

Chapter 4

Using Technology to Increase Fairness in Hiring

Kelly Trindel, *pymetrics*¹

Frida Polli, *pymetrics*

Kate Glazebrook, *Applied*

SUMMARY

Traditional recruitment and hiring practices are plagued with bias. While technological advances do not offer a silver bullet, when properly designed they can work to reduce discrimination. We offer concrete steps toward fairness for technology-enabled employment selection tools. Although the current state of public conversation around technology in employment selection highlights the potential danger and recent missteps, it is important to keep

in mind that traditional analog recruitment and hiring approaches have resulted in a situation that is not working for women and racial/ethnic minority group members. Carefully designed technological solutions cannot be ignored as viable alternatives to a biased human approach. We describe the approaches of two organizations that are grounded in principles of fairness and achieve optimal results.

KEY FINDINGS

- The current state of employment selection is biased.
- Fairness is defined here as the lack of disparate treatment and disparate impact.
- Technology-enabled employment selection tools offer a viable alternative to biased human selection decisions.
- Technology must be designed mindfully in order to avoid pitfalls and reach its potential for fair hiring.

Despite countless studies that have shed light on the inevitability of human bias, for more than 50 years HR professionals have relied on methods that introduce these biases to talent pipelines and employment selection procedures. Jobseekers consequently face systems that significantly disadvantage particular groups because of their demographic characteristics. In the following section, we investigate an important question: can technology reduce systematic discrimination in employment procedures?

1 Special thanks to Sara Kassir, Senior Policy and Research Analyst at pymetrics, who assisted with drafting and editing.

The Current State of Employment Selection

To motivate this exploration, it is worth noting a few economic and labor trends that demonstrate the extent of bias in the current order.

For almost 50 years, researchers have used audit and correspondence studies to measure rates of discrimination in the labor market.² Studies have consistently found that resumes submitted by equally qualified candidates receive differential outcomes that can be directly traced to changes in the name or other demographic signals. In one seminal study, candidates with “white-sounding” names received 50 percent more requests for interviews than their equally qualified black counterparts. Put another way, black candidates would need to have approximately five more years of work experience than white candidates to reach the same rate of interviews per job application.³

Despite increases in awareness of unconscious bias and discrimination over recent decades, meta analyses have shown no evidence of a decline in discrimination against black job-seekers in the labor market since the late 1980s, and only slight declines for Hispanic and Latino applicants.⁴ Similar effects persist across genders, with one study indicating that even top-tier academic faculty inadvertently rate female candidates for STEM positions lower than identically-qualified male candidates.⁵

This degree of bias in the labor market is not simply an academic trend; it represents the daily experience of thousands and thousands of individuals. In 2018 alone, the U.S. Equal Employment Opportunity Commission (EEOC) received 76,418 charges of employment discrimination, 64 percent of them on the basis of race or sex.⁶ Furthermore, while black workers make up 13 percent of the U.S. workforce, as a group they file 26 percent of claims with the agency and its partners.⁷ In 2017 the aggregate national unemployment rate was 4.4 percent; however for blacks it was 7.5 percent and for Hispanics it was 5.1 percent.⁸ Clearly, the current state of employment selection is plagued with discrimination.

2 Marianne Bertrand and Esther Duflo, “Field Experiments on Discrimination,” NBER Working Paper Series, February 2016, <https://www.nber.org/papers/w22014.pdf>.

3 Marianne Bertrand and Sendhil Mullainathan, “Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination,” *American Economic Review* 94, no. 4 (2004): 991–1013.

4 Lincoln Quillian, Devah Pager, Ole Hexel, and Arnfinn H. Midtboen, “Meta-Analysis of Field Experiments Shows No Change in Racial Discrimination in Hiring Over Time,” *Proceedings of the National Academy of Sciences* (2017), <https://www.pnas.org/content/early/2017/09/11/1706255114>.

5 Corinne A. Moss-Racusin, John F. Dovidio, Victoria L. Brescoll, Mark J. Graham, and Jo Handelsman, “Science Faculty’s Subtle Gender Biases Favor Male Students,” *Proceedings of the National Academy of Sciences* (2012), <https://doi.org/10.1073/pnas.1211286109>.

6 “EEOC Releases Fiscal Year 2018 Enforcement and Litigation Data,” U.S. Equal Employment Opportunity Commission, April 10, 2019, <https://www1.eeoc.gov/eeoc/newsroom/release/4-10-19.cfm>.

7 Maryam Jameel and Joe Yerardi, “Despite Legal Protections, Most Workers Who Face Discrimination Are on Their Own.” The Center for Public Integrity, February 28, 2019, <https://publicintegrity.org/workers-rights/workplace-inequities/injustice-at-work/workplace-discrimination-cases/>.

8 “Labor Force Characteristics by Race and Ethnicity, 2017,” U.S. Bureau of Labor Statistics, Updated August 2018, <https://www.bls.gov/opub/reports/race-and-ethnicity/2017/home.htm>.

While recruiting and hiring practices are not the only driver of racial inequality in this country, they inevitably contribute to a system that makes it more difficult for certain groups of people to achieve their socioeconomic and professional potential. This is a major motivation for developing new technologies to increase fairness in hiring across our society.

What Is Fairness?

To better understand how technology might improve the employment selection process, it is important to establish what a “fair” process looks like. In the U.S., two legal theories are commonly used to describe workplace discrimination: “disparate treatment” and “disparate impact.” For HR risk and compliance experts, a hiring procedure is deemed fair only if it is absent of *both* types of discrimination.

Disparate treatment occurs when a candidate is affected by an intentional act of overt discrimination. For example, a recruiter or hiring manager who discards a resume because they believe the applicant is black based on the name is engaging in disparate treatment. Note that, in this case, the discrimination is overt and intentional.

Disparate impact, on the other hand, occurs when a test or selection procedure disproportionately excludes candidates based on a protected characteristic (race, color, religion, sex [including pregnancy, sexual orientation, and gender identity], national origin, age, disability, and genetic information).⁹ Here, intent to discriminate is not necessary. As an example, consider the landmark case decided by the U.S. Supreme Court in 1971, *Griggs v. Duke Power*. Following the passage of the Civil Rights Act of 1964, North Carolina’s Duke Power company adopted the requirement that employees in all departments (except its lowest-paying labor department) have a high school diploma and a minimum score on two paper and pencil tests, the Bennett Mechanical Comprehension Test and the Wonderlic Cognitive Ability Test, for work in any department other than its lowest-paying labor department. While the requirement of a high school diploma and a passing score on two tests appears neutral on its face (and hence does not constitute disparate treatment), at the time these requirements demonstrably excluded Black employees. Specifically, the passing rate for the Bennett Mechanical and the Wonderlic tests was 58 percent for whites but only 6 percent for blacks. Further, the 1960 Census showed that 34 percent of white males had high school diplomas while only 12 percent of black males did.¹⁰ In contrast to the overt and conscious decisions that lead to disparate treatment, disparate impact discrimination is typically covert and seemingly unintentional.

Putting the above definitions together, a hiring process is considered fair when candidates are not intentionally singled out for discriminatory treatment and when the overall effect of the selection process does not disproportionately disadvantage members of any one demographic group. While other

9 “Frequently Asked Questions,” U.S. Equal Employment Opportunity Commission, accessed January 27, 2020, https://www.eeoc.gov/employers/smallbusiness/faq/who_is_protected.cfm.

10 *Griggs v. Duke Power Co.*, 401 U.S. 424 (1971), <https://www.law.cornell.edu/supremecourt/text/401/424>.

academic and philosophical definitions of fairness are available, these legal standards are typically the default for practitioners.

Using Technology to Increase Fairness

In reflecting on the state of the world in the 1970s, when most labor laws in the U.S. were written, it is easy to imagine that equal employment opportunity advocates did not view technology as a tool for their cause. On the contrary, over the past several decades the idea of using automated systems to score tests and check backgrounds have likely alarmed labor law practitioners far more than the harms of manually-applied human prejudices that pervade traditional approaches to selection. However, in recent years, advancements in data collection and processing have fundamentally changed the prospects of using technology to overcome bias, and lawmakers are growing receptive to this potential. For example, in September 2019, the California State Assembly passed ACR 125, Bias and Discrimination in Hiring Reduction through New Technology, which affirms that artificial intelligence (AI) may be used to promote fairer employment practices than the status quo. Importantly however, this resolution calls for ethical standards to be established to inform development and use of AI.¹¹ As of February 2020, both the New York City Council¹² and the California State Senate¹³ have introduced bills to actually amend outdated regulations around employment selection tools to account for advancements in technology.

Of course, technology in and of itself is not a silver bullet that will end discrimination. One need not dig deep before unearthing examples of technology gone awry.¹⁴ When utilizing historic datasets to train algorithms, technologists must be mindful to avoid codifying human biases. For example, in an organization where successful incumbents are mostly white and male, a tool that is modeled on them is likely to disadvantage non-white and female candidates, unless that tool is proactively inspected and stripped of such biases. Here, the benefit of an automated tool is that it *can* be stripped of such biases and trained to focus on truly job-relevant signals rather than the “noise” associated with demographic indicators or proxies for such indicators. The brains of human resume reviewers cannot be similarly stripped of such biases. When developers understand the importance of fairness and agree on the goals for the technology, novel approaches to employment selection can reduce discrimination and increase fairness.

In simple terms, the success of these approaches lies in how the technology is designed and how it interacts with humans. Essential steps towards fairness for any technology-enabled assessment include:

11 Bias and Discrimination in Hiring Reduction through New Technology, ARC-125, September 4, 2019, https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201920200ACR125.

12 <https://www.wsj.com/articles/new-york-city-bill-aims-to-regulate-ai-hiring-tools-11582801200>.

13 <http://sacobserver.com/2020/02/new-california-bill-aiming-to-eliminate-racial-and-other-biases-in-hiring/>.

14 See e.g., Jeffrey Dastin, “Amazon Scraps Secret AI Recruiting Tool that Showed Bias Against Women,” Reuters, October 9, 2018, <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>.

- Start by utilizing meaningful data points that evidence fairness across demographic groups in the aggregate.
- Design technology-enabled assessments and selection procedures that are objectively job-relevant and predictive of success in-role.
- Proactively test for and address disparate impact in selection algorithms before they are deployed on jobseekers.
- Hide demographic indicators from decision makers and allow objective assessment results to guide the decision-making process.
- Where human decision-making takes place, design tools that mitigate the risk of human bias from influencing outcomes.
- Audit procedures for disparate impact on candidates after deploying the assessment and revisit the solutions to develop improvements as necessary.

Two Examples of Fair Hiring Technology

pymetrics and Applied are two companies that put the above-mentioned principles for fair hiring technology into practice. The authors of this chapter have substantial experience at these companies developing, deploying, validating, and back-testing such solutions. Kelly Trindel is the Head of I/O Science & Diversity Analytics and Frida Polli is Co-Founder and CEO at pymetrics. Kate Glazebrook is the Co-Founder and CEO of Applied. Below we review each approach in greater detail.

pymetrics

pymetrics¹⁵ has gamified well-known behavioral science assessments adopted from the peer-reviewed academic literature and utilizes the data points collected from these exercises to build custom success profiles for clients. The behavioral assessments were chosen because they measure cross-culturally relevant cognitive, social, and emotional traits reliably and in such a way that minimizes demographic differences in performance.

Each time pymetrics builds a custom success profile based on the performance of locally successful incumbents, the algorithm behind the profile is proactively audited for disparate impact before it is deployed for candidate selection. This is done using a diverse hold-out set of individuals who previously completed the exercises and voluntarily provided their demographic information.

¹⁵ To learn more about pymetrics, visit <https://www.pymetrics.com/employers/>.

The performance of this hold-out group against each custom algorithm indicates objectively whether the algorithm results in disparate impact. The tests deployed by pymetrics to check for disparate impact are open-sourced¹⁶ and based on federal guidelines.¹⁷ Where disparate impact is uncovered, pymetrics is able to identify the data point(s) in the custom algorithm that cause the difference in performance across demographic groups and de-weight or remove those data points prior to use on candidates. The custom algorithm is then tested iteratively against the hold-out set until all demographic groups pass the assessment at a substantively similar rate. The company refers to this as a proactive debiasing process. It could also be considered a technology-enabled search for the least discriminatory alternative, as specified in guidance from the U.S. Equal Employment Opportunity Commission.¹⁸

In the end, pymetrics identifies the version of the algorithm that is best able to differentiate the locally successful incumbent group from a baseline population while being least likely to result in disparate impact on candidates. The cognitive, social, and emotional factors most heavily weighted in the custom algorithm are then compared to the knowledge, skills, and abilities required for the job to confirm and document rational job relevance. All custom algorithms are back-tested to determine whether disparate impact occurred in candidate selection and the degree to which the algorithm predicted success in role longitudinally (i.e., predictive validity). Algorithms are rebuilt annually if back-testing indicates room for improvement.

Applied

Applied is a technology platform that focuses on building tools that guardrail human decisions from bias.¹⁹ Applied redesigns the hiring environment so that only the most relevant (predictive) information is made available to assessors and removes distractions so that candidates are assessed based on skill not demography. Specifically, candidates are assessed based on their responses to work-based scenarios, not their resumes. Their answers are then anonymized to remove all candidate details, chunked up to allow for comparative assessment and reduce “halo” effects, randomized to mitigate ordering effects, and then scored and averaged across multiple independent assessors to harness the wisdom of the crowd and reduce idiosyncratic bias.

A similar methodology is applied to interview processes to increase predictive validity and reduce bias. This methodology codifies behavioral and data science research into a decision tool, and the experiments that underpin the technology are publicly shared.²⁰ Data from the hiring process is also collected and

16 The tests are available to the public at <https://github.com/pymetrics/audit-ai>.

17 The federal guidelines followed by pymetrics are available online from the U.S. Department of Labor. See “Validation of Employee Selection Procedures: Frequently Asked Questions,” Question 4, <https://www.dol.gov/ofccp/regs/compliance/faqs/ValidationEmployeeSelectionFAQs.htm#Q4> and “Practical Significance in EEO Analysis: Frequently Asked Questions,” <https://www.dol.gov/ofccp/regs/compliance/faqs/PracticalSignificanceEEOFAQs.htm>.

18 U.S. Equal Employment Opportunity Commission, “Employment Tests and Selection Procedures,” accessed January 27, 2020, https://www.eeoc.gov/policy/docs/factemployment_procedures.html.

19 To learn more about Applied, visit <https://www.beapplied.com/>.

20 The tests are available to the public at <https://www.beapplied.com/blog>.

analyzed for any latent disparate impact that may affect success rates. Only the most predictive and unbiased assessments are used.

While we have reviewed here the specific examples of pymetrics and Applied, the design principles utilized to ensure fairness in these technologies can be adopted by other technology-enabled assessments and selection devices. It is the opinion of these authors that indeed, they should be adopted.

Conclusion

Can technology work to reduce discrimination in recruiting and hiring? The answer is yes, so long as the developers of such technologies optimize these tools for fairness, transparency, and validity. Technology that is built by, trained on, and utilized by humans must be designed with an eye to avoiding the typical shortcomings human bias produces. When we agree on clear goals for the technology, these prosocial approaches can be coded and adopted with minimal effort. If however the creators and users of technology do not commit to ethical principles in both their processes and procedures, such systems will unfortunately mask unfair practices under the guise of automated objectivity.

In this chapter we have identified the value of technology for reducing discrimination and increasing fairness, we have provided a straightforward definition of fairness, and we have laid out essential steps that developers should follow and users should demand. It is our hope that moving forward, practitioners will continue to iterate on the steps described here with the goal of creating and using the fairest and most predictive technology-enabled approaches to employment selection and human capital development.

References

Bertrand, Marianne, and Esther Duflo. "Field Experiments on Discrimination." NBER Working Paper Series. February 2016. <https://www.nber.org/papers/w22014.pdf>.

Bertrand, Marianne, and Sendhil Mullainathan. "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review* 94, no. 4 (September 2004): 991–1013. <https://www.pnas.org/content/early/2017/09/11/1706255114>.

Bias and Discrimination in Hiring Reduction through New Technology. ARC-125, September 4, 2019. https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201920200ACR125

Dastin, Jeffrey. "Amazon Scraps Secret AI Recruiting Tool that Showed Bias Against Women." Reuters, October 9, 2018. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>.

Griggs v. Duke Power Co. 401 U.S. 424 (1971). <https://www.law.cornell.edu/supremecourt/text/401/424>.

Jameel, Maryam, and Joe Yerardi. "Despite Legal Protections, Most Workers Who Face Discrimination Are on Their Own." The Center for Public Integrity, February 28, 2019. <https://publicintegrity.org/workers-rights/workplace-inequities/injustice-at-work/workplace-discrimination-cases/>.

Moss-Racusin, Corinne A., John F. Dovidio, Victoria L. Brescoll, Mark J. Graham, and Jo Handelsman. "Science Faculty's Subtle Gender Biases Favor Male Students." *Proceedings of the National Academy of Sciences* (October 2012). <https://www.pnas.org/content/109/41/16474>.

Quillian, Lincoln, Devah Pager, Ole Hexel, and Arnfinn H. Midtboen. "Meta-Analysis of Field Experiments Shows No Change in Racial Discrimination in Hiring Over Time." *Proceedings of the National Academy of Sciences* (September 2017). <https://www.pnas.org/content/early/2017/09/11/1706255114>.

U.S. Bureau of Labor Statistics. "Labor Force Characteristics by Race and Ethnicity, 2017." August 2018, <https://www.bls.gov/opub/reports/race-and-ethnicity/2017/home.htm>.

U.S. Department of Labor. Office of Federal Contract Compliance Programs. "Validation of Employee Selection Procedures: Frequently Asked Questions." <https://www.dol.gov/ofccp/regs/compliance/faqs/ValidationEmployeeSelectionFAQs.htm#Q4>.

U.S. Department of Labor. Office of Federal Contract Compliance Programs. "Practical Significance in EEO Analysis: Frequently Asked Questions." <https://www.dol.gov/ofccp/regs/compliance/faqs/PracticalSignificanceEEOFAQs.htm>

U.S. Equal Employment Opportunity Commission. "Employment Tests and Selection Procedures." Accessed January 27, 2020, https://www.eeoc.gov/policy/docs/factemployment_procedures.html.

U.S. Equal Employment Opportunity Commission. "EEOC Releases Fiscal Year 2018 Enforcement and Litigation Data." Updated April 10, 2019. <https://www1.eeoc.gov/eeoc/newsroom/release/4-10-19.cfm>.

U.S. Equal Employment Opportunity Commission. "Frequently Asked Questions." Accessed January 27, 2020. https://www.eeoc.gov/employers/smallbusiness/faq/who_is_protected.cfm.