
Rankings for Two-Sided Market Platforms

Yi Su
Cornell University
Ithaca, NY
ys756@cornell.edu

Thorsten Joachims
Cornell University
Ithaca, NY
tj@cs.cornell.edu

Abstract

Rankings have become the standard interface for presenting results to customers in online systems. Traditional online systems connect customers with items (e.g. books, music, news), where only the customers have preference that the system needs to model. This has led to a broad array of learning-to-rank algorithms for optimizing customer's preference. However, new applications of ranking in multi-sided market platforms (e.g., job search, online dating, college admission), where both sides (e.g. job candidates and employers) explore the market through rankings, are gaining in importance. We find that in such ranking-based matching markets, the scarcity on the supplier side and competition among customers make the naive preference-based rankings highly sub-optimal. To address this problem, we explore methods for globally and simultaneously optimizing all rankings in a two-sided market platform, and we propose an optimization-based framework for maximizing social welfare in the market. Under a specific examination model, our objective can be solved effectively using any off-the-shelf convex optimization algorithms. Further, we discuss potential trade-offs between stability and the overall social welfare, and incorporate various constraints to achieve this balance.

1 Introduction

Most online systems (e.g., search engine, recommender system, music streaming) rely on rankings as the prevalent way of presenting results to the customers. A fundamental principle guiding the design of ranking algorithms is called the Probability Ranking Principle (PRP) [12], which argues that the ideal ranking should rank the items based on their relevance to the customer, since this will provide maximum utility to the customers under virtually any reasonable utility measure. However, rankings are increasingly used in new and emerging multi-sided market platforms [1, 9, 10, 17] (e.g., job search, online dating, property renting). In these markets, the online system works as an intermediary facilitating the interaction between the two-sides of the market: suppliers and customers. Different from traditional systems, in which an item could be consumed an (almost) unlimited amount of times by different customers, the items to be ranked in the two-sided markets are jobs, people, houses, which are scarce by nature. Furthermore, both sides of the market have preferences, the customers have preferences over the suppliers, but the suppliers also have preferences over the customers. This results in the fact that the utility to a customer on either side of the market is not determined by their own preference assessment, but also depends on the preferences of potentially every other customer in the market.

In these market platforms, ranking solely based on individual customer preference will give suboptimal results, for both an individual customer as well as for the effectiveness of the platform overall. For example, in the online networking recommendation systems for virtual conferences, a famous/senior researcher who works broadly on multiple topics might be of high preference to many attendees. Hence, a naive preference-based system will rank this person at the top to a large amount of people - overloading this individual. Consequently, many attendees will get no any reply from this person,

leading to low networking engagement and hence low satisfaction with the system. To alleviate this problem, globally optimized rankings could serve an important role in reducing collision between customer choices, for example, by presenting ranking that balance the load and maximize the number of successful matches. This calls for new frameworks for efficiently design ranking algorithms in these two-sided market platforms.

In this work, we study the setting in which the interaction of customers and suppliers are happening in a sequential manner. When a customer enters the market, the online system presents specific rankings to this customer. Based on this specific ranking, the customer initiates the matches with different suppliers probabilistically according to a behavioral model that depends on where the supplier appears in the ranking. This reflects the position bias that rankings induce on the probability that a rank position will be observed. From the supplier side, after receiving the initiated matches from the customers, the supplier responds back to the customers based on the preference of these customers in his perspective. This two-stage interaction is highly relevant to various practical applications, like job search, renting platforms, and online markets, and they are also used in the assortment planning literature [1, 16]. Under this setting, each individual’s utility is characterized by the expected matches he will get, and this not only depends on the two-sided preferences between him and suppliers, but also highly depends on how others act in the market. Motivated by this, we work with the objective of maximizing the expected number of matches in the markets as a form of social welfare function.

Our main contribution is as follows. We formulate the problem of ranking in these two-sided markets. Followed by this, we propose a tractable lower bound of the social-welfare objective. We show that, for certain examination models, this objective can be effectively solved iteratively based on any off-the-shelf convex optimization solvers. To further discuss the trade-off between overall social welfare and individual utilities, we show our framework can incorporate constraints related to the fairness of the globally optimized rankings without hurting the computational complexity. This need for fairness was pointed out in multiple recent works [15, 18, 2, 19] that discuss the social responsibility that the ranking systems should have. This includes fairness to all participants, interpretability, and transparent in the guarantees it provides. Our work also has implications on the societal role of these online systems, as these platforms are likely to take important roles in mediating consequential processes that affect not only the individual, but also the effectiveness of the market as a whole.

2 Problem Formulation

For simplicity of exposition, we discuss the two-sided platform in the specific setting of a job application platform. However, many other settings fit into the same schema. We denote the set of candidates in the markets as \mathcal{C} , and the set of employers as \mathcal{J} . In these markets, we use $r(c, j)$ to denote the candidate c ’s preference for employer j and $r(j, c)$ to denote the employer j ’s preference for candidate c , we assume all the preferences in $[0, 1]$ and known. By using different notations for the two-side preference, we also incorporate the scenario when the preference is asymmetric in the two-sided markets.

The interaction in the market goes as follows: for each candidate $c \in \mathcal{C}$, the system presents a ranking $r(c)$ of all $|\mathcal{J}|$ employers to the candidate. The candidate then browses the ranking, assuming a position-based model [11, 7] to account for presentation bias. This means the probability of applying to a specific employer j (i.e., initiate the matching) depends on the preference $r(c, j)$, as well as the examination probability $\mathbb{P}(e(j) = 1 | \text{rank}(j | r(c)))$. In the position-based model, the examination probability is a function of the position of employer j in the presented ranking and it models how much attention employer j gets at rank $\text{rank}(j | r(c))$. Typically v is an application-dependent function; popular choices include $v(x) = 1/x$ and $v(x) = \frac{1}{\log(1+x)}$, which is used in the Discounted Cumulative Gain (DCG) [6] measure. Therefore we have that

$$\mathbb{P}(c \text{ initiates matching to } j) := P_{c,j} = r(c, j) \left(\sum_{r(c) \in R^c} \mathbb{P}(r(c) | c) v(\text{rank}(j | r(c))) \right). \quad (1)$$

Here we use R^c to denote the set of all possible rankings over $|\mathcal{J}|$ employers for candidate c , and $\mathbb{P}(\cdot | c)$ is our *stochastic* ranking policy, which is a probability distribution maps from $R^c \rightarrow \Delta$ with Δ denotes the probability simplex over all the possible rankings over $|\mathcal{J}|$ employers. However, the number of possible rankings is exponential in nature. In order to avoid any computation over this large combinatorial space, we utilize the fact that the term $\sum_{r(c) \in R^c} \mathbb{P}(r(c) | c) v(\text{rank}(j | r(c)))$

captures the marginal rank distributions of employer j under stochastic ranking policy $\mathbb{P}(\cdot|c)$. It can therefore be rewritten as

$$P_{c,j} = r(c,j) \left(\sum_{k=1}^{|\mathcal{J}|} \mathbb{P}^M(\text{rank}(j) = k|c)v(k) \right) = r(c,j) \mathbf{e}_j^T P^c \mathbf{v},$$

where $\mathbb{P}^M : \{1, 2, \dots, |\mathcal{J}|\} \rightarrow \Delta_{|\mathcal{J}|}$ captures the marginal probability of placing employer j at rank k , under stochastic policy \mathbb{P} . This reduces the representation of the stochastic ranking policy to $O(|\mathcal{J}|^2)$, and it is uniquely characterized by a doubly stochastic matrix P^c , where $[P^c]_{j,k}$ denotes the probability that stochastic ranking policy \mathbb{P} places employer j at rank k .¹ The second equality further simplifies the expression by using matrix P^c and vector $\mathbf{v} \in \mathbb{R}_+^{|\mathcal{J}|}$, where $\mathbf{v}(rk) = v(rk)$. For brevity of notation, let $\mathbf{P}_C = \{P^c\}_{c \in \mathcal{C}}$ denote the rankings we provide for all candidates in the market, and the probability of initiating a match is a function of \mathbf{P}_C , hence $P_{c,j}(\mathbf{P}_C) = r(c,j) \mathbf{e}_j^T P^c \mathbf{v}$.

After all candidates initiate their matchings, then the employers start to act. For each employer j , the ranking system ranks the candidates who apply to employer j by the employers preference, i.e., $r(j,c)$. Let \mathcal{C}_j denote the set of candidates who initiates the match to job j and it is worth noting that \mathcal{C}_j is a random set since it depends on how each candidate acts probabilistically. Then after the employer j sees this ranked list, he will reply back to the matching (i.e., interview the candidate) based on a similar position-based model.

$$\mathbb{P}(j \text{ replies back to } c|c \text{ applied to } j, \mathbf{P}_C) := P_{j,c}(\mathbf{P}_C) = r(j,c)v(\text{rk}(c|\mathcal{C}_j(\mathbf{P}_C))) \quad (2)$$

$\text{rk}(c|\mathcal{C}_j)$ denotes the rank of candidate c over \mathcal{C}_j candidates, where the ranking is determined by employer j 's preference only. We say a match between candidate employer pair (c,j) is successful if candidate c initiates the match/applies to employer j , and employer j replies back/interviews candidate c . Our goal is to design ranking functions for all candidates $\mathbf{P}_C = \{P^c\}_{c \in \mathcal{C}}$ that maximize the social welfare: the total number of expected matches in the whole market.

$$\mathbf{SW}(\mathbf{P}_C) = \sum_{c \in \mathcal{C}} \sum_{j \in \mathcal{J}} \mathbb{P}(c \text{ and } j \text{ matches}) = \sum_{c \in \mathcal{C}} \sum_{j \in \mathcal{J}} P_{c,j}(\mathbf{P}_C) P_{j,c}(\mathbf{P}_C) \quad (3)$$

3 Rankings in two-sided markets

In this section, we analyze the social welfare objective and derive a tractable lower bound. Under some specific examination function $v(\cdot)$, we show this lower bound can be solved effectively by using iterative methods, where each iteration can be solved using off-the-shelf convex optimization methods.

Recall the expected number of match between each candidate employer pair (c,j) is $P_{c,j}(\mathbf{P}_C)P_{j,c}(\mathbf{P}_C)$, with $P_{c,j}(\mathbf{P}_C) = r(c,j) \mathbf{e}_j^T P^c \mathbf{v}$. Now we are going to analyze $P_{j,c}(\mathbf{P}_C)$, the probability of employer j replying back to candidate i given that all candidates act according to the examination model underlying \mathbf{P}_C .

Before going into the detail, we will first introduce the following notation. Let $\sigma_j(c) \in [|\mathcal{C}|]_n$ be the rank of candidate c among all candidates in the set \mathcal{C} , where the ranking is in descending order of $r(j,c)$ for any fixed employer j (we break ties arbitrarily). Similarly, we denote with $\sigma_c(j) \in [|\mathcal{J}|]_n$ the rank of employer j among all employers, where the ranking is based on candidate c 's preference function $r(c,j)$. Inversely, we define $\sigma_j^{-1}(s) := \{k \in \mathcal{C} | \sigma_j(k) = s\}$ be the candidate who is listed at rank s in employer j 's preference list. Define the priority set for employer j w.r.t. candidate i as $\Sigma_j(c) := \{k \in \mathcal{C} | \sigma_j(k) < \sigma_j(c)\}$, which include the candidates who receive higher preference to employer j than candidate i . Based on this, we let $F_{(j,c)}^k$ be the set of all subsets of k items that can be selected from $\Sigma_j(c)$, and $A^c := \mathcal{C} \setminus A$ be the complement of A for any $A \subseteq \mathcal{C}$. Let $R_{j,c}$ denotes the rank of candidate c in employer j 's ranking (depends on $r(j,c)$ and the random set \mathcal{C}_j) given c applied to j , then we have the following lemma w.r.t. the distribution of $R_{j,c}$.

Lemma 1. $R_{j,c}$ is a random variable that lies in $\{1, 2, \dots, \sigma_j(i)\}$. The distribution of $R_{j,c} - 1$ is a Poisson Binomial Distribution with parameters $[P_{\sigma_j^{-1}(1)}(\mathbf{P}_C), P_{\sigma_j^{-1}(2)}(\mathbf{P}_C), \dots, P_{\sigma_j^{-1}(\sigma_j(i)-1)}(\mathbf{P}_C)]$,

¹Regarding the question of how to sample rankings given the doubly stochastic matrix P^c , the Birkhoff-von Neumann (BvN) decomposition provides a transformation to decompose a doubly stochastic matrix into a convex sum of permutation matrices [3, 15].

which is a vector of probabilities corresponding to the probability of candidate $k \in \Sigma_j(c)$ initiating the match with to employer j .

$$\mathbb{P}(R_{j,c} = k) = \sum_{A \in F_{(j,c)}^{k-1}} \prod_{s \in A} P_{s,j}(\mathbf{P}_C) \prod_{r \in A^c} (1 - P_{r,j}(\mathbf{P}_C)) \quad (4)$$

Given Lemma 1, it is easy to see that

$$P_{j,c}(\mathbf{P}_C) = r(j,c) \mathbb{E}[v(R_{j,c})], \quad (5)$$

where the randomness is over the probabilistic way of interacting in the market as we introduced in the previous section. However, the PMF of the Poisson binomial distribution involves $n!/((n-k)!k!)$ (with $n = |\Sigma_j(c)|$ and $k \in \{1, 2, \dots, \Sigma_j(c)\}$) terms of summation. Though some recursive formula exists [13], the complicated form w.r.t the $P_{\sigma_j^{-1}(s)}(\mathbf{P}_C)$'s poses significant challenges when we treat these probabilities as a function of \mathbf{P}_C , especially when we eventually perform optimization. Instead, we work with a lower bound, derived in the following lemma.

Lemma 2. *Assume the examination model $v(x)$ is convex, then the lower bound for $P_{j,c}(\mathbf{P}_C)$ is:*

$$P_{j,c}(\mathbf{P}_C) \geq r(j,c)v(1 + \sum_{k \in \Sigma_j(c)} P_{k,j}(\mathbf{P}_C)) \quad (6)$$

Proof. The proof is based on the convexity of v and the simple form of the expectation of Poisson Binomial R.V.

$$P_{j,c}(\mathbf{P}_C) := r(j,c) \mathbb{E}[v(R_{j,c})] \geq r(j,c)v(\mathbb{E}[v(R_{j,c})]) = r(j,c)v(1 + \sum_{k \in \Sigma_j(c)} P_{k,j}(\mathbf{P}_C)) \quad (7)$$

□

It is worth noting that convexity of v is not a restrictive condition, since many popular examination models (e.g. $v(x) = 1/x$ and $v(x) = 1/\log_2(1+x)$) satisfy this assumption. Now we are ready to present our main theorem, which gives a valid and tractable lower bound of the social welfare objective in Equation 3.

Theorem 3. *Assume the examination model $v(x)$ is convex, a lower bound for the SW objective is given by*

$$\mathbf{SW}(\mathbf{P}_C) \geq \mathbf{SW}_{lower}(\mathbf{P}_C) := \sum_{c \in \mathcal{C}} \sum_{j \in \mathcal{J}} r(c,j)r(j,c)v(1 + \sum_{k \in \Sigma_j(c)} P_{k,j}(\mathbf{P}_C)) e_j^T P^c \mathbf{v}. \quad (8)$$

Sub-optimality of the naive preference-based ranking. In the introduction section, we briefly introduce an example about how naively ranking by each participants preferences will induce collision and hence reduce the successful matching rate in the networking recommendation example. Here, we quantify the sub-optimality of naive preference-based ranking using the following theorem.

Theorem 4. *In a two-sided market with $|\mathcal{C}| = |\mathcal{J}| = n$, there exists an instance of the interaction model (which is characterized by the choice of preference functions $r(c,j)$, $r(j,c)$, and the examination model v), such as the gap measured by $\mathbf{SW}_{lower}(\mathbf{P}_C)$ between the optimal ranking \mathbf{P}_C^* and the naive preference-based ranking \mathbf{P}_C^{naive} is larger than $\Theta(\sqrt{n})$.*

The sub-optimality of the naive preference-based ranking motivates optimizing the social welfare objective directly, or more tractably our lower bound of the social welfare. In this paper, we work with a specific and widely used examination model $v(x) = 1/x$ afterwards, and show how the lower bound on \mathbf{SW} leads to an efficient optimization procedures. In particular, our *social welfare aware* optimization objective becomes:

$$\begin{aligned} & \underset{P^c, c \in \mathcal{C}}{\text{maximize}} && \sum_{c \in \mathcal{C}} \sum_{j \in \mathcal{J}} \frac{r(c,j)r(j,c)e_j^T P^c \mathbf{v}}{1 + \sum_{k \in \Sigma_j(c)} r(k,j)e_j^T P^k \mathbf{v}} \\ & \text{subject to} && P^c \text{ is doubly stochastic, } c \in \mathcal{C} \end{aligned}$$

Recall that $\mathbf{v} \in \mathbb{R}_+^{|\mathcal{J}|}$ with $\mathbf{v}(i) = 1/i$. This problem is non-convex. However, the objective is a sum of ratios, where both numerator and denominator are linear in P^c . For single-ratio fractional programming, the Charnes-Cooper transformation [4] gives an efficient formulation to decouple the numerator and denominator, and jointly optimizing them. For multi-ratio fractional programming, [14] gives an equivalent formulation using a quadratic transform. Using the quadratic transform, the optimization problem is equivalent to:

$$\text{maximize}_{P^c, \mathbf{y}} \sum_{j \in \mathcal{J}} \sum_{c \in \mathcal{C}} \left(2y_{cj} \sqrt{r(j, c)r(c, j)e_j^T P^c \mathbf{v}} - y_{cj}^2 \left(1 + \sum_{k \in \Sigma_j(c)} r(k, j)e_j^T P^k \mathbf{v} \right) \right) \quad (9)$$

subject to P^c is doubly stochastic, $y_{cj} \in \mathbb{R}$, $\forall c \in \mathcal{C}, j \in \mathcal{J}$

For this equivalent formulation, due to the linearity inside the square-root function and the otherwise linear terms, it is a function concave in P^c given fixed y_{cj} . The problem can thus be solved via an iterative concave-convex procedure:

- Step 1: update y_{cj} using

$$y_{cj} = \frac{\sqrt{r(j, c)r(c, j)e_j^T P^c \mathbf{v}}}{1 + \sum_{k \in \Sigma_j(c)} r(k, j)e_j^T P^k \mathbf{v}}. \quad (10)$$

- Step 2: update P^c for $c \in \mathcal{C}$ by solving Equation 9 for fixed y_{cj} , which can be solved by convex optimization algorithms.

This procedure is guaranteed to find the stationary points. The benefits of using this equivalent form is that it is easy to incorporate additional objectives (see Section 4) in the form of linear constraints without worrying about projection. The problem can still be solved with the same iterative procedure, without adding substantial complexity.

4 Social Welfare vs. Individual Utility

While we so far focused on optimizing the overall effectiveness of the matching process, this clearly raises questions about whether the optimum is fair and desirable to each individual candidate. To begin to understand the trade-off between social welfare and individual utility, we begin with taking a game theoretic perspective [8]. In particular, we can view each candidate in the market as an independent player, where his choice of strategy lies in the selection of his ranking. Consider each individual's utility as the expected number of matches over all employers he could get – the higher the better. In this setting, it is natural to ask: are the action (i.e., the rankings for all candidates) maximizing the social welfare is also a Nash equilibrium? The answer for this question is no. Social welfare is not a valid potential function in this game, since in general some player's best response could hurt others exposure on the employer's side, reduce other players' expected matching rate, and hence make the social welfare lower. The lack of equilibrium poses an interesting questions: how strong is the incentive for people to deviate from the presented ranking, and may they even be better off not participating and using their naive ranking according to their own preferences alone? We envision that we can control both both aspects by adding constraints into the original optimization problem.

Constraint 1: No individual is much worse off comparing with not deploying the proposed ranking. In this constraint, we want to ensure that no individual lose too much in their individual utility, compared with the case that we do not deploy the *social welfare aware* ranking in the market. When we do not deploy the proposed social welfare aware rankings in the market, each individual will be given the naive greedy preference-based ranking. Under examination model $v(x) = 1/x$, the examination probability of candidate c to employer j in the naive preference-based ranking is $\frac{1}{\sigma_c(j)}$, then the probability to initiate the ranking for candidate c to employer j is $\frac{r(c, j)}{\sigma_c(j)}$. Then we are ready to write our constraints as:

$$\frac{r(j, c)r(c, j)}{\sigma_c(j) \left(1 + \sum_{k \in \Sigma_j(c)} \frac{r(k, j)}{\sigma_k(j)} \right)} - \frac{r(j, c)r(c, j)e_j^T P^c \mathbf{v}}{1 + \sum_{k \in \Sigma_j(c)} r(k, j)e_j^T P^k \mathbf{v}} \leq \epsilon \quad \forall c \in \mathcal{C}, j \in \mathcal{J} \quad (11)$$

Note that these constraints are linear in the P^c matrix we want to optimize, therefore we could incorporate this easily in the aforementioned optimization framework.

Constraint 2: Limit the incentive of any individual to deviate. We aim to add a constraint so that at the solution of the global optimization problem, no individual has a strong incentive to deviate (without access to other individual’s information, such as preference function and rankings given by the recommender system). This means that given other players playing the recommended strategy/ranking, the individual will not gain a lot by switching to his personal greedy strategy (without knowing information about others), and the greedy strategy is just naively ranking this candidate’s preference list.

$$\frac{r(j, c)r(c, j)}{\sigma_c(j)(1 + \sum_{k \in \Sigma_j(c)} r(k, j)e_j^T P^k \mathbf{v})} - \frac{rel(j, c)r(c, j)e_j^T P^c \mathbf{v}}{1 + \sum_{k \in \Sigma_j(c)} r(k, j)e_j^T P^k \mathbf{v}} \leq \epsilon \quad \forall c \in \mathcal{C}, j \in \mathcal{J} \quad (12)$$

These constraints ensure that each individual does not have a strong incentive to deviate, or no each individual will deviate if the cost of deviating is greater than ϵ . Similarly, these constraints are also linear in P^c and could be effectively incorporated in our optimization framework. This enables us to find the ranking with the highest social welfare, while ensuring a certain stability in of the solution.

5 Experiments

We evaluate our proposed method in a semi-synthetic two-sided ranking environment, and we compare its performance with the naive ranking based on individual preference. Our secondary goal is to examine how individual utility is affected by maximizing social welfare, and empirically study the possible values of ϵ we could enforce in the constraints.

Dataset and Environment. We simulate preferences in a two-sided markets using the relevance judgments in the Jester Recommender dataset [5], which contains 4.1 million ratings (ranges from -10.00 to +10.00) of 100 jokes from 73,496 users. We treat movies as employers, and users as candidates in the markets. From this dataset, we treat the rating as the candidate’s preference for the employers. We also need to manually construct employer’s preference for the candidates. This semi-synthetic framework enables us to examine our proposed method in different markets: (1) the preferences of the two sides are similar; (2). the preferences of the two sides are reversed (3). the preferences of the two sides are independent. To construct similar two-sided preferences, we take $r(j, c) = \min\{\max\{r(c, j) + e, -10\}, 10\}$ with $e \sim \mathcal{N}(0, 1)$; for the reversed preferences, we take $r(j, c) = \min\{\max\{-r(c, j) + e, -10\}, 10\}$ with $e \sim \mathcal{N}(0, 1)$; for the independent preferences, we just take $r(j, c) \sim \mathcal{U}(-10, 10)$. Finally we do the min-max normalization for all the preferences to be in $[0, 1]$.

Evaluation Metric and Results. Due to the intractability of the original utility $\mathbf{SW}(\mathbf{P}_C)$, we compare the utility of our proposed method and the naive ranking by individual preference using the lower bound $\mathbf{SW}_{lower}(\mathbf{P}_C)$, which is easy to calculate and can be expected to be a reasonable surrogate. For the naive ranking of each candidate, we just rank each employer deterministically based on its preference to the candidate. For our proposed method, we use the iterative concave-convex procedure proposed in Section 3. In Figure 1, we compare the performance of our proposed method with the naive rankings, over 6 scenarios: 2 different market sizes (one with $|\mathcal{J}| = 20$ and $|\mathcal{C}| = 50$ and one with $|\mathcal{J}| = 50$ and $|\mathcal{C}| = 100$) and 3 different two-sided preferences settings defined above. It can be seen that among all scenarios, the globally optimized rankings provide substantially higher overall utility than the naive rankings. The difference is largest when the preferences on the two sides are uncorrelated.

We also examine the difference between the individual utility for each candidate under our proposed ranking and the naive ranking, which corresponds to the first constraint in Section 4. The histogram is shown in Figure 2. We find that most candidates gain in utility under the globally optimized rankings. However, in the worst case, we find that some candidates could lose around 0.2, while the gain for others can be as high as 0.4. This suggests that explicitly constraining the fairness of the solution as suggested in the previous section is indeed worth considering.

6 Discussion and Future work

In this paper, we have formulated the problem of optimizing rankings in a two-sided matching market, incorporating a behavioral models of how the ranking focuses the attention of the customers. We find




Figure 1: Left: Comparison between proposed method and preference-based ranking when $|\mathcal{J}| = 20$ and $|\mathcal{C}| = 50$; Right: Comparison between proposed method and preference-based ranking when $|\mathcal{J}| = 50$ and $|\mathcal{C}| = 100$.




Figure 2: Histogram of individual utility gain under 3 different environments (same reward, reverse reward and random reward) when $|\mathcal{J}| = 50$ and $|\mathcal{C}| = 100$.

empirically that globally optimizing the rankings can provide substantial gains in social welfare. To make this optimization tractable, we identify a lower bound that can be optimized effectively. For a class examination functions in the position-based model, we propose an efficient iterative algorithm that reduces to a sequence of convex optimization problems. We also explore the relationship between social welfare and individual utility, viewing the problem from the game theoretic perspective and proposing additional linear constraints that can control the balance between maximizing social welfare and achieving stability and aspects of individual fairness.

This paper a number of interesting questions for future work: (1) It will be interesting to empirically investigate to which extent the collision among candidates happens in the current preference-based recommendation systems, and evaluate the performance of the proposed methods using real-world dataset. (2) We assume known preference functions in the proposed framework, it will be important to extend to the setting with unknown preference function. One direct way may be the use the click data on both sides to learn the preference functions $\hat{r}(c, j)$ and $\hat{r}(j, c)$. However, we can also view this as a policy optimization that directly learns the ranking policies using the two-sided click data. (3) In this paper, we view the "stability" of the market as no individual wanting to deviate for each employer. It may also be interesting to view this from an aggregated perspective, i.e., the benefit of deviation is measured by the individual's utility across all employers, and how to effectively incorporate this in our framework.

References

- [1] Itai Ashlagi, Anilesh K Krishnaswamy, Rahul Makhijani, Daniela Saban, and Kirankumar Shiragur. Assortment planning for two-sided sequential matching markets. *arXiv preprint arXiv:1907.04485*, 2019.
- [2] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international acm sigir conference on research & development in information retrieval*, pages 405–414, 2018.
- [3] Garrett Birkhoff. *Lattice theory*, volume 25. American Mathematical Soc., 1940.
- [4] Abraham Charnes and William W Cooper. Programming with linear fractional functionals. *Naval Research logistics quarterly*, 9(3-4):181–186, 1962.
- [5] Ken Goldberg, Theresa Roeder, Dhruv Gupta, and Chris Perkins. Eigentaste: A constant time collaborative filtering algorithm. *information retrieval*, 4(2):133–151, 2001.
- [6] Kalervo Järvelin and Jaana Kekäläinen. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)*, 20(4):422–446, 2002.
- [7] Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. Unbiased learning-to-rank with biased feedback. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 781–789, 2017.
- [8] Fuhito Kojima and Parag A Pathak. Incentives and stability in large two-sided matching markets. *American Economic Review*, 99(3):608–27, 2009.
- [9] Lydia T Liu, Horia Mania, and Michael Jordan. Competing bandits in matching markets. In *International Conference on Artificial Intelligence and Statistics*, pages 1618–1628. PMLR, 2020.
- [10] Martin Mladenov, Elliot Creager, Omer Ben-Porat, Kevin Swersky, Richard Zemel, and Craig Boutilier. Optimizing long-term social welfare in recommender systems: A constrained matching approach. *arXiv preprint arXiv:2008.00104*, 2020.
- [11] Matthew Richardson, Ewa Dominowska, and Robert Ragno. Predicting clicks: estimating the click-through rate for new ads. In *Proceedings of the 16th international conference on World Wide Web*, pages 521–530, 2007.
- [12] Stephen E Robertson. The probability ranking principle in ir. *Journal of documentation*, 1977.
- [13] BK Shah. Distribution of sum of independent integer valued random-variables. *American Statistician*, 27(3):123–124, 1973.
- [14] Kaiming Shen and Wei Yu. Fractional programming for communication systems—part i: Power control and beamforming. *IEEE Transactions on Signal Processing*, 66(10):2616–2630, 2018.
- [15] Ashudeep Singh and Thorsten Joachims. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2219–2228, 2018.
- [16] Alfredo Torrico, Margarida Carvalho, and Andrea Lodi. Provable guarantees for general two-sided sequential matching markets. *arXiv preprint arXiv:2006.04313*, 2020.
- [17] Kun Tu, Bruno Ribeiro, David Jensen, Don Towsley, Benyuan Liu, Hua Jiang, and Xiaodong Wang. Online dating recommendations: matching markets and learning preferences. In *Proceedings of the 23rd international conference on world wide web*, pages 787–792, 2014.
- [18] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th international conference on world wide web*, pages 1171–1180, 2017.
- [19] Honglei Zhuang, Xuanhui Wang, Michael Bendersky, Alexander Grushetsky, Yonghui Wu, Petr Mitrichev, Ethan Sterling, Nathan Bell, Walker Ravina, and Hai Qian. Interpretable learning-to-rank with generalized additive models. *arXiv preprint arXiv:2005.02553*, 2020.

7 Appendix

In this Appendix, we provide proof for Theorem 4, the sub-optimality of the preference-based ranking.

Sub-optimality of naive ranking. Here we show that the gap of the naive ranking algorithm and the optimal solution by our algorithm can be $O(\sqrt{n})$, and it is achieved by identifying some specific instance (particular choice of preference model, examination model). Consider the following instance: there are n employers and n candidates in the market. On the candidate side, they have highly correlated preference over employers, while employers have different preference over candidates. The employers' preference ordering is given by:

j_1	c_1	c_2	c_3	\cdots	c_n
j_2	c_2	c_3	\cdots	c_n	c_1
j_3	c_3	\cdots	c_n	c_1	c_2
\vdots	\vdots				
j_n	c_n	c_1	\cdots	c_{n-2}	c_{n-1}

while the candidates' preference ordering is:

c_1	j_1	j_2	j_3	\cdots	j_n
c_2	j_1	j_2	j_3	\cdots	j_n
c_3	j_1	j_2	j_3	\cdots	j_n
\vdots	\vdots				
c_n	j_1	j_2	j_3	\cdots	j_n

Given this, the preference function is defined by:

$$rel(c_i, j_k) = \frac{1}{\sqrt{k}} \quad rel(j_k, c_i) = \frac{1}{\sqrt{\sigma_{j_k}(c_i)}}$$

The examination model E_m in this case is given by:

$$v(x) = \begin{cases} 1/x, & \text{if } x \leq m \\ 0, & \text{otherwise} \end{cases}$$

This mimic the scenario that even the market size n grows, people tends to only examine the top m recommendations due to time and resource constraints. The naive ranking in this case corresponds to ignoring the two-side perspective, and only recommending the employer to candidates by their preference to the candidates. Therefore, the optimal doubly stochastic matrix should be deterministic and it is given by

$$P^{c_i} = I_n$$

for all the candidates c_i , since they have the same preference judgment. For ease of notation, we let $c_{j_k(i)}$ denotes the candidate who is in rank i in employer j_k 's preference list. The utility of the naive ranking is given by:

$$\begin{aligned} \mathbf{SW}_{lower}(\mathbf{P}^{naive}) &= \sum_{k=1}^n \sum_{i=1}^n \frac{rel(j_k, c_{j_k(i)})rel(c_{j_k(i)}, j_k)}{1 + \sum_{q=1}^{i-1} rel(c_{j_k(q)}, j_k)e_j^T P^{c_{j_k(q)}} v} e_j^T P^{c_{j_k(i)}} v \\ &= \sum_{k=1}^m \sum_{i=1}^m \frac{\frac{1}{\sqrt{k}} \frac{1}{\sqrt{i}}}{1 + (i-1)(\frac{1}{\sqrt{k}} \frac{1}{k})} \frac{1}{k} \\ &\leq \sum_{k=1}^m \frac{1}{k^{3/2}} \sum_{i=1}^m \frac{1}{\sqrt{i}} := f(m) \end{aligned} \quad (13)$$

Here we consider a different ranking \mathbf{P}^s , which is also deterministic and given by:

$$P_{kl}^{c_i} = \begin{cases} 1, & \text{if } rk(c_i) = l \text{ for employer } j_k \\ 0, & \text{otherwise} \end{cases}$$

P^{c_i} is a doubly stochastic matrix from our design, note this comes from the fact $rk(c_i)$ is different for different j_k . The utility for this specific ranking \mathbf{P}^s is given by:

$$\begin{aligned}
\mathbf{SW}_{lower}(\mathbf{P}^s) &= \sum_{k=1}^n \sum_{i=1}^n \frac{rel(j_k, c_{j_k(i)})rel(c_{j_k(i)}, j_k)}{1 + \sum_{q=1}^{i-1} rel(c_{j_k(q)}, j_k)} e_j^T P^{c_{j_k(i)}} v \\
&= \sum_{k=1}^n \sum_{i=1}^m \frac{\frac{1}{\sqrt{k}} \frac{1}{\sqrt{i}}}{1 + \frac{1}{\sqrt{k}} \sum_{q=1}^{i-1} (\frac{1}{q})} \frac{1}{i} \\
&\geq (m^{-3/2}) \sum_{k=1}^n \frac{1}{\sqrt{k} + \sum_{i=1}^m \frac{1}{i}} \\
&= \Theta(\sqrt{n})
\end{aligned} \tag{14}$$

Therefore, we have $\mathbf{SW}_{lower}(\mathbf{P}_{\mathcal{C}}^*) - \mathbf{SW}_{lower}(\mathbf{P}_{\mathcal{C}}^{naive}) \geq \mathbf{SW}_{lower}(\mathbf{P}_{\mathcal{C}}^s) - \mathbf{SW}_{lower}(\mathbf{P}_{\mathcal{C}}^{naive}) = \Theta(\sqrt{n})$ for fixed m .