

Statistical Discrimination in Ratings-Guided Markets

Yeon-Koo Che*

Kyungmin Kim[†]

Weijie Zhong[‡]

December 2019

Preliminary and Incomplete

Abstract

We study statistical discrimination of individuals based on payoff-irrelevant social identities in markets where ratings/recommendations facilitate social learning among users. Despite the potential promise and guarantee for the ratings/recommendation algorithms to be fair and free of human bias and prejudice, we identify possible vulnerability of the ratings-based social learning to discriminatory inferences on social groups. In our model, users' equilibrium attention decision may lead data to be sampled differentially across different groups so that differential inferences on individuals may emerge based on their group identities. We explore policy implications in terms of regulating trading relationships as well as algorithm design [to be added].

JEL CLASSIFICATION NUMBERS:

KEYWORDS: ratings-based social learning; statistical discrimination; directed search; algorithmic fairness.

1 Introduction

Discrimination of individuals based on their race, gender, ethnicity and other social identities, is a pervasive problem. While the problem is as old as humanity, it has taken on a new meaning and form at digital marketplaces and social media where our social and economic interactions increasingly take place. Evidence suggests that discrimination is prevalent in popular online platforms such as airbnb ([Edelman et al.](#)

*Che: Columbia University, yc2271@columbia.edu.

[†]Kim: Emory University, kyungmin.kim@emory.edu.

[‡]Zhong: Stanford University, weijie.zhong@stanford.edu.

(2017) and Cui et al. (2019)), freelancing worksites (Hannák et al. (2017)), ride-sharing platforms (Ge et al. (2016)), and math Stack exchange (Bohren et al. (2019)).

At first glance, online platforms are an unlikely place for discrimination to occur, given the widely-used ratings system that facilitates social learning among users. Platforms for ride-sharing, house-sharing, freelancing, credit, and insurance collect information about drivers, customers, guests, workers, and loan and insurance applicants based on performers' past records and user experiences. Platforms then aggregate the information into simple *ratings*, and make *recommendations* based on these ratings. It is now routine that machine algorithms pre-screen résumés of job applicants, evaluate loan or insurance applicants and freelance workers, recommend their promotion or firing, and rates recidivism of parolees.

By limiting subjective human judgment and replacing it with accurate information and objective recommendation, one would think that data-driven social learning should limit the scope for statistical discrimination. Intuitively, there should be simply less room for statistical inference on an individual based on his/her group identity—and therefore discrimination based on it—if one is guided by more accurate information about his/her *individual* characteristics. The logic of this reasoning is at first glance plausible, and indeed, if full information were available, discrimination, except based on tastes, should disappear.¹ However, it is not at all clear that more information and social learning should necessarily lead to less discrimination.

Most importantly, it is not clear that social learning mechanisms at the heart of these platforms work fairly and unbiasedly to mitigate discrimination. At a high level, social learning involves a feedback of two processes: (1) *the sampling of data (or experience)* and (2) *the informing (or recommending) of user decisions*. The latter process is fair or unbiased, or one can at least guarantee it to be so—in keeping with the recent call for algorithmic fairness. However, the former process is neither random nor unbiased. Data is sampled when transaction occurs, and this process is dictated by the economic interests of parties: users seek reliable, trust-worthy, high-value partners with favorable ratings, not random or representative ones—a far cry from the idealized notion of statistical sampling. Without a deeper understanding of this feedback process, particularly the selective nature of sampling, one cannot truly understand the fairness of social learning and its implications for discrimination.

The purpose of the current paper is to build a model of social learning that accounts for this feedback process and investigate its implications for statistical discrimination. Specifically, we study the possibility of, and the extent to, which social learning through ratings can ameliorate or exacerbate statistical discrimination of social groups.

¹ There is some empirical evidence that reputation/ratings ameliorate discrimination in some contexts (see Cui et al. (2019) and Bohren et al. (2019)).

Our model features directed search/matching between two sides, (masses of long-lived) buyers and sellers, guided by user-contributed ratings. Each seller is indexed by her social group identity $j = 1, 2$ and her productive type, H (“high”) or L (“low”): a seller’s group identity is unchanging, but her productive type changes over time, according to a continuous time Markov chain. The group identity is payoff irrelevant since it has no effect on sellers’ productive types. Buyers seek to match and trade with sellers; the surplus generated from trade is higher if the seller is of high type rather than a low type.

The search-matching process is frictional, and is guided by imperfect information about sellers’ types, called ratings. The rating is binary, either G or B , and is updated after each trade. Although it is impossible for the ratings to perfectly reveal the sellers’ types due to the ever changing nature of seller type (except in the limit), we parameterize the effectiveness of social learning, or the extent to which the ratings “track” sellers’ true types, by $\alpha \in (0, 1)$ —the probability that an incorrect rating gets corrected (e.g., type- H seller having G -rated) after each transaction. Buyers direct their search attention to sellers based on their ratings and possibly on their group identities: a seller of a given rating $j = G, B$ and group $\ell = 1, 2$ matches with probability that depends on the number of buyers directing attention to the sellers with (j, ℓ) . We study the steady state of this system.

Absent any ex ante bias in belief updating, discrimination may still arise from the selective nature of data sampling mentioned earlier. In our model, data on a seller is sampled whenever the seller gets matched, and matching (thus sampling) arises from buyers’ search for sellers with favorable beliefs. A positive feedback loop can then ensue since sellers with a favorable posterior belief will get sampled more, and those believed to have been sampled more often but maintained G rating would enjoy even more favorable belief than those with the same G -rating but believed to have been sampled less often. We identify a positive feedback loop that could lead to a systematic statistical discrimination of social groups.

To begin, the payoff irrelevance of group identity means that there always exists a “non-discriminatory” equilibrium in which sellers of the two groups are treated identically. In this equilibrium, a seller of any given rating enjoys the same amount of attention regardless of her group identity; a G -rated seller enjoys a higher match/employment rate than a B -rated seller, due to the more favorable signal contained in the former, but identically across the group identities. Identical treatment by buyers across the social groups means that there is no bias in the sampling process between the groups, and this in turn leads to the identical updating of beliefs based on the group identity. Hence, non-discrimination survives and perpetuates in steady state.

However, a non-discriminatory equilibrium need not be the unique or even stable

steady state.² There could be another steady state that is discriminatory in the following sense. Suppose buyers direct search attention toward $G1$ sellers (G -rated sellers in group 1) away from $G2$ sellers (G -rated sellers in group 2). This leads to $G1$ sellers being sampled more often than $G2$ in the steady state and to a more intensive weeding out of type L sellers from $G1$ sellers than from $G2$ sellers. This in turn leads to $G1$ sellers enjoying more favorable posterior beliefs than $G2$ sellers and validates the more intensive attention $G1$ sellers are receiving. Hence, the feedback loop is now complete, supporting a discriminatory equilibrium in which buyers favor $G1$ sellers over $G2$ sellers. Interestingly, the same feedback loop does not exist between $B1$ and $B2$ sellers, the B -rated sellers in groups 1 and 2, respectively. Say $B1$ sellers receive more attention from buyers than $B2$ ones. This makes the ratings of former sellers more accurate, which, however, leads to a *less favorable* belief for those sellers. Hence, the initial shift of the attention toward $B1$ gets self-corrected. In the steady state, therefore, $B1$ and $B2$ sellers are treated identically.³

In sum, group 1 sellers are favored than group 2 sellers in this discriminatory equilibrium, despite there being no payoff relevance of the group identity and no bias in either the algorithmic rating/recommendation and belief updating. A discriminatory equilibrium need not always exist, but interestingly, when it does, the non-discriminatory steady state may become unstable; a small perturbation in terms of buyers shifting their attention toward one group may break non-discriminatory equilibrium and trigger a shift that leads to a discriminatory equilibrium. Our analysis shows that a discriminatory equilibrium exists if the matching friction is small and social learning friction is of intermediate value. This suggests that the advance of the online marketplace, as measured by the reduction of these frictions, may have a non-monotonic effect on discrimination. The economy could very well begin with high enough frictions on both accounts that support only the non-discriminatory equilibrium. With an advance in matching and social learning, frictions may diminish and a discriminatory equilibrium may emerge.

The current paper joins the long-line of research on discrimination. In particular, our research follows the literature of statistical discrimination originated by [Phelps \(1972\)](#) and [Arrow \(1973\)](#).⁴ Unlike the taste-based theories of discrimination (see [Becker \(1957\)](#)), this literature explains group inequality and stereotype as resulting

² As will be seen, stability is defined in the usual manner, by the robustness of an equilibrium to small perturbations.

³ Note that this is observationally equivalent to firms paying more attention to the workers in one group than those in the other, resulting in *differential* rewarding across the groups for (perceived) high quality. Interestingly, this behavior is consistent with the empirical findings of [Bertrand and Mullainathan \(2004\)](#) and [Bartoš et al. \(2016\)](#): [Bertrand and Mullainathan \(2004\)](#) performed field experiments by sending out fictitious résumés to help-wanted ads under white-sounding and black-sounding names. They find that not only résumés by white-sounding names receive more call-backs for interviews than those by black-sounding names, but the call-back rates gap between high-quality and low-quality résumés is significantly higher for the former group than for the latter. Similar field experiments were performed by [Bartoš et al. \(2016\)](#) on the pre-screening behavior in job application and apartment application contexts; they find that the advantaged group receives more scrutiny in the former (“cherry-picking”) context whereas a disadvantaged group receives more scrutiny in the latter (“lemon-dropping”) context.

⁴ The earlier theories focus on taste-based discrimination. See [Becker \(1957\)](#).

from rational statistical inferences on groups’ characteristics. In Phelps (1972) and its modern incarnations, discrimination originates from exogenous differences in group characteristics,⁵ whereas Arrow (1973) derives average group differences, and the associated differential treatment, as an endogenous equilibrium behavior. In the same spirit, Coate and Loury (1993b) and the subsequent literature focus on workers’ skill acquisition as a source of discriminatory group stereotyping:⁶ if employers view a certain group as less skilled and thus become more selective against them for assigning higher-paying positions, the affected group will indeed lose incentives for acquiring skills, thus fulfilling the employers’ adverse beliefs on that group.

Our paper is also related to the burgeoning literature in computer science on ethical algorithm (see Dwork et al. (2012), Corbett-Davies et al. (2017), Kearns et al. (2018), and Kleinberg et al. (2018), among others). This literature explores ways to ensure that decision/recommendation algorithms satisfy a variety of fairness standards. The current paper qualifies the effectiveness of this approach, by identifying the possibility that algorithmic fairness alone may not be enough to accomplish a fairness goal. In our model, a discriminatory equilibrium may arise even when the rating algorithm treats both groups identically, as long as agents (buyers in our model) can *interpret* the ratings in a way that can lead to discriminatory sampling. The debate on algorithmic fairness must keep this aspect of social learning into consideration, so that either the interpretational scope is totally eliminated to guarantee outcome fairness, an approach that appears to be in line with the prescription of Kleinberg et al. (2018), or in case that is impossible, the ratings system may be designed to counteract the interpretational response by the users.

2 Model

We consider a frictional search market in which buyers (or firms) search for sellers (or workers) of unknown types.

Players. There is a unit mass of sellers in the market. The sellers are indexed by two characteristics, *type* and *group*. The type of a seller represents any payoff-relevant information, such as the productivity or quality of the seller. At a given moment, a seller is either of high type (H) or low type (L). Each seller’s type, however, changes according to a continuous-time Markov process. Specifically, each type turns into the other type at rate $\delta > 0$. The group of a seller describes her payoff-irrelevant identity,

⁵Cornell and Welch (1996) explains discrimination and its inter-generational persistence from group-specific evaluational familiarity. Bohren et al. (2019) and Monachou and Ashlagi (2019) focus on the “mis-specified” prior beliefs as a source of discrimination.

⁶See Coate and Loury (1993a), Mailath et al. (2000), Moro and Norman (2003), Norman (2003), and a survey by Fang and Moro (2011).

such as her gender, ethnic or racial identity. Each seller belongs to either group 1 or 2, with $\ell = 1, 2$ being used as the generic index. Unlike her type, a seller's group does not change over time. For strong symmetry between the two groups, we assume that both groups have the same total size (i.e., each group has mass 1/2).

On the other side of the market, there is mass $Q(> 0)$ of buyers. They search for sellers based on public information about sellers. Specifically, they condition their search on sellers' two observable characteristics, rating $j = G, B$ and group identity $\ell = 1, 2$.

The sellers who share the same observable characteristics, (j, ℓ) , and the buyers that search for them constitute a "submarket." Clearly, each submarket can be indexed by (j, ℓ) . Sellers are assigned to those submarkets according to their (perfectly persistent) group identity and (evolving) ratings, while buyers choose which submarket to enter.

Matching. We adopt the canonical search-and-matching framework to model an interaction between sellers and buyers. We assume that matching technology is common across all submarkets and exhibits constant returns to scale. The latter assumption implies that all agents' matching rates in each submarket depend only on the ratio λ of buyers to seller in the submarket. We let $\psi(\lambda)$ denote a seller's matching rate and $\phi(\lambda)$ denote a buyer's matching rate. Note that consistency requires that $\psi(\lambda) = \lambda\phi(\lambda)$ for all $\lambda > 0$: matching is one-to-one, and thus the number of matched sellers should be identical to that of matched buyers at each point in time.

For expositional clarity, we focus on the parametric case where $\psi(\lambda) = \lambda^k$ for some $k \in (0, 1)$. This corresponds to the constant-returns-to-scale Cobb-Douglas matching function and, therefore, satisfies various natural and desirable properties.⁷ In particular, $\psi(0) = 0$, $\lim_{\lambda \rightarrow \infty} \psi(\lambda) = \infty$, $\psi'(\lambda) > 0$, and $\psi''(\lambda) < 0$. In addition, $\phi(0) = \infty$, $\lim_{\lambda \rightarrow 0} \phi(\lambda) = 0$, $\phi'(\lambda) < 0$, and $\phi''(\lambda) > 0$. As becomes clear later, most of our results require only these standard properties of the matching function and, therefore, easily generalize beyond our parametric case.

Trade. Once a buyer and a seller meet, they transact instantaneously and go back to the market.⁸ The transaction yields surplus u_H if the seller's type is H and u_L if the seller's type is L , where $u_H > u_L \geq 0$. If a buyer transacts with a seller, the buyer pays $p(\geq 0)$ to the seller. In order to exclude trivial cases, we assume $u_H > p$ (so that

⁷ Formally, let $f(b, s)$ denote the measure of matched formed at each instant when there are mass b of buyers and mass s of sellers. If $f(b, s) = b^k s^{1-k}$, then the rate at which each individual seller is matched with a buyer is given by

$$\psi(\lambda) \equiv \psi\left(\frac{b}{s}\right) = \frac{f(b, s)}{s} = \frac{b^k s^{1-k}}{s} = \left(\frac{b}{s}\right)^k = \lambda^k.$$

⁸ The assumption on instantaneous transaction makes the model more naturally applicable for one-shot relationship such as free-lance independent work and insurance or loan purchase. It is conceptually straightforward, but technically complicated, to extend our model to capture the canonical persistent employment relationship.

there are gains from trade when a seller is of type H), but consider both the case when $u_L > p$ and the case when $u_L \leq p$.

Ratings. Market accumulates information about sellers through simple summary indices, called “ratings.” There are two possible ratings: G (as in “good”) and B (as in “bad”). When searching for sellers, buyers can only observe the current rating of the sellers; no information about sellers’ underlying types or their past ratings history is available to them. After each transaction, the seller’s rating may be updated to reveal her type. Specifically, with probability $\alpha \in (0, 1]$, a B -rated seller with type H receives G rating, and a G -rated seller with type L receives G rating. A seller with correct rating keeps the same rating after a transaction. With remaining probability $1 - \alpha$, the seller’s rating remains unchanged. Note that due to the changing environment (or changing type), a correct rating may turn inaccurate.

Buyers’ beliefs over a seller’s type will depend on the rating, and the (equilibrium) behavior of all players in the system. In particular, the belief may depend on the group identity. If the two groups of agents are treated differently, the inference a buyer makes on a seller with a given rating depend nontrivially on her group identity. This will be made clear in our analysis.

Solution Concept. We consider a steady state of the economy in terms of the distribution of sellers of different types, ratings and group identity, and the beliefs that the buyers hold for each submarket. Specifically, an equilibrium is a tuple $\{(P_{ij}^\ell, \lambda_j^\ell, \mu_j^\ell)\}_{i=H,L, j=G,B}^{\ell=1,2}$ in the stationary distribution, where P_{ij}^ℓ is the mass of sellers of type i with rating j and group ℓ , λ_j^ℓ is the ratio of buyers to sellers in submarket (j, ℓ) , and $\mu_j^\ell \in [0, 1]$ is the public belief on the sellers in submarket (j, ℓ) , i.e., the probability that they are of type H . Note that the masses of buyers who participate in alternative submarkets are determined by the first components.

3 Non-Discriminatory Equilibrium

We begin by studying equilibria in which buyers search for sellers only based on their ratings and do not distinguish between the two groups. As in other statistical-discrimination models, this equilibrium always exists and provides a benchmark for discriminatory equilibria studied in Section 4. Several technical results in this section will also be useful in Section 4.

3.1 Steady-State Distribution

In non-discriminatory equilibrium, buyers condition their search strategies only on sellers' ratings. Therefore, effectively, there are only 2 submarkets indexed by rating $j = G, B$. For each $j = G, B$, let q_j denote the measure of buyers that join submarket j and P_{ij} denote the measure of type $i = H, L$ sellers with rating $j = G, B$. Then, the ratio of buyers to sellers ("queue length") in submarket $j = G, B$ is given as follows:

$$\lambda_j \equiv \frac{q_j}{P_{Hj} + P_{Lj}}.$$

Our matching technology implies that a j -rated seller is hired by a buyer at rate $\psi_j \equiv \psi(\lambda_j)$, while a buyer in submarket j successfully finds a seller at rate $\phi_j \equiv \psi_j/\lambda_j$.

In steady state, P_{ij} 's must satisfy the following system of equations:

$$\begin{aligned} P_{HG}\delta &= P_{LG}\delta + P_{HB}\psi_B\alpha, \\ P_{LG}(\delta + \psi_G\alpha) &= P_{HG}\delta, \\ P_{HB}(\delta + \psi_B\alpha) &= P_{LB}\delta, \text{ and} \\ P_{LB}\delta &= P_{HB}\delta + P_{LG}\psi_G\alpha. \end{aligned}$$

In each equation, the left-hand side represents the outflow of sellers from status (i, j) , while the right-hand side quantifies the corresponding inflow. For example, consider P_{HG} . Due to our rating technology (with no false negative), a type H seller with rating G changes his status only when his type changes to L , which occurs at rate δ (thus, $-P_{HG}\delta$). On the other hand, a type L seller with rating G becomes type H at rate δ (thus $+P_{LG}\delta$). In addition, a type H seller with rating B improves his rating to G once he meets a buyer and receives a good rating (thus $+P_{HB}\psi_B\alpha$). In steady state, net flow must be equal to 0, which yields the first equation.

The following result is straightforward from the above equations.

Lemma 1 (Steady-state Distribution). *In steady state, the measure of sellers with type $i = H, L$ and rating $j = G, B$ is given as follows:*

$$\begin{aligned} P_{HG} &= \frac{\psi_B(\delta + \psi_G\alpha)}{2(\delta(\psi_G + \psi_B) + \alpha\psi_G\psi_B)}, & P_{LG} &= \frac{\psi_B\delta}{2(\delta(\psi_G + \psi_B) + \alpha\psi_G\psi_B)}, \\ P_{HB} &= \frac{\psi_G\delta}{2(\delta(\psi_G + \psi_B) + \alpha\psi_G\psi_B)}, & P_{LB} &= \frac{\psi_G(\delta + \psi_B\alpha)}{2(\delta(\psi_G + \psi_B) + \alpha\psi_G\psi_B)}. \end{aligned}$$

Letting $\mu_j \equiv P_{Hj}/(P_{Hj} + P_{Lj})$ for each $j = G, B$,

$$\mu_G = 1 - \frac{\delta}{2\delta + \psi_G\alpha} \text{ and } \mu_B = \frac{\delta}{2\delta + \psi_B\alpha}.$$

Proof. See the appendix.

Q.E.D.

Notice that for each $j = G, B$, the fraction μ_j of type H sellers in submarket j depends only on ψ_j . This is due to the fact that ψ_j has the same proportional effect on P_{Hj} and P_{Lj} , and thus P_{Hj}/P_{Lj} is independent of ψ_j . More concretely, our rating technology involves no type-I errors, and thus a seller's rating improves from B to G only when her type is H , while her rating falls from G to B only when her type is L . This implies that in each submarket, the mass of sellers with wrong rating (P_{LG} in submarket G , and P_{HB} in submarket B) is fully determined by the mass of sellers with correct rating (P_{HG} in submarket G , and P_{LB} in submarket B) and the matching rate of the submarket (see the second and the third equations in the above system of equations). This drives the convenient independence property of μ_j .

In addition, μ_G increases in ψ_G , while μ_B decreases in ψ_B . This is intuitive: an increase in ψ_G or ψ_B can be interpreted as better/faster screening of sellers. Therefore, a seller with rating G becomes more likely to be type H as ψ_G increases. Similarly, a seller with rating B becomes less likely to be type H (more likely to be type L) as ψ_B increases.

3.2 Buyers' Expected Payoffs

Let u_j denote a buyer's flow expected payoff when he targets j -rated sellers (i.e., searches in submarket j). Given the steady-state queue length λ_j and the fraction μ_j of type H sellers, u_j is given by

$$u_j = \phi_j(\mu_j u_H + (1 - \mu_j)u_L - p).$$

There are the following two cases to consider.

- (i) $u_j \leq 0$: This case arises if and only if⁹

$$\mu_j u_H + (1 - \mu_j)u_L - p \leq 0 \Leftrightarrow \mu_j \leq \underline{\mu} \equiv \frac{p - u_L}{u_H - u_L}.$$

This is when a buyer's utility from a type L seller falls short of the price p (i.e., $u_L < p$) and the probability that a buyer meets a type L seller is sufficiently large. In this case, clearly, no buyers search in submarket j , that is, $\lambda_j = 0$.

- (ii) $u_j > 0$: Opposite to (i), this arises if and only if $\mu_j > \underline{\mu}$. In addition, it must be that a positive measure of buyers search in submarket j and, therefore, $\lambda_j > 0$: if $\lambda_j = 0$, then a participating buyer would meet sellers infinitely frequently, each

⁹ If $\lambda_j = \infty$ then $\phi_j = 0$, and thus $u_j = 0$ even if $\mu_j u_H + (1 - \mu_j)u_L - p > 0$. However, this case clearly cannot be sustained in equilibrium.




Figure 1: The blue solid curves shows buyers' expected payoffs in submarket G , as a function of λ_G . The common parameter values used for this figure are $\delta = \alpha = 0.1$, $u_H = 2$, and $u_L = w = 1$ (which leads to $\underline{k} = 0.8536$). $k = 0.7682$ in the left panel, while $k = 0.8828$ in the right panel.

of whom gives the buyer a positive expected payoff $\mu_j u_H + (1 - \mu_j) u_L - p > 0$, and thus his expected payoff becomes unbounded.

Recall that $\phi_j = \phi(\lambda_j)$ and in steady state, μ_j is also a function only of λ_j (see Lemma 1). Therefore, u_j also can be interpreted as a function of λ_j . Interestingly and importantly, whereas $u_B(\lambda_B)$ is always monotone, $u_G(\lambda_G)$ may not be monotone, as formally reported in the following lemma (and shown in Figure 1).

Lemma 2. *Suppose that $(y_H + y_L)/2 > w$. For both $j = G, B$, $\lim_{\lambda_j \rightarrow 0} u_j(\lambda_j) = \infty$, $\lim_{\lambda_j \rightarrow \infty} u_j(\lambda_j) = 0$, and $u_j(\lambda_j)$ is continuous. $u_B(\lambda_B)$ is always strictly monotone (decreasing), while $u_G(\lambda_G)$ is monotone if and only if*

$$k \leq \underline{k} \equiv \frac{1 + \sqrt{1 - \frac{u_H - u_L}{2(u_H - w)}}}{2}.$$

If $k > \underline{k}$, then there exist $\underline{\lambda}_G (> 0)$ and $\bar{\lambda}_G (> \underline{\lambda}_G)$ such that $u_G(\lambda_G)$ is strictly increasing if and only if $\lambda_G \in (\underline{\lambda}, \bar{\lambda})$.

Proof. See the appendix. Q.E.D.

This result is due to the fact that $\mu_B(\lambda_B)$ is decreasing in λ_B , while $\mu_G(\lambda_G)$ is increasing in λ_G (Lemma 1). Since a seller's matching rate $\phi(\lambda)$ always decreases with relatively more sellers (i.e., higher λ), the utility effect of increasing λ is always negative in submarket B but ambiguous in submarket G . Lemma 2 shows that the quantity effect (through $\phi(\lambda_G)$) always dominates the quality effect (through $\mu_G(\lambda_G)$) if λ_G is sufficiently small or sufficiently large: the former is because the quantity effect

$(\phi'(\lambda_G) = (k-1)\lambda_G^{k-2})$ is arbitrarily large when λ_G is close to 0, while the latter is because as λ_G tends to infinity, the quality effect vanishes faster than the quantity effect. Given these observations, it is intuitive that the quality effect can outweigh the quantity effect, and thus $u_G(\lambda_G)$ decreases, over an interval if and only if k is sufficiently close to 1 (so that given $\lambda_G > 0$, $\phi'(\lambda_G) = (k-1)\lambda_G^{k-2}$ is close to 0).

3.3 Equilibrium Characterization

We now characterize non-discriminatory steady-state equilibria of our model. The following result provides a necessary and sufficient condition for there to be no trade in equilibrium.¹⁰

Proposition 1. *If $(u_H + u_L)/2 \leq p$, then it is the unique non-discriminatory steady-state equilibrium outcome that buyers do not search for sellers, regardless of their ratings (i.e., $\lambda_G = \lambda_B = 0$). Conversely, if $(u_H + u_L)/2 > p$, then buyers search for both ratings of sellers (i.e., $\lambda_G, \lambda_B > 0$).*

Proof. Suppose that $\lambda_j = 0$. Since $\psi_j = \psi(\lambda_j) = 0$, by Lemma 1, $\mu_j = 1/2$. But then,

$$\mu_j u_H + (1 - \mu_j) u_L - p = \frac{u_H + u_L}{2} - p.$$

Therefore, $\lambda_j = 0$ (no trade in submarket j) can be an equilibrium if and only if $(u_H + u_L)/2 \leq p$. Since this condition is independent of rating j , it is either $\lambda_G = \lambda_B = 0$ or $\lambda_G, \lambda_B > 0$. Q.E.D.

This result is fairly intuitive. Since both types of sellers change their types at an identical rate, the unconditional proportion of type H sellers is $1/2$ in steady state: recall that in Lemma 1, $(P_{HG} + P_{HB}) / \sum_{i,j} P_{ij} = 1/2$. If there is no trade, then there is also no market learning, that is, a seller's rating becomes uninformative of his type (observe that $\mu_G = \mu_B = 1/2$ if $\lambda_G = \lambda_B = 0$). Therefore, no trade outcome can be sustained if and only if $(u_H + u_L)/2 \leq p$.

Now consider the case where $(u_H + u_L)/2 > p$, so that $\lambda_G, \lambda_B > 0$. In this case, searching buyers receive positive surplus (i.e., $\mu_G, \mu_B > 0$), and thus it is necessarily the case that all buyers search. Therefore, in equilibrium, the following "market clearing" condition must hold:

$$q_G + q_B = Q.$$

Since $\lambda_j = q_j / (P_{Hj} + P_{Lj})$ for each $j = G, B$, this condition can be rewritten, in terms of

¹⁰Note that Proposition 1 argues only uniqueness of no trade outcome, not that of symmetric steady-state equilibrium. This is because in no-trade equilibrium, there is no seller movement between submarkets G and B , and thus any distribution of sellers can be sustained as long as $\mu_G = \mu_B = 1/2$.

(λ_G, λ_B) , as follows:

$$\begin{aligned} Q &= \lambda_G(P_{HG} + P_{LG}) + \lambda_B(P_{LG} + P_{LB}) \\ &= \frac{\lambda_G \psi_B(2\delta + \alpha \psi_G)}{2(\delta \psi_G + \delta \psi_B + \alpha \psi_G \psi_B)} + \frac{\lambda_B \psi_G(2\delta + \alpha \psi_B)}{2(\delta \psi_G + \delta \psi_B + \alpha \psi_G \psi_B)}. \end{aligned} \quad (1)$$

In addition, in equilibrium, a buyer must be indifferent between submarket G and submarket B , that is, the following buyer-indifference condition should hold:

$$\phi_G(\mu_G u_H + (1 - \mu_G)u_L - p) = u_G(\lambda_G) = u_B(\lambda_B) = \phi_B(\mu_B u_H + (1 - \mu_B)u_L - p). \quad (2)$$

This equation shows a buyer's trade-off between ratings G and B . By Lemma 1, it is always the case that a seller is more likely to be type H when his rating is G than when his rating is B (i.e., $\mu_G \geq 1/2 \geq \mu_B$). This makes rating G attract relatively more buyers than rating B , which reduces a buyer's chance to hire a G -rated seller (i.e., $\lambda_G > \lambda_B$, and thus $\phi_G < \phi_B$). In equilibrium, λ_G and λ_B are such that a buyer is indifferent between the two submarkets.

Combining the two conditions leads to the following result.

Proposition 2. *If $(u_H + u_L)/2 > w$, then there always exists a non-discriminatory steady-state equilibrium in which $\lambda_G > \lambda_B > 0$.*

Proof. Let $\lambda_B^{MC}(\lambda_G)$ be the implicit function defined by equation (1). It is straightforward that $\lim_{\lambda_G \rightarrow 0} \lambda_B^{MC}(\lambda_G) = \infty$ and $\lim_{\lambda_G \rightarrow \infty} \lambda_B^{MC}(\lambda_G) = 0$. In addition, since the right-hand side of equation (1) increases in both λ_G and λ_B , $\lambda_B^{MC}(\lambda_G)$ is monotone.¹¹ Similarly, let $\lambda_B^{BI}(\lambda_G)$ be the implicit function defined by equation (2). By Lemma 2, $\lim_{\lambda_G \rightarrow 0} \lambda_B^{BI}(\lambda_G) = 0$, $\lim_{\lambda_G \rightarrow \infty} \lambda_B^{BI}(\lambda_G) > 0$, and $\lambda_B^{BI}(\lambda_G)$ is continuous. Therefore, there exists $\lambda_G^* > 0$ such that $\lambda_B^* \equiv \lambda_B^{MC}(\lambda_G) = \lambda_B^{BI}(\lambda_G)$. By construction, the pair $(\lambda_G^*, \lambda_B^*)$ constitute a non-discriminatory steady-state equilibrium. *Q.E.D.*

Figure 2 explains the argument behind Proposition 2. The red dashed curves represent the market clearing condition (1). They are always decreasing from infinity to zero as λ_G increases, which is because the right-hand side is increasing in both λ_G and λ_B . The blue solid curves capture the buyer-indifference condition (2). As shown in Figure 2, they are not necessarily monotone, which is because, whereas $u_B(\lambda_B)$ always decreases, $u_G(\lambda_G)$ may increase over an interval (Lemma 2). Nevertheless, they are always continuous, start from 0 and eventually stay away from 0, and thus the two curves always intersect.

¹¹ For the result with λ_G , one can express the right-hand side as follows:

$$\frac{(\lambda_G/\psi_G)\psi_B(2\delta + \alpha\mu_G)}{2(\delta + \delta\psi_B/\psi_G + \alpha\psi_B)} + \frac{\lambda_B\psi_G(2\delta + \alpha\psi_B)}{2(\delta + \delta\psi_B/\psi_G + \alpha\psi_B)}.$$

The desired result follows from the fact that $\phi(\lambda) = \psi(\lambda)/\lambda$ is decreasing in λ , while $\psi(\lambda)$ is increasing in λ .




Figure 2: The blue solid curves depict the buyer-indifference condition (2), while the red dashed curves depict the market clearing condition (1). The common parameter values used for this figure are $\delta = 1$, $\alpha = 0.1$, $u_H = 2$, and $u_L = w = 1$ (which leads to $\underline{k} = 0.8536$). $k = 0.8682$ in the left panel, while $k = 0.9121$ in the right panel.

As exemplified by the right panel of Figure 2, there may exist multiple equilibria. This is, again, because the buyer-indifference condition may produce a non-monotone relationship between λ_G and λ_B (i.e., the implicit function $\lambda_B^{BI}(\lambda_G)$ may not be monotone). Clearly, if $k \leq \underline{k}$ (in which case $\lambda_B^{BI}(\lambda_G)$ is monotone), then there always exists a unique non-discriminatory equilibrium. Even if $k > \underline{k}$, it is often the case that non-discriminatory equilibrium is unique (see the left panel). However, there is a non-negligible set of parameter values that yield multiple equilibria.¹²

4 Discriminatory Equilibrium

In this section, we investigate discriminatory equilibria in which buyers condition their search strategies on sellers' identities as well as their ratings. We first derive a condition under which such equilibria exist and then compare them to non-discriminatory equilibria studied in Section 3.

4.1 Notation and Assumption

We use the same notation as in Section 3 but distinguish the two groups with superscripts $\ell = 1, 2$. For example, we denote by q_j^ℓ the measure of buyers who are targeting j -rated sellers in group ℓ , and P_{ij}^ℓ the measure of sellers with type i , rating j , and group ℓ .

¹² If there are multiple equilibria, then they are ranked in terms of buyer surplus: since the market-clearing condition is always decreasing in λ_G , if there are two equilibria, (λ_G, λ_B) and (λ'_G, λ'_B) , and $\lambda_G < \lambda'_G$, then $\lambda_B > \lambda'_B$. In this case, buyers' expected payoffs are necessarily higher with (λ'_G, λ'_B) , because μ_B is decreasing in λ_B (see Lemma 1), and thus $u_B(\lambda'_B)$ has not only a higher value of $\phi(\lambda_B)$ but also a higher value of μ_B .

Within each group $\ell = 1, 2$, sellers follow the same transition dynamics as in the non-discriminatory case. Therefore, Lemma 1 applies unchanged to each group. In particular, for each $\ell = 1, 2$, the proportion of type H sellers in submarket $j\ell$ is given as follows:

$$\mu_G^\ell \equiv \mu_G(\lambda_G^\ell) = 1 - \frac{\delta}{2\delta + \psi(\lambda_G^\ell)\alpha} \text{ and } \mu_B^\ell \equiv \mu_B(\lambda_B^\ell) = \frac{\delta}{2\delta + \psi(\lambda_B^\ell)\alpha}.$$

In addition, buyers' expected payoffs are determined as follows:

$$u_j^\ell(\lambda_j^\ell) = \phi(\lambda_j^\ell)(\mu_j^\ell u_H + (1 - \mu_j^\ell)u_L - p).$$

Proposition 1 also applies unchanged. Specifically, by the same logic and proof as for non-discriminatory equilibria, there is no trade in any submarket if and only if $(u_H + u_L)/2 \leq p$. In addition, if $(u_H + u_L)/2 > p$, then trade must take place in all four submarkets. Since the analysis is trivial for the no-trade case, from now on, we maintain the following assumption:

Assumption 1. $(u_H + u_L)/2 > p$, and thus $\lambda_j^\ell > 0$ for all $j = G, B$ and $\ell = 1, 2$.

4.2 Existence of Discriminatory Equilibrium

We begin by presenting a necessary condition for the existence of discriminatory equilibria, namely that the function $u_G(\lambda)$ must be non-monotone, which is the case if and only if $k > \bar{k}$.

Proposition 3. *If $u_G(\lambda)$ is monotone (i.e., $k \leq \bar{k}$), then there does not exist a discriminatory equilibrium in which $\lambda_B^1 \neq \lambda_B^2$ or $\lambda_G^1 \neq \lambda_G^2$.*

Proof. In equilibrium, buyers must be indifferent over all 4 submarkets, that is,

$$u_G(\lambda_G^1) = u_B(\lambda_B^1) = u_B(\lambda_B^2) = u_G(\lambda_G^2).$$

Since $u_B(\lambda)$ is monotone, it is always the case that $\lambda_B^1 = \lambda_B^2$. If $u_G(\lambda)$ is also monotone, then it is also the case that $\lambda_G^1 = \lambda_G^2$. Therefore, a discriminatory equilibrium cannot exist. Q.E.D.

Let us first explain how the non-monotonicity of $u_G(\lambda)$ can lead to the existence of discriminatory equilibria. Suppose that $u_G(\lambda)$ is non-monotone, and fix λ_B such that $u_B(\lambda_B) = u_G(\lambda_G) = u_G(\lambda'_G)$ for two distinct values, $\lambda_G < \lambda'_G$: in Figure 3, it suffices to choose $\lambda_B \in [\underline{\lambda}_B, \bar{\lambda}_B]$. Given λ_B and λ_G , let Q^1 be the value that supports (λ_B, λ_G) as a




Figure 3: The blue solid curve represents the buyer-indifference condition (2). The parameter values used for this figure are $\delta = 0.2$, $\alpha = 0.5$, $u_H = 3$, $u_L = 1$, $w = 1.5$ (which leads to $\underline{k} = 0.7887$), and $k = 0.8204$.

non-discriminatory equilibrium, that is,

$$Q^1 \equiv \frac{\lambda_G \psi(\lambda_B)(2\delta + \alpha \psi(\lambda_G))}{2(\delta \psi(\lambda_G) + \delta \psi_B + \alpha \psi(\lambda_G) \psi(\lambda_B))} + \frac{\lambda_B \psi(\lambda_G)(2\delta + \alpha \psi(\lambda_B))}{2(\delta \psi(\lambda_B) + \delta \psi(\lambda_B) + \alpha \psi(\lambda_G) \psi(\lambda_B))}.$$

Similarly, let Q^2 be the corresponding value for (λ_B, λ'_G) , that is,

$$Q^2 \equiv \frac{\lambda'_G \psi(\lambda_B)(2\delta + \alpha \psi(\lambda'_G))}{2(\delta \psi(\lambda'_G) + \delta \psi_B + \alpha \psi(\lambda'_G) \psi(\lambda_B))} + \frac{\lambda_B \psi(\lambda'_G)(2\delta + \alpha \psi(\lambda_B))}{2(\delta \psi(\lambda_B) + \delta \psi(\lambda_B) + \alpha \psi(\lambda'_G) \psi(\lambda_B))}.$$

Suppose that $Q = Q^1 + Q^2$, and consider a steady state in which $\lambda_B^1 = \lambda_B^2 = \lambda_B$, $\lambda_G^1 = \lambda'_G$, and $\lambda_G^2 = \lambda_G$. By construction, buyers are indifferent over all 4 submarkets, that is,

$$u_B(\lambda_B^1) = u_B(\lambda_B^2) = u_G(\lambda_G^1) = u_G(\lambda_G^2).$$

In addition, for both $\ell = 1, 2$, we have

$$Q^\ell = \lambda_G^\ell (P_{HG}^\ell + P_{LG}^\ell) + \lambda_B (P_{HB}^\ell + P_{LB}^\ell).$$

Therefore, this (discriminatory) strategy profile is a steady-state equilibrium.

This discriminatory equilibrium gives higher expected payoffs to group 1 sellers than to group 2 sellers. Intuitively, this is because the two groups have the same matching rate $\psi(\lambda_B)$ with rating B , but the former have a higher matching rate than the latter with rating G (i.e., $\psi(\lambda_G^1) > \psi(\lambda_G^2)$).¹³ Let us emphasize that this discriminatory out-

¹³ To be precise, the following composition effect needs to be taken into account: more group 2 sellers are G -rated than group 1 sellers (i.e., $P_G^1 < P_G^2$). However, it can be shown that this composition effect never outweighs the effect due to the difference in matching rates: the unconditional matching rate $P_G^\ell \psi_G(\lambda_G^\ell) + P_B^\ell \psi_B(\lambda_B^\ell)$ is increasing in both λ_G and λ_B .

come arises even if the two groups have no fundamental differences (including the size of the group) and, perhaps more importantly, despite the fact that there may exist a unique non-discriminatory equilibrium. In other words, unlike in some related papers, the possibility of discriminatory equilibria relies neither on any intrinsic differences between the two groups nor on the equilibrium multiplicity of our underlying environment.

The mechanism behind our discriminatory equilibria is the novel feedback loop between ratings and trade. More buyers search for group 1 sellers because their ratings are more informative (i.e., $\mu_G^1 > \mu_G^2$). Conversely, group 1 ratings are more accurate because they are hired and reviewed more frequently (i.e., $\psi(\lambda_G^1) > \psi(\lambda_G^2)$).

As shown by Proposition 3, this feedback effect never results in the existence of discriminatory equilibria if $u_G(\lambda)$ is monotone. It may not work even if $u_G(\lambda)$ is not monotone, depending on Q . Nevertheless, the following result shows that whenever $u_G(\lambda)$ is not monotone, there is a positive measure of Q 's that give rise to discriminatory equilibria.

Theorem 1. *If $u_G(\lambda)$ is not monotone (i.e., $k \in (\underline{k}, 1)$), then there exist $\underline{Q}(> 0)$ and $\overline{Q}(> \underline{Q})$ such that a discriminatory equilibrium exists if and only if $Q \in (\underline{Q}, \overline{Q})$.*

Proof. See the appendix.

Q.E.D.

In order to see why a discriminatory equilibrium does not exist if $Q < \underline{Q}$ or $Q > \overline{Q}$, note that in our model, buyers use group identities only to improve informational content of ratings. If Q is sufficiently small, then $\lambda_j^\ell \approx 0$ for all $j\ell$, and thus ratings do not contain much information (i.e., $\mu_j^\ell \approx 1/2$ for all $j\ell$). To the contrary, if Q is sufficiently large, then each λ_j^ℓ is large, and thus ratings convey precise information about sellers' types (i.e., for both $\ell = 1, 2$, $\mu_B^\ell \approx 0$, while $\mu_G^\ell \approx 1$). In both cases, group identities do not add much more information to ratings, and thus it is unlikely that buyers condition their search strategies on them.

4.3 Stability of Ratings-based Discrimination

The analysis so far has demonstrated that non-discriminatory and discriminatory equilibria can coexist. The possibility of discriminatory equilibria is intriguing by itself, but they would be unlikely to matter in practice if they could not be sustained in any robust manner. We now examine whether (and when) discriminatory equilibria exhibit desirable stability properties.

Since discriminatory equilibria can exist only when $u_G(\lambda)$ is not monotone, we mainly focus on the case where $k \in (\underline{k}, 1)$ throughout this subsection. In addition, so as to streamline the analysis, we also restrict attention to the parameter space that yields a unique non-discriminatory equilibrium.

Assumption 2. For all $Q \in \mathcal{R}_+$, there exists a unique non-discriminatory equilibrium.

Since the market clearing condition is monotone in both λ_G and λ_B , this assumption holds whenever the function $u_G(\lambda)$ does not increase too sharply, which is guaranteed if k is close to \underline{k} .

In order to provide a relevant stability concept, let $U(q)$ denote buyers' equilibrium expected payoffs in the non-discriminatory equilibrium when the total measure of buyers is given by $2q$. In other words, if (λ_B, λ_G) is a non-discriminatory equilibrium with $2q$ measure of buyers, then $U(q) = u_B(\lambda_B) = u_G(\lambda_G)$. Proposition 2 and Assumption 2 ensure that $U(Q)$ is well-defined for all Q .

Both non-discriminatory and discriminatory equilibria can be summarized by a pair (Q^1, Q^2) such that $Q^1 + Q^2 = Q$ (market clearing) and $U(Q^1) = U(Q^2)$ (buyer indifference among all 4 submarkets). The only difference between them is that an equilibrium is non-discriminatory if $Q^1 = Q^2$, while it is discriminatory if $Q^1 \neq Q^2$. Based on this observation, we make use of the following tractable notion of stability.

Definition 1. An equilibrium with (Q^1, Q^2) is stable if $U'(Q^1) + U'(Q^2) \leq 0$ and unstable otherwise.

In order to understand this definition, fix a steady-state equilibrium with (Q^1, Q^2) , and suppose that a small measure of buyers move from group 1 to group 2, so that $Q^1 - \Delta$ measure of buyers search for group 1 sellers and $Q^2 + \Delta$ measure of buyers search for group 2 sellers. After the change, if buyers targeting group 1 receive a higher expected payoff than those targeting group 2, then those buyers who left group 1 will move back to group 1, restoring the original equilibrium with (Q^1, Q^2) . If it is the opposite, then even more buyers would move to group 2, making the economy drift further away from the equilibrium (Q^1, Q^2) . Clearly, the (original) equilibrium is *stable* in the former case and *unstable* in the latter case. Our stability definition captures this idea in a particularly simple fashion. For the simple condition in the definition, observe that, since $U(Q^1) = U(Q^2)$,

$$U(Q^1 - \Delta) \geq U(Q^2 + \Delta) \Leftrightarrow -\frac{U(Q^1) - U(Q^1 - \Delta)}{\Delta} \geq \frac{U(Q^2 + \Delta) - U(Q^2)}{\Delta},$$

which reduces to $U'(Q^1) + U'(Q^2) \leq 0$ in the limit as Δ tends to 0.

We first apply our stability notion to non-discriminatory equilibria.

Proposition 4. Fix a non-discriminatory equilibrium in which $\lambda_B^\ell = \lambda_B$ and $\lambda_G^\ell = \lambda_G$ for both $\ell = 1, 2$. The equilibrium is stable if and only if $u_G(\cdot)$ is decreasing at λ_G .

Proof. In a non-discriminatory equilibrium, $Q^1 = Q^2 = Q/2$. Therefore, it is stable if and only if $U'(Q/2) \leq 0$. Since the market clearing condition (1) always expands as

Q increases, this is equivalent to the equilibrium value of λ_B increasing in Q (so that $u_B(\lambda_B)$ decreases), which in turn holds if and only if $u_G(\cdot)$ is decreasing at λ_G .

Q.E.D.

Recall that if $k \leq \underline{k}$, then there exists a unique steady-state equilibrium, which is non-discriminatory (Propositions 2 and 3). Proposition 4 suggests that the unique equilibrium is stable, which is desirable. It further suggests that even if a non-discriminatory equilibrium coexists with discriminatory equilibria, the former may be stable, but not always. If an equilibrium lies on a decreasing region of $u_G(\cdot)$ (either $\lambda_G \leq \underline{\lambda}_G$ or $\lambda_G \geq \bar{\lambda}_G$), then it is stable. Otherwise (i.e., $\lambda_G \in (\underline{\lambda}_G, \bar{\lambda}_G)$), the equilibrium is unstable.

Can a discriminatory equilibrium be stable? The following result argues that (there is a sense in which) discriminatory equilibria are more likely to be stable than non-discriminatory equilibria.

Theorem 2. *Whenever there exists a discriminatory equilibrium, there exists a stable discriminatory equilibrium.*

Proof. Suppose that a discriminatory equilibrium exists. By Theorem 1, this is the case if and only if $u_G(\lambda)$ is non-monotone (i.e., $k \in (\underline{k}, 1)$) and $Q \in (\underline{Q}, \bar{Q})$. Now consider the function $g : [0, Q/2) \rightarrow \mathcal{R}$ such that

$$g(x) = U\left(\frac{Q}{2} + x\right) - U\left(\frac{Q}{2} - x\right).$$

Since $U(Q)$ is finite, while $\lim_{q \rightarrow 0} U(q) = \infty$, $g(x) < 0$ if x is sufficiently close to $Q/2$. In addition, the existence of discriminatory equilibrium implies that there exists $x^* \in [0, Q/2)$ such that $g(x^*) = 0$ (i.e., $x^* = |Q^1 - Q^2|/2$). Combining these with continuity of g , it follows that there exists $x^{**} \in [x^*, Q/2)$ such that

$$g(x^{**}) = 0 \text{ and } g'(x^{**}) \leq 0.$$

The first condition implies that $(Q/2 - x^{**}, Q/2 + x^{**})$ is a discriminatory equilibrium, while the second condition implies that the equilibrium is stable. *Q.E.D.*

Note that Theorem 2 does not claim that all discriminatory equilibria are stable. There may exist an unstable discriminatory equilibrium. Theorem 2 implies that if such an unstable equilibrium exists, then there exists another discriminatory equilibrium that is stable. Of course, it also implies that if there is a unique discriminatory equilibrium, then it is necessarily stable.

4.4 Rating Quality and Discrimination

As ratings have become increasingly more prevalent, the associated technology also has improved, extracting more (accurate) information from more users (buyers). Will such a technological advance bring more fairness by weakening the role of prejudice in decision making, or can it actually worsen discrimination? The following result shows that the effect is non-monotone in general: an improvement in rating quality, measured by $\beta \equiv \alpha/\delta$, may create discrimination initially.¹⁴ However, if the technology becomes sufficiently effective, then discrimination becomes unsustainable.

Proposition 5. *Fix $k > \underline{k}$ and $Q > 0$. There exist $\underline{\beta}(> 0)$ and $\bar{\beta}(> \underline{\beta})$ such that a discriminatory equilibrium exists if and only if $\beta \in (\underline{\beta}, \bar{\beta})$.*

Proof. Let $\lambda' \equiv \lambda\beta^{1/k}$. Then, the two equilibrium equations, (1) and (2), can be written as follows:

$$Q\beta^{1/k} = \frac{\lambda'_G \psi'_B (2 + \psi'_G)}{2(\psi'_G + \psi'_B + \psi'_G \psi'_B)} + \frac{\lambda'_B \psi'_G (2 + \psi'_B)}{2(\psi'_G + \psi'_B + \psi'_G \psi'_B)}$$

and

$$\phi(\lambda'_G) \left(\frac{1 + \psi'_G}{2 + \psi'_G} (u_H - u_L) + u_L - p \right) = \phi(\lambda'_B) \left(\frac{1}{2 + \psi'_B} (u_H - u_L) + u_L - p \right).$$

Notice that β appears only on the left-hand side of the first equation, together with Q . This implies that if there is an equilibrium (whether discriminatory or non-discriminatory) with Q and β , then effectively the same equilibrium exists with $Q\beta^{1/k}$ and 1 as well, and vice versa. Combining this observation with Theorem 1, it follows that for a fixed value of Q , a discriminatory equilibrium exists if and only if $\beta \in (\underline{\beta}, \bar{\beta})$, where $Q\underline{\beta}^{1/k} = \underline{Q}$ and $Q\bar{\beta}^{1/k} = \bar{Q}$. Q.E.D.

For an intuition, recall that discrimination arises in our model because a seller's group identity may provide extra information about her productive type. Such extra information is of little value when ratings are sufficiently uninformative (i.e., β is close to 0) or sufficiently informative (i.e., β is large). It can make a large enough difference that can sustain a discriminatory equilibrium only when rating quality belongs to an intermediate range.

5 Comparison to Coate and Loury (1993a)

Just as in our paper, Coate and Loury (1993a) (CL, hereafter) demonstrated that discriminatory outcomes can arise despite no exogenous differences between two groups.

¹⁴ Note that α and δ always enter the steady-state system together (see Lemma 1), and thus they cannot be separately identified. Intuitively, α measures how fast ratings get corrected, while δ measures how fast ratings become obsolete. Therefore, their ratio, $\beta = \alpha/\delta$, is the proper measure of rating quality.

They considered a labor market model in which workers invest in their human capital and employers assign a job to each worker based on the worker's group identity and a noisy signal about the worker's human capital. They showed that discrimination can result from a "self-fulfilling prophecy": employers believe that group 2 workers are less likely to invest in their human capital than group 1 workers and, therefore, assign less group 2 workers to a better job than group 1 workers. Expecting lower returns, group 2 workers indeed invest in their human capital less than group 1 workers.

The underlying mechanism behind statistical discrimination in our model differs from that of CL in several important ways. First, the discrimination in our model is based on rational inference on sellers' transaction histories ("ratings"), and thus it does not require endogenous human capital acquisition as in CL. The perspectives on the cause of discrimination also differ. In the latter, "future" anticipated discrimination discourages the candidates in the discriminated group from acquiring skills, thus validating the discrimination in equilibrium. In our theory, "past" discrimination disadvantages the discriminated sellers/workers in the inference formed by prospective buyers/employers about their good ratings, thus perpetuating discrimination.

Second, unlike CL, the existence of discriminatory equilibrium in our model does not rely on multiplicity of non-discriminatory equilibria. In CL, discrimination arises when (and because) different groups coordinate on different equilibria, making multiplicity of non-discriminatory equilibria a necessary and sufficient condition for discrimination. As explained in Section 4, in our model, a discriminatory equilibrium can exist even if there is a unique non-discriminatory equilibrium.

Finally, while the non-discriminatory equilibrium is *stable* in CL, it is often *unstable* in our model: as shown in Proposition 4 and Theorem 2, when a discriminatory equilibrium exists, the non-discriminatory equilibrium can be unstable, while there always exists a stable discriminatory equilibrium. This means that discrimination is a more robust prediction in our model than in CL.

It is worth noting that, despite the above differences, the effect recognized in our theory is consistent with, and further reinforced by, the force that Coate and Loury identified. Namely, the payoff gap between H and L types is higher for group 1 than for group 2 in the asymmetric equilibrium, suggesting that the incentive for becoming type H will be higher for the former group if indeed the type is endogenous, as envisioned by Coate and Loury (1993a).

References

Arrow, Kenneth, "The theory of discrimination," *Discrimination in Labor Markets*, 1973, 3 (10), 3–33. 4, 5

- Bartoš, Vojtěch, Michal Bauer, Julie Chytilová, and Filip Matějka**, “Attention discrimination: Theory and field experiments with monitoring information acquisition,” *American Economic Review*, 2016, 106 (6), 1437–75. 4
- Becker, Gary S**, *The Economics of Discrimination*, The University of Chicago Press, 1957. 4
- Bertrand, Marianne and Sendhil Mullainathan**, “Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination,” *American economic review*, 2004, 94 (4), 991–1013. 4
- Bohren, J Aislinn, Alex Imas, and Michael Rosenberg**, “The dynamics of discrimination: Theory and evidence,” *American Economic Review*, 2019, 109 (10), 3395–3436. 2, 5
- Coate, Stephen and Glenn Loury**, “Antidiscrimination enforcement and the problem of patronization,” *The American Economic Review*, 1993, 83 (2), 92–98. 5, 19, 20
- and —, “Will affirmative-action policies eliminate negative stereotypes?,” *The American Economic Review*, 1993, pp. 1220–1240. 5
- Corbett-Davies, Sam, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq**, “Algorithmic decision making and the cost of fairness,” in “Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining” ACM 2017, pp. 797–806. 5
- Cornell, Bradford and Ivo Welch**, “Culture, information, and screening discrimination,” *Journal of political Economy*, 1996, 104 (3), 542–571. 5
- Cui, Ruomeng, Jun Li, and Dennis J Zhang**, “Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on Airbnb,” *Management Science*, 2019. 2
- Dwork, Cynthia, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel**, “Fairness through awareness,” in “Proceedings of the 3rd innovations in theoretical computer science conference” ACM 2012, pp. 214–226. 5
- Edelman, Benjamin, Michael Luca, and Dan Svirsky**, “Racial discrimination in the sharing economy: Evidence from a field experiment,” *American Economic Journal: Applied Economics*, 2017, 9 (2), 1–22. 1
- Fang, Hanming and Andrea Moro**, “Theories of statistical discrimination and affirmative action: A survey,” in “Handbook of social economics,” Vol. 1, Elsevier, 2011, pp. 133–200. 5

Ge, Yanbo, Christopher R Knittel, Don MacKenzie, and Stephen Zuepf, “Racial and gender discrimination in transportation network companies,” in “Technical Report” National Bureau of Economic Research 2016. 2

Hannák, Anikó, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson, “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” ACM 2017, pp. 1914–1933. 2

Kearns, Michael, Seth Neel, Aaron Roth, and Zhiwei Steven Wu, “Preventing fairness gerrymandering: Auditing and learning for subgroup fairness,” in “International Conference on Machine Learning” 2018, pp. 2569–2577. 5

Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Cass R Sunstein, “Discrimination in the Age of Algorithms,” *Journal of Legal Analysis*, 2018, 10. 5

Mailath, George J, Larry Samuelson, and Avner Shaked, “Endogenous inequality in integrated labor markets with two-sided search,” *American Economic Review*, 2000, 90 (1), 46–72. 5

Monachou, Faidra Georgia and Itai Ashlagi, “Discrimination in Online Markets: Effects of Social Bias on Learning from Reviews and Policy Design,” in “Advances in Neural Information Processing Systems” 2019, pp. 2142–2152. 5

Moro, Andrea and Peter Norman, “Affirmative Action in a Competitive Economy,” *Journal of Public Economics*, 2003, 87, 567–594. 5

Norman, Peter, “Statistical Discrimination and Efficiency,” *The Review of Economic Studies*, 2003, 70 (3), 615–627. 5

Phelps, Edmund, “The Statistical Theory of Racism and Sexism,” *American Economic Review*, 1972, 62, 659–661. 4, 5

Appendix: Omitted Proofs

Proof of Lemma 1. Arranging the equations for P_{LG} , P_{HB} , and P_{LB} , we get

$$P_{LG} = \frac{\delta}{\delta + \psi_G \alpha} P_{HG}, \quad P_{HB} = \frac{\delta}{\delta + \psi_B \alpha} P_{LB}, \quad \text{and} \quad P_{LB} = P_{HB} + \frac{\psi_G \alpha}{\delta} P_{LG}.$$

Combining the latter two equations yields

$$P_{LB} = \frac{\delta}{\delta + \psi_B \alpha} P_{LB} + \frac{\psi_G \alpha}{\delta} P_{LG} = \frac{\psi_G (\delta + \psi_B \alpha)}{\delta \psi_B} P_{LG}.$$

Since the total measure of workers is always equal to 1, we have

$$1 = P_{HG} + P_{LG} + P_{LB} + P_{HB} = P_{HG} \left[1 + \frac{\delta}{\delta + \psi_G \alpha} \left(1 + \frac{\psi_G (\delta + \psi_B \alpha)}{\delta \psi_B} \left(1 + \frac{\delta}{\delta + \psi_B \alpha} \right) \right) \right].$$

Arranging the terms, we get the expression for P_{HG} . From there, we can also find the expressions for P_{LG} , P_{LB} , and P_{HB} as well. The results for μ_G and μ_B are immediate from the solutions to P_{HG} , P_{LG} , P_{LB} , and P_{HB} . Q.E.D.

Proof of Lemma 2. Recall from Lemma 1 that

$$\mu_G = 1 - \frac{\delta}{2\delta + \psi_G \alpha} \text{ and } \mu_B = \frac{\delta}{2\delta + \psi_B \alpha}.$$

Therefore,

$$u_G(\lambda) = \phi(\lambda)(\mu_G y_H + (1 - \mu_G)y_L - w) = \frac{\psi(\lambda)}{\lambda} \left((u_H - w) - (u_H - u_L) \frac{\delta}{2\delta + \psi(\lambda)\alpha} \right),$$

and

$$u_B(\lambda) = \phi(\lambda)(\mu_B y_H + (1 - \mu_B)y_L - w) = \frac{\psi(\lambda)}{\lambda} \left(\frac{\delta}{2\delta + \psi_B \alpha} (y_H - y_L) + y_L - w \right).$$

The continuity of $u_j(\lambda)$ follows from the same property of $\phi(\lambda)$ and $\mu_j(\lambda)$.

If λ tends to 0, then μ_j approaches 1/2 for both $j = G, B$. The result that $\lim_{\lambda \rightarrow 0} u_j(\lambda) = \infty$ then follows from the fact that $(y_H + y_L)/2 - w > 0$ and $\lim_{\lambda \rightarrow 0} \phi(\lambda) = \infty$. For the case when λ tends to infinity, observe that

$$\mu_j(\lambda) = \phi(\lambda)(\mu_G y_H + (1 - \mu_G)y_L - w) \leq \phi(\lambda)(y_H - w).$$

The desired result is immediate because $\lim_{\lambda \rightarrow \infty} \phi(\lambda) = 0$.

The monotonicity of $u_B(\lambda)$ follows from the fact that both $\phi(\lambda)$ and μ_B are strictly decreasing in λ . For $u_G(\lambda)$, observe that

$$\begin{aligned} \frac{du_G(\lambda)}{d\lambda} &= \left(\frac{\psi'}{\lambda} - \frac{\psi}{\lambda^2} \right) \left(u_H - w - (u_H - u_L) \frac{\delta}{2\delta + \psi\alpha} \right) + \frac{\psi}{\lambda} (u_H - u_L) \frac{\delta \alpha \psi'}{(2\delta + \psi\alpha)^2} \\ &= \frac{1}{(2\delta + \psi\alpha)^2} \left(\left(\frac{\psi'}{\lambda} - \frac{\psi}{\lambda^2} \right) ((u_H - w)(2\delta + \psi\alpha)^2 - (u_H - u_L)\delta(2\delta + \psi\alpha)) + \frac{\psi}{\lambda} (u_H - u_L)\delta\alpha\psi' \right). \end{aligned}$$

When $\psi(\lambda) = \lambda^k$, $du_G(\lambda)/d\lambda$ has the same sign as

$$\begin{aligned} h(\lambda) &= -(1 - k)(u_H - w)(2\delta + \psi\alpha)^2 + (1 - k)(u_H - u_L)\delta(2\delta + \psi\alpha) + k(u_H - u_L)\delta\alpha\psi \\ &= -(1 - k)\alpha^2(u_H - w)\psi^2 + (-4(1 - k)(u_H - w) + u_H - u_L)\delta\alpha\psi \\ &\quad - 2(1 - k)\delta^2(u_H + u_L - 2w). \end{aligned}$$

$h(\lambda)$ is a quadratic equation of ψ , and its maximal value is equal to

$$\frac{(u_H - u_L - 4(1-k)(u_H - w))^2 \delta^2}{4(1-k)(u_H - w)} - 2(1-k)\delta^2(u_H + u_L - 2w),$$

which has the same sign as

$$\begin{aligned} & (u_H - u_L)^2 - 8(1-k)(u_H - u_L)(u_H - w) + 16(1-k)^2(u_H - w)^2 \\ & - 8(1-k)^2(u_H - w)(u_H + u_L - 2w) \\ = & 8(u_H - u_L)(u_H - w) \left((1-k)^2 - (1-k) + \frac{u_H - u_L}{8(u_H - w)} \right). \end{aligned}$$

Let $\underline{k} \in (0, 1)$ be the unique value that equates the maximal value of $h(\lambda)$ to 0:

$$1 - \underline{k} = \frac{1 - \sqrt{1 - \frac{u_H - u_L}{2(u_H - w)}}}{2} \Rightarrow \underline{k} = \frac{1 + \sqrt{1 - \frac{u_H - u_L}{2(u_H - w)}}}{2}.$$

If $k \leq \underline{k}$, then $h(\lambda) \leq 0$ for any λ , which implies that $u_G(\lambda)$ is monotone (decreasing). If $k > \underline{k}$, then $h(\lambda) \leq 0$ has two solutions, $\underline{\lambda}_G$ and $\bar{\lambda}_G$. Then, $u_G(\lambda)$ is strictly increasing if and only if $\lambda \in (\underline{\lambda}_G, \bar{\lambda}_G)$. Q.E.D.

Proof of Theorem 1. Suppose that $u_G(\lambda)$ is not monotone, that is, $k \in (\underline{k}, 1)$. Then, as shown in Lemma 2, there exist $\underline{\lambda}_G (> 0)$ and $\bar{\lambda}_G (> \underline{\lambda}_G)$ such that $u_G(\lambda)$ is strictly decreasing if and only if $\lambda \in (\underline{\lambda}_G, \bar{\lambda}_G)$. Let $\bar{\lambda}_B$ and $\underline{\lambda}_B$ be the values such that $\bar{\lambda}_B = \lambda_B^{BI}(\underline{\lambda}_G)$ (i.e., $u_B(\bar{\lambda}_B) = u_G(\underline{\lambda}_G)$) and $\underline{\lambda}_B = \lambda_B^{BI}(\bar{\lambda}_G)$ (i.e., $u_B(\underline{\lambda}_B) = u_G(\bar{\lambda}_G)$), respectively (see Figure 3). Then, for each $\lambda_B \in [\underline{\lambda}_B, \bar{\lambda}_B]$, there exist $h_1(\lambda_B)$, $h_2(\lambda_B)$, and $h_3(\lambda_B)$ such that $h_1(\lambda_B) \leq h_2(\lambda_B) \leq h_3(\lambda_B)$, with at least one inequality holding strictly if $\lambda_B = \underline{\lambda}_B$ or $\lambda_B = \bar{\lambda}_B$ and both inequality holding strictly otherwise, and $u_B(\lambda_B) = u_G(h_m(\lambda_B))$ for all $m = 1, 2, 3$. By construction, all $h_i(\lambda)$'s are continuous over $(\underline{\lambda}_B, \bar{\lambda}_B)$. In addition, $h_1(\lambda_B)$ and $h_3(\lambda_B)$ are strictly increasing, while $h_2(\lambda_B)$ is strictly decreasing (again, see Figure 3).

By the explanation given just before Theorem 1, every discriminatory equilibrium can be represented (produced) by two distinct points, λ_G and λ'_G , that correspond to the same value of $\lambda_B \in [\underline{\lambda}_B, \bar{\lambda}_B]$. Combining this with the fact that for each $\lambda_B \in [\underline{\lambda}_B, \bar{\lambda}_B]$, there are three such values, $h_1(\lambda_B)$, $h_2(\lambda_B)$, and $h_3(\lambda_B)$, it follows that each λ_B yields three possible values of Q 's that lead to a discriminatory equilibrium. In other words, for any pair (m, m') such that $m, m' = 1, 2, 3$ and $m \neq m'$, it is an equilibrium that one group trades according to $(h_m(\lambda_B), \lambda_B)$ and the other group trades according to $(h_{m'}(\lambda_B), \lambda_B)$ if the total measure of firms is given by

$$Q = \Theta(h_m(\lambda_B), \lambda_B) + \Theta(h_{m'}(\lambda_B), \lambda_B),$$

where $\Theta(\lambda_G, \lambda_B)$ is the measure of buyers necessary to support a steady-state equilibrium in which one group of sellers trade according to (λ_G, λ_B) , that is,

$$\Theta(\lambda_G, \lambda_B) \equiv \frac{\lambda_G \psi(\lambda_B)(2\delta + \alpha\psi(\lambda_G))}{2(\delta\psi(\lambda_G) + \delta\psi_B + \alpha\psi(\lambda_G)\psi(\lambda_B))} + \frac{\lambda_B \psi(\lambda_G)(2\delta + \alpha\psi(\lambda_B))}{2(\delta\psi(\lambda_B) + \delta\psi(\lambda_B) + \alpha\psi(\lambda_G)\psi(\lambda_B))}.$$

In order to show the interval structure of Q 's, for each $m = 1, 2, 3$, let I_m denote the set of all Q 's that are associated with $h_m(\lambda_B)$ and $h_{m+1}(\lambda_B)$. Formally, define

$$I_m = \{\Theta(h_m(\lambda_B), \lambda_B) + \Theta(h_{m+1}(\lambda_B), \lambda_B) | \lambda_B \in [\underline{\lambda}_B, \bar{\lambda}_B]\}.$$

Since $\Theta(\lambda_G, \lambda_B)$ is continuous, each I_m is an interval. Furthermore, $I \equiv \bigcup_m I_m$ is also an interval, because

$$\lim_{\lambda_B \rightarrow \underline{\lambda}_B} h_2(\lambda_B) = \lim_{\lambda_B \rightarrow \underline{\lambda}_B} h_3(\lambda_B) \text{ and } \lim_{\lambda_B \rightarrow \bar{\lambda}_B} h_1(\lambda_B) = \lim_{\lambda_B \rightarrow \bar{\lambda}_B} h_2(\lambda_B).$$

Q.E.D.