FEDERAL RESERVE BANK OF RICHMOND®
Richmond ▪ Baltimore ▪ Charlotte

June 2020

SaMMF Workshop: Discrimination in Labor Markets

Summary of Presented Research

June 26, 2020

1:00 p.m.
Peter Norman
Economist, University of North Carolina at Chapel Hill
A Search Model of Statistical Discrimination

1:40 p.m.
J. Aislinn Bohren
Economist, University of Pennsylvania
Inaccurate Statistical Discrimination: An Identification Problem

2:30 p.m.
Suqin Ge
Economist, Virginia Tech
Testing for Asymmetric Employer Learning and Statistical Discrimination

3:20 p.m.
Arjada Bardhi
Economist, Duke University
Early-Career Discrimination: Spiraling or Self-Correcting?

4:00 p.m.
Kyungmin (Teddy) Kim
Economist, Emory University
Statistical Discrimination in Ratings-Guided Markets

The Search and Matching in Macro and Finance (SaMMF) Virtual Seminar Series was created to facilitate the exchange of ideas in macroeconomics and finance with an emphasis on search and matching, networks, intermediation, market structure and design, information frictions, and other issues surrounding decentralized markets. Given recent events and discussions about race and discrimination in the United States, SaMMF held a virtual seminar on June 26 to discuss research on discrimination in labor markets. The papers presented highlighted insights from search and matching, networks, and employer learning.

While this seminar was not organized by the Richmond Fed, Bruno Sultanum, an economist at the Richmond Fed, is one of the organizers of the SaMMF series. The workshop was co-organized by Hanming Fang of the University of Pennsylvania.

Within this publication, you will find summaries of the research discussed at the workshop. Papers, slides, and a video recording of the entire workshop can be found at https://sammf.com/sammf-workshop-discrimination/.

A Search Model of Statistical Discrimination

By Jiadong Gu (University of North Carolina at Chapel Hill) and Peter Norman (University of North Carolina at Chapel Hill)

Much of the gender wage gap can be explained by the different occupational choices made by men and women. One common interpretation of this difference in choice is that women place a higher value on certain job amenities than men. Another common explanation is that women gravitate toward low-skill jobs as a result of lower confidence or higher risk aversion than men. Gu and Norman develop a model to show how job choice, search externalities, and signal extraction problems can combine to create unequal representation in jobs.

In their model, firms are either high-tech or low-tech, and workers are either qualified or unqualified. All workers are equally productive in low-tech firms, but only qualified workers are productive in high-tech firms. As a result, high-tech firms wish to hire only qualified workers, but they cannot directly observe whether a worker is qualified. They must extract that information from noisy signals, such as a job interview, test, or other screening device, as well as the probability that a worker is qualified. In the model, this probability is simply the proportion of qualified workers in the total pool of potential workers. If qualified female workers accept low-tech positions at a higher rate than qualified men, it reduces the proportion of qualified women in the overall pool of job seekers, making it less likely that a high-tech firm will hire a woman. This creates a self-fulfilling prophecy — women are more incentivized to accept low-tech positions because high-tech firms are less likely to hire them.

Gu and Norman find that affirmative action policies can be effective at eliminating this asymmetry. Whereas in traditional statistical discrimination models affirmative action may reduce the incentive of the targeted group to obtain human capital, generating perverse welfare effects, this unintended consequence does not occur in Gu and Norman's model.

Inaccurate Statistical Discrimination: An Identification Problem

By J. Aislinn Bohren (University of Pennsylvania), Kareem Haggag (Carnegie Mellon University), Alex Imas (Carnegie Mellon University), and Devin G. Pope (University of Chicago)

Economists typically distinguish between two types of discrimination: taste-based and statistical. In the first type, an individual or company discriminates against a group of people due to animus. In the second type, discrimination arises when an individual or firm uses observable characteristics (e.g., race or gender) as a proxy for unobservable characteristics (e.g., productivity).

Bohren, Haggag, Imas, and Pope argue that in many cases, beliefs about the productivity of different groups may not be accurate. These inaccuracies may stem from biases or simply a lack of information. The extent to which discrimination is driven by inaccurate beliefs affects the types of policy responses one might consider to reduce discrimination.

Many studies that attempt to identify the source of discrimination assume that beliefs are accurate. As a result, if researchers rule out statistical discrimination, the underlying cause of the observed discrimination is classified as taste-based.

Using a theoretical model and stylized experimental setting, Bohren, Haggag, Imas, and Pope show how statistical discrimination due to inaccurate beliefs can be misclassified as taste-based discrimination. They also show how either eliciting beliefs or manipulating information can separate taste-based sources from accurate or inaccurate belief-based sources. Further, the authors show that employers who received information about productivity averages by group adjusted their behavior to reduce discrimination. This suggests that policy interventions that provide accurate information may be an effective way to reduce discrimination that is driven by inaccurate beliefs.

Testing for Asymmetric Employer Learning and Statistical Discrimination

By Suqin Ge (Virginia Tech), Andrea Moro (Vanderbilt University), and Beibei Zhu (Slack)

Many studies have demonstrated the presence of statistical discrimination in the labor market. This is when employers use observable demographic characteristics, such as race, as proxies for unobservable worker characteristics, such as productivity. Other studies have also shown that employers learn about the unobservable characteristics of their employees over time. This implies that as a worker gains tenure, firms will know more about that worker's true productivity and will rely less on statistical discrimination. A key question is whether this accumulated information follows a worker to different firms or whether it stays within individual firms. In the latter case, workers in discriminated groups will face renewed discrimination whenever they change jobs as the information about their true skills obtained by their previous employer is lost.

Ge, Moro, and Zhu find evidence that employers learn asymmetrically about workers' skills and statistically discriminate against non-college-educated black workers. Black workers without a college education suffer an initial wage penalty, but as their tenure with the same employer increases, their wages become more correlated with their skills. Employers learn about the true productivity of their workers over time and no longer rely on imperfect proxies. However, this information is asymmetrically shared — other employers do not have access to it. As a result, workers face renewed discrimination when they change jobs.

In contrast with non-college-educated workers, Ge, Moro, and Zhu find no evidence of asymmetric learning or statistical discrimination for college-educated workers, suggesting that employers can more easily observe those workers' productivity upon their initial entry into the labor market. This implies that policies that are able to rectify information asymmetry in the labor market could be effective at reducing statistical discrimination, particularly for non-college-educated workers.

Early-Career Discrimination: Spiraling or Self-Correcting?

By Arjada Bardhi (Duke University), Yingni Guo (Northwestern University), and Bruno Strulovici (Northwestern University)

Do the wage and employment effects of early-career discrimination in the labor market fade over time as employers learn about workers' productivities, or do workers who face early-career discrimination suffer intensifying wage losses throughout their careers because they miss out on early opportunities critical for career advancement? Bardhi, Guo, and Strulovici find that the answer depends on the environment in which employers learn about workers.

They consider two job environments: breakthrough and breakdown. In a breakthrough environment, employers reward successes. Examples of this type include researchers and salespeople. In a breakdown environment, employers reward avoiding failures. Examples of this include airline pilots and surgeons. In the breakthrough environment, high-productivity workers generate breakthroughs at randomly distributed times, and low-productivity workers generate none. In the breakdown environment, low-productivity workers generate failures at randomly distributed times, while high-productivity workers do not. They show that in the breakthrough environment, early-career discrimination is self-correcting. If an employer hires a worker from a group because he or she believes this group is slightly more productive but the worker fails to produce a breakthrough, the employer will quickly start to assign tasks to workers from the other group. Because the discriminated group does not wait long to get a chance, the initial statistical discrimination is corrected quickly. In contrast, discrimination in a breakdown environment can lead to spiraling inequality. Employers in this environment will assign tasks to workers until they demonstrate low productivity by making a mistake. But mistakes are rare, even when the worker has low productivity. As long as no mistakes are made, workers from groups that face initial discrimination will not get an opportunity to perform tasks, even if the productivity difference between the two groups is arbitrarily small. This suggests that the way employers learn about workers — not just the speed at which they learn — is key to understanding the long-lasting effects of early-career discrimination.

Statistical Discrimination in Ratings-Guided Markets

By Yeon-Koo Che (Columbia University), Kyungmin (Teddy) Kim (Emory University), and Weijie Zhong (Stanford University)

Discrimination is prevalent in many markets, but in theory, online sharing platforms like Airbnb or Uber should be less prone to statistical discrimination. This is because many online marketplaces employ a rating system, allowing users to share information about their experiences and make recommendations. Ratings enable social learning among users, which should reduce reliance on biases based on characteristics like race or gender. However, Che, Kim, and Zhong argue that the availability of more information through ratings and recommendations does not necessarily lead to less discrimination. In fact, ratings can foster discrimination.

Che, Kim, and Zhong present a model in which buyers seek to match and trade with sellers. The search and matching process has frictions, and buyers' decisions are guided by imperfect information about sellers presented through ratings. Which sellers the buyers pick is not random — buyers gravitate toward the highest-rated sellers. This can create a feedback loop, where one group of sellers gets sampled more by buyers, generating more ratings and information and increasing the likelihood that future buyers will also transact with that group. If initial buyer preferences are driven by some bias in favor of one group over another, the feedback loop from the ratings system can solidify this discriminatory equilibrium.

Che, Kim, and Zhong plan to extend their research to consider potential policy responses to correct this feedback loop. One policy they hope to test is making it easier or harder for certain groups of sellers to obtain and/or retain a good rating, which might encourage buyers to explore a more diverse set of sellers. Traditional affirmative action, such as hiring quotas, might also be effective.