Fair Sharing for Sharing Economy Platforms

Abhijnan Chakraborty
Max Plank Institute for Software Systems, Germany
Indian Institute of Technology Kharagpur, India

Asia J. Biega
Max Planck Institute for Informatics
Saarland Informatics Campus, Germany

Aniko Hannak
Central European University, Hungary

Krishna P. Gummadi
Max Plank Institute for Software Systems, Germany

ABSTRACT

Sharing economy platforms, such as Airbnb, Uber or eBay, are an increasingly common way for people to provide their services to earn a living. Yet, the focus in these platforms is either on the satisfaction of the customers of the service, or on boosting successful business transactions. However, recent studies provide a multitude of reasons to worry about the providers in the sharing economy ecosystems. The concerns range from bad working conditions and worker manipulation to discrimination against minorities. This is worsened by the fact that the algorithms used for matching customers and providers, that de facto decide the amount of exposure each provider receives, are proprietary and non-transparent.

In this position paper, we propose a novel framework to think about fairness in the matching mechanisms of online sharing economy platforms. Specifically, we focus on various fairness guarantees from the providers' perspective. Our notion of fairness relies on the idea that, spread *over time*, all providers should receive the amount of exposure proportional to their relevance or the utility they provide. We postulate that by not requiring every match to be fair, but rather distributing the fairness over time, we can (i) give better guarantees in terms of the overall benefit for the providers and the customers, (ii) make use of implementations from a long line of research concerned with fair division of constrained resources. Overall, our work takes the first step towards rethinking fairness in online sharing economy systems with an additional emphasis on the well-being of providers, and provides insights into parallels with well-established practical implementations in other domains.

1 INTRODUCTION

While the *sharing economy* or *two-sided market* has traditionally been thought of as a movement towards more democratized marketplaces, increasing number of studies and articles are concerned with the potential discriminatory effects of some sharing economy *giants* such as Uber or AirBnB [8, 21]. Sharing economy is loosely defined as *peer-to-peer-based activity of obtaining, giving, or sharing the access to goods and services*, which is coordinated through a

web-based platform or a mobile-app [12, 28]. The rise of social technological systems and online platforms has enabled it to become a major competitor to the traditional (B2C) economic model. According to a report by the United States Department of Commerce from 2016, the spending in the most common areas of sharing economy was 5% of the total economy, and this is predicted to grow to 50% by 2025 [23]. Recognizing the opportunity, startups that build sharing economy platforms have mushroomed and become a mega trend among Silicon Valley investors, attracting millions of dollars in venture capital funds [26].
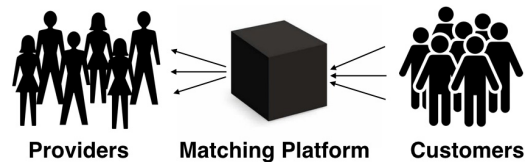

Figure 1: Different stakeholders in sharing economy framework: providers, customers and the platform.

In the sharing economy framework, there are three stakeholders: (i) **providers** of goods and services, (ii) **customers** who pay for them, and (ii) the **platform** which provides the *matching* between the providers and customers. As depicted in Figure 1, the platform lies in the center of the framework, enabling the providers and the customers to connect and do business. Crucially, the platform has a control over the exposure of service providers to potential customers and vice-versa. For example, Uber matches drivers with passengers "under-the-hood" [8]. In case of Airbnb or different freelance websites, customers have more control over the choice of the provider, but the platform still decides how much exposure and attention each provider gets and which customers they are shown to. Hence, a pertinent question to ask in this context is ***what should be the objectives of the matching platform?***

Most matching platforms today try to maximize the *utility* for the customers, driven by the *customer is always right* philosophy [16]. The underlying idea is that the party paying the platform for using the service (passengers in Uber, renters in Airbnb, employers in Freelancer, or buyers on Amazon) should be most satisfied, yielding the most revenue for the platform. However, the providers may get squeezed in the process. For example, Uber drivers complain that they have to work long hours for little pay [11]. Similarly, sellers in ecommerce marketplaces like Amazon complain about increasing participation costs and declining profits [25].

Moreover, recent studies have shown that inequalities from the offline world easily transfer to these online platforms, making it harder for minorities to succeed. This is, for instance, partly due to

biased hiring choices of the employers (job platform customers) [1, 10, 14, 24], or biased evaluations and ratings of providers, based on their gender, race, etc. [13, 15]. However, even more worrying are the effects of the big data algorithms that the platforms use to match customers with providers. A lot of prior works have raised concerns regarding the biases algorithmic decisions carry [5–7, 9, 18, 22]. Likewise, the algorithms deployed by the matching platforms may reproduce or even reinforce biases that are present among the customers, leading to a *rich-get-richer* effect for the providers. Thus, while optimizing for the most profitable matching, unfortunately the minorities among the providers may not receive their deserved attention.

In this position paper, we argue for establishing the notion of **fairness of matching mechanisms in sharing economy platforms**. The related research questions can be threefold:

RQ1. What notions, measures and criteria of fairness should be applied on these platforms?

RQ2. Do the existing sharing economy platforms satisfy them?

RQ3. How can we (re)design platforms to satisfy the fairness criteria identified in RQ1?

In this work, we focus on answering RQ1, and leave the other research questions for future work. Fairness of the matching can be ensured by optimizing for: (a) fairness for providers, (b) fairness for customers, or (c) fairness for both customers and providers. In this work, we focus on goal (a), investigating the trade-off between utility for customers and fairness to providers, and leaving other scenarios as subjects of future work.

In a marked departure from past efforts on defining fairness in search and recommendation systems [17, 19], we add a **temporal dimension to the notion of fairness**. We argue that a *fair matching platform* would be something that distributes the exposure of providers to customers *over time in a fair way*. The exact fairness notions are discussed in Section 2.

The advantage of introducing the time dimension into the notion of fairness in sharing economy platforms is two fold:
(i) The fairness notion becomes more relaxed than the constrained requirement of being fair with respect to every matching, and
(ii) Abstractly, the problem maps to one of fair division of a constrained resource (in this context, the exposure of providers to customers) over time. There is a long line of work on bringing fairness in generalized processor sharing algorithms [2, 29], which could be applied in a temporally fair matching system.

In summary, in this position paper, we make two contributions: (i) we provide a systematic way to think about fairness in sharing economy platforms, and (ii) by introducing the temporal dimension to the notion of fairness in such platforms, we enable reuse of the existing techniques for fairness in processor scheduling. In future work, we would go beyond fairness for providers, and consider fairness for all members of the sharing ecosystem. Subsequently, we would like to investigate RQ2 and RQ3 as outlined earlier, to provide a comprehensive understanding and potentially required addition of fairness in today's sharing economy platforms.

2 NOTION OF FAIRNESS IN SHARING ECONOMY PLATFORMS

In this section, we first present the system model of a matching platform. Then, we introduce the notions of fairness for such platforms.

2.1 System Model

We present the system model for the matching platform in Figure 2 to conceptualize the notions of fairness. The platform produces the sequence of matches between customers and providers over time. The matching decision at any time instant t can be represented as a tuple $\{C_i, P_j, t\}$ involving the customer-provider pair $\{C_i, P_j\}$.
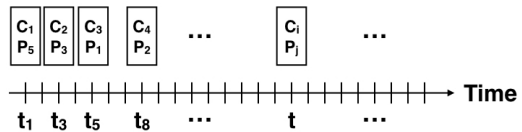
Figure 2: The matching platform produces the matching between different customers and providers over time.

Then, we define two functions over the tuple $\{C_i, P_j, t\}$ at time t:

(i) Utility function $U(\{C_i, P_j, t\})$: A measure of utility derived by the customer C_i for a matching $\{C_i, P_j, t\}$, and

(ii) Benefit function $B(\{C_i, P_j, t\})$: A measure of benefit received by the provider P_j for the matching $\{C_i, P_j, t\}$.

The exact form of utility function may vary from domain to domain. For example, for a renter in Airbnb, utility will depend on the size of the rooms, cleanliness, the location, or the price of the rental property; whereas, for an employer in Freelancer, the utility will be decided based on the qualification of the provider to carry out the task. Similarly, the exact benefit function for the providers may also vary depending on the domain, although the benefit will often map to the revenue of the provider earned by finishing the task. The customer utility may not be as explicit as the provider benefit.

For a customer request at any given time, the platform can produce a list of available providers ranked according to their relevance. The relevance corresponds to the predicted utility the matched provider would generate for the customer. Starting from the top ranked provider, the platform can then move down the list, deciding on the match by additionally considering other simultaneous and outstanding demands. While such an approach would maximize the utility for the customers (i.e., maximize $U(\{C_i, P_j, t\}), \forall t$), in platforms where a provider becomes available as soon as the formerly assigned task is complete, often only a few providers (e.g., the ones that happen to have the highest rating) would be matched again and again, resulting in the *starvation* of other providers in a rich-get-richer scheme. Such a scenario might lead to amplification of small differences in ratings between equally relevant providers, and would make it especially difficult for new, equally qualified providers, to join the system.

Since there is no fairness for the providers in the standard matching scheme we described, we propose two notions of provider fairness that the matching platform can attempt to guarantee.

2.2 Parity Fairness

A simple notion of fairness would be to maintain parity between the providers, i.e., **over time, the sum of received benefits of all providers should be similar.** This can be achieved by maintaining a virtual queue of all providers, and exposing them to the customers in a round-robin manner.

Advantages and limitations: Although the above notion guarantees that no provider is starved and they all accumulate similar benefits over the long term, it assumes all matches are equally relevant and is thus oblivious to customer utility. However, different customers have different service requirements, and hence they might gain no utility if a non-relevant provider is matched. For example, Uber offers the option to hire different cabs like hatchbacks, SUVs etc. Therefore, a passenger asking for a SUV would get zero utility if a hatchback driver is assigned to her. Similarly, in Freelancer some employers might need people with experience in Machine Learning, whereas others may want to hire people with background in Networking. In such cases, the employers would gain little utility if the platform does not match them to people with relevant expertise.

2.3 Proportional Fairness

Actual and deserved benefits

To control for the missed benefit opportunities of all relevant drivers, we introduce two types of benefits for a provider P_j:
(i) **Actual benefit**, which is the sum of the benefits received by P_j from the matches made by the platform till time t, and
(ii) **Deserved benefit**, which is the sum of benefits provider P_j would have made on all platform matches made until time t: $\{C_i, P_j, t_k\}$, $t_k < t$, scaled by the utility the customer C_i would have received on that match.

On each customer request, the platform predicts the customer utility (relevance) and provider benefit for every provider-customer pair. If a provider is matched to the customer, the platform updates the *actual benefit* of the provider. Otherwise, the missed benefit of the non-match scaled by the relevance (customer utility) is added to the *deserved benefit* of the provider. In other words, the deserved benefit accumulates all benefits which could have been accrued by the provider if she was matched to all requesting customers, scaled by the chance of matching expressed by the match relevance.

Proportionally fair sharing-economy platforms

Using the definitions of deserved and actual benefits provided above, we propose a new notion of fairness in sharing economy platforms: A platform is fair if **over time, the actual benefit is proportional to the deserved benefit for every provider.**

The proposed notions have a number of useful properties:
(i) the deserved benefit of non-relevant providers does not increase on missed matches, (ii) the deserved benefit of equally relevant providers will increase at the same rate, (iii) the deserved benefit will increase proportionately to the potential benefit of each match, which enables one to control the actual and deserved benefit ratios.

In Uber, for example, we can think of the deserved benefit for drivers as the accumulation of potential benefits if they were matched with every passenger requesting a ride. However, the deserved benefit for a hatchback driver will not increase if the passenger asked for SUV. The deserved benefit will grow reverse proportionately

to the distance between the passenger's location and the driver's location, when location is used to determine relevance. When a driver is logged out of the system, her deserved benefit does not change. Moreover, as the deserved benefit for a provider will increase during her waiting time when logged-in to the system, the proposed definition ensures there is no starvation.

3 MAPPING TO PROCESSOR SCHEDULING

When looking at the problem of fairness in two-sided markets in a temporal way, we can draw an analogy to problems of fair sharing of constrained resource, such as a processor or network bandwidth. In our scenario, the customers requesting a match correspond to incoming packets, and the providers correspond to processors with limited capacity. The goal of the resource sharing algorithms is to spread the workload among the processors in a fair way.

The problem of fairness has been extensively studied in realtime systems and networking literature, where the ideal notion of fair scheduling is captured by **generalized processor sharing** (GPS) algorithms [29]. However, GPS algorithms make certain assumptions which do not hold in practice. For example, they assume that the traffic in a network is *fluid* and can be split at any arbitrary point. In reality, traffic comes in discrete packets. Analogously, in the context of sharing economy platforms, the matching between a provider and a consumer constitutes a discrete event.

To escape the unrealistic assumptions of GPS, alternative techniques such as **Weighted Fair Queuing** [2] have been proposed. An implementation of fair matching in sharing economy platforms could directly utilize the weighted fair queuing techniques. An especially promising direction is to use a priority based scheduler [27], which would generate the matching between providers and consumers depending on the priority value

$$PR_j = 1 - \frac{AB_j}{DB_j}$$

where AB_j and DB_j denote the actual and desired benefits of provider j. At time t, the matching platform should pick a provider with high relevance and high priority value.

Finally, utilizing previous work on **Hierarchical Fair Scheduling** [3], the proposed solution could be generalized to tackle the problem of group fairness. Recent works have shown that consumers tend to give biased evaluations and ratings based on providers' gender, race, and other protected attributes [13, 15]. This means that the utility function can be biased against some socially salient groups (e.g., women, blacks) and recreate inequality in the matching. Hierarchical Fair Scheduling would allow the platforms to implement group fairness on top of the individual fairness, and control for factors such as equal exposure of men and women, minimum exposure for a minority group, or closing the wage gap between different demographic groups.

4 CONCLUDING DISCUSSION

In this position paper, we have provided a systematic way to think about fairness in sharing economy platforms. The proposed definition of fairness incorporating the temporal dimension enables implementations that could draw from past works on fairness in processor scheduling.

While our focus in this paper was to consider fairness for the providers, the are remaining questions regarding the incentives of different stakeholders to participate in a fair matching framework.

Incentives for the providers: The obvious incentive for providers is the guarantee of getting a fair share of exposure, which translates to equal income opportunities. Additionally, there is an intrinsic feedback loop in the two-sided market systems – with more providers joining a fair system, the pool of customers can increase, guaranteeing a more steady demand for the providers.

Incentives for the customer: In many two-sided markets, matches of customers with a single most relevant provider do not occur. Rather, it is often the case that many providers are equally relevant to a request, and that otherwise the relevance scale is discrete (e.g., relevant, somewhat relevant, not relevant). For example, in cab-riding services like Uber, there is not much difference between the skillset of the providers, therefore all equally relevant providers should have an equal chance of being matched to customers.

Customers should also care about having a wide pool of providers available for the times when the demand is higher, and having fair matching algorithms may help keep providers inside the system. Alternatively, the platform might explicitly ask the customers to participate in the fair matching scheme by offering additional monetary incentives. For example, Uber might offer a cheaper fare if a passenger would be willing to wait longer to get the cab. There are research efforts on designing incentive strategies [20], which may be applicable in this context.

Incentives for the platform: The matching platforms in the two-sided markets thrive on attracting both providers and customers. With a wide pool of users on both sides, they will be more resilient to the loss of user interest. Interestingly, recent incidents like the *Delete Uber* movement [4] show that the users are motivated to boycott platforms due to unfairness, among other reasons. Therefore, building user trust and loyalty by guaranteeing fairness for different stakeholders is a reasonable strategy towards a long-term success of the company.

Overall, our work provides a novel way of thinking about fairness in sharing economy platforms. We hope that it will spark the research on fair matching implementations, and investigations into the guarantees that could be provided for different ecosystem participants.

REFERENCES

[1] Ian Ayres, Mahzarin Banaji, and Christine Jolls. 2015. Race effects on eBay. *The RAND Journal of Economics* 46, 4 (2015).

[2] Jon CR Bennett and Hui Zhang. 1996. WF/sup 2/Q: worst-case fair weighted fair queueing. In *IEEE INFOCOM*.

[3] Jon CR Bennett and Hui Zhang. 1997. Hierarchical packet fair queueing algorithms. *IEEE/ACM Transactions on Networking (TON)* 5, 5 (1997).

[4] Emma Brockes. 2017. Is it OK to use Uber now that Travis Kalanick has resigned? theguardian.com/technology/2017/jun/28/uber-travis-kalanick-should-i-delete-app. (2017).

[5] Abhijnan Chakraborty, Saptarshi Ghosh, Niloy Ganguly, and Krishna P Gummadi. 2015. Can trending news stories create coverage bias? on the impact of high content churn in online news media. In *Computation and Journalism Symposium*.

[6] Abhijnan Chakraborty, Saptarshi Ghosh, Niloy Ganguly, and Krishna P Gummadi. 2016. Dissemination Biases of Social Media Channels: On the Topical Coverage of Socially Shared News. In *AAAI ICWSM*.

[7] Abhijnan Chakraborty, Johnnatan Messias, Fabricio Benevenuto, Saptarshi Ghosh, Niloy Ganguly, and Krishna P Gummadi. 2017. Who Makes Trends? Understanding Demographic Biases in Crowdsourced Recommendations. In *AAAI ICWSM*.

[8] Le Chen, Alan Mislove, and Christo Wilson. 2015. Peeking Beneath the Hood of Uber. In *ACM IMC*.

[9] Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2015. Automated Experiments on Ad Privacy Settings: A Tale of Opacity, Choice, and Discrimination. In *PETS*.

[10] Benjamin G. Edelman, Michael Luca, and Dan Svirsky. 2017. Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment. *American Economic Journal: Applied Economics* (2017).

[11] Carla Green and Sam Levin. 2017. Homeless, assaulted, broke: drivers left behind as Uber promises change at the top. theguardian.com/us-news/2017/jun/17/uber-drivers-homeless-assault-travis-kalanick. (2017).

[12] Juho Hamari, Mimmi Sjöklint, and Antti Ukkonen. 2016. The sharing economy: Why people participate in collaborative consumption. *Journal of the Association for Information Science and Technology* 67, 9 (2016).

[13] Anikó Hannák, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson. 2017. Bias in Online Freelance Marketplaces: Evidence from TaskRabbit and Fiverr. In *ACM CSCW*.

[14] Tamar Kricheli-Katz and Tali Regev. 2016. How many cents on the dollar? Women and men in product markets. *Science Advances* 2, 2 (2016).

[15] Michael Luca. 2016. Reviews, reputation, and revenue: The case of Yelp. com. (2016).

[16] Hughston McBain. 1944. Are customers always right. *The Rotarian* (1944).

[17] Rishabh Mehrotra, Ashton Anderson, Fernando Diaz, Amit Sharma, Hanna Wallach, and Emine Yilmaz. 2017. Auditing Search Engines for Differential Satisfaction Across Demographics. In *ACM WWW Companion*.

[18] Christian Sandvig, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. 2014. Auditing algorithms: Research methods for detecting discrimination on internet platforms. *Data and discrimination: converting critical concerns into productive inquiry* (2014).

[19] Dimitris Serbos, Shuyao Qi, Nikos Mamoulis, Evaggelia Pitoura, and Panayiotis Tsaparas. 2017. Fairness in Package-to-Group Recommendations. In *ACM WWW*.

[20] Adish Singla, Marco Santoni, Gábor Barták, Pratik Mukerji, Moritz Meenen, and Andreas Krause. 2015. Incentivizing Users for Balancing Bike Sharing Systems.. In *AAAI*.

[21] Ana Swanson. 2015. *Racial bias in everything: Airbnb edition.* The Washington Post. washingtonpost.com/news/wonk/wp/2015/12/12/racial-bias-in-everything-airbnb-edition.

[22] Latanya Sweeney. 2013. Discrimination in online ad delivery. (2013).

[23] Rudy Telles Jr. 2016. Digital Matching Firms: A New Definition in the "Sharing Economy" Space. *ESA Issue Brief* (2016).

[24] Jacob Thebault-Spieker, Loren G. Terveen, and Brent Hecht. 2015. Avoiding the South Side and the Suburbs: The Geography of Mobile Crowdsourcing Markets. In *ACM CSCW*.

[25] Siddharth Tiwari. 2017. 'Underhand tactics' boost sales in e-commerce. sundayguardianlive.com/investigation/9881-underhand-tactics-boost-sales-e-commerce. (2017).

[26] Matt Vella. 2012. The mega trend that swallowed Silicon Valley. fortune.com/2012/10/03/the-mega-trend-that-swallowed-silicon-valley. (2012).

[27] Ji Yang, Zhang Yifan, Wang Ying, and Zhang Ping. 2004. Average rate updating mechanism in proportional fair scheduler for HDR. In *IEEE GLOBECOM*.

[28] Georgios Zervas, Davide Proserpio, and John W Byers. 2014. The rise of the sharing economy: Estimating the impact of Airbnb on the hotel industry. *Journal of Marketing Research* (2014).

[29] Zhi-Li Zhang, Don Towsley, and Jim Kurose. 1995. Statistical analysis of the generalized processor sharing scheduling discipline. *IEEE Journal on Selected Areas in Communications* 13, 6 (1995).