Non-Price Discrimination by a Prejudiced Platform*

Ming Gao

Tsinghua University

Email: gaom@sem.tsinghua.edu.cn

Travis Ng

The Chinese University of Hong Kong

Email: TravisNg@cuhk.edu.hk

August 22, 2017

Abstract

Recent lawsuits and anecdotal evidence suggest that some platforms discriminate against certain users through non-price practices, discouraging their participation without directly increasing revenue. We show that a monopolist two-sided platform with a prejudice against certain users - modeled as more costly to serve - chooses to discriminate only if the cost savings from reducing such users' participation outweigh the network benefits they create. Surprisingly, user surpluses may increase under discrimination because the platform often voluntarily lowers price(s) - sometimes on both sides - to attract other users. Therefore, tightening anti-discrimination policies for platforms can increase price and decrease welfare.

Key Words: discrimination; prejudice; regulation; policies on platforms; two-sided market; non-price strategies.

JEL **Classification**: K20, L50, J14, J15, J16

*Part of this article has previously been circulated under the title "The Economics of App Review." We thank Bruno Jullien, Martin Osborne, Andrew Rhodes, Alexander White, Heiko Karle, Chiu Yu Ko, Jie Zheng, Junjie Zhou, Lin Zhu and the seminar participants at the Toulouse 2016 E-commerce, Digital Economy and Delivery Services Conference, the 2017 International Industrial Organization Conference at Boston, the Nanjing University 2016 International Conference on Innovation and Industrial Economics, Shanghai Jiao Tong University and Shanghai University of Finance and Economics. Ming Gao acknowledges the financial assistance of Project 20151080389 supported by the Tsinghua University Initiative Scientific Research Program and Project 71203113 supported by the National Natural Science Foundation of China.

1 Introduction

What are efficient policies for tackling discrimination by platforms? Following the pioneering work by Becker (1957), discrimination has accumulated much research attention in the economics field. Why would discrimination by platforms deserve yet another study? There are four reasons. First, platforms are becoming more important in our daily lives (Hagiu and Rothman, 2016). Second, a platform market can tip (Hossain, Minor, and Morgan, 2011 and Hossain and Morgan, 2013), inhibiting market competition from curbing discrimination à la Becker (1957). Third, there has been an increasing number of court cases and allegations of discrimination by platforms. Fourth, platforms are now capable of collecting much more data from users; if a platform does discriminate against certain users, its data allow it to better identify these users and discriminate against them in more sophisticated ways.

We focus on non-price discrimination because outright price discrimination by platforms is both much easier to detect - and hence to prosecute if illegal - and much more thoroughly researched in the economics literature. In contrast, non-price discriminatory practices are usually subtler and do not take a specific form, and therefore can elude regulatory scrutiny. For instance, they may involve prolonged application processes, additional documentation requirements, or reduced responsiveness in service. (Section 1.1 provides more examples.) Discrimination makes victims suffer, in turn discouraging them from joining the platforms. Unlike price discrimination, however, non-price discriminatory actions do not merely transfer value from the users to the platforms. Little is known about the factors that incentivize a platform to discriminate against certain users in non-price ways and about the ensuing welfare implications.

Suppose some users allege that a platform has discriminated against them. A regulator or a court first investigates whether such discriminatory practices exist. If they do, the platform must give an explanation. Many potential explanations boil down to the following: “We are in the business of ensuring quality matches for both sides of our users.”

Such a claim is a subtle way to say that the platform discriminates against certain users because *its users* on the other side want it to do so.[1]

A competing story, of course, is that the platform discriminates against certain users because *it* wants to.[2] Perhaps most would think that the platform is guilty in this case. It is even worse if the platform is the dominant one in a tipped market, because the users will have nowhere else to go. If we tick both the "discrimination" and "tipped market" checkboxes, the conclusion that regulators should step in and sanction the platform seems irresistible.

Our analysis, however, does not lead us to this conclusion. In our model, the platform discriminates because it wants to, not because its users want it to. We find that tightening anti-discrimination policies for platforms can lead to price rises and welfare losses. These consequences occur even if the platform is a monopoly in a tipped market. Our argument does not rely on the potentially high costs of administering anti-discrimination rules (such as the costs of investigating whether a discriminatory practice exists, and whether the platform or the users want the discrimination). In fact, assuming zero enforcement costs and no market competition, the counterintuitive result that it is efficient to leave the platform alone still holds.

Specifically, we study a two-sided market where users on each side create positive network externalities through cross-side interactions. A monopolist platform has a prejudice against a subset of users on one side, because they are more costly to serve. All users are, however, unprejudiced and view everyone on the opposite side as equally valuable. Although the platform cannot price-discriminate against high-cost users, it can choose to discourage them from joining it through non-price practices.

We find several results. First, despite its prejudice, the platform does not always

[1]For instance, a lending platform may discriminate against borrowers who are likely to default from joining the platform, because its lenders would prefer not to lend money to them. A ride-sharing platform may discourage certain types of drivers from joining the network because riders would feel unsafe in their cars.

[2]It is certainly difficult to find out which story is the truth, as we will see in the real cases in section 1.1.

choose to discriminate. It has an incentive to discriminate only if the cost savings from reducing high-cost users' participation outweigh the network benefits they create. Second, and perhaps most surprisingly, discrimination may increase user surpluses, as when the platform discriminates it often voluntarily lowers price(s) - sometimes on both sides - to retain sufficient levels of user participation on both sides and to profit from the network benefits created. Third, when the platform chooses to discriminate, and when this practice increases users' surplus, total social welfare clearly increases. In such cases, anti-discrimination policies lead to higher prices and welfare losses.

The driving force of these results is the two-sidedness of a platform, whose business relies on attracting users from both sides. Therefore, even when it is optimal to discourage a subset of users on one side, the platform still needs to appeal to the other users on both sides. This need incentivizes the platform to cut prices, which increases user surpluses. This mechanism is absent in one-sided markets.

There are certainly alternative policy objectives concerning discrimination, such as fairness, morality, friendship, and other value judgments. We do not address these concerns in this article. Our model, however, enables us to explore other alternative policies. We find that if prohibiting discrimination is either impossible or expensive to administer, mandatory disclosure requirements and/or user privacy protection policies that are well crafted can reduce discrimination.

Two policy implications follow. First, it is not always socially efficient to prohibit discrimination even if a two-sided market has tipped. As a result, it is hard to make a case for stricter anti-discrimination policies for platforms. Second, although antitrust regulators may be concerned about platform markets for a number of reasons, the fear that a dominant platform can better discriminate does not appear to be a major cause for alarm.

1.1 Recent allegations of non-price discrimination by platforms

Non-price discrimination allegations against platforms abound. Some borrowers have complained that Lending Club has discriminated against them by requesting them to submit more documents than others are required to submit.[3] Some app developers have complained that Apple App Store either rejected their apps or approved them with substantial delays, without giving clear reasons. Shopping malls connecting businesses and shoppers have also been accused of discrimination by both sides. In *Radek v. Henderson Development (Canada)*, an aboriginal woman successfully sued a mall for discriminating against her.[4] A Singaporean mall issued an open apology for denying the rental application of a Malaysian businesswoman, which appeared to be due to racial bias.[5] In a published article, Match.com requested users not to state "young for my age" in their online dating profile. The article irritated some older women.[6] An atheist has alleged that eHarmony has discriminated against him by not turning up any matches for him.[7] In *Fair Housing Council of San Fernando Valley v. Roommates.com*, Roommates.com was convicted of extracting information from potential customers as a condition of accepting them as clients.[8] Some customers find it offensive for Roommates.com to ask about gender, sexual orientation, number of children, and whether the children live with the customer.

Table 1 summarizes these cases. In each case, we briefly describe the allegation. We also highlight the alleged prejudice of and possible counter arguments by the platform.

[3]The two links are here: https://lending-club.pissedconsumer.com/scam-and-discriminate-20160422834627.html and http://www.ripoffreport.com/r/lending-club/internet/lending-club-lendingclubcom-discrimination-in-the-utmost-against-women-scam-to-get-your-78008

[4]The Canadian case was *Radek v. Henderson Development (Canada) Ltd. (No. 3) (2005), 52 C.H.R.R. D/430, 2005 BCHRT 302*. http://www.cdn-hr-reporter.ca/hr_topics/systemic-discrimination/shopping-mall-discriminates-against-aboriginal-people.

[5]A Singaporean mall has been alleged of discriminating against a Malaysian businesswoman. http://www.straitstimes.com/singapore/tampines-1-says-sorry-after-customer-complains-of-racial-discrimination-in-e-mail-exchange.

[6]The article can be seen on Match.com http://www.match.com/magazine/article/6793/Over-50-And-Online/.

[7]The article can be seen here: http://www.patheos.com/blogs/friendlyatheist/2007/06/16/eharmony-saying-no-to-atheists/.

[8]*Fair Housing Council of San Fernando Valley v. Roommates.com, LLC*, 521 F.3d 1157 (9th Cir. 2008)

A few observations emerge from these cases. First, it is not impossible for a platform to give preferential treatment to the kind of users with whom the platform "likes" to associate. Second, it is possible that a platform may discriminate purely for business reasons, with no prejudice. Third, it appears quite difficult to determine what the true motives are behind these allegedly discriminatory practices.

Table 1: Recent Allegations of Non-Price Discrimination by Platforms

	Platform	Description	Suspected Prejudice	Possible Counter Argument
1	LendingClub.com	Requiring more information than usual for loan approval and requesting unnecessary changes to personal profile	Discriminated against female loan applicants	Additional information from certain users lowers the risk of fake identities and facilitates internal procedures, which ultimately improve applicants' reliability and creditworthiness.
2	Apple Store	Rejecting some apps or approving them with substantial delays, without clear reasons	Discriminated against apps that Apple, rather than users, does not like to appear	The apps are not well built or have insufficient security for the intended functionality. They either appear incomplete or pose threats to users.
3	Canadian shopping mall (court case)	Telling an aboriginal disabled person to leave the mall	Discriminated against an aboriginal disabled person	Aboriginal and disabled people usually require more services and support from the mall, which takes away dis-proportionately what would otherwise be available for other customers. The mall only invites suspicious people of this background to leave.
4	Singaporean shopping mall	Rejecting rental application for a mall fair	Discriminated against Malay businesswomen	The fair targets Chinese customers only. Space also runs out rapidly and becomes available only later in the year.
5	Match.com	Discouraging users from stating "young for my age" in their online profiles	Discriminated against seniors	If almost all users above 50 years old claim this, then some of them must be making an inaccurate statement. Lacking the ability to verify the statement, Match.com universally discourages the usage of such phrase to avoid dating targets feeling cheated.
6	eHarmony	Returning no match for a user who has no religion	Discriminated against atheists	The pool at the time did not have a good suggestion for that particular atheist's profile. eHarmony serves people with or without religious beliefs.
7	Roommate.com (court case)	Requiring information beyond the basics, such as gender, sexual orientation, and family status, before one can search or post housing opportunities	Discriminated against customers based on conditions irrelevant to housing	The information helps both sides to identify desirable roommates. As such, it improves the value both parties gain from Roommate.com and hence benefits society.

2 Literature Review

Discrimination among platform users. While our paper studies discrimination *by a platform*, many notable papers have studied a related kind of discrimination - discrimination *among users*. A growing and exciting area of research focuses on discrimination among users of *digital platforms*. Does discrimination persist in online platforms? Do online platforms even exacerbate discrimination?

Fisman and Luca (2016) highlight some papers addressing these questions and draw many insightful managerial implications.[9] Edelman, Luca, and Svirsky (2017) also suggest several managerial implications faced by online platforms; they conduct experiments via Airbnb and find that applications from guests with distinctively African-American names are significantly less likely to be accepted by hosts relative to control groups with distinctively white names. On the other side of the Airbnb market, Edelman and Luca (2014) find that African-American hosts ask and get significantly lower prices than otherwise similar white hosts. Pope and Sydnor (2011) look into peer-to-peer lending at Prosper.com and find that loan listings with blacks in the attached picture are significantly less likely to receive funding than those of whites with similar credit profiles. Duarte, Siegel, and Young (2012) also examine Prosper.com; they rate borrowers trustworthiness only by viewing their photos and find that those who look trustworthy are significantly more likely to have their loan requests granted. However, they are also more likely to eventually repay their loans. Experimenting with Uber and Lyft, Ge, Knittel, MacKenzie, and Zoepf (2016) find that African-American passengers suffer from longer waiting times and their orders are more likely to be canceled by drivers. Doleac and Stein (2013) examine Craigslist and find that the same iPod receives significantly fewer responses from potential buyers if it is held by a black hand than a white hand.

[9] Platforms such as Airbnb and Uber have published formal statements and policies addressing these types of discrimination. For Airbnb, please read: http://www.inc.com/tess-townsend/airbnb-hires-eric-holder.html. For Uber, please read: http://flavorwire.com/581580/ubers-evolving-relationship-with-discrimination.

Discrimination by platforms. There is a growing interest in discrimination by digital platforms, especially in addressing the question of whether their use of big data and algorithms exacerbates discrimination.[10] For instance, would a search engine show more expensive items to Internet users living in certain neighborhoods because its algorithms suggest that they are more likely to buy expensive items? These Internet users may feel discriminated against by the platform for making it more difficult to locate bargains online.

Computer scientists have developed various tools and methodologies that help regulators and researchers to examine black-box algorithms to detect discrimination (Sandvig, Hamilton, Karahalios and Langbort, 2014). Sweeney (2013) finds that "Googling" for common African-American names is significantly more likely to result in ads offering criminal background checks than "Googling" for names common among whites. Datta, Tschantz, and Datta (2015) also find that Google does not show as many ads for high-paying jobs if the Google profile's setting is female rather than male.

Discrimination can also be carried out by platforms that are non-digital. We are not aware of any research that clearly delineates whether discrimination by digital platforms is taste-based or statistical. We are also not aware of any economic research that analyzes whether regulators should step in, and what policies they should adopt. While computer scientists have been accumulating interesting research, little is known about the economics of this topic. Our paper fills this gap by offering a theoretical framework within which we can study a platform's incentive to discriminate against some users because of prejudice, and examine the welfare implications of non-price discrimination.

[10]The media has also reported this concern. See Miller, Claire Cain. 2015. "When Algorithms Discriminate." *The New York Times.* https://www.nytimes.com/2015/07/10/upshot/when-algorithms-discriminate.html. See also Kirchner, Lauren. 2015. "When Discrimination is Baked Into Algorithms." *The Atlantic.* http://www.theatlantic.com/business/archive/2015/09/discrimination-algorithms-disparate-impact/403969/

3 Model

3.1 Set-up

Our model builds on the monopoly model of Armstrong (2006).[11] A platform facilitates interactions between buyers (on side 1) and sellers (on side 2). Each side has a total mass normalized to 1. The platform can only charge the same price to all users on each side. Let p_i be the price it charges to a user on side i.

A user cares about the number of users on the other side. If the platform attracts n_1 buyers and n_2 sellers, then the utilities of a buyer and a seller are

$$u_1 = \alpha_1 n_2 - p_1; u_2 = \alpha_2 n_1 - p_2, \tag{1}$$

respectively, where $\alpha_i (\geqslant 0)$ measures the network externality that each side-i user enjoys from interacting with a user on the other side. Parameter α_i is exogenous and known to everyone.[12] Denote $\alpha \equiv \alpha_1 + \alpha_2$ as the total network externalities between a buyer-seller pair.

This model has three stages.

- **Stage 1**: The platform posts the prices to both sides, p_1 and p_2, and decides whether to introduce discrimination.

- **Stage 2**: Every user observes the platform's actions and decides whether to join it.

[11] Rochet and Tirole (2003 and 2006) and Armstrong (2006) offer the canonical models of two-sided markets. In some two-sided markets, platforms use quality as a criterion to exclude some users. For instance, some night clubs do not admit patrons wearing sandals or jeans. Hagiu (2009) uses a model to study platform exclusion using quality as a criterion.

[12] We do not study the possibility that a user cares about the number of users on his own side. We also do not model any user's pricing decisions, but focus on the platform's incentives. If the seller's pricing decisions are explicitly modeled, in equilibrium all sellers would charge the same price to buyers, because no two sellers would be different in the eyes of buyers. One can interpret parameter α_2 as the network externalities minus the endogenously determined common price that all sellers charge a buyer. Such an interpretation does not change our results or the intuition. In reality, some platforms, such as Uber, do not allow users to set their own prices, while others do, such as Airbnb.

- **Stage 3**: Those users who join interact and realize their utilities; the platform realizes its profit.

Following Armstrong (2006), we specify the number of buyers who join as a function of their utility. Given utility u_1, the number of buyers is

$$n_1 = \phi_1(u_1),$$

for some increasing function ϕ_1 known to everyone.[13] Assume it costs the platform $f_1 \geqslant 0$ to serve each buyer on side 1. Assume ϕ_1 is twice differentiable.

3.2 Two types of sellers

In the eyes of any buyer, all sellers are equal. However, they are not equally costly for the platform to serve. The cost to serve low-cost (i.e. type-L) sellers, comprising a fraction λ of the totality, is f_L whereas the cost to serve high-cost (i.e. type-H) sellers, comprising the remaining $(1-\lambda)$ fraction, is f_H. Assume $0 \leqslant f_L \leqslant f_H$.

Denote n_k as the number of type-k sellers joining the platform ($k \in \{L, H\}$). Without discrimination, the utility a seller gets from joining the platform determines these numbers as follows:

type-L sellers:	$n_L = \lambda \cdot \phi_2(u_2),$
type-H sellers (without discrimination):	$n_H = (1-\lambda) \cdot \phi_2(u_2),$

where ϕ_2 is increasing and known to everyone. Assume ϕ_2 is twice differentiable. The total number of sellers on the platform becomes

$$n_2 = n_L + n_H.$$

[13] In Section 5, we specify a functional form of ϕ_1 in which the users derive private benefits from joining the platform beyond network externalities; the private benefits follow exogenous distributions known to the platform.

Denote Δf as the additional cost that a type-H seller brings to the platform relative to a type-L seller:

$$\Delta f \equiv f_H - f_L.$$

The parameter Δf can have several interpretations.

- **Prejudice (non-pecuniary):** The platform owner can simply dislike type-H sellers, even though the operational cost to serve the different sellers is the same. When interacting with a type-H seller, the platform owner suffers a psychological cost of Δf. This interpretation makes Δf a measure of the prejudice.[14]

- **Appeal (non-pecuniary):** Equivalently, the platform owner may simply favor type-L sellers, even though the operational cost to serve the different sellers is the same. When interacting with a type-L seller, the platform owner gains a psychological benefit of Δf. This interpretation makes Δf a measure of the appeal of type-L sellers to the platform.[15]

- **Additional operation cost (pecuniary):** Alternatively, Δf can just be the additional operational cost of serving a type-H seller.

Our model addresses discrimination due to a platform's prejudice, and therefore we adopt the first interpretation. However, the following analyses would not change even if we adopted the other two interpretations.

3.3 Non-price discrimination

It is more costly to serve type-H sellers, which may give the platform an incentive to discourage their joining. As we rule out price discrimination, the platform can only do so

[14]Note that the platform monopoly does not believe that type-H sellers are less valuable in the eyes of the buyers either. Neither do type-H sellers have a chance to under-invest. Therefore, the kind of discrimination in our model is not driven by the possibility that the beliefs of the platform can be self-fulfilling, as in Phelps (1972) and Arrow (1973).

[15]For example, a shopping mall owner is an environmentalist. While there is no difference in the operational costs of serving different shops, those that have zero carbon footprints appeal to the owner more than those that have high carbon footprints.

through certain non-price practices.[16]

Formally, we allow the platform to impose a disutility of $D \in [0, \bar{D}]$ on each type-H seller who joins the platform.[17] The magnitude of D, therefore, measures the extent of discrimination.

Under discrimination, each type-L seller gets u_2 from joining the platform but each type-H seller only gets $(u_2 - D)$ from joining. The corresponding number of type-H sellers who join the platform becomes

$$\text{type-H sellers (with discrimination): } n_H = (1 - \lambda) \cdot \phi_2(u_2 - D).$$

The platform's profit is then given by

$$\pi = (p_1 - f_1)n_1 + (p_2 - f_L)n_L + (p_2 - f_H)n_H, \tag{2}$$

4 Equilibrium analysis

When does the platform act on its prejudice and discriminate against type-H sellers? Only when doing so increases profit (including all benefits and costs, psychological or otherwise).[18] We consider the utilities that the platform offers to users (u_1 and u_2) and the extent of discrimination (D) as the choice variables.

Rewriting the profit function in (2) as a function of u_1, u_2 and D and substituting

[16]The non-price practices that are allegedly discriminatory in Section 1.1 motivate such a modeling approach. A platform can selectively expedite services (such as help desk support) for some users but delay those for others. A mall can order its security guards to ask minority customers to show their ID cards before entering the mall. A peer-to-peer lending website can ask a certain group of users for more documents than are normally requested from other users. More such non-price discriminatory practices are described in Section 1.1.

[17]There is no discrimination when $D = 0$. One can interpret the upper bound of D as given by the law that defines the extent of harassment, insult, or annoyance that is illegal.

[18]Alchian and Kessel (1962) and Becker (1962) give reasons to support the notion that a monopoly is not profit maximizing but utility maximizing. As we interpret Δf as a psychological cost due to prejudice, the goal of a platform in our model is consistent with utility maximizing. However, if Δf is interpreted as simply an operational cost, the platform in our model is just maximizing profit. In the main text, we still say that the platform maximizes profit to avoid confusion.

$f_L = f_H - \Delta f$, we get:

$$\pi(u_1, u_2, D) \equiv (p_1 - f_1)n_1 + (p_2 - f_H + \Delta f)n_L + (p_2 - f_H)n_H. \quad (3)$$

By (1) and the previous definitions of buyer and seller numbers, the demand from both sides and the prices are also functions of u_1, u_2 and D as follows:

$$n_1(u_1) = \phi_1(u_1); \quad (4a)$$

$$n_L(u_2) = \lambda\phi_2(u_2); \quad (4b)$$

$$n_H(u_2, D) = (1-\lambda)\phi_2(u_2 - D); \quad (4c)$$

$$n_2(u_2, D) = n_L(u_2) + n_H(u_2, D); \quad (4d)$$

$$p_1(u_1, u_2, D) = \alpha_1 n_2(u_2, D) - u_1; \quad (4e)$$

$$p_2(u_1, u_2) = \alpha_2 n_1(u_1) - u_2. \quad (4f)$$

The first-order condition of (3) with respect to u_1 gives

$$\underbrace{(\alpha n_2 - u_1 - f_1)}_{\text{economic profit from each buyer}} \cdot \underbrace{\frac{dn_1}{du_1}}_{\text{rise in } n_1} = \underbrace{n_1}_{\text{loss in revenue}} \quad (5)$$

where $\frac{dn_1}{du_1} = \phi_1'(u_1)$.

As the interaction of each buyer-seller pair creates total network externalities of α, each buyer brings to the platform a total "economic profit" of $(\alpha n_2 - u_1 - f_1)$. Offering each buyer one more unit of utility increases the number of buyers by $\frac{dn_1}{du_1}$. Therefore, the left-hand side of (5) represents the platform's marginal benefit of increasing u_1. The platform sets u_1 optimally when this marginal benefit equals the right-hand side, which is the marginal cost in terms of the total loss in revenue, n_1.

The first-order condition of (3) with respect to u_2 is

$$\underbrace{(\alpha n_1 - u_2 - f_H)}_{\substack{\text{economic profit from} \\ \text{each type-H seller}}} \cdot \underbrace{\frac{\partial n_2}{\partial u_2}}_{\substack{\text{rise in \# of} \\ \text{sellers}}} + \underbrace{\Delta f}_{\substack{\text{cost saved by} \\ \text{each type-L}}} \cdot \underbrace{\frac{dn_L}{du_2}}_{\substack{\text{rise in \#} \\ \text{of type-L}}} = \underbrace{n_2}_{\substack{\text{loss in} \\ \text{revenue}}} \tag{6}$$

where $\frac{\partial n_2}{\partial u_2} = \lambda \phi_2'(u_2) + (1-\lambda)\phi_2'(u_2 - D)$ and $\frac{dn_L}{du_2} = \lambda \phi_2'(u_2)$.

Each seller creates a total value of αn_1 for the platform, and therefore the economic profit that a type-H seller generates is $(\alpha n_1 - u_2 - f_H)$. Increasing u_2 by one unit, the platform can attract a total of $\frac{\partial n_2}{\partial u_2}$ new sellers, and the first term on the left-hand side of (6) represents the associated increase in economic profit if all new sellers were of type-H. However, $\frac{dn_L}{du_2}$ of these new sellers are of type-L, for each of whom the platform saves a cost of Δf; the total cost savings are equal to $\Delta f \cdot \frac{dn_L}{du_2}$. These cost savings also add to the platform's benefits from increasing u_2, and are represented by the second term on the left-hand side. On the right-hand side, the marginal loss in revenue due to the additional utility offered to all sellers equals their total number n_2. The optimal u_2 makes the two sides equal.

Given some $D \in [0, \bar{D}]$, denote the pair of u_1 and u_2 that satisfies (5) and (6) as

$$(u_1^*(D), u_2^*(D)) \equiv \arg\max_{(u_1, u_2)} \pi(u_1, u_2, D). \tag{7}$$

Denote the maximized profit as a function of D alone as

$$\Pi(D) \equiv \pi(u_1^*(D), u_2^*(D), D). \tag{8}$$

Assume the profit function (3) is well-behaved such that $u_1^*(D)$ and $u_2^*(D)$ are differentiable for $D \in [0, \bar{D}]$, which implies that $\Pi(D)$ is also differentiable. We have the following useful property of $\Pi(D)$:

Lemma 1 $\Pi(D)$ *is quasiconvex on* $[0, \bar{D}]$.

(All omitted proofs are provided in the Appendix.) Lemma 1 implies that the platform's maximized profit as a function of D must have no peak.[19] Therefore, no intermediate level of discrimination $D \in (0, \bar{D})$ is optimal, which gives the following result.

Lemma 2 (Binary Discrimination Choice) *If the platform can choose the extent of discrimination* $D \in [0, \bar{D}]$, *it will either choose not to discriminate at all* $(D = 0)$, *or choose to fully discriminate* $(D = \bar{D})$.

As the maximized profit without discrimination is equal to $\Pi(0)$, we can use $\Pi'(0)$ and Lemma 1 to determine the platform's incentive to introduce discrimination. For $D \in [0, \bar{D}]$, apply an envelope argument for (8), and we have

$$\Pi'(D) = \frac{\partial}{\partial D}\pi(u_1^*(D), u_2^*(D), D) \tag{9}$$

Therefore, we can find $\Pi'(0)$ by evaluating $\frac{\partial \pi}{\partial D}$ at the optimal utilities without discrimination.

4.1 Optimal pricing without discrimination

Without discrimination, Proposition 1 states the optimal prices.

Proposition 1 (Optimal Pricing without Discrimination) *Without discrimination, the optimal prices for both sides,* (p_1^0, p_2^0), *are given by*

$$\begin{cases} p_1^0 = f_1 - \alpha_2 n_2 + \frac{\phi_1(u_1^0)}{\phi_1'(u_1^0)}, \\ p_2^0 = (f_H - \lambda \Delta f) - \alpha_1 n_1 + \frac{\phi_2(u_2^0)}{\phi_2'(u_2^0)}. \end{cases} \tag{10}$$

[19] The property in Lemma 1 results from the fact that non-price discrimination represents pure value destruction - no market participant directly benefits from it.

The optimal utilities that the platform offers to two sides, (u_1^0, u_2^0)*, are given by*

$$\begin{cases} \phi_1(u_1^0) = \phi_1'(u_1^0)(\alpha\phi_2(u_2^0) - u_1^0 - f_1) \\ \phi_2(u_2^0) = \phi_2'(u_2^0)(\alpha\phi_1(u_1^0) - u_2^0 - (f_H - \lambda\Delta f)) \end{cases} \tag{11}$$

As expected, the platform's optimal price for each side in (10) marks up on the users' cost to the platform, after deducting the network benefits they generate for the opposite side. Without discrimination, the platform does not treat the two types of sellers differently. Therefore, the utility it offers to them is determined by the *average* cost of serving each seller, $(f_H - \lambda\Delta f)$.

4.2 Private incentive to discriminate

Taking the partial derivative of the profit function in (3) with respect to D gives

$$\frac{\partial}{\partial D}\pi(u_1, u_2, D) = \frac{\partial n_H}{\partial D}(\alpha n_1 - u_2 - f_H), \tag{12}$$

where $\frac{\partial n_H}{\partial D} = -(1-\lambda)\phi_2'(u_2 - D)$. Intuitively, increasing discrimination D has the exact opposite effect on type-H sellers that increasing u_2 does. The negative impact on the platform's profit is the product between the decrease in the number of type-H sellers and the economic profit created by each of them.

Without discrimination, $n_2 = \phi_2(u_2^0)$ and $\frac{\partial n_2}{\partial u_2} = \phi_2'(u_2^0)$, and the economic profit from type-H sellers can be represented as follows:

$$\begin{aligned} \alpha n_1 - u_2^0 - f_H &= p_2^0 - f_H + \alpha_1 n_1 \\ &= -\lambda\Delta f + \frac{\phi_2(u_2^0)}{\phi_2'(u_2^0)}, \end{aligned}$$

where the first equality is due to (1) and the second is due to (10).

Substituting the first equality in (12), and by the second equality and (9), we have

$$\begin{aligned}\Pi'(0) &= \frac{\partial}{\partial D} n_H(u_2^0, 0) \cdot (p_2^0 - f_H + \alpha_1 n_1) \\ &= -(1-\lambda)\phi_2'(u_2^0) \cdot [-\lambda \Delta f + \frac{\phi_2(u_2^0)}{\phi_2'(u_2^0)}].\end{aligned} \tag{13}$$

These equations show the determinants of the platform's incentive to introduce discrimination. As increasing discrimination always dissuades more type-H sellers from joining the platform (i.e. $\frac{\partial n_H}{\partial D} < 0$), the platform would only do so when they each create a negative economic profit, i.e., $(p_2^0 - f_H + \alpha_1 n_1) < 0$.

According to the platform's optimal pricing rule without discrimination in (10), its optimal markup for sellers is $\frac{\phi_2(u_2^0)}{\phi_2'(u_2^0)}$.[20] Rewriting (10), we thus find that the economic profit of each type-H seller is exactly equal to the difference between the optimal markup $\frac{\phi_2(u_2^0)}{\phi_2'(u_2^0)}$ and the fraction of cost that all type-L sellers save in the calculation of the average seller cost, $\lambda \Delta f$.

Therefore, the platform only finds an incentive to introduce discrimination when its optimal markup on sellers, $\frac{\phi_2(u_2^0)}{\phi_2'(u_2^0)}$, is smaller than the cost saved by type-L sellers, $\lambda \Delta f$, as shown in (13). This condition is summarized in the following result.

Proposition 2 (Incentive to Discriminate) *The platform has an incentive to introduce non-price discrimination if and only if*

$$\lambda \Delta f \geqslant \frac{\phi_2(u_2^0)}{\phi_2'(u_2^0)}, \tag{14}$$

where u_2^0 is the utility that each seller obtains from the platform without discrimination, as given by (11).

Fixing Δf, a higher fraction of type-L sellers (λ) makes it more likely that the platform will discriminate against type-H sellers. Doing so discourages some type-H sellers from joining and reduces the platform's cost. However, it also lowers the attractiveness of the

[20] In (10), the optimal price, p_2^0, is equal to this markup plus the average seller's cost to the platform, $(f_H - \lambda \Delta f)$, and minus the network benefits each seller generates for the opposite side, $\alpha_1 n_1$.

platform to buyers because there are fewer sellers under discrimination. The platform introduces discrimination against type-H sellers only when their proportion is small enough, such that even after discouraging some of them from joining, the platform will still remain fairly attractive to buyers.

Similarly, fixing a λ fraction of type-L sellers, a larger cost difference Δf makes it more likely that the platform will discriminate against type-H sellers. Therefore, if the platform is sufficiently prejudiced, it is more likely to discriminate against type-H sellers.

Discrimination enables the platform to discourage type-H sellers from joining it. When its prejudice is minor (i.e., small Δf), type-H sellers are not much different from type-L sellers; when their proportion $(1-\lambda)$ is large, their participation becomes crucial to the platform's business. In either case, introducing discrimination is less likely to be beneficial.

As the platform does not always act on its prejudice to discriminate against type-H sellers, we conclude that prejudice does not always lead to discriminatory actions even if the platform market has tipped.

4.3 Optimal pricing with discrimination

When condition (14) holds, Proposition 3 states the platform's optimal prices.

Proposition 3 (Optimal Pricing with Discrimination) *When the platform discriminates (with* $D = \bar{D}$*), its optimal prices for two sides are given by*

$$\begin{cases} p_1^* = f_1 - \alpha_2 n_2 + \frac{n_1}{dn_1/du_1}, \\ p_2^* = f_H - \Delta f \cdot \frac{dn_L/du_2}{\partial n_2/\partial u_2} - \alpha_1 n_1 + \frac{n_2}{\partial n_2/\partial u_2}, \end{cases} \tag{15}$$

where $\frac{dn_1}{du_1} = \phi_1'(u_1^*(\bar{D}))$, $\frac{dn_L}{du_2} = \lambda\phi_2'(u_2^*(\bar{D}))$ *and* $\frac{\partial n_2}{\partial u_2} = \lambda\phi_2'(u_2^*(\bar{D})) + (1-\lambda)\phi_2'(u_2^*(\bar{D}) - \bar{D})$.

We postpone the comparison between the optimal prices with and without discrimination to Proposition 6 in Section 5.3.2, where we derive more intuitive comparative statics under a uniform distribution.

5 Welfare analysis

To further examine the effects of discrimination on users and on welfare, we need to provide a "micro-foundation" of the model.

5.1 A general framework

Suppose joining the platform gives a buyer an idiosyncratic value t_1, while her outside option yields zero utility. If the platform offers each buyer u_1, the buyer joins the platform if and only if

$$u_1 + t_1 \geqslant 0.$$

Denote F_1 on $\mathbb{R}$ as the cumulative distribution of t_1. The accumulation process of buyers is

$$\phi_1(u_1) = \Pr[u_1 + t_1 \geqslant 0] = 1 - F_1(-u_1).$$

Consistent with our previous specification of ϕ_1, the function $1 - F_1(-u_1)$ is increasing in u_1. The buyers' total surplus is now

$$v_1(u_1) \equiv \mathbb{E}_{t_1}[\max(u_1 + t_1, 0)] = \int_{-u_1}^{+\infty} (u_1 + t) dF_1(t), \tag{16}$$

which implies

$$v_1'(u_1) = \phi_1(u_1) = n_1.$$

Similarly, suppose joining the platform gives a seller an idiosyncratic value t_2, while his outside option yields zero utility. If the platform offers each seller u_2, he joins the platform if and only if

$$u_2 + t_2 \geqslant 0.$$

Denote F_2 on $\mathbb{R}$ as the cumulative distribution of t_2. The accumulation process of sellers is

$$\phi_2(u_2) = \Pr[u_2 + t_2 \geqslant 0] = 1 - F_2(-u_2).$$

The function $1 - F_2(-u_2)$ is also consistent with our previous specification of ϕ_2.

As a fraction λ of the sellers are of type-L, the sellers' total surplus depends on both u_2 and D as follows:

$$\begin{aligned} \nu_2(u_2, D) &\equiv \lambda \mathbb{E}_{t_2}[\max(u_2 + t_2, 0)] + (1-\lambda)\mathbb{E}_{t_2}[\max(u_2 - D + t_2, 0)], \\ &= \lambda \int_{-u_2}^{+\infty} (u_2 + t) dF_2(t) + (1-\lambda) \int_{D-u_2}^{+\infty} (u_2 - D + t) dF_2(t), \end{aligned} \tag{17}$$

which implies

$$\begin{aligned} \frac{\partial}{\partial u_2}\nu_2(u_2, D) &= \lambda\phi_2(u_2) + (1-\lambda)\phi_2(u_2 - D) = n_2; \\ \frac{\partial}{\partial D}\nu_2(u_2, D) &= -(1-\lambda)\phi_2(u_2 - D) = -n_H. \end{aligned}$$

The social surplus in this market is

$$w(u_1, u_2, D) = \pi(u_1, u_2, D) + \nu_1(u_1) + \nu_2(u_2, D).$$

5.2 The effects of discrimination

How does discrimination change the surplus for sellers, buyers, and society?

We use the optimal $u_1^*(D)$ and $u_2^*(D)$ as in (7). Denote $V_1(D) \equiv \nu_1(u_1^*(D))$ as the aggregate buyer surplus in (16) as a function of D alone. Denote $V_2(D) \equiv \nu_2(u_2^*(D), D)$ as the aggregate seller surplus in (17) as a function of D alone. Denote the social surplus when the platform maximizes profit, given D, as

$$\begin{aligned} W(D) &\equiv \Pi(D) + V_1(D) + V_2(D), \\ &= \pi(u_1^*(D), u_2^*(D), D) + v_1(u_1^*(D)) + v_2(u_2^*(D), D). \end{aligned}$$

The marginal impact of discrimination on social surplus is captured by:

$$\begin{aligned} W'(D) &= \Pi'(D) + V_1'(D) + V_2'(D), \\ &= \Pi'(D) + v_1'(u_1^*(D)) \cdot u_1^{*\prime}(D) + \frac{\partial}{\partial u_2} v_2(u_2^*(D), D) \cdot u_2^{*\prime}(D) + \frac{\partial}{\partial D} v_2(u_2^*(D), D), \\ &= \Pi'(D) + n_1 \cdot u_1^{*\prime}(D) + n_2 \cdot u_2^{*\prime}(D) - n_H. \end{aligned}$$

This equation shows that the platform's private incentive to discriminate ($\Pi'(D)$) may not always align with the interest of society ($W'(D)$). Starting from no discrimination (i.e., $D = 0$), the misalignment between the social and private incentive to introduce discrimination is:

$$\begin{aligned} \text{Social-private misalignment} &= W'(0) - \Pi'(0), \\ &= V_1'(0) + V_2'(0), \\ &= n_1 \cdot u_1^{*\prime}(0) + n_2 \cdot u_2^{*\prime}(0) - n_H. \end{aligned} \tag{18}$$

When calculating whether to introduce discrimination, the platform does not internalize all of the effects on users (i.e., $V_1'(0) + V_2'(0)$), which are included in the calculation of the social surplus.

The platform's interest is perfectly aligned with that of society if and only if $V_1'(0) + V_2'(0) = 0$. However, there is no reason for this condition to hold generally. When $V_1'(0) + V_2'(0) < 0$, the platform's incentive to discriminate is too large from society's perspective. When $V_1'(0) + V_2'(0) > 0$, however, society actually benefits from discrimination but the platform may lack the incentive to do so.

Next we show that it is perfectly possible for users' surplus to increase when the platform introduces discrimination, which in turn raises social welfare. It is therefore hard to make a case for tightening anti-discrimination policies on platforms, simply based on welfare concerns. Note that this is so even if we assume away the costs of enforcing anti-discrimination policies.

5.3 Welfare analysis under uniform distribution

As (18) shows, we need to know $u_1^{*\prime}(0)$ and $u_2^{*\prime}(0)$, or at least some of their properties to tell the sign of $V_1'(0) + V_2'(0)$, or $W'(0) - \Pi'(0)$. It is impossible to do so without specifying the accumulation processes of buyers and sellers (i.e. the F_1 and F_2 defined previously).

To illustrate the channels through which discrimination changes surpluses, in this section we assume that the accumulation processes of participants on both sides of the market follow the same uniform distribution.

Assumption 1 (Uniform Distribution) *Both* t_1 *and* t_2 *follow the same uniform distribution on* $[a, b]$, *such that*

$$F_1(x) = F_2(x) = \frac{x-a}{b-a}, \textit{ for } x \in [a, b];$$

$$\phi_1(y) = \phi_2(y) = Ay + B, \textit{ for } y \in [-b, -a],$$

where $A \equiv \frac{1}{b-a}$, $B \equiv \frac{b}{b-a}$.

5.3.1 Simulation results

Assuming uniform distribution allows us to run simulations. We run 503,119 simulations, where we use different sets of parameter values for the model. In addition to the parameters for the distributions of buyers and sellers, the other parameters include the strength of cross-group network externalities, the costs of serving different users, and the proportion of sellers that the platform is prejudiced against. These results can be summarized as follows. The Appendix shows more details.

Proposition 4 (Welfare Simulation) *Simulation results show that the platform's incentive to discriminate may be aligned or misaligned with the interest of society, depending on the distributions of buyers and sellers, and other model parameters.*

In particular, we check if there exist parameter values that support the four cases summarized in Table 2, under the condition that $\Pi(0) > 0$ (such that the platform exists).

Table 2: (Mis)Alignment between Social and Private Incentives to Discriminate

	$\Pi'(0) > 0$	$\Pi'(0) < 0$
$W'(0) > 0$	I	II
$W'(0) < 0$	III	IV

In both (I) and (IV), there is no incentive misalignment, in the sense that when the platform has an incentive to discriminate, society benefits too, and vice versa. In case (III), the platform has an incentive to discriminate but society suffers. In case (II), society benefits from discrimination but the platform has no incentive to discriminate. Only (III) justifies tightening anti-discrimination policies on platforms. Our simulation results show that all four cases are possible under uniform distribution.

Ensuring that it is socially efficient to prohibit discrimination is equivalent to making sure that the parameter values of that two-sided market fall exactly within the set that generates case (III). However, doing so requires the rather challenging task of carrying out empirical estimation of all the relevant parameters.

5.3.2 The channels of welfare improvement

Now we show that when the platform introduces discrimination, it may optimally lower its prices, and hence there is a potential to increase sellers' and buyers' surpluses.

One can view non-price discrimination as a tool for the platform to adjust the seller composition to economize the cost of serving sellers. As the platform's business depends on network externalities across the two sides, it still needs to attract sufficient numbers of users from both sides regardless of its choice of discrimination.

When the platform discriminates against type-H sellers, if it does not simultaneously reduce its prices, the number of sellers joining will decrease. To maintain its attractiveness to buyers, it usually has to lower the price for buyers. If doing so is itself insufficiently effective at persuading buyers to stay, the platform may also have to lower the price for sellers to induce more of them to join it. As long as the total cost savings are large enough to compensate for the losses from lower prices, the platform will optimally lower its prices.

We now derive the formulas for the model under uniform distribution, and characterize the conditions under which platform discrimination increases social surplus. By Assumption 1, given $u_1, u_2, (u_2 - D) \in [-b, -a]$, we can write the analytical formulas in (4). We then solve for the first-order conditions (5) and (6) for the (interior) optimal utilities offered to both sides (u_1^* and u_2^*) and their derivatives with respect to D.[21] The signs of $u_1^{*\prime}(D)$ and $u_2^{*\prime}(D)$ in turn help us find the following result.

Proposition 5 (Welfare Improvement) *Suppose Assumption 1 holds, and the optimal utilities that the platform offers to two sides, with or without discrimination, are all interior solutions. In this case:*

i) There exists a social incentive to discriminate whenever there exists a private incentive to discriminate, if and only if $\frac{\alpha}{b-a} > 2$*; or equivalently*

ii) The aggregate surpluses of all buyers and sellers in the market increase when the platform introduces discrimination, if and only if $\frac{\alpha}{b-a} > 2$.

Note that α represents the total network externalities created when each buyer-seller pair interacts, the main source of value creation in this two-sided market. The two results

[21] These formulas, conditions, and solutions are provided in the proof of Proposition 5 in the Appendix.

in Proposition 5 refer to the signs of $W'(0) - \Pi'(0) = V_1'(0) + V_2'(0)$ in (18), and the necessary and sufficient condition for them to be positive under Assumption 1 is $\frac{\alpha}{b-a} > 2$.

The condition $\frac{\alpha}{b-a} > 2$ requires that cross-side network externalities are strong compared with the range of idiosyncratic values (i.e., t_1 and t_2) that market users derive. In Section 5.3.3 we show a numerical solution in which this condition holds. In that case, discrimination increases buyers and sellers surpluses. Of course, the platform's profit increases too. Therefore, discrimination increases the social surplus.

Given $D \in [0, \bar{D}]$, rewrite (4e) and (4f) and denote the platform's optimal prices as functions of D alone:

$$
\begin{aligned}
p_1^*(D) &\equiv \alpha_1 \cdot n_2(u_2^*(D), D) - u_1^*(D), \\
p_2^*(D) &\equiv \alpha_2 \cdot n_1(u_1^*(D)) - u_2^*(D),
\end{aligned}
$$

and we have the following conclusion.

Proposition 6 (Voluntary Price Cuts Under Discrimination) *Suppose Assumption 1 holds, and the optimal utilities that the platform offers to two sides, with or without discrimination, are all interior solutions. Then, for $D \in [0, \bar{D}]$, we have:*

i) $p_1^{*\prime}(D) + p_2^{*\prime}(D) < 0$;

ii) $p_1^{*\prime}(D) < 0$ *if* $\frac{\alpha}{b-a} > 2$ *and* $\frac{\alpha_1}{\alpha} < \frac{1}{2}$;

iii) $p_2^{*\prime}(D) < 0$ *if* $\frac{\alpha}{b-a} > 2$ *and* $\frac{\alpha_1}{\alpha} > \frac{2(b-a)^2}{\alpha^2}$.

i) Under uniform distribution, whenever the platform discriminates, it *always* voluntarily cuts the total equilibrium prices that it charges to both sides. This self-correcting pricing mechanism alleviates the negative impact on user participation due to discrimination, and enables the platform to remain fairly attractive to both sides.

ii) ii) If the buyers enjoy a smaller fraction of the total network externalities from interactions than do sellers ($\frac{\alpha_1}{\alpha} < \frac{1}{2}$), the platform needs to make an extra effort to keep buyers when it chooses to discriminate. This is reflected in a lower equilibrium price for buyers under discrimination.

iii) If the sellers enjoy too small a fraction of the total network externalities ($1 - \frac{\alpha_2}{\alpha} = \frac{\alpha_1}{\alpha} > \frac{2(b-a)^2}{\alpha^2}$), the equilibrium price for all sellers must decrease under discrimination to prevent them from leaving the platform.

Depending on the model parameters, it is possible for the equilibrium prices for both sides to decrease under discrimination, i.e., when $\frac{\alpha}{b-a} > 2$ and $\frac{2(b-a)^2}{\alpha^2} < \frac{\alpha_1}{\alpha} < \frac{1}{2}$. We show such an example in section 5.3.3.

Propositions 5 and 6 show that the platform's voluntary price cuts that come with discrimination have a potential of increasing users' surpluses if and only if there exist sufficiently strong cross-side network externalities (i.e., $\frac{\alpha}{b-a} > 2$). In other words, the two-sidedness of a platform market is crucial. The following result formalizes this finding.

Proposition 7 (One-Sided Market) *Suppose Assumption 1 holds, and the optimal utilities that the platform offers to two sides, with or without discrimination, are all interior solutions. When* $\alpha_1 = \alpha_2 = 0$, *such that the market is one-sided, we must have:*

i) $W'(0) - \Pi'(0) < 0$;

ii) $p_1^{*\prime}(0) = 0, V_1'(0) = 0$;

iii) $p_2^{*\prime}(0) < 0, V_2'(0) < 0$.

When there exist no network externalities whatsoever, buyers and sellers are separated as if the platform simply operates in two distinct one-sided markets. In this case, under uniform distribution, the social incentive to introduce discrimination is always weaker than the platform's private incentive.

The platform's discrimination against sellers is completely irrelevant to buyers. Therefore, neither the buyers' surplus nor the platform's pricing for them is affected. On the seller side, the platform uses price and discrimination simultaneously to adjust the composition of the group of sellers. Discrimination is always accompanied by a price cut. Absent cross-side network externalities, however, such a price cut is insufficient to compensate for sellers' lost surplus under discrimination. Therefore, introducing discrimination in one-sided markets never benefits any market participants besides the platform itself.

The stark contrast between Propositions 5 and 7 and between Propositions 6 and 7 show that two-sidedness makes it possible for social and private incentives to discriminate to be aligned. Therefore, previous studies of discrimination without two-sidedness are not applicable to platform markets.

5.3.3 A numerical example

Let t_1 and t_2 both follow a uniform distribution on $[a, b] = [-10, 5]$, and let the parameters in the model take the following values.

$$\begin{aligned} f_1 &= 8, & \alpha_1 &= 20, & \lambda &= 0.8, \\ f_H &= 12, & \alpha_2 &= 60, & \bar{D} &= 0.5, \\ \Delta f &= 10, & \alpha &= 80. \end{aligned}$$

We find the following interior numerical solution.[22]

[22]We use Scientific WorkPlace 5.0 to generate our numerical solutions.

No Discrimination $(D=0)$		Discrimination $(D=\bar{D})$
$p_1^0 = 5.7364,$	$>$	$p_1^* = 5.7255,$
$p_2^0 = 4.5364,$	$>$	$p_2^* = 4.5155,$
$u_1^0 = -4.9727,$	$<$	$u_1^* = -4.9509,$
$u_2^0 = -4.4273,$	$<$	$u_2^* = -4.3191,$
$\Pi(0) = 1.6364 \times 10^{-2},$	$<$	$\Pi(D) = 6.5852 \times 10^{-2},$
$n_1^0 = 1.8182 \times 10^{-3}$	$<$	$n_1^* = 3.2727 \times 10^{-3},$
$n_L^0 = 3.0545 \times 10^{-2},$	$<$	$n_L^* = 3.6315 \times 10^{-2},$
$n_H^0 = 7.6364 \times 10^{-3},$	$>$	$n_H^* = 2.4121 \times 10^{-3},$
$n_2^0 = 3.8182 \times 10^{-2},$	$<$	$n_2^* = 3.8727 \times 10^{-2},$
$n_H/n_2^0 = 20\%(= 1-\lambda),$	$>$	$n_H/n_2^* = 6.2\%.$

Condition (14) in Proposition 2 holds, and therefore the platform indeed earns more profits when it discriminates. All the conditions in Propositions 5 and 6 also hold. It is clear from the comparison that when it discriminates, the platform indeed chooses to lower its prices for both buyers and sellers.

5.3.4 The "victims" of discrimination

Discrimination creates a wedge between the net utilities obtained by type-L and type-H sellers, which necessarily results in a reduction of the proportion of type-H sellers on the platform, regardless of the price adjustments.

Formally, given u_2 and D, denote $\tau(u_2, D)$ the proportion of type-H sellers among all sellers on the platform, i.e.,

$$\tau(u_2, D) \equiv \frac{n_H(u_2, D)}{n_2(u_2, D)},$$

and we have the following result.

Lemma 3 (Proportion of Type-H Sellers) *Under Assumption 1, we have*

$$\frac{d}{dD}\tau(u_2^*(0),0) < 0.$$

Nevertheless, it is still possible for type-H sellers' total number and aggregate surplus to rise in equilibrium. Given (u_2, D), denote type-H sellers' aggregate surplus as

$$\nu_H(u_2, D) \equiv (1-\lambda)\int_{D-u_2}^{+\infty}(u_2 - D + t)dF_2(t).$$

and let $V_H(D) \equiv \nu_H(u_2^*(D), D)$ denote their aggregate surplus in equilibrium. We have the following result.

Proposition 8 (Welfare of Type-H Sellers) *Suppose Assumption 1 holds, and the optimal utilities that the platform offers to two sides, with or without discrimination, are all interior solutions. Then, we have*

$$V_H'(D) > 0 \text{ and } \frac{\partial n_H}{\partial D} > 0 \text{ if and only if } \frac{2-\lambda(\alpha^2A^2-2)}{\alpha^2A^2-4} > 0.$$

The necessary and sufficient condition for both $V_H'(D) > 0$ and $\frac{\partial n_H}{\partial D} > 0$ is $u_2^{*\prime}(D) > 1$. Intuitively, if $u_2^{*\prime}(D) > 1$, the utility that the platform offers to all sellers, $u_2^*(D)$, increases faster than does the disutility from discrimination. Therefore, the platform's price cut for all sellers more than compensates for the disutility it imposes on type-H sellers through discrimination. The resulting net utility ($u_2^* - D$) that they obtain thus rises, attracting more of them to the platform and increasing their user surplus.

Given the uniform distribution on $[a, b]$ and when $\alpha A = \frac{\alpha}{b-a} > 2$, the condition in Proposition 8 reduces to $\lambda < \frac{2}{\alpha^2A^2-2}$, such that the proportion of type-H sellers $(1-\lambda)$ needs to be sufficiently large for type-H sellers to benefit from discrimination. In the proof of this result (in the Appendix), we provide a numerical example with $1-\lambda = 99.5\%$, where the platform's optimal discrimination against type-H sellers indeed benefits them in equilibrium.

6 Alternative policy: making discrimination imperfect

In reality, it is often difficult to detect discrimination. For instance, the developers of different smartphone apps may find it difficult to tell whether the difference in the amount of time they spend waiting for their apps to be approved by Apple is due to discrimination. It is even more difficult to distinguish between discriminatory practices based on prejudice and those based on actual operational costs (Heckman, 1998). It may not be feasible at all to invoke anti-discrimination policies in order to eliminate discrimination.

In this section, we show that making discrimination imperfect can reduce a platform's incentive to discriminate. Such an approach is therefore an alternative policy when enforcing anti-discrimination laws is difficult.

Suppose that discrimination aimed at one subset of users may have unintended adverse impacts on *other* users. For instance, the platform may mistake a type-L seller for a type-H seller. Such imprecision means type-L sellers may also suffer from discrimination aimed at type-H sellers. Another conceivable possibility is the following. Even though the platform can precisely target discrimination at type-H sellers, if type-L sellers know about these discriminatory practices, they may feel sympathetic and thus suffer utility loss.

To formalize these ideas, suppose that whenever the platform discriminates against type-H sellers by imposing a disutility D, each type-L seller also experiences a disutility θD, where $\theta \in [0, 1]$ represents either the "degree of sympathy" that type-L sellers hold towards type-H sellers, or the "degree of imprecision" in the platform's targeting when carrying out discrimination. Therefore, when it is easier to induce type-L sellers to feel sympathetic, or when the discriminatory targeting is more imprecise, type-L sellers experience more utility losses under discrimination. In either case, we say the discrimination is *imperfect*.

Under imperfect discrimination, while a type-H seller gets $u_2 - D$ from joining the platform, a type-L seller gets a utility of $u_2 - \theta D$ (instead of u_2) from joining. This imprecision changes the number of type-L sellers joining under discrimination to the following:

$$n_L(u_2, D) = \lambda \phi_2(u_2 - \theta D). \tag{19}$$

Denote $\frac{\partial n_L}{\partial u_2} = \lambda \phi_2'(u_2 - \theta D)$, and $\frac{\partial n_L}{\partial D} = -\theta \lambda \phi_2'(u_2 - \theta D)$.

This new set-up leads to a slightly modified Proposition 2 as the following:

Proposition 9 (Incentive to Discriminate with Imperfection) *When discrimination is imperfect, the platform has an incentive to introduce it if and only if*

$$\frac{(1-\theta)\lambda(1-\lambda)}{1-\lambda+\theta\lambda}\Delta f \geqslant \frac{\phi_2(u_2^0)}{\phi_2'(u_2^0)}, \tag{20}$$

where θ *represents the degree of imperfection, and* u_2^0 *is the utility that each seller obtains from the platform without discrimination, as given by (11).*

Propositions 9 and 2 share the same intuition. Moreover, because the left-hand side of (20) is strictly decreasing in θ, whereas the right-hand side does not depend on θ at all, we have the following result.

Proposition 10 (Reduced Discrimination due to Imperfection) *The more imperfect the discriminatory strategy is - i.e., the larger* θ *is - the less likely the platform is to have an incentive to discriminate.*

Which policies may make discrimination more imperfect? It is worth taking a look at the case of *Fair Housing Council of San Fernando Valley v. Roommates.com, LLC*. Roommates.com matches users renting rooms with those who need rooms. However, it requires users to disclose their sexual orientation, number of children, and whether the children live with the user. The 9th Circuit panel ruled that Roommates.com should not require users to choose from "potentially discriminatory options."

The decision makes it more difficult for Roommates.com to identify users with certain attributes, and can be understood as a form of privacy protection policy. Privacy protection policies forbidding a platform from requesting specific user information, such as pictures, race, and marital status, may increase the imprecision of its discriminatory targeting. Proposition 10 implies that the platform would then be less likely to discriminate based on such characteristics.[23]

Another potential policy concerns information disclosure. For instance, the regulator can require that smartphone app platforms disclose information about waiting times for all app review submissions. In the case of lending platforms, the regulator can require the platform to clearly list all documents required to apply for a loan, and to report, anonymously, cases where more than the required documents are requested. If the platform does engage in discriminatory practices (which may not be verifiable in court), by disclosing its treatment of different users, the users who are treated more favorably may feel sympathetic to those who are treated differently.

How to design these disclosure requirements and/or privacy protection policies remains an open question, and probably depends on the contexts of specific platform markets.

7 Conclusion

In a tipped platform market, the prejudice of the platform does not necessarily translate into actual discriminatory practices, as the platform may not always find it beneficial to do so. When a platform does discriminate, the social surplus does not necessarily

[23]Private firms adopt similar forms of anti-discriminatory policies. Uber, for instance, prevents drivers from seeing the complete profile of riders and their destination before accepting a ride request. It is therefore more difficult for drivers to discriminate against short-haul riders or riders of certain races. If Prosper.com prevents users from posting their pictures, lenders can less easily tell whether a borrower is black. However, these cases concern alleged discrimination among platform users rather than by platforms.

decrease. Even if we ignore the platform's own interests, surpluses of the platform's users may increase under non-price discrimination. Therefore, in a platform market with a tendency to tip, tightening anti-discrimination policies on platforms cannot be always socially efficient.

When a platform discourages participation by a certain subset of users through non-price discrimination, it may optimally choose to lower the prices it charges to all users, as its business depends crucially on attracting sufficient numbers of users from both sides. This self-correcting pricing mechanism is the fundamental source of welfare improvement under non-price discrimination. It only exists in two-sided markets.

When it is difficult or costly to enforce anti-discrimination laws, our model suggests that mandatory disclosure requirements or user privacy protection policies that are well crafted can make it less likely for the platform to discriminate. Identifying other forms of regulations to address non-price discrimination is a promising area for future research.

8 Appendix - proofs and derivations

Lemma 1 For any (u_1, u_2, D), using (4), the partial derivative of (3) with respect to D gives

$$\frac{\partial}{\partial D}\pi(u_1, u_2, D) = \frac{\partial n_H}{\partial D} \cdot (\alpha n_1 - u_2 - f_H),$$

where

$$\frac{\partial n_H}{\partial D} = -(1-\lambda)\phi_2'(u_2 - D) < 0.$$

And therefore

$$\frac{\partial^2}{\partial D^2}\pi(u_1, u_2, D) = \frac{\partial^2 n_H}{\partial D^2} \cdot (\alpha n_1 - u_2 - f_H).$$

Suppose there exists $D_0 \in [0, \bar{D}]$ such that $\Pi'(D_0) = 0$. By (7), (8) and the envelope theorem, we know

$$\begin{aligned}
\Pi'(D_0) &= \frac{\partial}{\partial D}\pi(u_1^*(D_0), u_2^*(D_0), D_0) \\
&= \frac{\partial}{\partial D}n_H(u_1^*(D_0), u_2^*(D_0), D_0) \cdot [\alpha \cdot n_1(u_1^*(D_0)) - u_2^*(D_0) - f_H] \\
&= 0.
\end{aligned}$$

Therefore, we must have

$$\alpha \cdot n_1(u_1^*(D_0)) - u_2^*(D_0) - f_H = 0,$$

which implies

$$\frac{\partial^2}{\partial D^2}\pi(u_1^*(D_0), u_2^*(D_0), D_0) = 0.$$

We now prove that we must have $\Pi''(D_0) > 0$.

First, we examine function $\pi(u_1, u_2, D)$ by fixing the utilities at $u_1 = u_1^*(D_0)$ and $u_2 = u_2^*(D_0)$, and for any $D \in [0, \bar{D}]$ we must have

$$\pi(u_1^*(D_0), u_2^*(D_0), D) \leqslant \pi(u_1^*(D), u_2^*(D), D) = \Pi(D),$$

with the first inequality holding with equation at $D = D_0$. Given that all of these functions are differentiable, we must have

$$\Pi''(D_0) > \frac{\partial^2}{\partial D^2}\pi(u_1^*(D_0), u_2^*(D_0), D)\bigg|_{D=D_0} = 0.$$

Therefore, we have proved the following claim:

$$\text{For any } D_0 \in [0, \bar{D}],\ \Pi''(D) > 0 \text{ whenever } \Pi'(D) = 0. \tag{21}$$

We now prove that this implies quasiconvexity.

Suppose $\Pi(D)$ is not quasiconvex. Then, there exist $[x, y] \subseteq [0, \bar{D}]$, and $k \in [0, 1]$ such that

$$\Pi(kx + (1-k)y) > \max\{\Pi(x), \Pi(y)\}.$$

As $kx + (1-k)y \in [x, y]$, and twice-differentiability implies that $\Pi(D)$ is continuous and differentiable on $[0, \bar{D}]$, there must exist $z \in [x, y]$ such that $z = \arg\max_{D \in [x,y]} \Pi(D)$, which in turn implies that $\Pi'(z) = 0$ and $\Pi''(z) < 0$, a contradiction. Therefore $\Pi(D)$ must be quasiconvex on $[0, \bar{D}]$.■

Lemma 2 This is immediately implied by quasiconvexity of $\Pi(D)$ on $[0, \bar{D}]$.■

Proposition 1 (10) and (11) are found by letting $D = 0$ in (5) and (6), and inverting (1).■

Proposition 2 Let $D = 0$ in (9) and we have

$$\Pi'(0) = (1-\lambda)[\lambda\Delta f \phi_2'(u_2^0) - \phi_2(u_2^0)]$$

As $\phi_2'(u_2^0) > 0$, we know $\frac{\phi_2(u_2^0)}{\phi_2'(u_2^0)} > 0$. Therefore, we have

$$\Pi'(0) > 0 \text{ if and only if } \lambda\Delta f > \frac{\phi_2(u_2^0)}{\phi_2'(u_2^0)}.$$

Now consider the case when $\lambda\Delta f = \frac{\phi_2(u_2^0)}{\phi_2'(u_2^0)}$, that is, when $\Pi'(0) = 0$. By (21) we know that $\Pi''(0) > 0$, therefore $D = 0$ is a minimizer of $\Pi(D)$. As introducing discrimination increases D from 0, it will make $\Pi'(D) > 0$ and is hence profitable.

Then, by Lemma 1, we know $\Pi(\bar{D}) > \Pi(0)$ must hold when (14) holds.■

Proposition 3 (15) is found by solving (5) and (6), and inverting (1).■

Proposition 4 Assume t_1 and t_2 follow the same uniform distribution on $[a, b]$, and let $A \equiv \frac{1}{b-a}$, and $B \equiv \frac{b}{b-a}$. Assume $\alpha_2 = 0$. (More details are provided in Section 5.3.2.) For example, the following four sets of parameter values show four different cases of (mis-)alignment between the private and social incentives to discriminate.

Example 1. Misalignment: $\Pi'(0) > 0$ and $W'(0) < 0$.

$\{A, B, \alpha, f_1, f_H, \Delta f, \lambda\} = \{0.5, 0.8, 2.4, 1, 2, 1.5, 0.1\}$

$\Pi'(0) = 0.0288281; W'(0) = -0.075542; \Pi(0) = 0.0473633.$

Example 2. Alignment: $\Pi'(0) > 0$ and $W'(0) > 0$.

$\{A, B, \alpha, f_1, f_H, \Delta f, \lambda\} = \{1/15, 1/3, 80, 8, 12, 10, 0.8\}$

$\Pi'(0) = 0.09903; W'(0) = 0.09973; \Pi(0) = 0.016364.$

Example 3. Misalignment: $\Pi'(0) < 0$ and $W'(0) < 0$.

$\{A, B, \alpha, f_1, f_H, \Delta f, \lambda\} = \{0.5, 0.8, 3, 1, 2, 1.5, 0.1\}$

$\Pi'(0) = -0.0353571; W'(0) = -0.334745; \Pi(0) = 0.0564286.$

Example 4. Misalignment: $\Pi'(0) < 0$ and $W'(0) > 0$.

$\{A, B, \alpha, f_1, f_H, \Delta f, \lambda\} = \{1/15, 1/3, 80, 8, 5, 4, 0.15\}$

$\Pi'(0) = -0.00031; W'(0) = 0.00391; \Pi(0) = 0.00065.$ ■

Proposition 5 With uniform distribution on $[a, b]$, given $u_1, u_2, (u_2 - D) \in [-b, -a]$, (4) becomes

$$
\begin{aligned}
n_1 &= \phi_1(u_1) = Au_1 + B, \\
n_L &= \lambda\phi_2(u_2) = \lambda(Au_2 + B), \\
n_H &= (1-\lambda)\phi_2(u_2 - D) = (1-\lambda)(Au_2 + B - AD), \\
n_2 &= n_L + n_H = Au_2 + B - (1-\lambda)AD, \\
p_1 &= \alpha_1 n_2 - u_1, \\
p_2 &= \alpha_2 n_1 - u_2.
\end{aligned}
$$

Given $D \in [0, \bar{D}]$, the first-order conditions for the interior optimal utilities offered to two sides u_1^* and u_2^* are

$$\begin{cases} \alpha(Au_2^* + B) - \alpha(1-\lambda)AD - 2u_1^* - f_1 - \frac{B}{A} = 0, \\ \alpha(Au_1^* + B) - 2u_2^* - f_H - \frac{B}{A} + (1-\lambda)D + f_3\lambda = 0, \end{cases}$$

whose solution is

$$\begin{cases} u_1^*(D) = \frac{2B - AB\alpha + 2Af_1 + A^2D\alpha - A^2D\alpha\lambda + A^2\alpha f_H - A^2\alpha\lambda f_3 - A^2B\alpha^2}{A^3\alpha^2 - 4A}, \\ u_2^*(D) = \frac{2B - 2AD - AB\alpha + 2AD\lambda + 2Af_H - 2A\lambda f_3 + A^2\alpha f_1 - A^2B\alpha^2 + A^3D\alpha^2 - A^3D\alpha^2\lambda}{A^3\alpha^2 - 4A}. \end{cases}$$

Therefore, we have

$$\begin{aligned} u_1^{*\prime}(D) &= \frac{(1-\lambda)\alpha A}{\alpha^2A^2 - 4}, \\ u_2^{*\prime}(D) &= \frac{(1-\lambda)(\alpha^2A^2 - 2)}{\alpha^2A^2 - 4}, \end{aligned} \tag{22}$$

which immediately imply the following:

i) If $\alpha A > 2$, we have $u_1^{*\prime}(D) > 0$ and $u_2^{*\prime}(D) > 0$;

ii) if $\sqrt{2} < \alpha A < 2$, we have $u_1^{*\prime}(D) < 0$ and $u_2^{*\prime}(D) < 0$; and

iii) if $\alpha A \leqslant \sqrt{2}$, we have $u_1^{*\prime}(D) < 0$ and $u_2^{*\prime}(D) \geqslant 0$.

With uniform distribution on $[a, b]$, we have

$$V_1'(0) = n_1 \cdot u_1^{*\prime}(0) = n_1 \cdot \frac{(1-\lambda)\alpha A}{\alpha^2A^2 - 4},$$

and

$$V_2'(0) = n_2 \cdot u_2^{*\prime}(0) - n_H = \frac{2(1-\lambda)(Au_2 + B)}{\alpha^2A^2 - 4}.$$

And finally by (18), we have $W'(0) - \Pi'(0) = V_1'(0) + V_1'(0) > 0$ if and only if $\alpha A > 2$.■

Proposition 6 Because $p_1 = \alpha_1 n_2 - u_1$, and $p_2 = \alpha_2 n_1 - u_2$, by (22), we have

$$p_1^{*\prime}(D) = \frac{(1-\lambda)A}{\alpha^2 A^2 - 4} \cdot (2\alpha_1 - \alpha),$$
$$p_2^{*\prime}(D) = \frac{(1-\lambda)}{\alpha^2 A^2 - 4} \cdot (2 - \alpha A^2 \alpha_1),$$

which immediately imply

i) $p_1^{*\prime}(D) + p_2^{*\prime}(D) = -\frac{(1-\lambda)(\alpha_1 A+1)}{\alpha A+2} < 0$;

ii) $p_1^{*\prime}(D) < 0$ if $\frac{\alpha}{b-a} > 2$ and $\frac{\alpha_1}{\alpha} < \frac{1}{2}$;

iii) $p_2^{*\prime}(D) < 0$ if $\frac{\alpha}{b-a} > 2$ and $\frac{\alpha_1}{\alpha} > \frac{2(b-a)^2}{\alpha^2}$.■

Proposition 7 Let $\alpha_1 = \alpha_2 = 0$ in the basic model. This immediately produces the results.■

Lemma 3 $\frac{\partial}{\partial u_2}\tau(u_2, D) = \frac{\lambda(1-\lambda)A^2 D}{(n_2)^2}$, which implies $\frac{\partial}{\partial u_2}\tau(u_2^*(0), 0) = 0$. And $\frac{\partial \tau}{\partial D} = -\frac{\lambda(1-\lambda)A(Au_2+B)}{(n_2)^2} < 0$.■

Proposition 8 $V_H'(D) = n_H(u_2^{*\prime}(D) - 1)$, and $\frac{\partial n_H}{\partial D} = A(1-\lambda)(u_2^{*\prime}(D) - 1)$. Therefore, they are both positive if and only if $u_2^{*\prime}(D) - 1 = \frac{2-\lambda(\alpha^2 A^2 - 2)}{\alpha^2 A^2 - 4} > 0$.

The following is a numerical example for this case. Let $[a, b] = [-10, 5]$, and let the parameters in the model take the following values.

$$f_1 = 3.2, \quad \alpha_1 = 20, \quad \lambda = 0.005,$$
$$f_H = 10, \quad \alpha_2 = 60, \quad \bar{D} = 0.05,$$
$$\Delta f = 4, \quad \alpha = 80.$$

We find the following interior numerical solution.

No Discrimination $(D=0)$		Discrimination $(D=\bar{D})$
$p_1^0 = 4.0804,$	$>$	$p_1^* = 4.0749,$
$p_2^0 = 8.7424,$	$>$	$p_2^* = 8.7320,$
$u_1^0 = -4.0607,$	$<$	$u_1^* = -4.0499,$
$u_2^0 = -4.9853,$	$<$	$u_2^* = -4.9315,$
$\Pi(0) = 5.3912 \times 10^{-2},$	$<$	$\Pi(D) = 5.3922 \times 10^{-2},$
$n_1^0 = 6.2618 \times 10^{-2}$	$<$	$n_1^* = 6.3342 \times 10^{-2},$
$n_L^0 = 4.9091 \times 10^{-6},$	$<$	$n_L^* = 2.2849 \times 10^{-5},$
$n_H^0 = 9.7691 \times 10^{-4},$	$<$	$n_H^* = 1.2303 \times 10^{-3},$
$n_2^0 = 9.8182 \times 10^{-4},$	$<$	$n_2^* = 1.2532 \times 10^{-3},$
$n_H/n_2^0 = 99.5\% (= 1-\lambda),$	$>$	$n_H/n_2^* = 98.2\%.$

Note that condition (14) in Proposition 2 holds, and therefore the platform indeed earns more profits when it discriminates. All of the conditions in Propositions 5, 6, and 8 also hold. It is clear from the comparison that the platform indeed chooses to lower its prices for both buyers and sellers when it discriminates. The result is an increase in the number of buyers, type-L sellers, and type-H sellers.■

Proposition 9 The first-order conditions for the optimal u_1 and u_2 are still given by (5) and (6), and the platform's incentive to discriminate still depends on $\Pi'(0)$ derived from (8), except that because n_L also depends on D according to (19), we now have

$$\begin{aligned}\Pi'(D) &= \frac{\partial}{\partial D}\pi(u_1^*(D), u_2^*(D), D) \\ &= \frac{\partial n_L}{\partial D}\cdot(\alpha n_1 - u_2 - f_L) + \frac{\partial n_H}{\partial D}\cdot(\alpha n_1 - u_2 - f_H).\end{aligned}$$

By (5) and (6), and let $D=0$, we have

$$\Pi'(0) = (1-\theta)\lambda(1-\lambda)\Delta f \cdot \phi_2'(u_2^0) - (1-\lambda+\theta\lambda)\cdot\phi_2(u_2^0),$$

which implies (20).■

Proposition 10 We can take the first-order derivative of the left-hand side of (20) with respect to θ to find the impact of imperfection on the incentive to discriminate:

$$\frac{\partial \frac{(1-\lambda)\lambda(1-\theta)}{1-\lambda+\theta\lambda}}{\partial\theta} = \frac{\lambda(\lambda-1)}{(1-\lambda+\theta\lambda)^2} < 0.$$

which immediately implies Proposition 10.■

References

[1] **Alchian, Armen A., and Reuben A. Kessel**. 1962. "Competition, Monopoly and the Pursuit of Money." *Aspects of Labor Economics*, 157-175. Princeton: Princeton University Press (for NBER).

[2] **Armstrong, Mark**. 2006. "Competition in Two-sided Markets." *RAND Journal of Economics*, 37(3): 668-691.

[3] **Arrow, Kenneth J.** 1973. "The Theory of Discrimination." In *Discrimination in Labor Markets* edited by Orley Ashenfelter and Albert Rees, 3-33. Princeton NJ: Princeton University Press.

[4] **Becker, Gary S.** 1957. *The Economics of Discrimination*. Chicago, IL: University of Chicago Press, 2nd edn.

[5] **Becker, Gary S.** 1962. "Comments: Competition, Monopoly and the Pursuit of Money." *Aspects of Labor Economics*, 175-183. Princeton: Princeton University Press (for NBER).

[6] **Datta, Amit, Michael Carl Tschantz, and Anupam Datta.** 2015. "Automated Experiments on Ad Privacy Settings." *Proceedings on Privacy Enhancing Technologies*, 2015(1): 92-112.

[7] **Doleac, Jennifer L., and Luke C. D. Stein**. 2013. "The Visible Hand: Race and Online Market Outcomes." *Economic Journal*, 123(572): 469-492.

[8] **Duarte, Jefferson, Stephan Siegel, and Lance Young**. 2012. "Trust and Credit: The Role of Appearance in Peer-to-Peer Lending." *Review of Financial Studies*, 25(8): 2455-2484.

[9] **Edelman, Benjamin, and Michael Luca.** (2014). "Digital Discrimination: The Case of Airbnb.com." *Harvard Business School Working Paper*, 14-054.

[10] **Edelman, Benjamin, Michael Luca, and Dan Svirsky.** 2017. "Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment." *American Economic Journal: Applied Economics*, 9(2): 1-22.

[11] **Evans, David, and Richard Schmalensee.** 2015. "The Antitrust Analysis of Multisided Platform Businesses." In *The Oxford Handbook of International Antitrust Economics, Volume 1,* edited by Roger D. Blair and D. Daniel Sokol, 404-450. New York, NY: Oxford University Press.

[12] **Fisman, Raymond, and Michael Luca.** 2016. "Fixing Discrimination in Online Marketplaces." *Harvard Business Review*, 94(12): 88-95.

[13] **Ge, Yanbo, Christopher R. Knittel, Don MacKenzie, and Stephen Zoepf**. 2016. "Racial and Gender Discrimination in Transportation Network Companies." *NBER Working Paper*, 22776.

[14] **Hagiu, Andrei**. 2009. "Quantity vs. Quality and Exclusion by Two-sided Platforms." *Harvard Business School Working Paper*, 09-094.

[15] **Hagiu, Andrei, and Simon Rothman**. 2016. "Network Effects Aren't Enough." *Harvard Business Review*, April, 6471.

[16] **Heckman, James**. 1998. "Detecting Discrimination." *Journal of Economic Perspectives*, 12(2): 101-116.

[17] **Hossain, Tanjim, Dylan Minor, and John Morgan**. 2011. "Competing Matchmakers: An Experimental Analysis." *Management Science*, 57(11): 1913-1925.

[18] **Hossain, Tanjim, and John Morgan**. 2013. "When Do Markets Tip? A Cognitive Hierarchy Approach." *Marketing Science*, 32(3): 431-453.

[19] **Phelps, Edmund S.** 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review*, 62(4): 659-661.

[20] **Pope, Devin G., and Justin R. Sydnor**. 2011. "What's in a Picture?: Evidence of Discrimination from Prosper.com." *Journal of Human Resources*, 46(1), 53-92.

[21] **Rochet, Jean-Charles, and Jean Tirole**. 2003. "Platform Competition in Two-sided Markets." *Journal of the European Economic Association*, 1(4): 990-1029.

[22] **Rochet, Jean-Charles, and Jean Tirole**. 2006. "Two-sided Markets: A Progress Report." *Rand Journal of Economics*, 37(3): 645-667.

[23] **Sandvig, C., Hamilton, K., Karahalios, K., and Langbort, C.** 2014. "Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms." In *Proceedings of "Data and Discrimination: Converting Critical Concerns into Productive Inquiry," a Preconference at the 64th Annual Meeting of the International Communication Association*.

[24] **Sweeney, Latanya**. 2013. "Discrimination in Online Ad Delivery." *Communications of the Association of Computing Machinery*, 56(5): 44-54.