

**Project Title:**

*Is it the picture or a thousand words? How do job applicant photos and messaging language affect hiring outcomes on a technology-mediated labor market*

**Project Description:**

It is widely acknowledged that discrimination against African American or female job applicants persists in the labor market today. While we are beginning to unravel where in the job seeking process bias may occur, we still lack a detailed understanding of the mechanisms underlying these outcomes (Pager and Karafin 2009). In particular, it may be that a job applicant's race and gender are correlated with other factors that affect their employment chances, such as their behaviors or traits. If so, then these factors, unobservable to a researcher but observable to an employer, may explain these outcomes. We seek to advance our understanding of what information employers rely on to make hiring decisions. We plan to address this challenge by bringing to bear computationally sophisticated techniques on unprecedentedly detailed data of a hiring process and, ultimately, by designing and testing interventions in a field experiment.

Employers today have access to a breadth of technologically enabled information with which to judge their job applicants (Bachrach 2015). Social platforms like Facebook and online job platforms like LinkedIn provide a copious amount of publicly available applicant data to potential employers; applicants, perhaps unwittingly, build personal profiles filled with a variety of information, ranging from short personal biographies to pictures to portfolios of past projects. Consequently, these details are likely to alter employers' perceptions of job applicants. Therefore, we believe that there is a critical need to understand more precisely if and how employers are utilizing photo and text information to hire and whether that may lead to discrimination. We propose that by bringing computationally sophisticated analytical tools to novel data, we can shed light on processes we could not study before.

We theorize on the effects of a job applicant's self-presentation on hiring outcomes through two technologically-enabled artifacts of his or her photo and linguistic behaviors. Specifically, we plan to investigate whether and how these artifacts affect an employer's perception of important characteristics, such as the job applicant's professionalism, warmth, and trustworthiness. We suspect that a job applicant's race and gender may be correlated with how they present themselves to employers and thereby affect their likelihood of being hired through how they are perceived. For example, if African American job applicants present themselves as being less professional, this likely reduces their chances of being hired but not because of their race.

Although securing this type of data is not trivial, we have been fortunate to gain access with a local company that runs a technology-mediated labor market for low-wage jobs. Employers use

this platform to post their jobs (there are 20 categories of jobs, of which the most common are fulfillment/warehousing, delivery, and general labor). After a job is posted, push notifications with job details are sent to potential workers, who then sign up for the job. Once workers have signed up, the employer can view each worker's profile picture, contact information, overall rating on past jobs, etc. Employers can also message with potential workers over this platform. Since the platform's inception, there have been over 180,000 jobs posted, with over 800 employers and over 9,000 active job seekers. In preliminary analysis, we find that rejected application rates significantly differ by race and gender of the job seeker (see Table 1 in the Appendix).

We plan to analyze the transactions on this labor market along two, interrelated, paths. Our first path is to extract perceived characteristics from a job applicant's photo. A photo can inform an employer of both the applicant's objective characteristics as well as the applicant's behavioral characteristics. We plan to code the job applicant's photo for objective characteristics - such as age, gender, and race - using available APIs that automate this process. Then, with the aid of the Fisher Center grant, we plan to code for more behavioral characteristics - such as their professionalism, trustworthiness, and competence - by surveying individuals on Amazon Mechanical Turk. By surveying individuals on their perceptions of a subset of job applicant photos, we can build a training data set for a machine learning model that can then automate coding of the remainder of the photos. To do this, we will use novel technologies in the field of computer vision, namely convolutional neural networks (CNN); results from the 2015 ImageNet Large-Scale Visual Recognition Challenge, a world-wide computer vision competition, has demonstrated phenomenal accuracy of these models (He, Zhang, Ren, and Sun 2016). We hope to be able to detect a systematic relationship between photo characteristics and employer judgments. With a working model, we can not only reduce the cost of coding pictures, but also design interventions by precisely generating photos with specific characteristics that could alter an employer's perceptions of the person.

Our second path is to code the linguistic text messaging data. Because language is highly contextual and varied, it is difficult to predict what aspects of a text message may affect hiring perceptions. Therefore, we plan to employ several different natural language processing techniques in examining the variation in the linguistic data. First, with the help of the Fisher Center grant in acquiring an updated license, we plan to use the popular Linguistic Inquiry and Word Count (LIWC) dictionary to code each individual's messaging data (Pennebaker, Boyd, Jordan, and Blackburn 2015), providing insight into the socio-cognitive processes occurring. Second, given that the in-app messaging data is similar to text messaging, we plan to train word embeddings on similar datasets (i.e. Twitter data) for sentiment analysis. The sentiment an employee expresses to his or her employer is likely to be increasingly important in a technology-mediated setting, where there are few other means of interaction and ways to build trust. Third, we plan to use slang dictionaries to capture additional patterns of behavior for the employees. While there are a great number of natural language processing methodologies we could additionally try, we believe the three here present a strong start in understanding the linguistic data and how employers may discriminate as a result of it.

## References

- Bachrach, Y. 2015. “Human judgments in hiring decisions based on online social network profiles.” In *2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 1–10.
- He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. “Deep residual learning for image recognition.” In *2016 IEEE Conference on Computer Vision and Pattern Recognition*.
- Pager, Devah and Diana Karafin. 2009. “Bayesian Bigot? Statistical Discrimination, Stereotypes, and Employer Decision Making.” *The ANNALS of the American Academy of Political and Social Science* 621:70–93.
- Pennebaker, J.W., R. L. Boyd, K. Jordan, and K. Blackburn. 2015. “The development and psychometric properties of LIWC2015.” .

## Appendix

Table 1: Application Rejection Rate

	<i>Dependent variable:</i>
	Rejected
Age	-0.002*** (0.0001)
Gender (Male)	-0.019*** (0.003)
Distorted Photo	0.057*** (0.007)
Race (Black)	0.080*** (0.006)
Race (Hispanic)	0.005 (0.009)
Race (White)	0.029*** (0.007)
Log Avg. Prior Rating	-0.460*** (0.008)
Observations	79,490

*Note:* \*p<0.1; \*\*p<0.05; \*\*\*p<0.01. Linear regression. Difference between 1st and 99th percentile for log avg. prior rating is 0.097. Additional controls: job category and number of jobs worker previously completed on platform. Omitted race is Asian.