An Overview of Crowd-based Markets and Racial Discrimination

Completed Research

Lauren Rhue
Wake Forest University School of Business
rhuela@wfu.edu

Abstract

Crowd-based markets are gaining in prominence as providers of goods and services (Airbnb, Uber), sources of employment (TaskRabbit, Mechanical Turk), and sources of capital (Kickstarter, IndieGogo). Although these markets have an open and egalitarian premise, multiple studies have found evidence of racial discrimination. This paper provides an overview of the existing research in the racial discrimination of crowd-based markets and examines their measurement choices, recommendations to reduce discrimination, and the ethical values embedded in the anti-discrimination suggestions. Based on the survey of the current literature, this analysis suggests new research avenues and highlights automation as the most promising new avenue to reduce racial discrimination in the crowd.

Keywords

Crowdfunding, sharing economy, crowd-based platforms, race, discrimination, ethnicity.

Introduction

After struggling to book a room on Airbnb, Quirtina Crittenden blogged about her experience on Twitter using the hashtag #AirbnbWhileBlack. Her friends recommended that she shorten her name to "Tina" and change her profile photo to a cityscape photo, and subsequently Crittenden had no problems booking a room (Penman, Vedantam, and Nesterak, 2016). Although she launched the hashtag #AirbnbWhiteBlack, Crittenden's experience is a single narrative and there could be other, non-discriminatory explanations for her difficulties on Airbnb. Are there other anecdotes of racial discrimination?

Dyne Suh, a UCLA law student, experienced overt racial discrimination on Airbnb. Her host, Tami Baker, canceled Suh's reservation for a Big Bear cabin saying that she would not rent to an Asian (Wang, 2017). Suh complained, and eventually the California Department of Fair Employment and Housing reached an agreement with Baker that she would pay $5,000 in damages and complete a course in Asian-American studies (Martin, 2017).

Incidents like these are a public relations problem for a platform like Airbnb, which has grown into a multi-billion dollar company and the largest provider of short-term accommodations (Sundararajan, 2018). More generally, crowd-based capitalism is a new hybrid between the firm and the market (Sundararajan, 2018) in which a platform facilitates economic transactions among participants. Crowd-based platforms are expected transform our economy by building a "gig economy" of on-demand workers (Greenwood, Burtch, and Carnahan, 2017), by creating a "sharing economy" of access to goods and services (Sundararajan, 2018), and by opening additional avenues for fundraising. The egalitarian and open nature of these platforms is thought to democratize access to funding, goods, services, and workers, ushering in an era of economic opportunities.

Despite the egalitarian promise of crowd-based platforms like Airbnb, experiences of Crittenden and Suh highlight an ongoing challenge with respect to peer-to-peer markets. Members of the crowd can exhibit bias, which decreases the attractiveness of the platform to participants and exposes the platform to legal liability. This paper discusses the evidence of racial discrimination across different peer-to-peer markets in a number of recent studies. Although the final takeaways are consistent among all studies—certain racial groups experience worse outcomes on crowd-based platforms—these studies differ in their operationalization of race, their success measures, and their recommendations to reduce discrimination.

This study contributes to the information systems literature by coalescing the existing research on discrimination in crowd-based markets, examining the research methodologies, and discussing the recommendations. Based on this analysis, automation is discussed as a promising means to reduce racial discrimination while balancing the ethical values of fairness and transparency. The information systems literature is starting to consider the design challenges in bias reduction, and this study focuses on the opportunities and challenges in regulating the crowd.

Peer-to-Peer Markets

Crowd-based, peer-to-peer platforms are characterized by their ability to connect at least two distinct sets of participants and facilitate transactions among those disparate parties. Crowd-based markets encompass multiple categories of peer-to-peer platforms: the sharing economy, crowdwork, and crowdfunding.

Crowd Categories

Sharing economy. Sharing economy platforms enable individual sellers to "share" (provide access to a good or service) with buyers for a fee (Sundarajan, 2016).[1] In the sharing economy, the participant terminology differs based on the context. For home-sharing platforms like Airbnb, the sellers are the "hosts" and the buyers are the "guests". For ride-sharing platforms like Uber and Lyft, the sellers are the drivers and the buyers are the riders. To maintain consistency across contexts, this paper will use the terms buyers and sellers.

Crowdwork. Crowdwork platforms enable employers to broadcast job opportunities and find workers to complete these various tasks. In this instance, workers are sellers because they sell their labor, and employers are consumers because they purchase this labor.

Crowdfunding. Crowdfunding platforms enable fundraisers to broadcast a request for funding for a particular project or campaign, and backers can browse the requests and contribute funds. In this context, the fundraisers are the sellers because they sell the potential access to a product, and the project contributors are the buyers because they purchase the potential access or rewards.

Crowd Challenges

Regulation and discrimination are important challenges in all crowd-based markets. For creative work, workers are sensitive to the division of revenue between the worker and the platform, feeling that it is "unfair" for the platform to reap more than a certain amount of the revenue (Franke et al., 2013). For rote tasks such as those on Amazon Mechanical Turk (AMT), there is an increasing call for labor protections, noting that social harms and regulatory failures can occur (Felstiner, 2011). Airbnb and Uber are fighting the classification of hotel and taxi-service respectively, but there are questions about how to regulate these platforms to ensure consumer protections. Airbnb recently allowed government regulators to check hosts for discrimination (Martin, 2017). Regulations already exist to protect consumers from discrimination by businesses; it is less clear how government entities can protect consumers from other market participants. There are a number of recommendations such as regulating crowdsourcing vendors and changing the employee classification of crowd workers on platforms like MTurk and TaskRabbit (Felstiner, 2011).

Platforms can also intervene to promote labor protections such as a default minimum wage, reputation portability (Felstiner, 2011), and payment transparency (Bederson and Quinn, 2011). There are numerous pending lawsuits regarding the state of the workers on these platforms (Codagnone, 2016). To increase overall fairness, Borromeo et al. (2017) suggest increasing fairness in the task assignment process and limiting the anonymity of both workers and requesters. Since fairness is defined according to the attributes of the individuals (Borromeo et al., 2017), this idea of unfair task assignment and opacity in payment choices could lead to race-based discrimination, as requesters prefer particular groups of workers.

[1] Sharing economy platforms that control the product supply, such as Rent-the-Runway or Zipcar, are not peer-to-peer markets and are excluded.

Crowd-based Platforms: A Foundation of Trust

In crowd-based markets, participants may not have a prior relationship, so platforms must engender sufficient trust among the participants to create economic transactions. If participants do not feel that there is sufficient information about the other participants in the market, then high quality participants could leave and the market would disintegrate (Tomboc, 2013). In order to facilitate economic transactions among participants, crowd-based platforms aim to produce trust among participants (Sundararajan, 2016), e.g., driver and rider in ride-sharing; guest and host in home-sharing; worker and employer in crowdsourcing; fundraiser and backer in crowdfunding. To that end, crowd-based platforms include elements such as social network information (Tomboc, 2014) as well as user reputations and detailed user profiles with pictures (Ert et al., 2016). Each of these elements can convey information about the racial identity of the market participant (Rhue and Clark, 2016); therefore, any discussion about the discrimination in crowd-based capitalism literature needs to consider the conceptualization of race.

Discrimination in the Crowd

Methodology

This paper reviews the literature of racial discrimination in crowd-based markets, as described by Webster and Watson (2002). Papers were discovered by examining the incoming and outgoing citation network for the most recent crowd-based platforms papers: Younkin and Kuppuswamy (2017), Ye et al. (2017), Edelman et al. (2017), and Edelman and Luca (2014). To find potentially unconnected papers, Google Scholar keyword searches were performed for race-related and crowd-related papers (e.g., "racial sharing economy", "discrimination sharing economy", "discrimination crowdfunding"). There were two main criteria for inclusion in this review. First, the studies must examine crowd-based markets. Racial discrimination in a broader context, such as racial discrimination in labor markets, is outside of the scope of this analysis. Second, the studies must discuss racial discrimination, so studies that examine bias due to gender, religion, nationality or ethnicity are excluded from this review. One of the key components of this study is the measurement and operationalization of race, so this paper focuses on studies that have grappled with defining race and measuring its impact in a peer-to-peer context.

Given these criteria, racial discrimination has been found in numerous crowd-based platforms: the crowd-lending platform Prosper (Pope and Sydnor, 2011), the home-sharing platform Airbnb (Edelman and Luca, 2014), Craigslist (Ghoshal and Gaddis, 2015), the auction platform eBay (Ayers et al., 2015), the ride-sharing platform Uber (Ye, Pierce, Alahmad, Robert, 2017; Greenwood, Adjerid, and Angst, 2017) and the crowdfunding platform Kickstarter (Younkin and Kuppuswamy, 2017). These studies differ in their operationalization of race, in their methodology, in the measured discriminatory outcomes, and in the recommendation to reduce bias in the future.

Measuring and Operationalizing Race

Race is a social construct that is considered to be a combination of self-identification constrained by external perception (Nagel, 1994). Race can be measured in a number of ways. In the social sciences, it is common for researchers to use questionnaires and ask their subjects to identify their affiliation with particular racial and ethnic groups (Davis and Engel, 2010). In the information systems literature, this approach is less common, and most studies determine whether a participant belongs to a particular racial category by analyzing his/her name or his/her perceived race in a photo. By using a single method to determine race, researchers may find race-based differences when race is actually a proxy for another variable such as socioeconomic status (Winker, 2004).

The use of distinctively racial names to measure a subject's race is problematic as it can conflate race with socioeconomic status and education. Distinctively African American names are often proxies for socioeconomic status of the parents (Fryer and Levitt, 2004). Plus, Gaddis (2017) finds that the names used by highly educated African-American mothers such as Jalen and Nia are perceived as less black than names used by less educated African-American mothers such as DaShawn and Tanisha, highlighting the confluence of race and parental education when names are used to measure race.

In addition to names, research on crowd-based discrimination measures perceived race by users' pictures, typically using skin tone or facial recognition software to identify race. In studying discrimination, it is appropriate to measure race based on user perception and not the participant's self-identification because it mirrors the process by which other users infer a person's race and act on that inference (Rhue and Clark, 2016). One challenge with defining race based on visual features, particularly skin tone, is the disregard of skin color differences within particular races and the reduction of race to mutually exclusive categories.

Sharing economy				
Study	*Groups*	*Race Concept*	*Finding*	*Recommendation*
Ge et al., 2017	• African-American (AA) • White	Name (Field Study)	AA riders experienced longer wait for their trip to be accepted AA riders had more cancelations	• Concealment • Transparency • Data analysis • Fixed fares
Cui et al., 2016	• AA • White	Name (Field Study)	AA guests had acceptance rates of 29% (vs. 49%) AA guests with a positive review saw no difference (56% vs. 58%)	• Reputation systems
Ye et al., 2017	• AA • White	Photo (Lab Exp.)	Booking intention on Airbnb	*None*
Edelman and Luca, 2014	• AA • White	Photo	AA hosts seek and receive lower prices than white hosts	• Concealment
Edelman et al., 2017	• AA • White	Name (Field Study)	Guests with AA names receive 42% positive response (vs. 50%)	• Concealment • Instant bookings • Data analysis
Airbnb	N/A	Photo (Complaints)	Complaints about racial and gender discrimination	• Less photo prominence • Penalties
Crowd funding				
Study	*Groups*	*Race Concept*	*Finding*	*Recommendation*
Rhue and Clark, 2016	• AA • Asian • White	Photo + Text (Observational)	AA are half as likely to be funded as white or Asian fundraisers Fundraisers with AA-sounding text are less successful	• Concealment • Promotions
Younkin and Kuppuswamy, 2017	• AA • White	Photo (Observational)	AA fundraisers experience lower success	• Concealment
Pope and Sydnor, 2011	• AA • White	Photo (Observational)	Loans with AA are less likely to receive funding	• Concealment
Economic markets				
Study	*Groups*	*Race Concept*	*Finding*	*Recommendation*
Doleac and Stein, 2013	• AA • White	Photo (Field Exp.)	AA receive 13% fewer responses and 18% fewer offers on iPods	• None; competition will stop discrimination
Ayres et al., 2015	• AA • White	Photo (Field Exp.)	AA receive 20% lower price on baseball cards	• Concealment

Table 1. Studies by the Racial Signal, Outcome, and Recommendation

Because most studies focus on race as a dichotomy between white and black, there is a need to study the multifaceted racial landscape and potential interactions among groups. Rhue and Clark (2016) compare three racial groups and find that African-Americans have distinctly lower success rates than Asian and white fundraisers, whereas the difference in the success rates between Asian and white fundraisers is not significant. The observed success of Asians is not given, as previous studies in offline labor markets have indicated the potential for racial discrimination against Asians (Sun Young et al., 2016).

Summary of Findings

Study Results

Table 1 summarizes the research along several important dimensions. Across all these studies, African-Americans fare worse in crowd-based markets. In the ride-sharing context, African-American riders experience longer wait times and more cancelations (Ge, Knittel, MacKenzie, and Zoepf, 2017). In field studies on economic platforms, African-Americans receive less interest and lower final offers for identical goods (Ayers et al., 2015; Doleac and Stein, 2013). In observational analysis of crowdfunding, African-American fundraisers are less likely to reach their target goal (Younkin and Kuppuswamy, 2017; Rhue and Clark, 2016). In experiments and field studies on home-sharing platforms, African-Americans hosts seek and receive lower prices (Edelman and Luca, 2014) and are less likely to receive responses as guests (Edelman et al., 2017; Cui, Li, and Zhang, 2017). There is a consensus that African-Americans are less successful in crowd-based platforms, regardless of racial measurement, economic objective, and methodology. There is a surprising dearth of information about racial discrimination in online labor markets, potentially due to a reluctance of companies to investigate digital discrimination (Fisman and Luca, 2016).

The studies included in this review disagree on two important dimensions: the mechanisms that lead to discrimination and the recommendations to reduce discrimination. These studies typically identify three potential mechanisms for discrimination: statistical discrimination, taste-based discrimination, and in-group preferences. Across these studies, researchers find evidence for each of these forms of discrimination, and based on their analysis of the discriminatory mechanism, there are several recommendations to reduce discrimination. The next subsection discusses the three discriminatory mechanisms in more depth.

Mechanisms for Discrimination

Three potential discriminatory mechanisms are identified in these crowd-based platforms: statistical discrimination, taste-based or animus-based discrimination, and in-group preferences (Doleac and Stein, 2013). Statistical discrimination occurs when the buyer or seller uses race as a proxy for other quality-related attributes about the other transacting party, so additional information about the quality of that party would reduce future discrimination (Cui, Li, and Zhang, 2016). Taste-based discrimination occurs when the buyer or seller has a dislike of a particular race and avoids transacting or interacting with that race (Edelman and Luca, 2014). Statistical discrimination is often cited as the primary mechanism leading to racial discrimination (Pope and Sydnor, 2011; Cui et al., 2017); however, researchers observe that multiple discriminatory mechanisms can co-exist. Statistical discrimination can inform animus-based discrimination and vice versa (Doleac and Stein, 2013), suggesting the need for a more nuanced understanding of discriminatory mechanisms (Edelman et al., 2017). Lastly, people can prefer to transact with racial groups that mirror their own, forming an in-group preference (Rhue and Clark, 2016; Edelman et al., 2017). In this case, the platform would need to attract a sufficiently diverse set of participants in all roles in order to see improvement in the racial disparities of success outcomes. By identifying the mechanisms that underlie the observed discrimination, researchers aim to produce recommendations to reduce discrimination on the platform.

Recommendations for Bias Reduction

There are several common recommendations to reduce discrimination in crowd-based platforms: concealment, transparency, data analysis, and automation, and these recommendations are not necessarily consistent with each other. For example, Cui et al. (2016) recommend increasing the amount of reputation information and disagree with the de-emphasis of photos, as recommended by Younkin and Kuppuswamy (2017) and promoted by Airbnb in its revised guidelines (2016).

As a guiding principle, the community should consider *why* anti-discrimination measures are important in crowd-based markets. Discrimination is illegal in particular contexts, like short-term accommodations (Martin, 2017), but not in all contexts. There are no laws or regulations against racial discrimination in crowdfunding or hiring independent contractors. The prevalence of studies into discrimination and recommendations to reduce digital discrimination (Fisman and Luca, 2016) suggests that the imperative to reduce discrimination arises from an ethical notion of fairness and not from legal requirements.

Fairness matters in peer-to-peer interactions. Individuals are less likely to engage with a platform that engages in "unfair" behaviors like profiting from their users' creativity or not intervening in disputes (Franke et al., 2013). Researchers are recommending ways to ensure fair wage compensation for crowd workers (Bederson and Quinn, 2013) and axioms to espouse fairness and transparency in crowd-based marketplaces (Borromeo et al. 2017). Similarly, the recommendations to mitigate racial bias on peer-to-peer platforms can be analyzed through the goal of fairness, considering effectiveness and transparency as auxiliary goals and means to achieve fairness.

Concealment

The most consistent recommendation to reduce discrimination on crowd-based platforms is that users should conceal any markers that indicate their racial identity (Younkin and Kuppuswamy, 2017; Ge et al., 2017; Edelman et al., 2017). The removal of all racial identifiers may be possible in some situations, such as not showing the seller's hand when listing an iPod for sale online (Doleac and Stein, 2013). For platforms like Uber and Airbnb that require a name and picture, users would need to choose generic nicknames and generic profile photos in order to effectively conceal their race. In addition to anecdotal evidence (Penman et al., 2016), evidence from Kickstarter indicates that success rates do not drop with generic profile photos (Rhue and Clark, 2016), indicating that this recommendation may be an effective short-term solution to reduce discrimination. Furthermore, concealment is an effective strategy regardless of the discriminatory mechanism, preventing other users from incorporating race into their decisions.

Despite its effectiveness, this recommendation raises multiple important ethical questions. Who shoulders the burden for anti-discrimination measures (Younkin and Kuppuswamy, 2017)? This is the sole recommendation that puts the burden of bias reduction on the discriminated party instead of the platform or the government. Crowd-based platforms espouse user transparency to promote trust, so this suggestion that some users obscure or conceal aspects of their identity runs counter to user transparency. Also, should African-Americans be required to remove racial markers to fully participate in crowd-based markets? This recommendation, in the logical conclusion, leads to the complete removal of African-American faces and vernacular from crowd-based platforms.

Reputational and Attribute Transparency

Some studies take a different approach and recommend that platforms increase transparency as a means to provide a better experience for discriminated parties and/or to reduce discrimination. Ge et al. (2017) propose giving decision-makers more information, including racial identifiers, earlier in the decision-making process to forestall rejections later in the transaction. An earlier non-acceptance may create a better overall experience for the discriminated party than a series of later rejections. Cui et al. (2016) recommend more visible user reputation systems as a means to reduce racial discrimination, suggesting that market participants will place higher weight on the user reputation than on the racial characteristics in the user name. One caveat, this increased reputational transparency relies on an unbiased rating system, which may not occur in the real world as certain groups are more heavily penalized for particular mistakes (Greenwood et al., 2017).

As a recommendation, reputational and identifier transparency adhere to our notion of fairness. This recommendation places the burden of anti-discrimination measures on the platform, not the user, leading to a fairer division of the bias reduction labor. However, this recommendation only addresses racial discrimination if the underlying mechanism is statistical discrimination. This reputational transparency may reduce uncertainty about user quality and encourage participants to view each other as individuals and not as stereotypes. If racial discrimination is occurring due to taste-based or similarity-based preferences, however, then it would persist even with positive reviews and reduced uncertainty.

Policies and Data Analysis

There is evidence that only certain subsets of participants discriminate (Edelman et al., 2017), so platforms and regulators can perform data analysis to identify participants who engage in racial discrimination. There are already anti-discrimination laws and policies in many platforms, so data analysis would increase the compliance to these policies by identifying the violators (Murphy, 2016) and penalizing them (Martin, 2017). In this manner, anti-discrimination policies and data analysis must occur together to promote individual compliance and policy enforcement.

As a recommendation, anti-discriminatory data analysis promotes fairness and transparency. The burden to reduce discrimination is placed on platforms and regulators, not the affected users. Plus, data analysis is agnostic to the discriminatory mechanism because it would identify any form of racial discrimination. Its effectiveness depends on the quality of the data and the expertise of the analysts, so regulators and platforms should work in concert to identify participants engaging in discriminatory behaviors. Platforms would need to share their transactional data to identify participants engaging in discriminatory transactions, and this openness to the government may not be embraced by all companies (Fisman and Luca, 2016).

Automation

This final category of recommendations is a catch-all for the set of recommendations that remove choice from the individual and place it into an automated system. For example, Airbnb is promoting "Instant Book" as a means to reduce discrimination (Murphy, 2016; Edelman and Luca, 2014). As described in Murphy (2016), Instant Book is an expert system that automatically books any potential guest who satisfies the host's pre-defined criteria. By forcing hosts to lay out the acceptance criteria in advance, Airbnb is reducing the potential for hosts to engage in discriminatory behaviors for qualified guests. This same recommendation is extended to other crowd-based platforms, such as calling for automated worker acceptance for qualified workers on Mechanical Turk (Bederson and Quinn, 2013).

Ranking and matching algorithms are another avenue to promote automation to reduce discrimination. If Kickstarter increased the visibility of specific campaigns to reduce the racial disparities in success rates (Rhue and Clark, 2016), then this algorithm could incorporate racial identifiers to counter existing racial bias. There is a growing interest in understanding the values embedded in algorithms, and platforms could invest in matching and recommendation algorithms that incorporate fairness into their design.

Automated recommendations are most effective for bias caused by statistical discrimination. If people are primarily concerned about the quality of the applicants, then setting pre-defined criteria could alleviate their concerns and reduce the reliance on race as a proxy for quality. However, automation does not necessarily address taste-based or similarity preferences. Often, there is another variable that is highly correlated with race, so participants could set that as a criterion in their preferences, creating seemingly race-neutral policies with discriminatory effects. In general, these recommendations for automation adhere to fairness, but incorporating fairness into algorithms is an emerging and interesting field.

Responsibility	**Transparency**	**Fairness**
Individual	Conceal race identifiers (-) Use reputation systems (+)	
Platform	Share and analyze data to find discrimination (+) Promote transparency in qualification for automation (+) De-emphasize user photos (-) Enhance reputation systems (+)	Use automation to augment and/or replace participant choices (+) Enact anti-discriminatory policies (+)
Government	Analyze data to find discrimination (+)	Enact anti-discrimination regulations (+)

Table 2. Recommendations by Responsible Party and Ethical Principle

Table 2 shows a summary of the recommendations, divided by the responsible party and the adherence to the ethical goals of fairness and transparency. This table highlights the expectation that platforms shoulder most of the burden for anti-discriminatory measures, although there are some recommendations for individuals and governments. It also notes that most of the recommendations are concerned with transparency as a means to either reduce or increase information asymmetry among the transacting parties.

Discussion

Emerging research has tackled the notion of racial discrimination in crowd-based platforms. Although there is a consensus that African-Americans fare worse in crowd-based markets, there is not a consensus on how to measure race, on the underlying mechanism for racial discrimination, or on the recommendations to reduce discrimination on the platforms. Plus, there is a clear need to study racial discrimination for other racial groups aside from African-Americans.

One of the primary challenges in studying racial discrimination is isolating racial discrimination from discrimination based on correlated variables such as socioeconomic status. The distinctively African-American names that are often used to study race are also correlated with parental education, so an interesting avenue for future research is to examine distinctive names for other races and their correlation with other variables. For example, it would be interesting to create a 2x2 experimental design comparing the names used by highly educated white mothers and less educated white mothers with those used by highly educated African-American mothers and less educated African-American mothers. This framework would explore whether race or socioeconomic factors like education are driving the findings of discriminatory behavior. On a related note, the research on discrimination in crowd-based platforms could expand to examine name-based discrimination against other groups with distinctive names such as Asian-Americans, Latinx, and Indian-Americans.

The disagreement about the mechanisms for racial discrimination could indicate a need to create a more nuanced discussion about the origins of discrimination. These mechanisms are probably not mutually exclusive, as the extensive history of racial divisions has led to statistical differences across racial groups as well as ongoing racial prejudice and an in-group preference. Arguably, distilling the particular mechanism for discrimination is less critical than understanding the effects of racial discrimination, which party (if any) shoulders the burden of ameliorating them, and the potential effectiveness of anti-discriminatory recommendations. Individuals, platforms, and governments are all attempting to reduce racial discrimination on crowd-based platforms to varying degrees of effectiveness. Participants like Crittenden are posting generic names and photos to conceal their race; Airbnb has enacted new policies; and regulators are responding to discrimination claims. These existing methods to eliminate racial discrimination are insufficient, so the most promising new avenue is the potential for automation to remove some of these decision-making power from bias individuals into machines.

Although automation may reduce racial discrimination, automation can also reproduce discrimination (Borromeo et al., 2017) without careful consideration of the relevant values, so researchers must examine the implications of these automated anti-discrimination technologies. For example, are people using Instant Book to weed out minorities based on a correlated variable or is it reducing racial discrimination as intended? Is it fair and/or legal to include demographic indicators like race and gender into the recommendation or matching algorithms to reduce racial disparities? Also, automation must be balanced with autonomy and agency. Although automation may remove racial bias, individual autonomy is the essence of crowd-based markets. Although Airbnb could automatically assign a house for each guest and Kickstarter could assign a campaign to each backer, a recommendation for automation, in its logical extreme, would remove participants' agency, thus removing the crowd from crowd-based markets. When platforms promote automated systems, they must consider agency along with the fairness and bias reduction.

At the heart of this research agenda is the notion of fairness. At its root, we reject discrimination because of the presumption of unfairness–people with similar characteristics are treated differently because of irrelevant information. The push to understand racial discrimination in crowd-based platforms could lead to more equitable and inclusive crowd-based marketplaces.

REFERENCES

Ayres, I., Banaji, M., & Jolls, C. 2015. "Race effects on eBay," The RAND Journal of Economics (46:4), pp. 891-917.

Bederson, Benjamin B., and Alexander J. Quinn. 2011. "Web workers unite! addressing challenges of online laborers," CHI'11 Extended Abstracts on Human Factors in Computing Systems. ACM, 2011.

Bertrand, M., and Mullainathan, S. 2004. "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," The American Economic Review (94:4), pp. 991–1013. Retrieved from: www.jstor.org/stable/3592802.

Borromeo, R.M., Laurent, T., Toyama, M., and Amer-Yahia, S. 2017. "Fairness and Transparency in Crowdsourcing," Proceedings of the 20th International Conference on Extending Database Technology (EDBT), March 21-24, 2017 - Venice, Italy. ISBN 978-3-89318-073-8.

Codagnone, C., Abadie, F., and Biagi, F. 2016. "The Future of Work in the 'Sharing Economy'. Market Efficiency and Equitable Opportunities or Unfair Precarisation?" Institute for Prospective Technological Studies, Science for Policy report by the Joint Research Centre, May 27, 2016. Available at SSRN: http://dx.doi.org/10.2139/ssrn.2784774.

Cui, R., Li, J., and Zhang, D. 2016. "Discrimination with Incomplete Information in the Sharing Economy: Evidence from Field Experiments on Airbnb," (December 8, 2016). Available at SSRN: https://ssrn.com/abstract=2882982.

Cui, R., Li, J., and Zhang, D. 2017. "A Better Way to Fight Discrimination in the Sharing Economy," Harvard Business Review 2017. Available online at: https://hbr.org/2017/02/a-better-way-to-fight-discrimination-in-the-sharing-economy

Davis, L. E., and Engel, R.J. 2010. *Measuring race and ethnicity*. Springer Science & Business Media. DOI: 10.1007/978-1-4419-6697-1.

Doleac, J. L. and Stein, L.C.D. 2013. "The Visible Hand: Race and Online Market Outcomes," Economics Journal (123:572) pp. F469–F492. doi:10.1111/ecoj.12082.

Edelman, B. G., and Luca, M. 2014. "Digital Discrimination: The Case of Airbnb.com," *Harvard Business School NOM Unit Working Paper* (:14-054).

Edelman, B., Luca, M., and Svirsky, D. 2017. "Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment," American Economic Journal: Applied Economics (9:2), pp. 1-22. DOI: 10.1257/app.20160213.

Ert, E., Fleischer, A., & Magen, N. 2016. "Trust and reputation in the sharing economy: The role of personal photos in Airbnb," Tourism Management (55), pp. 62-73.

Felstiner, A. 2011. "Working the Crowd: Employment and Labor Law in the Crowdsourcing Industry," Berkeley Journal of Employment and Labor Law (32:1), pp. 143-203. Available at: http://www.jstor.org/stable/24052509

Fisman, R. and Luca, M. 2016. "Fixing Discrimination in Online Marketplaces," Harvard Business Review December 2016. Retrieved online from: https://hbr.org/2016/12/fixing-discrimination-in-online-marketplaces.

Franke, N., Keinz, P., and Klausberger, K. 2013. "Does This Sound Like a Fair Deal? Antecedents and Consequences of Fairness Expectations in the Individual's Decision to Participate in Firm Innovation," Organization Science (24:5), pp. 1495–1516.

Fryer, R.G. and Levitt, S.D. 2004. "The Causes and Consequences of Distinctively Black Names." Quarterly Journal of Economics (119:3), pp. 767–805.

Gaddis, S.M. 2017. "How Black Are Lakisha and Jamal? Racial Perceptions from Names Used in Correspondence Audit Studies.," Working Paper, available online at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2807176

Ge, Y., Knittel, C. R., MacKenzie, D., and Zoepf, S. 2016. "Racial and gender discrimination in transportation network companies," National Bureau of Economic Research, No. w22776.

Ghoshal, R., & Gaddis, S. M. 2015. "Finding a Roommate on Craigslist: Racial Discrimination and Residential Segregation." Working paper, Available on: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2605853.

Greenwood, B. N., Adjerid, I., & Angst, C. M. 2017. "How Unbecoming of You: Gender Biases in Perceptions of Ridesharing Performance," Working paper, available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2949511

Greenwood, B. Burtch, G., and Carnahan, S. 2017. "Unknowns of the gig-economy." Communications of the ACM (60:7), pp. 27-29.

Martin, H. 2017. "Airbnb host must pay $5,000 for canceling reservation based on race," Los Angeles Times, July 13, 2017. Available at: http://www.latimes.com/business/technology/la-fi-airbnb-discrimination-20170407-story.html

Murphy, L.W. 2016. "Airbnb's Work to Fight Discrimination and Build Inclusion," A Report Submitted to Airbnb. September 8, 2016. Available at: https://blog.atairbnb.com/wpcontent/uploads/2016/09/REPORT_Airbnbs-Work-to-Fight-Discrimination-and-Build-Inclusion.pdf?3c10be

List, J. A. 2004. "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field," The Quarterly Journal of Economics (119:1), pp. 49–89. Available at: www.jstor.org/stable/25098677.

Nagel, J. 1994. "Constructing Ethnicity: Creating and Recreating Ethnic Identity and Culture," Social Problems (41:1), pp. 152–176. www.jstor.org/stable/3096847.

Penman, M., Vedantam, V., and Nesterak, M. 2016. "#AirbnbWhileBlack: How Hidden Bias Shapes the Sharing Economy," NPR. Available at: https://www.npr.org/2016/04/26/475623339/-airbnbwhileblack-how-hidden-bias-shapes-the-sharing-economy.

Pope, D. G., and Sydnor, J. R. 2011. "What's in a Picture? Evidence of Discrimination from Prosper.com," Journal of Human Resources (46:1), pp. 53–92.

Rhue, L., and Clark, J. 2016. "The Cost of Authenticity: Quantifying the Racial Signals. Working paper, available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2837042.

Silberman, M., Irani, L., and Ross, J. 2010. "Ethics and Tactics of Professional Crowdwork," XRDS: Crossroads, The ACM Magazine for Students (17:2) pp. 39-43.

Smith, S. S. 2010. "Race and Trust," Annual Review of Sociology (36), pp. 453-475.

Sun Young, L., Pitesa, M., Thau, S., and Pillutla, M.M. 2015. "Discrimination in Selection Decisions: Integrating Stereotype Fit and Interdependence Theories," Academy Of Management Journal, (58:3), pp. 789-812. doi:10.5465/amj.2013.0571

Sundararajan, A., 2016. *The Sharing Economy: The End of Employment and the Rise of Crowd-Based Capitalism*. MIT Press, Cambridge, MA.

Sundararajan, A. 2018. "Crowd-Based Capitalism, Digital Automation, and the Future of Work," University of Chicago Legal Forum: Vol. 2017, Article 19. Available at: https://chicagounbound.uchicago.edu/uclf/vol2017/iss1/19

Tomboc, G. 2014. "The Lemons Problem in Crowdfunding," The John Marshall Journal of Information Technology & Privacy Law (30:2) pp. 253 (2013-2014).

Wang, A. B. 2017. "'One word says it all. Asian': Airbnb host banned after allegedly canceling guest because of her race," Los Angeles Times, April 7, 2017. Available at: http://www.latimes.com/business/technology/la-fi-airbnb-discrimination-20170407-story.html

Webster, J., & Watson, R. T. 2002. "Analyzing the past to prepare for the future: Writing a literature review," MIS quarterly, xiii-xxiii.

Winker, M.A. 2004. "Measuring Race and Ethnicity: Why and How?" JAMA October 6, 2004. (292:13) pp. 1612-14.

Ye T., Pierce, C., Alahmad, R., Robert, L. P. 2017. "Race and Rating on Sharing Economy Platforms: The Effect of Race Similarity and Reputation on Trust and Booking Intention in Airbnb," Proceedings of the 38th International Conference on Information Systems (ICIS 2017) in Seoul, Korea. Available: http://aisel.aisnet.org/icis2017/Peer-to-Peer/Presentations/4/.

Younkin, P., and Kuppuswamy, V. 2017. "The Colorblind Crowd: Founder Race and Performance in Crowdfunding," Management Science, *forthcoming*.