithms go through frequent changes and updates. Such updates are often very rapid and immediate, leaving no room for the producers to adjust to them. For example, with every change in the Facebook Newsfeed curation algorithm, news media outlets (i.e., the producers of news stories) complain about immediate drop in traffic to their websites (Lab 2019; AdExchanger 2018). Similar complaints have been reported also for other two-sided platforms like Amazon (Vox 2019).

While maximizing customer’s utility may be paramount, we argue that the platform also needs to be fair to the producers while updating. Multiple works in behavioral economics have shown that *human perceptions of fairness* of a new decision making system are influenced by how far the decision outcomes change from the *status quo* (i.e., the existing outcomes) (Bediou and Scherer 2014; Kahneman, Knetsch, and Thaler 1991). Motivated by this line of work, in this paper, we propose the notion of **Egocentric Fairness** for the producers, which requires that the impact of the changed system should be limited/small. We argue that a simple way of being fair would be to implement the change in a phased manner. This also has its practical advantage whereby a producer gets time to adjust to the change in demand.

One naive way of incremental update would be to change the platform for only a subset of customers and then gradually cover everyone. However, such an approach may be unfair to the customers. Since the change is supposed to provide higher satisfaction (utility) for the customers, those who experience the changed platform earlier will get higher utility than the customers covered later. To ensure fairness

to the customers, we formulate a constrained optimization problem whereby at every stage **every customer is guaranteed a minimum utility**, while the average change in exposure of the producers should be minimal. To model this update mechanism, in this paper, we focus on the recommendation systems deployed on two-sided platforms (we consider Amazon products and Google Local datasets), and consider three common types of updates: (i) addition of new features to better estimate customer preferences, (ii) deployment of a new recommendation algorithm reflecting technological advances, or (iii) addition of more data/customer feedback to account for the ever-changing choices of the customers.

The paper progresses in the following fashion, in §4 we introduce the datasets and update types, and perform a detailed experiment to show the impact of immediate update on the producers. The findings help us to succinctly define fairness from the perspective of producers and customers. Based upon this understanding of both-sided fairness, the constrained optimization formulation is developed in §5. The formulation takes into consideration several practical details – for example, optimization has to be performed at the level of each customer arrival and one may or may not have an estimate of the amount of changes which would happen if an update is applied. The experimental results show that both efficiency and fairness are ensured to the producers as well as the customers; the experiments bring forward the lacunae of updating algorithms popular in software engineering domain (used as baselines). To our knowledge, this is the first paper which focuses on issues associated with updates on two-sided platforms, and we believe that this work will be an important addition to the growing literature on fairness of algorithmic decision making systems.

2 Background and Related work

Fairness in Multi-Sided Platforms: Recently, few works have looked into the issues of unfairness and biases in platforms with multiple stakeholders. Disparity in customer utilities has led to the concerns of both individual and group fairness for customers. For example, studies have found instances of *group unfairness* – gender-based discrimination in career ads (Lambrecht and Tucker 2019), or racial bias in guest acceptance by Airbnb hosts (Edelman, Luca, and Svirsky 2017). On the other hand, (Serbos et al. 2017) have looked into *individual customer fairness* by studying the problem of *envy free* tour package recommendations on travel booking sites. Similarly, producer fairness relates to the disparity in producer utilities, and touches both group and individual fairness. For example, (Hannák et al. 2017) found racial and gender bias in ratings of freelancers on freelance marketplaces, (Chakraborty et al. 2019) proposed methods to ensure fair representation to different user groups in social media, (Geyik, Ambler, and Kenthapadi 2019) proposed fair exposure to candidates from different age and gender groups in LinkedIn. (Biega, Gummedi, and Weikum 2018) considered individual producer fairness in ranking in gig-economy platforms.

Few recent works have also explored fairness for both producers and customers. For example, (Abdollahpouri and Burke 2019; Burke 2017) categorized different types of

multi-stakeholder platforms and their desired fairness properties, (Sühr et al. 2019) presented a mechanism for two-sided fairness in matching problems, (Sürer, Burke, and Malthouse 2018) used minimum guarantee constraints for producers and diversity constraints for customers while recommending. However, these works have assumed that the underlying customer-item relevance model remains unchanged, whereas in reality, the algorithms go through frequent updates. In this paper, we focus on fairness issues arising out of such platform updates in multi-sided platforms.

Egocentric Perceptions of Fairness: Multiple research works have documented the existence of *egocentric biases* in what people perceive as fair. Through experiments in game theory (more specifically, Dictator Games and Ultimatum Games), researchers have observed that individuals take fairness concerns (such as equality) into account while distributing goods among players, and such concerns often originates from one’s *sense of endowment* (Bediou and Scherer 2014). Such endowment effect has also been studied in behavioral economics (Morewedge and Giblin 2015), where researchers found that individuals perceive a new system to be fair if the new outcomes are similarly beneficial as their *status quo* outcomes from the existing system (Kahneman, Knetsch, and Thaler 1991). Following this line of work, in this paper, we define the notions for *egocentric fairness* for producers in two-sided platforms and propose mechanisms to achieve the same.

Incrementalism: Incrementalism is a well-studied discipline in public policy making, which advocates for creating policies in iterations where new policy will build upon past policies, incorporating incremental rather than wholesale changes (Hayes 1992). Similar to policy issues, we argue for incremental algorithmic changes in two-sided platforms to limit large disruptive changes.

Minimum Utility Guarantee: (Pollin et al. 2008; Green and Harrison 2010; Falk, Fehr, and Zehnder 2006) proposed minimum wage guarantee as a fairness standard, and (Lin and Yun 2016; Engbom and Moser 2018) showed evidences of how minimum wage guarantee decreases income inequality. Inspired by these works, we propose notion of minimum utility guarantee for customer fairness.

3 Notations and Terminology

In this paper, U , P , S denote the sets of customers, producers, and items respectively. S_p represents the set of all items listed by a producer p such that $\bigcup_{p \in P} S_p = S$. R_u represents the set of k items recommended to customer u ; $R_u \subset S$, $|R_u| = k$. We assume k to be the same for every customer. Next, we define the terms used in the paper.

Relevance of Items: Relevance of an item s to a customer u represents the likelihood that u would prefer s . Formally, we can define relevance as a real function of customer and item; $V : U \times S \rightarrow \mathbb{R}$, and $V(u, s)$ denotes the relevance of item s to customer u . Alternatively, we can consider $V(u, s)$ as the amount of utility gained by customer u if item s is recommended to her.

Customer Utility: The utility of recommendation R_u to u w.r.t. a particular relevance function V can be written as;

$\phi(R_u, V) = \sum_{s \in R_u} V(u, s)$. u will get the maximum possible utility if k most relevant items – R_u^* , is recommended to her; $\phi_u^{\max}(V) = \phi(R_u^*, V) = \sum_{s \in R_u^*} V(u, s)$. A normalized form of customer utility from a recommendation R_u would be: $\phi^{\text{norm}}(R_u, V) = \frac{\phi(R_u, V)}{\phi_u^{\max}(V)} = \frac{\phi(R_u, V)}{\phi(R_u^*, V)}$.

Producer Exposure: The utility of a producer is the total amount of exposure/visibility its items get through recommendations. The exposure of an item s is the total amount of attention it receives from all the customers to whom s has been recommended. In an online scenario, U can be thought of as the sequence of customer-visits to the platform where some customers may visit multiple times. If $U([t_1, t_2])$ is the sequence of customer-logins in the interval $[t_1, t_2]$, then the exposure of an item s in the same interval will be $E_s([t_1, t_2]) = \frac{1}{k} \sum_{u \in U([t_1, t_2])} \mathbb{1}_{R_u}(s)$, and that of a producer p will be $E_p([t_1, t_2]) = \frac{1}{k} \sum_{s \in S_p} \sum_{u \in U([t_1, t_2])} \mathbb{1}_{R_u}(s)$ ¹.

Note that, in this work, we assume that customers pay similar attention to all k recommended items, and leave the consideration of position bias (i.e., top-ranked items may get more attention than low-ranked ones) for future work. We further assume an one-to-one correspondence between producers and items, and **henceforth use the terms ‘item’ and ‘producer’ interchangeably**. This assumption is valid for multiple platforms such as restaurant reservation/food delivery (Yelp, Google Local, Uber Eats), freelance marketplaces (Fiverr), human resource matchmaking (LinkedIn, Naukri), and so on. Even for e-commerce platforms where a producer can list multiple items, ensuring fair exposure to individual items would also ensure fairness at the producer level.

The **distribution of exposure** received by the items can be written as $D = \{D_s = \frac{E_s}{\sum_{s \in S} E_s} \forall s\}$. Given two exposure distributions D^{old} and D^{new} , we use L1-norm to calculate overall change in exposure: $EC(\text{old}, \text{new}) = \sum_{s \in S} |D_s^{\text{new}} - D_s^{\text{old}}|$.

In this paper, we assume that there is no change in overall demand of any item during the update. Although this assumption may not exactly replicate reality in some situations, but it helps us to focus on the main issue of two-sided fairness and bring out the nuances associated with it, rather than the general issue of unpredictability of demand.

4 Updating Recommendations in Two-Sided Platforms

In this section, we discuss the impact of platform updates on exposure of the producers. To concretely highlight the impact, we consider certain datasets, as well as different types of updates that are undertaken in real-world platforms.

4.1 Datasets and Types of Updates

In this work, we use the following datasets and test different types of updates on them.

Amazon Reviews dataset: We use the dataset released by (He and McAuley 2016), which comprises of customer re-

views and ratings for different Amazon products from the grocery category. From this dataset, we shortlist 1,000 most active customers (i.e., who have reviewed most number of products) and 1,000 most reviewed products, and only consider their corresponding ratings. Note that the rating act as a proxy to relevance score and the ratings of all the customer-item pairs are not available. Data-driven models are used to calculate the missing relevance values of other customer-item pairs. We test two kinds of updates on this dataset.

A. Changing the Model (Amazon-M): We test updating the recommender system (or the relevance scoring model) from a user-based collaborative filtering (Breese, Heckerman, and Kadie 1998) (it works on the assumption that similar users like same set of items) to a more sophisticated latent factorization based model (Koren, Bell, and Volinsky 2009).

B. Updating Training Data (Amazon-D): The most common type of update is the addition of new training data points. Here we calculate the relevance scores using a latent factorization method. At first the model is trained on the ratings received in the year 2013, and then trained on 2013 and 2014 rating data taken together.

We assume that since a platform is adopting a new recommendation algorithm, implicitly that means improved accuracy, otherwise there is no reason for the adoption. As a sanity check, our evaluation on held out ratings data shows improvements of 21.78% and 32.46% in root-mean-square-error by updating in Amazon-M and Amazon-D respectively.

Google Local dataset: We use data from Google Local, released by (He, Kang, and McAuley 2017), containing data about customers, local businesses and their locations (geographic coordinates), ratings, reviews etc. At any point in time, we consider each customer’s last reviewed location as a proxy for her location. We consider all active customers located in New York City and the business entities listed there, with more than 10 reviews. The dataset contains 45,305 customers, 3,029 businesses and 89,737 reviews.

C. Addition of New Feature(s) (GoogleLoc-F): Sometimes a new feature (e.g., customers’ location) is added to improve the relevance prediction model. We test an update from a purely ratings-based recommendation, $V^{\text{old}}(u, s) = \text{rating}(s)$, to a rating-cum-location based recommendation, $V^{\text{new}}(u, s) = \frac{\text{rating}(s)}{\text{distance}(u, s)}$.

4.2 Modeling Customer Arrivals

As we do not have the temporal customer arrival/login data, we model customer login events as Poisson point processes (Chiu et al. 2013), where we consider every customer’s logins to be independent of each other. The mean inter-arrival time (time interval between two consecutive arrivals on the platform) of each customer is sampled from a truncated Gaussian distribution (range $[0, 2]$) with a mean of 1 period (exact definition of a period may vary from platform to platform) and variance 0.2.

4.3 Impact of Immediate Updates

With the described customer arrival process, we implement updates listed in §4.1 in an immediate manner and report the distribution of the percentage changes in item exposures in

¹ $\mathbb{1}_{R_u}(s)$ is 1 if $s \in R_u$, and 0 otherwise.

Dataset	% of items with change of exposure		
	< 50%	50 – 100%	100 + %
Amazon-M	13.7	2.5	83.8
Amazon-D	24.1	6.7	69.2
GoogleLoc-F	0.12	1.17	98.71

Table 1: Percentage change in the exposure of individual items due to immediate update of recommendations.

Table 1. It is clear from the table that across different types of updates, 69 – 98% of the items experience more than 100% change (gain or loss) in their exposure values. Exposure or visibility often correlates with sales or economic opportunities on which the livelihood of many individuals depends (Wu and Bolivar 2009; Ghandour et al. 2008). An abrupt change (loss) in exposure could translate into economic loss or even shutdown; an abrupt gain may lead to degeneration of quality due to demand pressure. To capture the unfairness associated with such abrupt changes, we formalize the fairness notions for both producers and customers, as discussed next.

4.4 Formalizing Fairness in Two-Sided Platforms

Egocentric Fairness for Producers: As mentioned earlier, egocentric perception of fairness (Bediou and Scherer 2014; Kahneman, Knetsch, and Thaler 1991) depends on the change from the *status quo*. We define a platform update to be fair to the producers if the difference between the exposure distribution in the new system and that in the old system is minimal. More formally, if the new and previous exposure distributions are D^{new} and D^{old} respectively, then a platform is *egocentric fair* if $\sum_{s \in S} |D_s^{\text{new}} - D_s^{\text{old}}| < \epsilon$, where ϵ is a small positive number.

Minimum Guarantee for Customers: While being fair to the producers, the platform should not compromise on the satisfaction of the customers. We define a platform to be *fair* to the customers if it guarantees a minimum utility for everyone; $\phi^{\text{norm}}(R_u, V^{\text{new}}) \geq \theta, \forall u \in U$; where V^{new} is the new relevance scoring function to be implemented, and θ is the utility guarantee provided by the platform².

Table 1 clearly shows that updating recommendations immediately, violates the maxim of *egocentric fairness* for the producers. To ensure fairness, a phased update strategy can be undertaken. This is in line with research works in law, macroeconomics and business philosophy (Malerba 1992; Mintrom and Vergari 1996; Rabin 1997), where they have advocated for **incremental changes** for easy societal adaptation. However, updating recommendations incrementally in a two-sided market is challenging due to the dual task of

²Our proposal is comparable with the *fairness of minimum wage guarantee* (e.g., as required by multiple legislations in US, starting from *Fair Labor Standards Act 1938* to *Fair Minimum Wage Act 2007*) (Pollin et al. 2008; Green and Harrison 2010; Falk, Fehr, and Zehnder 2006). While ensuring minimum wage does not itself guarantee equality of income, it has been found to decrease income inequality (Lin and Yun 2016; Engbom and Moser 2018).

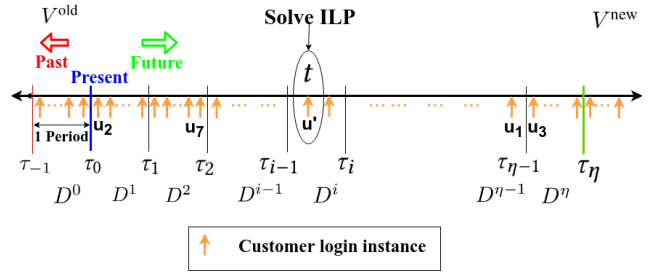


Figure 1: Timeline representation of incremental update: V^{old} and V^{new} are the current and new relevance scoring functions. The points on time axis are τ_{-1} (1 period in past), τ_0 (present or start time of update), τ_η (end time). $[\tau_0, \tau_\eta]$ is the time interval over which an incremental update will be implemented. D^0 is the observed exposure distribution in $[\tau_{-1}, \tau_0)$. D^i is the exposure distribution to be observed in step i ($[\tau_{i-1}, \tau_i)$). The objective of the proposed incremental update is to make each of the steps account for small changes in exposure, so that the transition is smooth for the items/producers, while ensuring fair customer utility.

protecting the producers, as well as ensuring a certain level of customer utility. We discuss this task in the next section.

5 Updating Recommendations Incrementally

In this section, we propose to update the recommendations in η steps (or η periods). We define points on the time axis like τ_{-1} (1 period in past), τ_0 (present or start point), τ_1, \dots, τ_η (end point) such that $[\tau_{i-1}, \tau_i)$ represents i^{th} period and the targeted update is achieved at τ_η . In each period, customers visit the platform (as modeled in §4.2) and personalized recommendation is provided to each of them. D^0 is the observed exposure distribution in $[\tau_{-1}, \tau_0)$. Let D^i represent the exposure distribution to be observed in step i ($[\tau_{i-1}, \tau_i)$). Figure 1 illustrates this online step-wise set up over time.

5.1 Incremental Update Formulation

Recommending items to minimize the exposure change can hurt customer utility; while maximizing customer utility can cause huge changes in exposure. To address this trade-off, we can formulate optimization problem in two ways: (i) one where we minimize exposure change constrained to a lower bound on customer utility, and (ii) another where we maximize customer utility constrained to an upper bound on exposure change. In this paper, we use the former one since the utility constraints are more interpretable and exposure objectives can be easily operationalized in an online scenario. We propose to come up with a target exposure distribution \overline{D}^i for step i , that is, we try to make the observed exposure distribution in step i as close as possible to \overline{D}^i , thus the objective can be written as below.

$$\text{minimize } \sum_{s \in S} \left| D_s^i - \overline{D}_s^i \right|, \forall i \in \{1, \dots, \eta\} \quad (1)$$

This objective needs to be transformed into an online version which deals with each individual customer logging in

at certain points of time. Assuming a specific customer u' logs into the platform at time t ($t \in [\tau_{i-1}, \tau_i]$), the objective transforms into below.

$$\operatorname{argmin}_{R_{u'}} \sum_{s \in S} \left| E_s^t + \frac{\mathbb{1}_{R_{u'}(s)}}{k} - (|U([\tau_{i-1}, t])| + 1) \cdot \overline{D}_s^i \right| \quad (2)$$

where E_s^t is the exposure of s in $[\tau_{i-1}, t]$, $\frac{\mathbb{1}_{R_{u'}(s)}}{k}$ is the attention to s from u' , $(|U([\tau_{i-1}, t])| + 1)$ is the total number of customer logins in $[\tau_{i-1}, t]$, and \overline{D}_s^i is the target exposure proportion for s in step i . As \overline{D}_s^i is a fraction, its multiplication with number of customer logins produces the total targeted exposure for s in $[\tau_{i-1}, t]$; the difference shows how far is the system from the target exposure.

Constraint for Minimum Utility: Along with the above objective, we also have to ensure a minimum utility to the customers, which would be a hard constraint. Thus, we use a constraint with lower bound on the normalized customer utility; For customer u' at time t ($t \in [\tau_{i-1}, \tau_i]$), we impose a constraint that the utility (based on new relevance scoring V^{new}) of the k -items chosen for the customer u' in step i must be above a threshold θ_i : $\phi^{\text{norm}}(R_{u'}, V^{\text{new}}) \geq \theta_i$ or $\phi(R_{u'}, V^{\text{new}}) \geq (\theta_i \cdot \phi_{u'}^{\text{max}}(V^{\text{new}}))$.

We formulate this optimization problem as an Integer Linear Program (ILP). For customer u' logging in at time t ($t \in [\tau_{i-1}, \tau_i]$), we introduce $|S|$ decision variables: $X_{u',s}$ which is set to 1 if s is recommended to u' , and set to 0 otherwise. Now we write the ILP as below.

$$\operatorname{argmin}_X \sum_{s \in S} \left| E_s^t + \frac{X_{u',s}}{k} - (|U([\tau_{i-1}, t])| + 1) \cdot \overline{D}_s^i \right| \quad (3a)$$

s.t.

$$X_{u',s} \in \{0, 1\} \quad \forall s \in S \quad (3b)$$

$$\sum_{s \in S} X_{u',s} = k \quad (3c)$$

$$\sum_{s \in S} X_{u',s} \cdot V^{\text{new}}(u', s) \geq (\theta_i \cdot \phi_{u'}^{\text{max}}(V^{\text{new}})) \quad (3d)$$

Here, constraint-3b ensures keeping the variables binary. Constraint-3c ensures selecting k items exactly. A minimum customer utility is guaranteed by constraint-3d.

5.2 Parameter Setting

There are two important parameters in the ILP formulation \overline{D}_s^i and θ_i which need to be fixed.

Setting \overline{D}_s^i : We propose two different ways to set the target exposure distributions for each step (\overline{D}_s^i for step i).

A. Estimated Steps: Using the current customer arrival frequency (as in $U([\tau_{i-1}, \tau_i])$) we can find an estimate of the final exposure distribution for the new relevance scoring (i.e., using top- k of V^{new} for $U([\tau_{i-1}, \tau_i])$) and let that be D^{pred} . Imagining D^0 and D^{pred} as points in multidimensional space (with $|S|$ dimensions), our proposition is to enforce certain level of change towards D^{pred} in each step. Thus we set the target exposure distribution for step i as: $\overline{D}_s^i = D_s^0 + i \cdot \delta$, $\forall s, i$, where $\delta = \frac{1}{\eta}(D_s^{\text{pred}} - D_s^0)$.

B. Preserving Steps: Here, we set target exposure distribution of a step to the observed one in last step (we try to preserve the observed exposure), i.e., $\overline{D}_s^i = D_s^{i-1}$, $\forall s, i$.

Setting θ_i : We use linearly increasing and geometrically increasing settings for θ_i .

A. Linear Steps: $\theta_i = \frac{i}{\eta}$, for $1 < i < \eta$,

B. Geometric Steps: $\theta_i = \theta_{i-1} + 2^{-i}$ for $1 < i < \eta$, while $\theta_0 = 0$, and $\theta_\eta = 1$.

5.3 Approximate Solution with Prefiltering

As our ILP operates on the whole item space, huge item space of some systems can be bottleneck for the ILP solvers. To deal with this issue, we propose to prefilter the item space, and then run the ILP on filtered (smaller) item space for an approximate solution. We prefilter in following two ways and merge the two filtered lists to get a smaller item space: (i) top- (k^2) ($k = \text{recommendation set size}$) items using new relevance scoring (V^{new}), which can help in satisfying the customer utility constraint, and (ii) top- (k^2) items based on $|\frac{E_s^t}{\sum_{s' \in S} E_{s'}^t} - \overline{D}_s^i|$, which can help in minimizing the objective function. We test the proposed ILP with both unfiltered and prefiltered item spaces in §6.

5.4 Baselines

Only a few prior works consider incremental changes; however they do not necessarily cater to two-sided platforms.

Baseline-1 (CanD): *Canary deployment* (Tseitlin and Sondow 2017) (also known as *phased roll out* or *incremental roll out*) is a popular approach traditionally used in software deployment, where a new software version is slowly rolled out in production for subsets of customers to reduce the risk of imminent failure in an unseen environment.

Baseline-2 (IRF): Another approach for incremental update would be to introduce *intermediate relevance functions* for each of the steps (gradually moving the relevance scores from V^{old} to V^{new}); relevance function for step i : $V^i(u, s) = (1 - \frac{i}{\eta}) \cdot V^{\text{old}}(u, s) + \frac{i}{\eta} \cdot V^{\text{new}}(u, s)$, $\forall u, s$. We can recommend the top- k according to $V^i(u, s)$ in step i .

6 Experimental Evaluation

For each customer u' logging into the platform at time t , we solve the proposed ILP with different settings which gives a set of items to be recommended. Using these results, we calculate and record item exposures and customer utilities in each step. We set the number of steps $\eta = 10$, and size of recommendation $k = 10$. We use *cvxpy* (*cvxpy.org*) paired with *Gurobi solver* (*gurobi.com*) for solving the ILP. In this section, we use the following abbreviations: **E- Estimated, P- Preserving steps in \overline{D}_s^i ; L- Linear, G- Geometric steps in θ_i ; PF- Prefiltering.**

6.1 Producer-Centric Metrics

First we define metric for change in exposure.

Exposure Change (EC): Given two exposure distributions D^{old} and D^{new} , exposure change (EC) is given by their L1 distance (also defined in §3); $EC(\text{old}, \text{new}) = \sum_{s \in S} |D_s^{\text{new}} - D_s^{\text{old}}|$ Based on EC, we define the following three metrics.

A. Transition Path Length (Υ) - Efficiency metric.

It is the sum of all exposure changes that the transition has

Method	Amazon-M ($EC(0, \eta) = 1.922$)			Amazon-D ($EC(0, \eta) = 1.557$)			GoogleLoc-F ($EC(0, \eta) = 1.967$)		
	Υ	π	Z	Υ	π	Z	Υ	π	Z
CanD	1.63	0.17	0.99	1.34	0.21	0.99	1.38	0.16	0.99
IRF	1.48	0.95	0.62	1.46	0.19	0.98	1.46	0.98	0.59
ILP-EL	1.03	0.12	0.99	1.04	0.13	0.99	1.04	0.11	0.99
ILP-EL(PF)	1.16	0.13	0.99	1.45	0.16	0.99	1.33	0.15	0.99
ILP-EG	1.18	0.28	0.93	1.12	0.22	0.97	1.16	0.19	0.96
ILP-PL	1.06	0.46	0.61	1.02	0.74	0.33	1.08	0.29	0.90
ILP-PG	1.10	0.30	0.86	1.08	0.23	0.86	1.24	0.29	0.90

Table 2: Producer-Centric Metrics for $\eta = 10$ and $k = 10$. CanD and IRF are the two baselines. For ILP-based methods we use abbreviations like; E: Estimated, P: Preserving steps in $\overline{D^i}$; L: Linear, G: Geometric steps in θ_i ; PF: Prefiltering item space.

gone through, relative to that of an immediate update.

$$\Upsilon = \frac{\sum_{i \in \{1 \dots \eta\}} EC(i-1, i)}{EC(0, \eta)} \quad (4)$$

The **lower** the path length, the **more efficient** is the transition.

B. Maximum Transition Cost (Π) - Fairness metric.

$$\Pi = \frac{\max_i [EC(i-1, i)]}{EC(0, \eta)} \quad (5)$$

Π checks the largest change during the incremental transition process relative to that of an immediate update; even a single big exposure change is undesirable as it is inherently disadvantageous (unfair) for producers.

C. Transition Inequality (Z) - Fairness metric.

Transition Inequality captures the dissimilarity in the quantum of transition among the steps, measured by entropy as defined below.

$$Z = - \sum_{i \in \{1 \dots \eta\}} \left(\frac{EC(i-1, i)}{M} \right) \log_{10} \left(\frac{EC(i-1, i)}{M} \right) \quad (6)$$

where $M = \sum_{i \in \{1 \dots \eta\}} EC(i-1, i)$.

An **ideal update** will have high efficiency (or **low Υ**), and small & equal sized changes (or **low π** and **high Z**).

6.2 Producer-Side Results

Table 2 reports the above metrics for all the baselines and proposed methods on different datasets.

Performance of Baselines: *CanD* ensures small (Π), performs very well in maintaining similar level of changes at each step (Z); but, *CanD* is less efficient due to high path lengths (Υ) as the changes (change from D^{i-1} to D^i in step i) it introduces may or may not be directed towards D^η . The performance of *IRF* is not stable; in Amazon-M and Google-F, it performs very poorly in both the fairness metrics Π and Z (the reason becomes clear when we look at the customer side results); It also shows poor efficiency (Υ) like *CanD*.

Performance of ILP-Estimated (EL/EG): ILP-EL shows the lowest Π and highest Z making it the most fair method for the producers; it is very efficient (Υ) too; ILP-EL with *pre-filtering*(PF) also performs better (in π and Z) than baselines and other ILP settings; however it incurs a loss in efficiency Υ due to its filtered item space. ILP-EG performs worse than ILP-EL in all metrics.

Performance of ILP-Preserving (PL/PG): ILP-PL shows very efficient (Υ) transitions; however it performs very poorly in π and Z which makes it even more unfair than

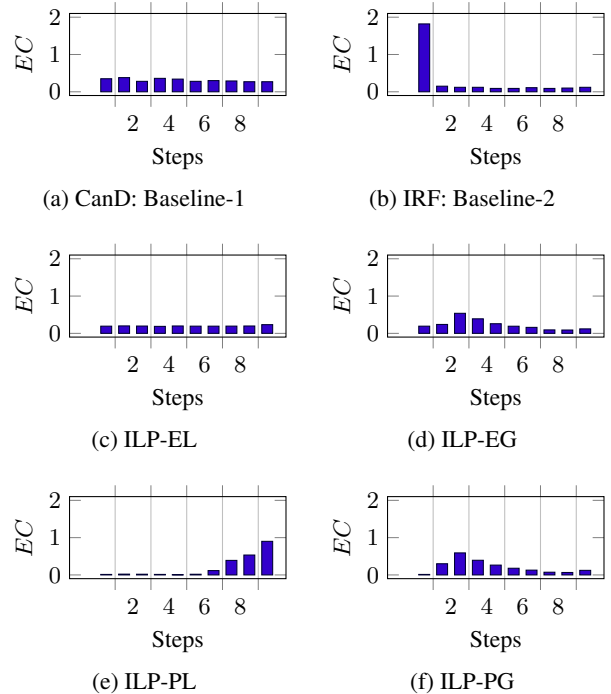


Figure 2: Exposure changes at all the steps ($EC(i-1, i) = \sum_{s \in S} |D_s^i - D_s^{i-1}|$) are plotted (x-axis: steps i , y-axis: EC) for **Amazon-M**. Hyperparameters: $\eta = 10$, $k = 10$.

CanD. On the other hand, ILP-PG shows good results comparable to ILP-EL and ILP-EG. The reasons become clear when we study the individual EC plots (elaborated next).

EC Exposure Change Plots: We plot exposure changes (EC) at each of the steps of updates for Amazon-M in Figure-2. We see small and equal sized changes for *CanD* and ILP-EL, however *CanD* produces slightly larger changes; this explains their similar Z performances, but different Υ and π . Both ILP-EG and ILP-PG show dissimilar changes in different steps; thus slightly larger Υ and π , and slightly lower Z . *IRF* causes huge change in the first step while ILP-PL shows large changes in the last few steps; this explains their high π and very low Z values.

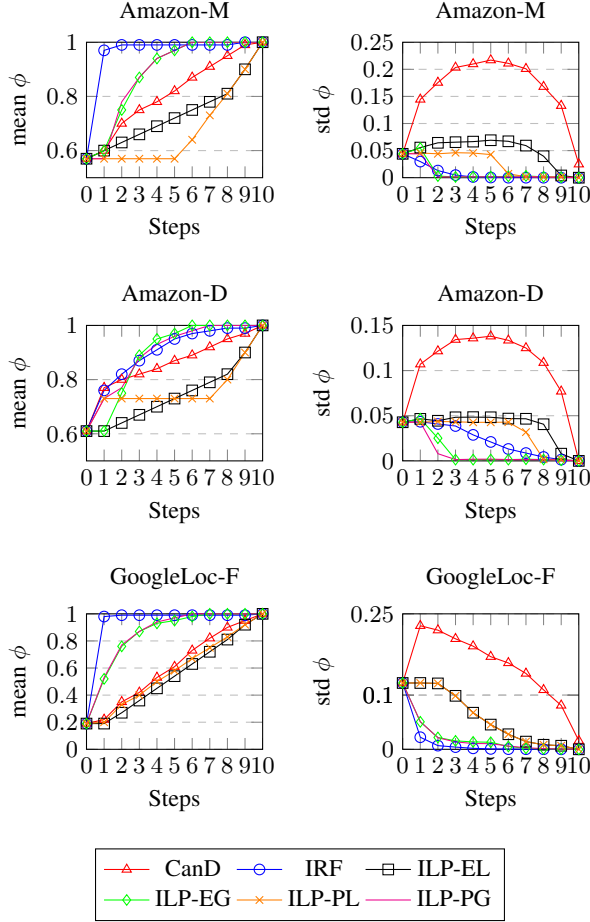


Figure 3: Mean (left column) and standard deviation (right) of $\phi^{\text{norm}}(R_u, V^{\text{new}}) \forall u \in U([\tau_{i-1}, \tau_i])$ for each step i .

6.3 Customer-Centric Metrics

In each step i for each customer $u \in U([\tau_{i-1}, \tau_i])$, we obtain the utility $\phi^{\text{norm}}(R_u, V^{\text{new}})$; We calculate their mean and standard deviation and plot them in Figure-3. The faster the mean utility grows, the faster the update applies. The standard deviation indicates the degree of unfairness.

6.4 Customer-Side Results

We explain the salient points of the results (in Figure-3).

Performance of Baselines: The rise in mean customer utilities for *CanD* is comparable to ILPs, however for *IRF* it is much faster. Note that *CanD* incurs large standard deviation, i.e., it introduces larger disparity in customer utilities, which is undesirable. The *IRF* shows a large increase in mean utility in the first step of Amazon-M and GoogleLoc-F which essentially means it fails to update incrementally; thus the producer fairness is severely compromised (correspondingly refer to Figure-3 and Table 2); As in *IRF*, we choose the top- k results using intermediate relevance functions, the intermediate function (V^1) at step 1 drastically changes (it could have happened at any other step too) the top- k set making it

close to the top- k of V^{new} ; This explains the large increase in mean utility and the large exposure change in the first step in those datasets (refer Figure-3,2 respectively). However, this is a very data-specific phenomenon as it doesn't happen in the Amazon-D. For the baseline methods, we see the above issues in ensuring producer and customer fairness; Reliability is also a major concern.

Performance of ILPs: By design, all ILPs ensure minimum utility guarantee to the customers in each of the steps. The ILPs (EG and PG) with geometric steps in θ_i increases the customer utilities quickly while the ILPs (EL and PL) with linear steps in θ_i show slower improvements. In Amazon-D, Amazon-M, Google-Loc, the status quo (period 0) mean utilities are near to 0.57, 0.61 and 0.19 correspondingly. Thus for ILP-Preserving (PL/PG), when there are scopes (θ_i becomes more than status quo utility), the ILPs show an update; This explains why ILP-PL generally shows significant updates (increase in utility Fig-3 and change in exposure Fig-2) only after some initial steps; while ILP-PG shows updates earlier due to geometric increase in θ_i and performs better. However for ILP-Estimated (EL/EG), such issue never comes as they enforce estimated changes in exposure (by setting \bar{D}^i) along with increase in θ_i in each step. The standard deviation of all the ILPs are small; the ILP-(EL/PL) have slightly higher values.

Summary: ILP-EL performs best in terms of producer fairness; its performance in maintaining customer utility is as per design; however, as the name suggests ILP-EL requires an estimation of change in producer exposure a priori. Whereas, ILP-PG performs a bit inferior to ILP-EL in terms of producer fairness but much better than baselines; the increase in customer utility is faster than ILP-EL. Most importantly, it doesn't require any estimation of the exposure change for designing each step which makes it an attractive choice. However, our aim has been to explore a whole range of possibilities, and leave it to the designer to choose one as per their requirement and available resources.

7 Conclusion

In this paper, we identified the adverse impact on the producers due to immediate updates in recommendations in two-sided platforms, and proposed an innovative ILP-based incremental update mechanism to tackle it. Extensive evaluations over multiple datasets and different types of updates show that our proposed approach not only allows smoother transition of producer exposures, but also guarantees a minimum customer utility in intermediate steps. In future, we plan to check the impact of updates in more complex settings, such as when the assumption of closed market (where neither new producers/customers enter the system nor the overall demand fluctuates) is relaxed. We also plan to consider *position/ranking bias* in customer attention.

Acknowledgements 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 European Union's Horizon 2020 Framework Programme (grant agreement no. 789373). G. K Patro is supported by TCS Research Fellowship.

References

- Abdollahpouri, H., and Burke, R. 2019. Multi-stakeholder recommendation and its connection to multi-sided fairness. *arXiv preprint arXiv:1907.13158*.
- AdExchanger. 2018. <https://adexchanger.com/data-driven-thinking/facebook-news-feed-changes-will-challenge-publishers-stay-relevant>.
- Bediou, B., and Scherer, K. R. 2014. Egocentric fairness perception: emotional reactions and individual differences in overt responses. *PLoS 9*(2).
- Biega, A. J.; Gummadi, K. P.; and Weikum, G. 2018. Equity of attention: Amortizing individual fairness in rankings. In *ACM SIGIR*.
- Breese, J. S.; Heckerman, D.; and Kadie, C. 1998. Empirical analysis of predictive algorithms for collaborative filtering. In *UAI*.
- Burke, R. 2017. Multisided fairness for recommendation. *arXiv preprint arXiv:1707.00093*.
- Chakraborty, A.; Hannak, A.; Biega, A. J.; and Gummadi, K. P. 2017. Fair sharing for sharing economy platforms.
- Chakraborty, A.; Patro, G. K.; Ganguly, N.; Gummadi, K. P.; and Loiseau, P. 2019. Equality of voice: Towards fair representation in crowdsourced top-k recommendations. In *ACM FAT**.
- Chiu, S. N.; Stoyan, D.; Kendall, W. S.; and Mecke, J. 2013. *Stochastic geometry and its applications*. John Wiley Sons.
- Edelman, B.; Luca, M.; and Svirsky, D. 2017. Racial discrimination in the sharing economy: Evidence from a field experiment. *American Economic Journal: Applied Economics* 9(2).
- Engbom, N., and Moser, C. 2018. Earnings inequality and the minimum wage: Evidence from Brazil. *Federal Reserve Bank of Minneapolis - Opportunity and Inclusive Growth Institute* 7.
- Falk, A.; Fehr, E.; and Zehnder, C. 2006. Fairness perceptions and reservation wages—the behavioral effects of minimum wage laws. *Qtrly. Joun. of Economics* 121(4).
- Geyik, S. C.; Ambler, S.; and Kenthapadi, K. 2019. Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In *ACM KDD*.
- Ghandour, A.; Deans, K.; Benwell, G.; and Pillai, P. 2008. Measuring ecommerce website success. *ACIS*.
- Graham, M.; Hjorth, I.; and Lehdonvirta, V. 2017. Digital labour and development: impacts of global digital labour platforms and the gig economy on worker livelihoods. *EU Review of Labour & Research* 23(2).
- Green, D. A., and Harrison, K. 2010. Minimum wage setting and standards of fairness. Technical report, Institute for Fiscal Studies.
- Hannák, A.; Wagner, C.; Garcia, D.; Mislove, A.; Strohmaier, M.; and Wilson, C. 2017. Bias in online freelance marketplaces: Evidence from taskrabbit and fiverr. In *ACM CSCW*.
- Hayes, M. T. 1992. *Incrementalism and public policy*. Longman.
- He, R., and McAuley, J. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *WWW*.
- He, R.; Kang, W.-C.; and McAuley, J. 2017. Translation-based recommendation. In *ACM RecSys*.
- Kahneman, D.; Knetsch, J. L.; and Thaler, R. H. 1991. Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic perspectives* 5(1).
- Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *IEEE Computer* (8).
- Lab, N. 2019. <https://www.niemanlab.org/2019/07/should-facebook-have-a-quiet-period-of-no-algorithm-changes-before-a-major-election/>.
- Lambrecht, A., and Tucker, C. 2019. Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads. *Management Science*.
- Lin, C., and Yun, M.-S. 2016. The effects of the minimum wage on earnings inequality: Evidence from China. In *Income Inequality Around the World*.
- Malerba, F. 1992. Learning by firms and incremental technical change. *The economic journal* 102(413).
- Mintrom, M., and Vergari, S. 1996. Advocacy coalitions, policy entrepreneurs, and policy change. *Policy studies journal* 24(3):420–434.
- Morewedge, C. K., and Giblin, C. E. 2015. Explanations of the endowment effect: an integrative review. *Trends in cognitive sciences* 19(6).
- Pollin, R.; Brenner, M.; Luce, S.; and Wicks-Lim, J. 2008. *A measure of fairness: The economics of living wages and minimum wages in the United States*. Cornell Press.
- Rabin, M. 1997. Fairness in repeated games.
- Serbos, D.; Qi, S.; Mamoulis, N.; Pitoura, E.; and Tsaparas, P. 2017. Fairness in package-to-group recommendations. In *WWW*.
- Sühr, T.; Biega, A. J.; Zehlike, M.; Gummadi, K. P.; and Chakraborty, A. 2019. Two-sided fairness for repeated matchings in two-sided markets: A case study of a ride-hailing platform. In *ACM KDD*.
- Sürer, Ö.; Burke, R.; and Malthouse, E. C. 2018. Multi-stakeholder recommendation with provider constraints. In *ACM RecSys*.
- Tseitlin, A., and Sondow, J. 2017. Progressive deployment and termination of canary instances for software analysis. US Patent 9,712,411.
- Vox. 2019. [vox.com/2019/3/8/18252606/amazon-vendors-no-orders-marketplace-counterfeits](https://www.vox.com/2019/3/8/18252606/amazon-vendors-no-orders-marketplace-counterfeits).
- Wu, X., and Bolivar, A. 2009. Predicting the conversion probability for items on c2c ecommerce sites. In *ACM CIKM*.