

Fairness Dynamics in Digital Economy Platforms with Biased Ratings

Martin Smit
University of Amsterdam
Amsterdam, Netherlands
j.m.m.smit@uva.nl

Fernando P. Santos
University of Amsterdam
Amsterdam, Netherlands
f.p.santos@uva.nl

ABSTRACT

The digital services economy consists of online platforms that facilitate interactions between service providers and consumers. This ecosystem is characterized by short-term, often one-off, transactions between parties that have no prior familiarity. To establish trust among users, platforms employ rating systems which allow users to report on the quality of their previous interactions. However, while arguably crucial for these platforms to function, rating systems can perpetuate negative biases against marginalised groups. This paper investigates how to design platforms around biased reputation systems, reducing discrimination while maintaining incentives for all service providers to offer high quality service for users. We introduce an evolutionary game theoretical model to study how digital platforms can perpetuate or counteract rating-based discrimination. We focus on the platforms' decisions to promote service providers who have high reputations or who belong to a specific protected group. Our results demonstrate a fundamental trade-off between user experience and fairness: promoting highly-rated providers benefits users, but lowers the demand for marginalised providers against which the ratings are biased. Our results also provide evidence that intervening by tuning the demographics of the search results is a highly effective way of reducing unfairness while minimally impacting users. Furthermore, we show that even when precise measurements on the level of rating bias affecting marginalised service providers is unavailable, there is still potential to improve upon a recommender system which ignores protected characteristics. Altogether, our model highlights the benefits of proactive anti-discrimination design in systems where ratings are used to promote cooperative behaviour.

KEYWORDS

Fairness; Cooperation; Evolutionary Game Theory; Recommender Systems; Digital Platforms; Indirect Reciprocity

ACM Reference Format:

Martin Smit and Fernando P. Santos. 2026. Fairness Dynamics in Digital Economy Platforms with Biased Ratings. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/CEJT6762>



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/CEJT6762>

1 INTRODUCTION

Evidence of discrimination against marginalised communities on online platforms is widespread in both the academic community and beyond. Research has found that guests on Airbnb with Black-sounding names are accepted 16% less often than guests with White-sounding names [11] and Black Airbnb hosts in the United States charge 5-7% less than White hosts for equivalent properties [16]. On Uber, users with Black-sounding names experience twice the cancellation rate than White-sounding names [12]. On eBay, after accounting for reviews, it was found that listings with photos showing a Black hand holding a baseball card sold for 20% less on average than when a White hand was photographed [4], and that women receive fewer and lower bids than men when selling identical items in new condition, leading to, again, a final sell price of 20% less on average [19]. Such biases do not require explicit group identifiers. Evidence shows that when information about a user or service provider's demographic characteristics are not directly available, they can be inferred [19]. Furthermore, other information can be used as a proxy, such as on Airbnb where the neighbourhood majority ethnicity is a statistically significant predictor of price after controlling for all observable features [14].

While research indicates that discrimination based on demographic characteristics can be partially alleviated by rating systems [3], these systems have, themselves, been shown to be susceptible to bias [14, 15] leading to lower demand, prices, and revenue for those discriminated against [36]. As argued by van Doorn et al. [38], while precise, cross-platform data on the socioeconomic, migratory, and residence status of those working on digital platforms is unavailable, data does indicate that migrant workers are hugely overrepresented on these platforms, even in countries with large domestic labour markets such as China and India [8, 28, 34].

Thus, a large number of economically (and often legally) vulnerable individuals are impacted by bias in the local populations for which they provide services. For many, this bias affects a significant proportion of their total earnings, as data from a 2022 Eurostat survey revealed that platform labour made up more than 75% of total income for almost a 25% of participants who reported using digital platforms for employment [1].

Given the importance placed on digital economies by various (inter-)governmental bodies such as the US Bureau of Labor Statistics [10] and the European Commission [2], there is a rich body of research analysing these problems from a legal, policy, and economic view. Although the potential sources of bias have been extensively identified, and many potential solutions proposed, there are few predictive models that explain the circumstances and mechanisms leading to unfair outcomes [7, 20]. Indeed, as highlighted in [20] as an open problem, it is unclear the extent to which the

decisions made by the platform designers can perpetuate or alleviate the biases of the platform’s users. Furthermore, besides one suggestion we discuss in Section 2.4, these works stop short of making explicit recommendations to platform designers on how to improve the situation.

In this paper we address this issue by analysing the role of the recommendation algorithm responsible for facilitating matches between users and providers on digital platforms. We introduce an evolutionary game-theoretical model of a digital platform in which providers adapt to user- and platform-based incentives to maintain high quality service, a trait essential to the longevity of a platform. This allows us to identify the emergent strategic and utility effects of platform designers’ decisions to tweak their algorithms by showing more or fewer providers of a certain rating or group.

Our contributions can be summarised as follows:

- (1) We formulate an evolutionary game theory (EGT) model that captures how different groups of agents respond to and are affected by recommender systems.
- (2) We numerically analyse the model’s strategy dynamics to define an intuitive necessary condition for a platform to sufficiently incentivise high-effort behaviour.
- (3) Using the model, we highlight the trade-off between user experience and group-level fairness faced by platform designers under the constraint of treating all providers equally despite rating bias.
- (4) We demonstrate how this trade-off vanishes when the group against which the population is biased can be algorithmically prioritised and that there is still a benefit in doing so under moderate uncertainty in the extent of the bias.

Structure of the paper. In Section 2, we summarise the related literature on fairness in two-sided recommender systems, evolutionary models of fairness more generally, and the two aforementioned mathematical models of bias in online platforms. In Section 3 we formulate the problem and introduce our model and its dynamics. Then, we present the experimental design of our simulations in Section 4 and the evaluation metrics we use to judge the effectiveness of platforms in Section 4.1. In Section 5.1 we analyse how the model’s dynamics respond to different parameter setups and define the necessary condition for a platform to succeed based on its strategy dynamics. Subsequently, we show in Section 5.2 how forcing algorithmic equal treatment creates a trade-off between provider fairness and user experience, and how is no longer the case when this constraint is removed in Section 5.3, and even when uncertainty is introduced in Section 5.4. Finally, in Section 6 we discuss the implications our results have for designers wanting to improve platforms, as well as avenues for future research. Our code (and the appendix) is available at [31].

2 RELATED LITERATURE

2.1 Fairness in two-sided recommender systems

Underpinning digital platforms are machine learning algorithms that attempt to reduce search friction for users by showing them relevant content, items, or service providers. A key challenge identified in such systems is the difficulty of simultaneously providing

accurate recommendations for both the consumer side which are fair for the producer side [5, 26].

One fundamental difficulty discussed extensively in [9] is that there are many, sometimes incompatible, definitions of fairness in recommender systems, each carrying an implicit set of normative judgements about the world. In [13] and [33], as well as our paper, we identify group-level fairness, as opposed to individual fairness, to be the most effective way to examine how protected characteristics can affect individuals.

Geyik et al. [13] provide specific algorithms which they show, both theoretically and in practice on LinkedIn data, are able to ensure the distribution of protected characteristics in the top k search results follows some desired distribution. Sühr et al. [33] then apply one of these algorithms to explore how the interactions between fair ranking algorithms and the inherent gender biases of employers can have a real impact on hiring decisions from online platforms. They conclude that the effectiveness of such algorithms is clear but that it may also be lessened by gender biases related to the task for which the candidate is being hired. Building on these compartmental studies, our paper shows the conditions under which such fair algorithms can be deployed without detrimentally harming the incentive structures that encourage high-effort behaviours. As we discuss in Section 4.1, we measure fairness by demographic parity ratio due to the ubiquity and simplicity of parity-based metrics.

Another concept which is important for two-sided recommender systems is the temporal nature of fairness [6, 26]. Given our use of stochastic processes to model digital platforms, our model explicitly includes an element of time, which allows us to extract both instantaneous and long-term measurements of fairness, the latter of which being the focus of this paper.

2.2 Fairness in multi-agent (reputation) systems

A large body of work studies how costly cooperation can be sustained in large populations through utilising reputations and **indirect reciprocity** [24, 29]. Recent work [32] shows how poorly designed reputation systems can lead to unfair outcomes, even when agents differ only in some arbitrary group label. The findings in [32] highlight that such biased reputation systems can prevent both fairness and universal (i.e. not based on group) cooperation. Other work considers how unfair outcomes can result from cooperation incentives and the emergence of different tasks in populations of learning agents [18]. In our paper we also consider the interplay between fairness and cooperation, yet we focus on digital economy platforms and consider a recommendation and choice system more realistic than [32], in which agents have an equal likelihood of interacting with anyone.

2.3 Correcting biased ratings

Some suggestions for handling biased ratings are summarised by Tushkevich et al. [37]. Beyond implementing structured reputation systems which encourage objective evaluation of interactions [25], adjusting or reinterpreting suspected biased reviews has been suggested as a way to achieve fairness in digital platforms [15]. This was later explored by Goel et al. [14], who show how platforms may be able to control for bias by performing post-aggregation bias correction

which aims to minimally alter the aggregated ratings of all individuals such that the correlation between sensitive attribute and score is at most some value δ . This paper acknowledges, however, that while the proposed mechanisms work in theory, the implementation of such policies faces issues in determining the acceptable loss in information and how much bias in the ratings is permissible. Given that rating biases may be impossible to eliminate entirely, and the line between what is considered a “good” and “bad” rating being thin, our study explores how platforms can address rating bias even when it cannot be fully corrected.

2.4 Modelling fairness in digital platforms

Efforts have been made to develop a mechanistic understanding of the effects of bias in online platforms. Monachou and Ashlagi [20] examine a digitally-mediated labour market where employers are matched with workers and must choose whether to hire or reject the candidate. Employers leave ratings for workers they hire which cumulatively reduce the uncertainty other employers have about the type of the worker, which can be either high- or low-skilled.

The authors exogenously introduce bias through *discriminating employers*, who make up a certain proportion of the market and hold the misspecified prior belief that the likelihood that a minority worker without any reviews is high-skilled is only β times as much as majority workers with $\beta \in (0, 1)$. Thus, even if ratings accurately reflect performance, bias in interpreting a group’s ratings can affect the rate at which employers acquire accurate information about worker ability, ultimately affecting their utility.

A similar conclusion is reached by Che et al. [7], who discuss how *statistical discrimination*, the inference of an individual’s characteristics through their group’s characteristics, can arise due to a self-fulfilling perception that the same numerical rating is a less reliable reflection of the underlying quality of an agent in one group than in another. In their model, ratings gradually become inaccurate over time as agents stochastically change their “type” (either high- or low-quality) at rate $\delta > 0$. Due to this, whichever group is chosen more often will have more up-to-date ratings and thus be trusted and favoured by potential users.

In their analysis, the authors of both papers mostly hold the design choices of the platform as exogenous, apart from one section in [20] where employers who are identified as “discriminating” through a pattern of discriminatory rejection are not matched with minorities until uncertainty about their skill is small enough. Even in this suggestion, however, they concede that their notion of fairness is one-sided, taking workers, but not employers, into account. Ultimately, while these models provide a good understanding of why rating bias is so pervasive, suggestions of how to address this issue have not been studied taking into account all the stakeholders’ perspectives i.e., providers, users, the platform itself.

3 PROBLEM FORMULATION

We model interactions on digital platforms as a four step process, visualised in Figure 1: When a user wants to find a service provider, 1) the platform recommends them a list of possible choices, taking into account their rating and possibly their group. Then, 2) the user chooses a provider in a boundedly rational way, 3) they subsequently interact, and then 4) the user truthfully reports how the

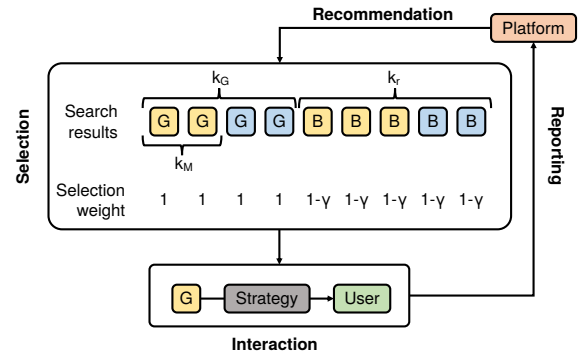


Figure 1: We model interactions on digital service economy platforms with four distinct steps. First, the platform recommends a list of k providers to a user, k_G of which are rated Good; out of these, k_M are part of a marginalised community. The user then chooses randomly from the list, giving those with a Good rating unit weight, and a Bad rating weight $1 - \gamma$. The chosen provider acts according to their strategy, either high or low-effort, rewarding the provider $b - c$ or b utility respectively. The user reports the action taken to the platform, who updates the rating of the provider. User reports are biased: with rate ϵ , they will report a marginalised user who played High as Bad even when they should be rated Good.

interaction went back to the platform. This process then repeats whenever a user looks for a provider.

Given that the number of users on the platform is typically multiple orders of magnitude larger than the number of providers,¹ we assume a finite population size Z of partners, Z_D of these being from some dominant social group and Z_M being marginalised, and assume for simplicity that the user population is infinitely large, as in Monachou and Ashlagi [20] and Che et al. [7]. This partial mean-field assumption, where one population is finite and the other infinite, allows us to neatly represent the user population using three parameters which express how users interact with different aspects of the interaction process previously described.

First, we use ϵ to measure the degree of **rating bias** experienced by marginalised users. We assume that all providers have an individual rating which is seen by users as either “good enough” or “not good enough” which we refer to as **Good (G)** and **Bad (B)** respectively. For simplicity, we assume that after every interaction, the involved user non-strategically reports the action taken by the provider to the platform. This action can be either high-effort (H) or low-effort (L), which, when reported to the platform, is translated into the corresponding rating G or B , overriding the provider’s current reputation and reflecting the “prosociality” of their most recent action. We introduce $\epsilon \in [0, 1]$ to be the likelihood that a provider of the marginalised group who plays H is falsely reported to have played L , which we refer to as the **rating bias**. We assume that, even if users do not know the group of their interaction partner before it takes place, they do know it by the time they write the

¹In 2020, when Airbnb most recently published their internal statistics, there were 14.1 million listings that had been booked in the last 365 days, with various sources reporting between 2.9 and 5 million hosts and 150 million users.

review given that they have communicated and sometimes met online or in person.

The plausibility of modelling ratings and actions with only two options is derived from platforms’ implementations of punishments for low ratings. The practice of removing partners with ratings below a certain threshold is a policy of both Airbnb² and Uber³, the latter even leading to a high-profile lawsuit alleging that this practice violated the Civil Right Act. Furthermore, because ratings are often so heavily skewed towards the positive end, it can be difficult to meaningfully distinguish between ratings beyond whether they are above or below some platform-specific threshold, something reflected in folk-advice written by/for users⁴. Because of this, a single bad rating can be enough to put a provider below the platform’s threshold for algorithmic preference and cause users to take caution before selecting the provider.

We measure how “hard” this rating threshold is, and in turn the amount of caution displayed by users, by $\gamma \in [0, 1]$, which we call the **rating sensitivity**. This parameter can alternatively be thought of as the trust users have in the rating system. When comparing between a list of providers, we assume that the platform obfuscates the demographic characteristics of the providers so users can only decide between providers using their rating. If $\gamma = 1$, then users will simply randomise amongst providers with rating G , but with $\gamma < 1$, users assign weight $1 - \gamma$ to providers with rating B and weight γ to providers with rate G , and randomise amongst *all* providers, selecting each with probability proportional to their weight. If $\gamma = 0$, users are indifferent between G or B providers.

Finally, we measure the number of providers that users will compare before making a selection with k , this is also known in business terms as **consumer involvement**. It has long been known that users selections are *heavily* skewed towards the first few results [17], and while this may be controlled to some extent by the platform (e.g. Uber’s fully automated matching), consumer involvement is more often explained by characteristics inherent to the market of which the platform is a part [22]. In the remainder of this paper, we refer to a user population by its rating bias, rating sensitivity, and consumer involvement written in a triple (ϵ, γ, k) .

3.1 Modelling the Recommender System

To encourage users to engage in cooperative, high-effort behaviour, platforms use recommender systems to prioritise partners with high ratings, making users more likely to interact with those deemed “high quality”. Rather than describing any particular algorithm or recommender system, we look at the output produced by whichever system is chosen by the platform. As providers who sign up to digital service platforms typically have to provide identification, we assume that platforms are able to infer the group of a provider.

Given some user population with consumer involvement k , we measure the extent of a platform’s “high rating prioritisation” by integer $k_G \leq k$. This value deterministically guarantees that at least k_G of the search results are rated G . Additionally, as some platform designers recognise that their ratings are biased and want to explicitly counteract this, we introduce $k_M \leq k_G$ which, again,

guarantees that at least k_M of the k_G providers will be from the marginalised group. To account for the fact that, depending on the number of providers playing each strategy, there may not always be k_G or k_M providers that fit the requirements to be prioritised, we let Z_G, Z_{GM} , and Z_{GD} be the number of providers, marginalised providers, and dominant providers rated G respectively and define $\hat{k}_G := \min(k_G, Z_G)$ and $\hat{k}_M := \min(k_M, Z_{GM})$.

Assuming that the remaining $k_R := k - \hat{k}_G$ providers are sampled from all providers *after* the \hat{k}_M , then $\hat{k}_G - \hat{k}_M$ providers are sampled from their respective sub-populations, we can calculate the probability that a provider in group $g \in \{M, D\}$ with rating $r \in \{B, G\}$ is included in any particular list of search results and subsequently chosen given that they were included. Because the likelihood that one is chosen depends on who else is included, and this in turn depends on at which stage (\hat{k}_M, \hat{k}_G , or k_R) one is chosen, we relegate the calculation of this probability to appendix Section 1 and subsequently refer to its value as $p(g, r | Z_{GD}, Z_{GM})$.

3.2 Utility calculation

In our model, the expected utility of a provider is their payoff from a single interaction multiplied by the likelihood that the provider is chosen for any given interaction. A high-effort (H) provider engages in behaviours such as ensuring cleanliness of a property and communicating in a timely manner with users, which we assume has per-interaction a utility cost $c > 0$. On the other hand, we assume that a low-effort (L) provider does not have to pay this cost, and that both types of provider have the same per-interaction revenue b with $b > c > 0$. If c and b are close, then margins in the market are tight and there is a large incentive to play low-effort (L).

As the likelihood that a partner is chosen depends on the strategies played by the rest of the population, our model’s dynamics consider every agent’s individual incentive to switch between their current strategy and the other, which we detail in Section 3.3. We write the state space of the dynamics as $\mathbf{h} = (h_M, h_D)$, where h_i represents the number of agents playing H in group i (so $0 \leq h_i \leq Z_i$), see how \mathbf{h} varies over time subject to stochastic dynamics, and then weigh each provider’s utility in each state by the time the dynamics spend in that particular state.

At any point in time, all of the h_M marginalised H -playing providers have a $1 - \epsilon$ probability of correctly being assigned rating G from their last interaction, and so $Z_{GM} \sim \text{Binomial}(h_M, 1 - \epsilon)$ with p.m.f $f_{Z_{GM}}$. Therefore, to calculate the likelihood of a provider in group g playing s being chosen, we sum over their rating r and over Z_{GM} given g and r (as when $(g, r) = (M, G)$, one partner is already accounted for), calculating $p(g, r | Z_{GD}, Z_{GM})$ for each case as outlined in appendix Section 1. $p(g, r | Z_{GD}, Z_{GM})$ represents the likelihood that a partner from group g with rating r is shown in the search results given the ratings of others, capturing competition among providers. As such, the utility u of a provider in group g , playing strategy $s \in \{L, H\}$ given the (competing) provider population is \mathbf{h} , is

$$u(g, s | \mathbf{h}) = \pi(s) \sum_{r \in \{B, G\}} p(r | (g, s)) \sum_{z=0}^{h_M} f_{Z_{GM}}(z) p(g, r | Z_{GD}, z), \quad (1)$$

²<https://www.airbnb.com/help/article/2895>

³<https://www.uber.com/us/en/drive/driver-app/deactivation-review/>

⁴<https://www.kenrockwell.com/tech/ebay/index.htm>

where $\pi(s) = (b - c \mathbb{I}_{\{s=H\}})$ and \mathbb{I} is the indicator function, and

$$p(G | (g, s)) = \begin{cases} 0, & \text{if } s = L \\ 1 - \epsilon, & \text{if } s = H \text{ and } g = M \\ 1, & \text{if } s = H \text{ and } g = D. \end{cases} \quad (2)$$

3.3 Population Dynamics

We model variations in $\mathbf{h}_t = (h_D, h_M)_t$ over time using a frequency-dependent Moran process [21] as in other models studying the emergence of cooperation [23, 35]. A Moran process is a stochastic evolutionary model used in genetics where the size of each “allele” (in our case strategy) remains constant, and where alleles with a fitness advantage tend to become more prevalent in the population.

We apply this model as such: each timestep t , we sample a random focal agent i and flip their strategy s_i (either from H to L or vice-versa) with probability μ , and otherwise compare their utility $u(g_i, s_i)$ to another randomly sampled agent from their group j using the **Fermi pairwise imitation rule** [35]: Let $\Delta u_{ij}^g(\mathbf{h}) := u(g, j | \mathbf{h}) - u(g, i | \mathbf{h})$ using the definition of $u(g, s | \mathbf{h})$ in (1) and let $\beta \in \mathbb{R}_+$ be the strength of selection, the probability that the number of agents playing H increases (f_+^g) or decreases (f_-^g) is given by

$$f_{\pm}^g(\mathbf{h}_t) = \frac{1}{1 + e^{\mp \beta \Delta u_{HL}^g(\mathbf{h}_t)}}. \quad (3)$$

This process of mutation and replacement, allows us to define a discrete-time Markov chain \mathbf{h}_t on the state space of the number of agents playing H in each group

$$\mathcal{S} = \{(h_M, h_D) : 0 \leq h_M \leq Z_M, 0 \leq h_D \leq Z_D\},$$

(ordered lexicographically), initial state \mathbf{h}_0 , mutation probability μ , and transition matrix $P^{\mathcal{S} \times \mathcal{S}}$ with entries

$$P_{\mathbf{h}_t, \mathbf{h}_{t+1}} = \begin{cases} \mu \frac{Z_D - h_D}{Z} + (1 - \mu) \frac{Z_D - h_D}{Z_D} \frac{h_D}{Z_D - 1} f_+^D(\mathbf{h}_t), & \text{if } \mathbf{h}_{t+1} = (h_D + 1, h_M) \\ \mu \frac{h_D}{Z} + (1 - \mu) \frac{h_D}{Z_D} \frac{Z_D - h_D}{Z_D - 1} f_-^D(\mathbf{h}_t), & \text{if } \mathbf{h}_{t+1} = (h_D - 1, h_M) \\ \mu \frac{Z_M - h_M}{Z} + (1 - \mu) \frac{Z_M - h_M}{Z_M} \frac{h_M}{Z_M - 1} f_+^M(\mathbf{h}_t), & \text{if } \mathbf{h}_{t+1} = (h_D, h_M + 1) \\ \mu \frac{h_M}{Z} + (1 - \mu) \frac{h_M}{Z_M} \frac{Z_M - h_M}{Z_M - 1} f_-^M(\mathbf{h}_t), & \text{if } \mathbf{h}_{t+1} = (h_D, h_M - 1) \\ 1 - \sum_{\mathbf{x} \neq \mathbf{h}_t} P_{\mathbf{h}_t, \mathbf{x}}, & \text{if } \mathbf{h}_{t+1} = \mathbf{h}_t \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

The Markov chain \mathbf{h}_t is irreducible if $\mu > 0$ as mutations draw the chain out of absorbing states, and aperiodic as the chain has a strictly non-zero probability of staying in the same place for all $\beta \neq \infty$. Therefore, we can find its stationary distribution \mathbf{h}^* defined as the limit of the recurrence relation $\mathbf{h}_{t+1} = \mathbf{h}_t P$ by solving the linear system $(I - P)^T \mathbf{h}^{*T} = \mathbf{0}$ subject to the constraint $\sum_i \mathbf{h}_i^* = 1$.

4 EXPERIMENTAL SETUP

In our experiments we aim to explore how different user populations (ϵ, γ, k) and recommender system parameters (k_G, k_M) affect 1) the incentive for each group to cooperate (play H), 2) the value users get out of the platform, and 3) the average utility of the two

groups of providers. In Section 4.1 we give a precise definition to each of these metrics, and below we detail the parameters common to every experiment we run.

To minimise differences in utility that are incidental to (relative) group size we set the size of both groups to be equal: $Z_D = Z_M = 20$ and keep this fixed from here on, referring readers to appendix Section 2 for results with $Z = 20$ and $Z = 80$. While this is certainly scaled back when compared to some real online platforms, it is large enough that k_G and k_M can take many values between 0 and $k \leq Z$, but also small enough such that parameter grid searches are computationally feasible (recall every simulation requires inverting a $(Z_D Z_M) \times (Z_D Z_M)$ matrix). Results for other population sizes, larger and smaller, are qualitatively the same, *mutatis mutandis*. As μ and β are both unitless parameters of evolution, we can scale one in terms of the other to achieve the same result. As such, we arbitrarily set the mutation rate $\mu = 1/Z = 1/40$, such that, on average, once in every Z strategic update steps the updating provider will randomise their strategy rather than imitate another provider.

To both emphasise the effect of platform design choices and simulate a case facing many service providers where profit margins are relatively tight, we set $c = 1$ and $b = 1.2$, raising b or lowering c simply makes cooperation easier to sufficiently incentivise. Given that our utilities are calculated as the likelihood of an interaction multiplied by the payoff from an interaction, if selections were done completely at random then each agent would have an average utility proportional to $1/40$. The selection strength β is set out of modelling convenience such that $\beta \times \Delta u$ has an order of magnitude close to 1, attempting keeping the output of the Fermi function away from its region of near-zero gradient, leading to our choice of $\beta = Z/(b - c) = 20$.

We initially hold $k_M = 0$, which represents the status-quo where, for implementation or political reasons, platforms do not explicitly prioritise marginalised providers. After exploring the dynamics of the model in Section 5.1 and introducing the “user-provider trade-off” subject to this restriction in Section 5.2, we allow platform designers to alter k_M while keeping k_G fixed, simulating a situation where platform designers want to minimally alter the user experience in Section 5.3. Finally, we remove all restrictions on k_G and k_M and introduce uncertainty on the precise value of ϵ , seeing how increasing uncertainty affects demographic parity in Section 5.4.

4.1 Evaluation Metrics

Mostly cooperative. A single group $g \in \{M, D\}$ is **mostly cooperative** if the stationary distribution of the strategy dynamics \mathbf{h}^* spends most of its time on the edge $h_g = Z_g$ i.e. the sum of \mathbf{h}^* over all states with $h_g = Z_g$ is greater than 0.5. A platform is mostly cooperative if both of its groups are.

User Experience. Given the stationary distribution \mathbf{h}^* , the **user experience (UX)** on the platform is defined as the likelihood that the service provider plays H . Let S be the state matrix whose entries are the number of agents playing H in each group

$$S = (S_1, S_2) := \begin{bmatrix} (0, 0) & (0, 1) & \cdots & (0, Z_D) \\ (1, 0) & (1, 1) & \cdots & (1, Z_D) \\ \vdots & \vdots & \ddots & \vdots \\ (Z_M, 0) & (Z_M, 1) & \cdots & (Z_M, Z_D) \end{bmatrix}$$

then define $\hat{S}_M := \frac{S_1}{Z_M}$, and $\hat{S}_D := \frac{S_2}{Z_D}$. Then, we can calculate the expected strategy in each group (σ^*) by taking the element-wise product (\odot) of \hat{S} with \mathbf{h}^* , before summing over the rows and columns, which is equivalent to taking the matrix 1-norm as S is non-negative:

$$\sigma^* = \|\hat{S} \odot \mathbf{h}^*\|_1 \quad (5)$$

Finally, we can calculate

$$\mathbf{UX} = \frac{\sigma_M^* Z_M + \sigma_D^* Z_D}{Z_M + Z_D} \quad (6)$$

Demographic Parity Ratio. As we discuss in Section 2, we measure fairness, for simplicity, through the **demographic parity ratio (DPR)**, which captures the ratio between the average utility of the better-off and worse-off group of service providers.

Define matrix $(U^g)_{\substack{0 \leq i \leq Z_M \\ 0 \leq j \leq Z_D}}$ with entries

$$U_{i,j}^g := u(g, L \mid \mathbf{h} = (i, j)) \frac{Z_g - h_g}{Z_g} + u(g, H \mid \mathbf{h} = (i, j)) \frac{h_g}{Z_g}$$

to be the average utility in group g at each state where h_g is the number of H -players in group g , taking either value i or j depending on which group is being considered. Then we can calculate

$$\mathbf{DPR} := \frac{\min(\|\mathbf{h}_M^* \odot U^M\|_1, \|\mathbf{h}_D^* \odot U^D\|_1)}{\max(\|\mathbf{h}_M^* \odot U^M\|_1, \|\mathbf{h}_D^* \odot U^D\|_1)} \quad (7)$$

Pareto Optimality. Given user population (ϵ, γ, k) , a platform designer’s choice of (k_G, k_M) , potentially subject to the constraint $k_M = 0$, is **Pareto optimal** if a different choice could not simultaneously increase the platform’s **UX** and **DPR**. The set of choices of (k_G, k_M) that are Pareto optimal is called the **Pareto front**.

5 RESULTS

5.1 Achieving a cooperative baseline

At a fundamental level, every platform must maintain an incentive for providers to be high-effort in order to ensure users have a good experience, given user parameters (ϵ, γ, k) . In Figure 2, we demonstrate the varying strategy dynamics that respective populations of providers have in three scenarios with varying user parameters and k_G , and who all have $k_M = 0$.

The left-most scenario (**A**) effectively has no reputation system ($k_G = \gamma = 0$), meaning neither the search algorithm nor the users use any information about the quality of a potential provider in their decision making. Unsurprisingly, without an incentive to be high-effort, no providers choose to do so. This outcome is characteristic of any platform that, given a certain user population, fails to sufficiently incentivise providers to cooperate.

The centre and right-most platform with $k = 10$, **B** and **C**, show how user populations with non-zero rating bias (here $\epsilon = 0.3$) and/or low sensitivity to ratings (here $\gamma = 0.6$) can lead to either a separating or pooling equilibrium depending on the platform’s choice of k_G . In the former case ($k_G = 5$), although the dominant group is “mostly cooperative” as previously defined in Section 4.1, the incentives are not sufficient for marginalised users to play H (high-effort) which we speculate is due to the difficulty maintaining a high enough rating to be prioritised by the recommender system and user population, and the insufficient prioritisation when their rating is high. This leads to the divergence of strategy between

dominant and marginalised providers that is not found in platform **C** ($k_G = 10 = k$), a platform that *does* feature sufficient incentives for both groups to be mostly cooperative.

5.2 The $(k_G, k_M = 0)$ Pareto front

In Figure 3, we show the **UX** and **DPR** of a platform with population parameters $k = 20$, $\epsilon = 0.15$, and $\gamma = 0.8$ and $k_M = 0$ subject to different choices of k_G . The dynamics of the three platforms **A**, **B**, and **C** from Figure 2 represent three qualitatively different regimes of cooperation (or lack thereof) which can be induced through choice of k_G and k_M . The size of these regimes in (k_G, k_M) -space depends on the population parameters. As in this case $k_M = 0$, we plot the k_G -change-points of the regimes as black vertical lines and label the regions correspondingly, also marking the maximum values attained by each metric in regime **C**.

Values of k_G between these maximum values make up Pareto front. Define $k_G^{\mathbf{UX}}$ and $k_G^{\mathbf{DPR}}$ to be the values of k_G maximise **UX** and **DPR** respectively while inducing regime **C** dynamics. When $k_M = 0$ we have $k_G^{\mathbf{UX}} \equiv k$ as **UX** is always monotonically increasing in k_G : the more highly-rated providers that users are shown, the more likely they are to choose a provider playing H .⁵ However, for $k_G > k_G^{\mathbf{DPR}}$, the increased pressure to cooperate is sufficiently small that the emergent effect is just showing marginalised providers less often due to the rating bias. This subsequently decreases the share of interactions they are involved in and lowers the **DPR**.

The decreasing nature of **DPR** for $k_G > k_G^{\mathbf{DPR}}$, combined with **UX** being monotonically increasing in the same region, implies that when $k_M = 0$, the Pareto front is the interval $[k_G^{\mathbf{DPR}}, k]$. In Figure 4, we show the size of the Pareto front when $k = 20$ (as in Figure 3), varying ϵ and γ in the open set $(0, 1)$. The reason we exclude the edges is that one of the metrics becomes completely flat (e.g. $\epsilon = 0 \implies \mathbf{DPR} \equiv 1$), but floating point error causes distracting fluctuations. We note that $k_G^{\mathbf{DPR}}$ monotonically increases with ϵ and γ , meaning biased and unreliable reputations require that platforms put more effort into incentivising cooperation, particularly in the marginalised group, to achieve maximum demographic parity.

This analysis shows that if platforms are unwilling or unable to introduce active countermeasures against inherent rating bias (ϵ) at the algorithmic level (i.e. $k_M = 0$), then they have to choose between a better user experience and a more equitable provider experience, where, even given optimal choices, one necessarily comes at the cost of the other.

5.3 Explicit anti-discrimination with $k_M \geq 0$

We first assume that the platform has a fixed k_G , which could be for many reasons including wanting to maintain the overall “feel” of the platform to users by showing the same variety of ratings, and can choose any $k_M \geq 0$. In Figure 5, we plot four different platforms. For each we arbitrarily choose k_G that maximises $\mathbf{UX} \times \mathbf{DPR}$, a value somewhere in the Pareto front, subject to $k_M = 0$, emulating the situation where a platform that cares about **DPR** wants to introduce an explicitly biased algorithm while making minimal changes to the platform from a user point of view.

⁵This would still be true if there were a chance for L -players to be G -rated as long as we are in regime **C**. In this case, no matter the false positive rate, because more agents are cooperating than defecting, showing more G -rated agents still increases the likelihood that a user selects an H -player

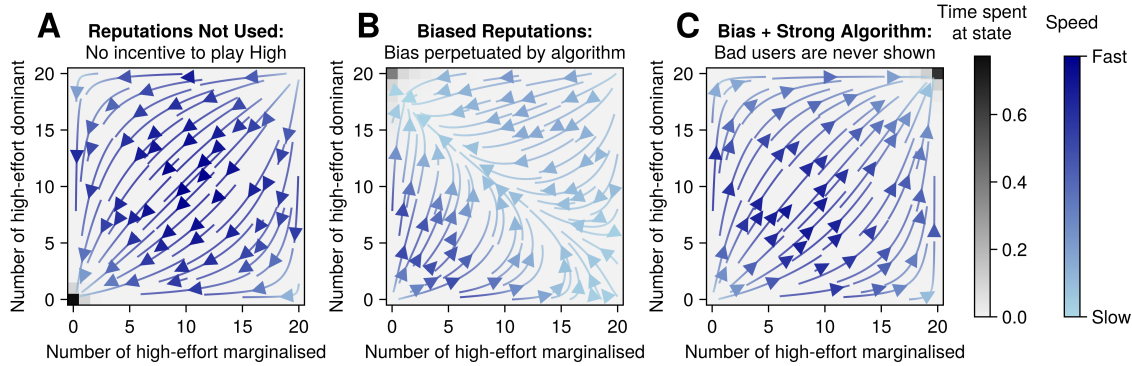


Figure 2: We demonstrate how changing platform design affects strategy dynamics for providers. The proportion of time the dynamics spend at each state under the stationary distribution is plotted as a greyscale heatmap, and the average direction of the dynamics are overlaid as a stream plot. All subplots have $k = 10$ and $\epsilon = 0.3$. For subplot A, we set $k_G = \gamma = 0$, which means that ϵ could take any value without affecting the dynamics. Subplot B and C have $k_G = 5$ and $k_G = 10 = k$ respectively. We observe that the platform’s design choices impact the providers’ strategic dynamics, leading them to adopt low-effort strategies (A), exacerbating biases by eliciting high-effort only from the dominant group (B) or inspiring high-effort by all groups (C).




Figure 3: In this Figure we fix the user population to have rating bias $\epsilon = 0.15$ and rating sensitivity $\gamma = 0.8$, and vary k_G while holding $k = 20$. We can partition the x-axis into three distinct regions defined by whether no groups, only the dominant group, or both groups are “mostly cooperative”. We indicate the maximum value attained by each of the three evaluation metrics defined in Section 4.1 using a dashed line.

We find that the higher the rating bias, the higher k_M is required to be to correct for it: when $\gamma = 0.8$, raising ϵ from 0.15 to 0.5 shifts the optimal value of k_M from 3 to 9. While changing k_M has a large effect on fairness, it has almost no effect on users, evidenced by the relatively flatness of the UX line compared to DPR. This is because raising k_M does not change the number of high-effort agents shown to users, simply the demographics of the agents shown. Only when ϵ is low and k_M is very high do we see a fall in UX. In this case the incentive for the dominant group to play H is severely impacted due to them almost exclusively being shown only as one of the providers in the random sampling stage which makes up only a small fraction k_R/k of all providers shown. Because of this, they are shown almost as often when playing L as when playing H .




Figure 4: For $k = 20$ (as in Figure 3), we vary ϵ and γ , calculating the value of k_G that maximises the demographic parity ratio (DPR) while leaving the dynamics in regime C.

While these results are promising, as any variable not controlled by the platform must be inferred from data, the designers may not have accurate values of k , ϵ , and γ . The sharp peak of DPR in Figure 5 with respect to k_M raises an issue for platform designers: by altering k_G and k_M you risk over- or under-correcting for bias. As inaction is less harshly judged than poor action, it is important to see that a $k_M > 0$ policy is still effective under uncertainty.

5.4 Optimising with uncertainty in rating bias

Suppose the true, unobserved amount of rating bias exhibited by a population of users was $\epsilon = 0.35$, and let $\gamma = 0.8$ and $k = 7$ be observed. For simplicity, we assume that the unobserved value of ϵ can be estimated to be within some interval $[\epsilon_{\min}, \epsilon_{\max}]$ with

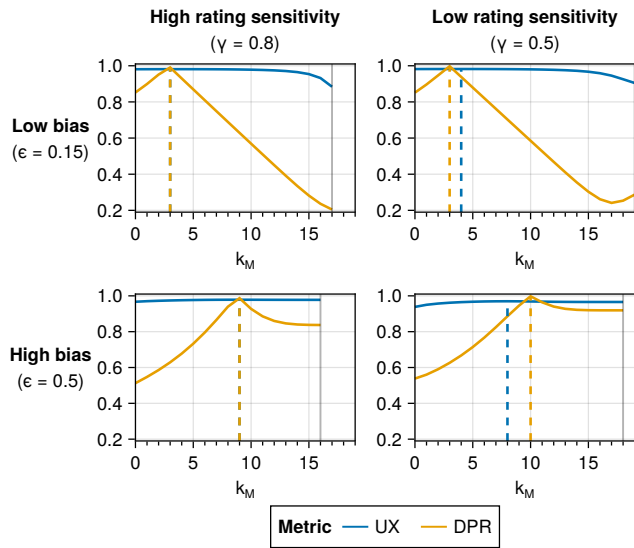


Figure 5: Consider a platform with $k = 20$ and vary ϵ and γ between low and high values. For each resulting scenario, find the value of k_G that maximises $UX \times DPR$ subject to $k_M = 0$, we plot this as a faint grey vertical line. As we now vary k_M , the dynamics stay inside regime C and the evaluation metrics peak at the same value of k_M when γ is high enough (0.8).




Figure 6: Let $k = 20$, $\epsilon = 0.35$, and $\gamma = 0.65$, and vary the uncertainty of rating bias. As previously established: forcing $k_M = 0$ leads to a low DPR. Even when uncertainty is 0.7, allowing $k_M \geq 0$ still results in a better worst-case outcome. Optimisation of max-min DPR and average DPR are both fine choices of metric to optimise over that only diverge significantly in outcome for very large uncertainty.

uniform likelihood over the interval. We refer to the width of this interval ($\epsilon_{\max} - \epsilon_{\min}$) as the **uncertainty** in ϵ .

Given some level of uncertainty, which values of k_G and k_M should the platform choose? One might expect this to depend considerably on what the platform tries to optimise for. Of course, the choice of values must leave the dynamics in regime C, but should you try to optimise the expected DPR or conservatively try to maximise its minimum value? Surprisingly, these two straightforward

metrics tend to yield very similar outcomes for low to moderate values of uncertainty, diverging only when uncertainty is very large. This can be seen in Figure 6, where until the uncertainty is 0.525, representing a range of $[0.0875, 0.6125]$, both the average and worst-case outcomes of applying either a maximin strategy or maximising the expected value are almost identical. After this point, the two diverge but importantly both still maintain better performance than a platform which keeps $k_M = 0$. In other words, even when uncertainty is very large, optimising for the worst-case outcome over $k_M \geq 0$ achieves better worst- and even *average*-case performance than leaving $k_M = 0$, highlighting how risk-free it is to allow k_M to vary.

6 CONCLUSION

Online marketplaces have become ubiquitous in modern society. In this paper we develop a model that shows how algorithmic design decisions, applied in such online platforms, affect utility and fairness of users and providers. Our findings reveal that recommender systems can play a decisive role in undoing the effects of rating bias, a phenomenon that is pervasive across platforms. However, if considerations of fairness take a back seat to the experience of the user population, then we show how seemingly fairness-neutral decisions can counterintuitively lead to less fair outcomes. The low barrier of entry to employment provided by these platforms means they are relied upon by many economically and legally vulnerable members of society who themselves are often part of the groups that are discriminated against. As such, designing these platforms with the goal of fair treatment of marginalised communities despite the inherent biases of users is a significant step towards fairness in the labour market more broadly.

By keeping the recommender system, user population, and the interaction itself abstract, we hope that this model might inform the application of complex adaptive systems, particularly evolutionary game theory, to more elaborate, tailored models in (digital) labour economics and industrial organisation: while we assume a the common assumption of a constant number of providers, in real markets participants are constantly arriving, and existing agents may leave during downturns in their utility which is modelled in [20]. Furthermore, our mean-field assumption on users and use of indirect as opposed to direct reciprocity may not be valid for smaller markets where consumers and producers are more equally numbered. We ignore providers’ forms of collective action, which might also affect fairness dynamics [27, 30], a topic of natural interest to be explored in future extensions of our model. Finally, we assume that providers’ demographic information is fully obfuscated, but this is obviously idealised and not typically the case [19].

Despite these limitations and suggestions for future work, our work already stresses the positive outcomes of anti-discrimination policies in digital economy platforms, and we test concrete interventions to improve fairness in systems where ratings are used to promote cooperative behaviour.

ACKNOWLEDGMENTS

FPS acknowledges funding through ERC grant (RE-LINK, <https://doi.org/10.3030/101116987>).

REFERENCES

- [1] 2023. Employment statistics - digital platform workers. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Employment_statistics_-_digital_platform_workers
- [2] 2024. Directive (EU) 2024/2831 of the European Parliament and of the Council of 23 October 2024 on improving working conditions in platform work (Text with EEA relevance). <http://data.europa.eu/eli/dir/2024/2831/oj/eng>
- [3] Bruno Abrahao, Paolo Parigi, Alok Gupta, and Karen S. Cook. 2017. Reputation offsets trust judgments based on social biases among Airbnb users. *Proceedings of the National Academy of Sciences* 114, 37 (Sept. 2017), 9848–9853. <https://doi.org/10.1073/pnas.1604234114>
- [4] Ian Ayres, Mahzarin Banaji, and Christine Jolls. 2015. Race effects on eBay. *The RAND Journal of Economics* 46, 4 (2015), 891–917. <https://www.jstor.org/stable/43895621>
- [5] Arpita Biswas, Gourab K. Patro, Niloy Ganguly, Krishna P. Gummadi, and Abhijnan Chakraborty. 2022. Toward Fair Recommendation in Two-sided Platforms. *ACM Transactions on the Web* 16, 2 (May 2022), 1–34. <https://doi.org/10.1145/3503624>
- [6] Abhijnan Chakraborty, Aniko Hannak, Asia J. Biega, and Krishna P. Gummadi. 2017. Fair Sharing for Sharing Economy Platforms. <https://doi.org/10.18122/B2BX2S> Institution: Boise State University.
- [7] Yeon-Koo Che, Kyungmin Kim, and Weijie Zhong. 2024. Statistical Discrimination in Ratings-Guided Markets. <https://doi.org/10.48550/arXiv.2004.11531> arXiv:2004.11531 [cs].
- [8] Julie Yujie Chen and Jack Linchuan Qiu. 2019. Digital utility: Datafication, regulation, labor, and DiDi's platformization of urban transport in China. *Chinese Journal of Communication* 12, 3 (July 2019), 274–289. <https://doi.org/10.1080/17544750.2019.1614964> eprint: <https://doi.org/10.1080/17544750.2019.1614964>
- [9] Yashar Deldjoo, Dietmar Jannach, Alejandro Bellogin, Alessandro Difonzo, and Dario Zanzonelli. 2024. Fairness in recommender systems: research landscape and future directions. *User Modeling and User-Adapted Interaction* 34, 1 (March 2024), 59–108. <https://doi.org/10.1007/s11257-023-09364-z>
- [10] Kristin Dell and Nicole Nestorik. 2020. *Assessing the Impact of New Technologies on the Labor Market: Key Constructs, Gaps, and Data Collection Strategies for the Bureau of Labor Statistics*. Congressional Report GS-00F-0078M. U.S. Department of Labor Bureau of Labor Statistics. <https://www.bls.gov/bls/congressional-reports/assessing-the-impact-of-new-technologies-on-the-labor-market.htm>
- [11] Benjamin Edelman, Michael Luca, and Dan Svirsky. 2017. Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment. *American Economic Journal: Applied Economics* 9, 2 (April 2017), 1–22. <https://doi.org/10.1257/app.20160213>
- [12] Yanbo Ge, Christopher R. Knittel, Don MacKenzie, and Stephen Zoepf. 2020. Racial discrimination in transportation network companies. *Journal of Public Economics* 190 (Oct. 2020), 104205. <https://doi.org/10.1016/j.jpubeco.2020.104205>
- [13] Sahin Cem Geyik, Stuart Ambler, and Krishnamurthy Venkatesh. 2019. Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, Anchorage AK USA, 2221–2231. <https://doi.org/10.1145/3292500.3330691>
- [14] Naman Goel, Maxime Rutagarama, and Boi Faltings. 2020. Tackling Peer-to-Peer Discrimination in the Sharing Economy. In *Proceedings of the 12th ACM Conference on Web Science (WebSci '20)*. Association for Computing Machinery, New York, NY, USA, 355–361. <https://doi.org/10.1145/3394231.3397926>
- [15] 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 *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. Association for Computing Machinery, New York, NY, USA, 1914–1933. <https://doi.org/10.1145/2998181.2998327>
- [16] Bastian Jaeger and Willem W. A. Sleegers. 2023. Racial Disparities in the Sharing Economy: Evidence from More Than 100,000 Airbnb Hosts across 14 Countries. *Journal of the Association for Consumer Research* 8, 1 (Jan. 2023), 33–46. <https://doi.org/10.1086/722700>
- [17] Mark T. Keane, Maeve O'Brien, and Barry Smyth. 2008. Are people biased in their use of search engines? *Commun. ACM* 51, 2 (Feb. 2008), 49–52. <https://doi.org/10.1145/1314215.1314224>
- [18] Woojun Kim and Katia P Sycara. [n.d.]. Fair Cooperation in Mixed-Motive Games via Conflict-Aware Gradient Adjustment. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- [19] Tamar Kricheli-Katz and Tali Regev. 2016. How many cents on the dollar? Women and men in product markets. *Science Advances* 2, 2 (Feb. 2016), e1500599. <https://doi.org/10.1126/sciadv.1500599>
- [20] Faidra Georgia Monachou and Itai Ashlagi. 2019. Discrimination in Online Markets: Effects of Social Bias on Learning from Reviews and Policy Design. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2019/hash/e00406144c1e7e35240afed70f34166a-Abstract.html
- [21] P. a. P. Moran. 1958. Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society* 54, 1 (Jan. 1958), 60–71. <https://doi.org/10.1017/S0305004100033193>
- [22] Andrea Niosi. 2021. Involvement Levels. (June 2021). <https://opentextbc.ca/introconsumerbehaviour/chapter/involvement-levels/> Book Title: Introduction to Consumer Behaviour.
- [23] Martin A. Nowak, Akira Sasaki, Christine Taylor, and Drew Fudenberg. 2004. Emergence of cooperation and evolutionary stability in finite populations. *Nature* 428, 6983 (April 2004), 646–650. <https://doi.org/10.1038/nature02414> Number: 6983.
- [24] Hisashi Ohtsuki and Yoh Iwasa. 2004. How should we define goodness?—reputation dynamics in indirect reciprocity. *Journal of Theoretical Biology* 231, 1 (Nov. 2004), 107–120. <https://doi.org/10.1016/j.jtbi.2004.06.005>
- [25] Amanda Pallais. 2014. Inefficient Hiring in Entry-Level Labor Markets. *The American Economic Review* 104, 11 (2014), 3565–3599. <https://www.jstor.org/stable/43495347>
- [26] Gourab K. Patro, Lorenzo Porcaro, Laura Mitchell, Qiuyue Zhang, Meike Zehlke, and Nikhil Garg. 2022. Fair ranking: a critical review, challenges, and future directions. In *2022 ACM Conference on Fairness, Accountability and Transparency*. ACM, Seoul Republic of Korea, 1929–1942. <https://doi.org/10.1145/3531146.3533238>
- [27] Gustavo R Pilatti, Flavio L Pinheiro, and Alessandra A Montini. 2024. Systematic literature review on gig economy: Power dynamics, worker autonomy, and the role of social networks. *Administrative Sciences* 14, 10 (2024), 267.
- [28] Noopur Atul Raval. 2020. *Platform-Living: Theorizing life, work, and ethical living after the gig economy*. Ph.D. Dissertation. UC Irvine. <https://escholarship.org/uc/item/4qc0x3mw>
- [29] Fernando P. Santos, Jorge M. Pacheco, and Francisco C. Santos. 2021. The complexity of human cooperation under indirect reciprocity. *Philosophical Transactions of the Royal Society B: Biological Sciences* 376, 1838 (Nov. 2021), 20200291. <https://doi.org/10.1098/rstb.2020.0291>
- [30] Dorothee Sigg, Moritz Hardt, and Celestine Mandler-Dünner. 2025. Decline now: A combinatorial model for algorithmic collective action. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [31] Jacobus Smit and Fernando P. Santos. 2026. Appendix and Code for Fairness Dynamics in Digital Economy Platforms with Biased Ratings. <https://doi.org/10.5281/zenodo.18600282> Language: eng.
- [32] Martin Smit and Fernando P. Santos. 2024. Learning Fair Cooperation in Mixed-Motive Games with Indirect Reciprocity. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Jeju, South Korea, 220–228. <https://doi.org/10.24963/ijcai.2024/25>
- [33] Tom Sühr, Sophie Hilgard, and Himabindu Lakkaraju. 2021. Does Fair Ranking Improve Minority Outcomes? Understanding the Interplay of Human and Algorithmic Biases in Online Hiring. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*. Association for Computing Machinery, New York, NY, USA, 989–999. <https://doi.org/10.1145/3461702.3462602>
- [34] Ambika Tandon and Aayush Rathi. 2024. Sustaining urban labour markets: Situating migration and domestic work in India's 'gig' economy. *Environment and Planning A: Economy and Space* 56, 4 (June 2024), 1245–1261. <https://doi.org/10.1177/0308518X221120822>
- [35] Arne Traulsen, Martin A. Nowak, and Jorge M. Pacheco. 2006. Stochastic dynamics of invasion and fixation. *Physical Review E* 74, 1 (July 2006), 011909. <https://doi.org/10.1103/PhysRevE.74.011909>
- [36] Timm Tuebner, Florian Hawlitschek, and David Dann. 2017. Price Determinants on Airbnb: How Reputation Pays Off in the Sharing Economy. *Journal of Self-Governance and Management Economics* 5, 4 (2017), 53–80. <https://www.ceol.com/search/article-detail?id=596092>
- [37] Miroslav Tushev, Fahimeh Ebrahimi, and Anas Mahmoud. 2022. A Systematic Literature Review of Anti-Discrimination Design Strategies in the Digital Sharing Economy. *IEEE Transactions on Software Engineering* 48, 12 (Dec. 2022), 5148–5157. <https://doi.org/10.1109/TSE.2021.3139961> Conference Name: IEEE Transactions on Software Engineering.
- [38] Niels van Doorn, Fabian Ferrari, and Mark Graham. 2023. Migration and Migrant Labour in the Gig Economy: An Intervention. *Work, Employment and Society* 37, 4 (Aug. 2023), 1099–1111. <https://doi.org/10.1177/09500170221096581>