

# Rawlsian Fairness in Online Bipartite Matching: Two-sided, Group, and Individual

Extended Abstract

Seyed A. Esmaeili, Sharmila Duppala, Vedant Nanda, Aravind Srinivasan, John P. Dickerson  
University of Maryland, College Park  
USA

## ABSTRACT

Online bipartite-matching platforms are ubiquitous and find applications in important areas such as crowdsourcing and ridesharing. In the most general form, the platform consists of three entities: two sides to be matched and a platform operator that decides the matching. The design of algorithms for such platforms has traditionally focused on the operator’s (expected) profit. Recent reports have shown that certain demographic groups may receive less favorable treatment under pure profit maximization. As a result, a collection of online matching algorithms have been developed that give a fair treatment guarantee for one side of the market at the expense of a drop in the operator’s profit. In this paper, we generalize the existing work to offer fair treatment guarantees to *both* sides of the market *simultaneously*, at a calculated worst case drop to operator profit. We consider group and individual Rawlsian fairness criteria. Moreover, our algorithms have theoretical guarantees and have adjustable parameters that can be tuned as desired to balance the trade-off between the utilities of the three sides. We also derive hardness results that give clear upper bounds over the performance of any algorithm.

## KEYWORDS

Online Matching; Fairness; Multi-Objective Optimization.

### ACM Reference Format:

Seyed A. Esmaeili, Sharmila Duppala, Vedant Nanda, Aravind Srinivasan, John P. Dickerson. 2022. Rawlsian Fairness in Online Bipartite Matching: Two-sided, Group, and Individual: Extended Abstract. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, 3 pages.

## 1 INTRODUCTION

Online bipartite matching has been used to model many important applications such as crowdsourcing [5, 10, 17, 18], rideshare [4, 12, 20] and online ad allocation [9, 13]. Although there has been some work which addresses fairness in this topic [8, 11, 14–16, 19], shortcomings still remain. In particular, online bipartite matching systems consist of three entities: the platform operator and the two sides of the market to be matched and the given utility guarantees have so far ignored at least one side. However, in this paper we provide algorithms with a guarantee for the operator’s profit as well fairness guarantees for each side of the market. We consider group as well as individual notions of fairness. Our algorithms have adjustable parameters that enable us to trade the utility of

each side as desired. We also provide impossibility results over the performance of any algorithm and complement our results with empirical verification over a real-world dataset.

## 2 ONLINE MODEL & OPTIMIZATION OBJECTIVES

Our model follows that of [1, 3, 7, 13] and others. We have a bipartite graph  $G = (U, V, E)$  where  $U$  represents the set of static (offline) vertices (workers) and  $V$  represents the set of online vertex types (job types) which arrive dynamically in each round. The online matching is done over  $T$  rounds. In a given round  $t$ , a vertex of type  $v$  is sampled from  $V$  with probability  $p_{v,t}$  with  $\sum_{v \in V} p_{v,t} = 1, \forall t \in [T]$  where  $p_{v,t}$  is known beforehand for each type  $v$  and each round  $t$ . This arrival setting is referred to as the known adversarial distribution (**KAD**) setting [1, 2, 4]. When the distribution is stationary, i.e.  $p_{v,t} = p_v, \forall t \in [T]$ , we have the arrival setting of the known independent identical distribution (**KIID**). Accordingly, the expected number of arrivals of type  $v$  in  $T$  rounds is  $n_v = \sum_{t \in [T]} p_{v,t}$ , which reduces to  $n_v = T p_v$  in the (**KIID**) setting. We assume that  $n_v \in \mathbb{Z}^+$  for (**KIID**) [3]. Every vertex  $u$  ( $v$ ) has a group membership, with  $\mathcal{G}$  being the set of all group memberships, for any vertex  $u \in U$ , we denote its group membership by  $g(u) \in \mathcal{G}$  (similarly, we have  $g(v)$  for  $v \in V$ ). Conversely, for a group  $g$ ,  $U(g)$  ( $V(g)$ ) denotes the subset of  $U$  ( $V$ ) with group membership  $g$ . A vertex  $u$  ( $v$ ) has a set of incident edges  $E_u$  ( $E_v$ ) which connect it to vertices in the opposite side of the graph. In a given round, once a vertex (job)  $v$  arrives, an irrevocable decision has to be made on whether to reject  $v$  or assign it to a neighbouring vertex  $u$  (where  $(u, v) \in E_v$ ) which has not been matched before. Suppose, that  $v$  is assigned to  $u$ , then the assignment is not necessarily successful rather it succeeds with probability  $p_e = p_{(u,v)} \in [0, 1]$ . This models the fact that an assignment could fail for some reason such as the worker refusing the assigned job. Furthermore, each vertex  $u$  has patience parameter  $\Delta_u \in \mathbb{Z}^+$  which indicates the number of failed assignments it can tolerate before leaving the system, i.e. if  $u$  receives  $\Delta_u$  failed assignments then it is deleted from the graph. Similarly, a vertex  $v$  has patience  $\Delta_v \in \mathbb{Z}^+$ , if a vertex  $v$  arrives in a given round, then it would tolerate at most  $\Delta_v$  many failed assignments in that round before leaving the system.

For a given edge  $e = (u, v) \in E$ , we let each entity assign its own utility to that match. In particular, the platform operator assigns a utility of  $w_e^O$ , whereas the offline vertex  $u$  assigns a utility of  $w_e^U$ , and the online vertex  $v$  associates a utility of  $w_e^V$ . This captures entities’ heterogeneous wants. For example, in ridesharing, drivers may desire long trips from nearby riders, whereas the platform operator would not be concerned with the driver’s proximity to the rider, although this maybe the only consideration the rider has. Similar motivations exist in crowdsourcing as well.

*Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online.* © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Letting  $\mathcal{M}$  denote the set of successful matchings made in the  $T$  rounds, then our optimization objectives for each entity in the system are as follows:

- **Operator’s Utility (Profit):** The operator’s expected profit is simply the expected sums of the profits across the matched edges, this leads to  $E[\sum_{e \in \mathcal{M}} w_e^O]$ .
- **Rawlsian Group Fairness:**
  - **Offline Side:** Denote by  $\mathcal{M}_u$  the subset of edges in the matching that are incident on  $u$ . Then our fairness criterion is equal to

$$\min_{g \in \mathcal{G}} \frac{E[\sum_{u \in U(g)} (\sum_{e \in \mathcal{M}_u} w_e^U)]}{|U(g)|}.$$

This value corresponds to the minimum average expected utility received by a group in the offline side  $U$ .

- **Online Side:** Similarly, we denote by  $\mathcal{M}_v$  the subset of edges in the matching that are incident on vertex  $v$ , and define the fairness criterion to be

$$\min_{g \in \mathcal{G}} \frac{E[\sum_{v \in V(g)} (\sum_{e \in \mathcal{M}_v} w_e^V)]}{\sum_{v \in V(g)} n_v}.$$

This value corresponds to the minimum average expected utility received throughout the matching by any group in the online side  $V$ .

- **Rawlsian Individual Fairness:**
  - **Offline Side:** The definition here follows from the group fairness definition for the offline side by simply considering that each vertex  $u$  belongs to its own distinct group. Therefore, the objective is  $\min_{u \in U} E[\sum_{e \in \mathcal{M}_u} w_e^U]$ .
  - **Online Side:** Unlike the offline side, the definition does not follow as straightforwardly. Here we cannot obtain a valid definition by simply assigning each vertex type its own group. Rather, we note that a given individual is actually a given arriving vertex at a given round  $t \in [T]$ , accordingly our fairness criterion is the minimum expected utility an individual receives in a given round, i.e.  $\min_{t \in [T]} E[\sum_{e \in \mathcal{M}_{v_t}} w_e^V]$ , where  $v_t$  is the vertex that arrived in round  $t$ .

### 3 MAIN CONTRIBUTIONS

*Performance Criterion:* We note that we are in the online setting, therefore our performance criterion is the competitive ratio. Denote by  $\mathcal{I}$  the distribution for the instances of matching problems, then  $\text{OPT}(\mathcal{I}) = E_{\mathcal{I} \sim \mathcal{I}}[\text{OPT}(I)]$  where  $\text{OPT}(I)$  is the optimal value of the sampled instance  $I$ . Similarly, for a given algorithm  $\text{ALG}$ , we define the value of the its objective over the distribution  $\mathcal{I}$  by  $\text{ALG}(\mathcal{I}) = E_{\mathcal{D}}[\text{ALG}(I)]$  where the expectation  $E_{\mathcal{D}}[\cdot]$  is over the randomness of the instance and the algorithm. The competitive ratio is then defined as  $\min_{\mathcal{I}} \frac{\text{ALG}(\mathcal{I})}{\text{OPT}(\mathcal{I})}$ .

*Main Contributions:* Our work provides three main contributions. First, we generalize the model and allow an edge to have different utilities with respect to different entities (two sides being matched and the platform operator). Second, we provide algorithms with competitive ratio guarantees for the operator, as well as fairness guarantees for the offline and online sides of the matching in the

form of group or individual fairness. We also consider two arrival settings: the **KIID** and the **KAD** settings. These theoretical guarantees are complemented by hardness results which show that we cannot achieve large competitive ratios over all objectives simultaneously and that group and individual fairness can conflict with one another, hence we have an upper bound on the competitive ratio we can achieve for both of them simultaneously. Below we concretely specify the results:

For the **KIID** arrival setting we have:

**THEOREM 3.1.** *For the **KIID** setting, algorithm  $\text{TSGF}_{\text{KIID}}(\alpha, \beta, \gamma)$  achieves a competitive ratio of  $(\frac{\alpha}{2e}, \frac{\beta}{2e}, \frac{\gamma}{2e})^1$  simultaneously over the operator’s profit, the group fairness objective for the offline side, and the group fairness objective for the online side, where  $\alpha, \beta, \gamma > 0$  and  $\alpha + \beta + \gamma \leq 1$ .*

The following two theorems hold under the condition that  $p_e = 1, \forall e \in E$ . Specifically for the **KAD** arrival setting we have:

**THEOREM 3.2.** *For the **KAD** setting, algorithm  $\text{TSGF}_{\text{KAD}}(\alpha, \beta, \gamma)$  achieves a competitive ratio of  $(\frac{\alpha}{2}, \frac{\beta}{2}, \frac{\gamma}{2})$  simultaneously over the operator’s profit, the group fairness objective for the offline side, and the group fairness objective for the online side, where  $\alpha, \beta, \gamma > 0$  and  $\alpha + \beta + \gamma \leq 1$ .*

Moreover, for the case of individual fairness whether in the **KIID** or **KAD** arrival setting we have:

**THEOREM 3.3.** *For the **KIID** or **KAD** setting, we can achieve a competitive ratio of  $(\frac{\alpha}{2}, \frac{\beta}{2}, \frac{\gamma}{2})$  simultaneously over the operator’s profit, the individual fairness objective for the offline side, and the individual fairness objective for the online side, where  $\alpha, \beta, \gamma > 0$  and  $\alpha + \beta + \gamma \leq 1$ .*

We also give the following hardness results. In particular, for a given arrival (**KIID** or **KAD**) setting and fairness criterion (group or individual), the competitive ratios for all sides cannot exceed 1 simultaneously:

**THEOREM 3.4.** *For all arrival models, given the three objectives: operator’s profit, offline side group (individual) fairness, and online side group (individual) fairness. No algorithm can achieve a competitive ratio of  $(\alpha, \beta, \gamma)$  over the three objectives simultaneously such that  $\alpha + \beta + \gamma > 1$ .*

It is natural to wonder if we can combine individual and group fairness. Though our algorithms can be extended to give such guarantees easily, we show the following hardness result which shows that individual and group fairness can conflict with one another. In fact, this is the case even when ignoring the operator’s profit and fairness on the other side of the graph.

**THEOREM 3.5.** *Ignoring the operator’s profit and focusing either on the offline side alone or the online side alone. With  $\alpha_G$  and  $\alpha_I$  denoting the group and individual fairness competitive ratios, respectively. No algorithm can achieve competitive ratios  $(\alpha_G, \alpha_I)$  over the group and individual fairness objectives of one side simultaneously such that  $\alpha_G + \alpha_I > 1$ .*

Please see the full manuscript [6] for the full details of the proofs as well as experimental results.

<sup>1</sup>Here,  $e$  denotes the Euler number, not an edge in the graph.

## REFERENCES

- [1] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. 2012. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. 18–35.
- [2] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. 2013. The online stochastic generalized assignment problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer, 11–25.
- [3] Nikhil Bansal, Anupam Gupta, Jian Li, Julián Mestre, Viswanath Nagarajan, and Atri Rudra. 2010. When lp is the cure for your matching woes: Improved bounds for stochastic matchings. In *European Symposium on Algorithms*. Springer, 218–229.
- [4] John P Dickerson, Karthik A Sankararaman, Aravind Srinivasan, and Pan Xu. 2018. Allocation problems in ride-sharing platforms: Online matching with offline reusable resources. In *AAAI*.
- [5] John P Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. 2019. Balancing relevance and diversity in online bipartite matching via submodularity. In *AAAI*.
- [6] Seyed A Esmaili, Sharmila Dupplala, Vedant Nanda, Aravind Srinivasan, and John P Dickerson. 2022. Rawlsian Fairness in Online Bipartite Matching: Two-sided, Group, and Individual. *arXiv preprint arXiv:2201.06021* (2022).
- [7] Jon Feldman, Aranyak Mehta, Vahab Mirrokni, and S. Muthukrishnan. 2009. Online Stochastic Matching: Beating 1-1/e. *arXiv:0905.4100* [cs.DS]
- [8] David García-Soriano and Francesco Bonchi. 2020. Fair-by-design matching. *Data Mining and Knowledge Discovery* (2020), 1–45.
- [9] Gagan Goel and Aranyak Mehta. 2008. Online budgeted matching in random input models with applications to Adwords. In *SODA*.
- [10] Chien-Ju Ho and Jennifer Vaughan. 2012. Online task assignment in crowdsourcing markets. In *AAAI*.
- [11] Nixie S Lesmana, Xuan Zhang, and Xiaohui Bei. 2019. Balancing efficiency and fairness in on-demand ridesourcing. In *NeurIPS*.
- [12] Meghna Lowalekar, Pradeep Varakantham, and Patrick Jaillet. 2018. Online spatio-temporal matching in stochastic and dynamic domains. *Artificial Intelligence (AIJ)* 261 (2018), 71–112.
- [13] Aranyak Mehta et al. 2013. Online Matching and Ad Allocation. *Foundations and Trends in Theoretical Computer Science* 8, 4 (2013), 265–368.
- [14] Vedant Nanda, Pan Xu, Karthik Abinav Sankararaman, John P Dickerson, and Aravind Srinivasan. 2020. Balancing the Tradeoff between Profit and Fairness in Rideshare Platforms During High-Demand Hours. In *AAAI*.
- [15] Govind S Sankar, Anand Louis, Meghana Nasre, and Prajakta Nimbhorkar. 2021. Matchings with Group Fairness Constraints: Online and Offline Algorithms. *arXiv preprint arXiv:2105.09522* (2021).
- [16] 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 *KDD*.
- [17] Yongxin Tong, Jieying She, Bolin Ding, Libin Wang, and Lei Chen. 2016. Online mobile micro-task allocation in spatial crowdsourcing. In *ICDE*.
- [18] Jinfu Xia, Yan Zhao, Guanfeng Liu, Jiajie Xu, Min Zhang, and Kai Zheng. 2019. Profit-driven Task Assignment in Spatial Crowdsourcing. In *IJCAI*.
- [19] Yifan Xu and Pan Xu. 2020. Trade the System Efficiency for the Income Equality of Drivers in Rideshare. In *IJCAI*.
- [20] Boming Zhao, Pan Xu, Yexuan Shi, Yongxin Tong, Zimu Zhou, and Yuxiang Zeng. 2019. Preference-aware task assignment in on-demand taxi dispatching: An online stable matching approach. In *AAAI*.