

Panel 1



U.S. Equal Employment Opportunity Commission

Testimony of Suresh Venkatasubramanian

Chair Burrows, and Commissioners of the EEOC, thank you for the opportunity to provide witness to the Commission today. My name is Suresh Venkatasubramanian, and I am a professor at Brown University and director of the Center for Technological Responsibility. I am a computer scientist who has for the last decade studied the ways in which automated systems, and especially those that use artificial intelligence, may produce discriminatory outcomes in employment and performance evaluation. Most recently, I served as the Assistant Director for Science and Justice in the White House Office of Science and Technology Policy in the Biden-Harris Administration and coauthored the Blueprint for an AI Bill of Rights^[1], a document that lays out five key protections for those meaningfully impacted by the use of automation, and a detailed technical companion for how these protections can be realized.

Automated systems, fueled by vast quantities of data, innovative machine learning algorithms, and fast computing resources, hold out the promise of faster, more efficient, and more accurate approaches to evaluating candidates for employment, whether it be algorithms based on natural language processing that can screen candidate resumes and identify salient factors, game-based interview sessions that seek to identify key cognitive factors that make an individual a good fit for a job, or multimedia analysis procedures that score candidates based on video interviews. These systems promise to make the interview process more seamless for candidates and recruiters, eliminate biases in judgement, and allow for a broader pool of candidates to be recruited and evaluated fairly.

The keyword in the above is 'promise': the tremendous hype surrounding the development of new technology, especially those that use artificial intelligence-based approaches, has obscured many documented problems that arise when these algorithms are deployed in actual employment settings. These include

differential outcomes for people from different demographics groups, inferences based on psychological premises (such as emotion recognition) that are unsound or unvalidated, and a lack of accountability arising from the shifting of responsibility between the vendors who develop such software and the companies that procure them for use in hiring.

Over the decades, any new technology that has been introduced into society – cars, medical treatments, airplanes, a host of consumer products – has been accompanied by rigorous testing regimes to ensure that the technologies work, are safe, and do not cause harm. These *guardrails* build trust in the technology and create an environment in which innovation flourishes without fear of liability. Indeed, we have already seen that in case of data-driven automated technologies such as machine learning, the insistence on guardrails to protect against discrimination and make the workings of systems more transparent has fostered a whole new area of innovation in the tech industry described as ‘Responsible AI’. Guardrails feed further innovation rather than hamper it: those who frame this as a zero-sum game are in effect advocating for sloppy, badly engineered and irresponsible technologies that would never be deployed in any other sector.

So what should these guardrails look like? The aforementioned Blueprint for an AI Bill of Rights, which I note was developed in consultation with agencies across the Federal government including the EEOC, as well as after extensive consultation with the private sector, civil society advocates, and academics, provides several relevant suggestions.

Firstly, a note about scope. The Commission correctly mentions both AI and automated systems in the title of this event, recognizing the varied nature of the systems that are used to assist in the employment process. As a computer scientist, I have seen the term ‘AI’ morph and evolve – going out of favor during AI winters^[2] and coming back into vogue as money and investments began to pour into the field. Therefore, it is important when the Commission provides guidance, that it focuses on the impact and harms on individuals rather than on the (rapidly evolving) technologies themselves and thus retain within scope any automated system as defined in the Blueprint.

Just like the Commission has done in the context of algorithms for employment and the Americans with Disabilities Act^[3], it should issue enforcement guidance and recommended questions that the designers and developers of such systems should answer as they develop their systems.

The Commission should direct the creators of automated systems used in employment to perform

- detailed validation testing that includes the specific technology being used as well as the system interaction with human operators or reviewers whose actions might impact overall system effectiveness. The results of this testing should be made available for review.
- risk identification and mitigation that can be based on the National Institute of Standard and Technology AI Risk Mitigation Framework[4].
- disparity assessments to determine how their systems might exhibit unjustified differential outcomes based on different protected characteristics and mitigate these differential outcomes as far as possible with the result of this assessment and mitigation made available for review.
- ongoing monitoring of the developed systems on a regular basis to ensure that the mitigations and validations continue to be maintained, since automated systems can “drift” away from their training over time especially if the underlying models are retrained based on new data.
- Evaluation of the *data* used to build models (in the case of AI or machine learning-based models) to ensure that only relevant, high quality data, tailored to the specific context of employment, is used. Relevancy itself should be determined based on research-backed demonstration of the causal influence of the data on the outcome, rather than via an appeal to historical practices.

The Commission should strongly encourage the following best practices by entities seeking to develop automated systems for use in employment contexts.

- **The use of transparent and explainable models.** Complex and opaque models make it difficult to understand why model predictions take the form that they do, and can render the system liable to make mistakes that are undetectable. Models that are simple enough to be easily explained, or that are augmented with procedures that can accurately explain the results of a prediction in a way that is tailored to the individual asking for the explanation are likely to be more accurate and less prone to unexpected errors or differential group outcomes.
- **The inclusion of human oversight.** Systems should provide timely human consideration and remedy by a fallback system to account for when the system

fails. This is important because automated systems are fallible especially when presented with scenarios far removed from the scenarios used to train them. This is important also to ensure that the use of the system does not prevent individuals with accessibility challenges from participating in the hiring process.

In conclusion, I once again would like to thank the Commissioners for giving me the opportunity to testify at this hearing and commend the Commission for taking up this complex and important civil rights challenge presented by modern technology.

[1] The White House Office of Science and Technology Policy. Blueprint for an AI Bill of Rights. Oct 2022. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>

[2] Wikipedia. AI Winter. https://en.wikipedia.org/wiki/AI_winter

[3] EEOC. The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees. <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence>

[4] NIST. AI Risk Management Framework. <https://www.nist.gov/itl/ai-risk-management-framework>



U.S. Equal Employment Opportunity Commission

Testimony of Pauline Kim

Chair Burrows, Commissioners, thank you for the opportunity to address the issue of artificial intelligence and employment discrimination.

I am the Daniel Noyes Kirby Professor of Law at Washington University School of Law in St. Louis. My research and scholarship center on the law of the workplace, with a particular focus on how emerging technologies are impacting anti-discrimination law and employee privacy rights.

When AI is incorporated into automated decision systems, or predictive algorithms, and used for hiring and promotion, these tools offers many advantages to employers, including efficiency and scalability. They also have the potential to remove some forms of human bias from these processes. However, as is now well recognized, these tools can operate in ways that are biased, and may discriminate along the lines of race, sex, and other protected characteristics.^[1] Computer based assessments can also create barriers to equal employment for individuals with disabilities. The technical assistance issued by the Chair last year provides crucial guidance to employers about how to comply with the Americans with Disabilities Act when using these tools, especially, the importance of providing individualized assessment and reasonable accommodation to individuals with disabilities. In these remarks, I will focus instead on issues of systemic bias that can arise when employers use algorithms to predict the suitability of workers for particular jobs.

Numerous studies and reports have documented the ways in which bias can creep into automated systems.^[2] When incomplete, unrepresentative, or error-ridden data are used to train a model, the resulting predictions can produce biased outcomes. Training data may encode biased human judgements, for example, when the data includes subjective scores assigned by humans, and the model takes them as objective measures of performance. And because predictive models extract patterns in past data to make future predictions, even highly accurate models may simply reproduce existing patterns of discrimination and occupational segregation.

In addition to data problems, the choice of the target variable can have a significant impact on who is given access to employment opportunities. The target variable is the outcome the system is designed to predict. Paying attention to how it is defined and measured is critical to avoiding bias. Good employees have many different traits and the designer of automated hiring systems must decide which attribute to focus on. Many of the most valuable qualities in an employee are difficult to define or to measure accurately with data. And so the designer may choose instead a target that is easily measurable, such as customer ratings or time on the job. That choice of the target variable can be highly consequential.

Take, for example, a model that predicts the best candidates by selecting those who most closely resemble applicants who were hired in the past. If past hiring decisions were infected by bias, the model's predictions will be as well. Another example is a model that rates highly applicants who are least likely to leave the paid labor force. Such a model will disproportionately screen out women of childbearing age or workers with disabilities, who are more likely to have breaks in employment, even though they are fully capable of performing the job. Thus, the initial step of problem formulation^[3]—deciding how the problem to be solved by the algorithm is defined—is crucial to avoiding discrimination.

Automated systems that rely on machine learning to constantly update can also create problematic feedback loops. Proponents of these systems argue that they can learn and improve continuously over time. However, unlike with online advertising, hiring tools cannot be subject to meaningful A/B testing. Low-ranked candidates will not be hired and their job potential cannot be observed. As a result, false negative outcomes cannot be detected and corrected, and erroneous assumptions about lack of ability may be reinforced over time.

Long before an employer makes its hiring decisions, predictive algorithms also play a critical role in matching job candidates with potential opportunities. Most employers today advertise job openings on social media sites like Facebook, or rely on job matching platforms to identify promising candidates. These new labor market intermediaries utilize algorithms to channel information about opportunities to different users, and their operation can determine which opportunities a job seeker learns about.^[4] Studies have documented that ad-targeting algorithms distribute job advertisements in racially- and gender- biased ways that reflect stereotypes about what kinds of people perform certain jobs.^[5] These effects occur even when the employer has requested race- and gender-

neutral ad targeting, and wants its job advertisements to be distributed to a broad and diverse pool.

-

The Uncertain Application of Existing Anti-Discrimination Law

Anti-discrimination laws apply to these automated decision tools and provide some leverage to prevent or redress discriminatory harms. However, current law is incomplete. There are a number of gaps and uncertainties about how the doctrine applies to automated decision systems.

Existing doctrine was developed with human decision-makers in mind and does not always fit the risks of discrimination posed by automated systems. For example, a formalistic view of disparate treatment discrimination might suggest that so long as a model does not take a protected characteristic into account, it does not violate disparate treatment. Conversely, it might assume that any model that does take a protected characteristic into account is discriminatory.

This interpretation of disparate treatment law is too simplistic. An employer could engage in disparate treatment without expressly relying on a protected characteristic like race or sex by using proxy variables to produce exactly the same effect.[6] On the other hand, in order to ensure that a model is fair for all groups, it may be necessary to take protected characteristics into account.[7] For example, the only way to audit for unintended bias is to make use of data about protected characteristics.[8]

Thus, simply prohibiting consideration of race or sex in a model would not only fail to prevent discrimination from occurring, it could be counterproductive as well.

Uncertainty also affects the application of disparate impact doctrine. Under current law, when an employer practice has a disparate impact on a protected group, the employer has a defense if it can show that the practice is “job related and consistent with business necessity.”[9] That defense, which was codified as part of disparate impact theory by the Civil Rights Act of 1991, is not explained in the statute. In order to interpret its meaning, many turn to the Uniform Guidelines on Employee Selection Procedures issued in 1978. The Guidelines set out methods for validating an employer test based on the industrial psychology literature at that time, and as a

result, they do not address some of the unique challenges posed by AI and predictive algorithms.

For example, some automated decision systems rely on data mining to extract patterns. They may uncover variables that are strongly predictive of the target variable, but have no clear connection to job performance. Some machine learning models are so complex that an employer that relies on them may not be able to explain its decision to reject candidates, making it difficult to apply concepts of “job relatedness” and “business necessity.” The Guidelines were not designed to address situations like these.

The third step of the disparate impact analysis allows a plaintiff to show that a less discriminatory alternative is available to the employer. Again, there is uncertainty how plaintiffs can show this when challenging predictive algorithms, given that there are many, potentially infinite, models that could be designed for a particular application.**[10]**

Another uncertainty surrounding the use of automated decision-tools relates to remedial efforts. If an employer detects that a predictive algorithm has a disparate impact on a disadvantaged group, what can it do in response? Some researchers have questioned whether efforts to *remove* discriminatory effects might themselves run afoul of anti-discrimination law by taking account of race, sex, or other protected characteristics. They have expressed concern that taking sensitive characteristics into account to prevent disparate impact might be construed as a form of disparate treatment.**[11]** Existing case law permits taking race and other sensitive characteristics into account in order to level the playing field and ensure equal access to opportunities;**[12]** however, the application of those principles needs to be clarified in the context of algorithmic decision-making.**[13]**

Finally, it is uncertain whether existing law reaches labor market intermediaries like online advertising and job matching platforms.**[14]** These entities play an increasingly important role in shaping the job market and access to opportunities, but it is unclear whether or when they would be considered “employment agencies” covered by Title VII, and what responsibilities employers have when relying on these platforms to recruit workers.**[15]**

Aside from these legal uncertainties, practical challenges exist as well. Title VII’s enforcement scheme relies primarily on retrospective liability to redress past discriminatory harms. Although the EEOC brings enforcement actions, individual

workers file the vast majority of employment discrimination suits and accessing remedies may be difficult for them. It has always been harder to detect and challenge discriminatory hiring decisions than firing decisions because of the difficulty obtaining evidence of discrimination when outside the firm. Individual workers will find it even more difficult to challenge biased hiring algorithms. Part of the problem is that applicants often do not know when or how employers are using automated systems. Even with greater transparency, they will typically lack the technical expertise and resources needed to assess the fairness of these tools or to bring a legal challenge.

The Limits of Self-Regulation

Before turning to some suggested reforms, I want to acknowledge that automated decision systems are not inevitably discriminatory. A well-designed and implemented system may help employers reduce the influence of human bias. Employers should be allowed some latitude to explore ways in which AI tools can help to remove bias and increase the diversity of their workforce. However, it is important not to get caught up in the rhetoric claiming that data-based tools are inherently neutral and objective.

If the goal is to create more equitable workplaces, relying on industry best practices and employer self-regulation is insufficient. While many firms care deeply about diversity, equity and inclusion, not all do. Robust regulatory tools remain important to address the bad actors. And even well-intentioned firms face significant constraints when trying to do the right thing. They may lack the expertise to understand the risks of discrimination, or the resources to engage in ongoing auditing and testing needed to prevent these harms. Detecting and removing bias requires close analysis and ongoing scrutiny of automated systems.

Another concern is that employers motivated primarily by liability risk avoidance will adopt pro forma, symbolic steps that do not meaningfully address discriminatory risks. Extensive research by sociologists has documented how many firms responded to civil rights laws, and in particular, the threat of sexual harassment liability, by creating procedures and checklists which signaled their good faith but did not address the root causes of discrimination and harassment.

[16] Given the experience with “best practices” that shaped firms’ responses to

sexual harassment, but often proved ineffective,[[17](#)] the EEOC should be cautious about allowing employers to rely on procedural checklists as evidence that their selection tools comply with anti-discrimination laws.

What Can the EEOC Do?

Given the ambiguity about how employment discrimination law applies to AI and other automated systems, it may be useful to clarify the law in a handful of discrete areas.

First, the EEOC should clarify that AI tools that produce a disparate impact cannot be defended solely on the basis of statistical correlations.[[18](#)] The employer should have to demonstrate the *substantive validity* of its selection tools. It should bear the burden of showing that the model was built using accurate, representative, and unbiased data, and that it actually measures job-relevant skills and abilities. This approach is consistent with the position taken by a coalition of civil rights organizations.[[19](#)] It would also create incentives for employers who purchase these systems from outside vendors to closely scrutinize these tools before deploying them.

Second, the EEOC should offer guidance on the duty of employers to explore less discriminatory alternatives. Researchers have demonstrated that there is no unique model for solving a given optimization problem.[[20](#)] Because there are typically multiple models that can be developed and used in any given application, designers should explore and document which options have the least discriminatory effect. If alternative, comparably effective models are available, then arguably, an employer's choice to use a model that has discriminatory impact is not consistent with business necessity.

Third, the EEOC should make clear that taking steps to correct or prevent a model from having a disparate impact is *not* a form of disparate treatment.[[21](#)] In order to de-bias models, designers will need to make use of data about sensitive characteristics. When building a model, they should examine proposed target variables for implicit bias. Avoiding discriminatory impacts may also require scrutinizing the representativeness and accuracy of training data, oversampling underrepresented groups, or removing features that encode human bias. And because AI tools may behave differently when applied to actual applicants

compared with training data, it is critical to audit their effects once deployed. Strategies like these require paying attention to race or other protected characteristics in order to avoid bias and build AI tools that are fair to all. Because these types of de-biasing strategies do not make decisions about individual workers turn on protected characteristics, they should not be considered a form of disparate treatment. By clarifying that it is permissible to take protected characteristics into account in order to remove disparate impact, the EEOC can encourage voluntary employer efforts to rigorously examine their practices and to avoid any discriminatory effects.^[22]

Fourth, the EEOC could offer guidance about the legal responsibilities of labor market intermediaries, such as job-matching platforms, that play a significant role in procuring workers for employers and employment opportunities for job seekers. Very little case law exists applying the statutory definition of an “employment agency” under Title VII, and as a result, the legal responsibility of online platforms to ensure that they provide a level playing field for all workers remains unclear. Even where these entities cannot be held directly liable, the EEOC could conduct research and educate employers about how the predictive algorithms these platforms use to distribute information can cause bias in the job advertising and recruiting process.

Beyond clarifying discrete issues in existing law, any regulatory steps should be taken cautiously. Because the technology at issue is complicated and rapidly evolving, it is important not to freeze into place standards that will quickly become obsolete. In particular, what appear to be “best practices” today may turn out to be sub-optimal solutions in the future. Locking them into place through legal doctrine or by recognizing employer defenses, could end up excusing or immunizing practices that are later determined to be harmful.

For these reasons, much of the EEOC’s efforts in this area should be forward-looking, aimed at building its capacity to audit automated hiring tools, to study their effects on workforce participation, and to research solutions, both technical and practical, that will ensure that these tools work to open opportunities to the widest possible pool of workers. An important part of this work will require increasing transparency by employers about when and how they are utilizing automated decision systems, how those systems were designed, what training data was used to build them, and the effects of these systems on workers.

Finally, the EEOC should consider developing data analytic tools to study employers and their human resources processes rather than workers. By leveraging data and

computational tools, these systems could help to diagnose where or why bias is occurring, or to predict which practices are more likely to broaden the diversity of employees who are hired and to support their success in the workplace.

Thank you again for your time and for focusing attention to the important issues of employment discrimination and AI.

[1] The report by the White House Office of Science and Technology Policy, *Blueprint for an AI Bill of Rights* (October 2022), documents many examples of algorithmic bias across a range of social applications.

[2] See, e.g., Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 Calif. L. Rev. 671 (2016).

[3] Samir Passi & Solon Barocas, *Problem Formulation and Fairness*, Proceedings of the Conference on Fairness, Accountability, and Transparency 39 (ACM 2019).

[4] Pauline T. Kim, *Manipulating Opportunity*, 106 Va. L. Rev. 69 (2020).

[5] Muhammad Ali et al., *Discrimination through Optimization: How Facebook's Ad Delivery Can Lead to Biased Outcomes*, 3 Proceedings of the ACM on Human-Computer Interaction 1 (2019); Piotr Sapiezynski et al., *Algorithms that "Don't See Color": Comparing Biases in Lookalike and Special Ad Audiences*, arXiv:1912.07579 [cs] (2019); Anja Lambrecht & Catherine Tucker, *Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads*, 65 Management Science 2966 (2019).

[6] See, e.g., Cynthia Dwork et al., *Fairness through Awareness*, Proceedings of the 3rd Innovations in Theoretical Computer Science Conference on - ITCS '12 214 (ACM Press 2012); Moritz Hardt et al., *Equality of Opportunity in Supervised Learning*, arXiv:1610.02413 [cs] (Oct. 2016).

[7] See, e.g., Talia B. Gillis & Jann L. Spiess, *Big Data and Discrimination*, 86 U. Chi. L. Rev. 459 (2019); Sam Corbett-Davies & Sharad Goel, *The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning*, arXiv:1808.00023 [cs] (Aug. 2018).

[8] Pauline T. Kim, *Auditing Algorithms for Discrimination*, 166 U. Pa. L. Rev. Online 189 (2017).

[9] 42 U.S.C. §2000e-2(k)(1)(A).

[10] Emily Black, Manish Raghavan, and Solon Barocas, *Model Multiplicity: Opportunities, Concerns, and Solutions*, 2022 ACM Conference on Fairness, Accountability, and Transparency 850 (ACM Jun. 2022); Charles T. Marx, Flavio du Pin Calmon, and Berk Ustun, *Predictive Multiplicity in Classification*, arXiv:1909.06677 [cs, stat] (Sep. 2020).

[11] This worry stems from what I argue is a misunderstanding of the holding in *Ricci v DeStefano*, 557 U.S. 557 (2009). See Pauline T. Kim, *Auditing Algorithms for Discrimination*, 166 U. Pa. L. Rev. Online 189, 200-202 (2017). Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 Wm. & Mary L. Rev. 857, 925–32 (2017).

[12] See, e.g., *Maraschiello v. City of Buffalo Police Department*, 709 F.3d 87, 89 (2d Cir. 2013); *Duffy v. Wolle*, 123 F.3d 1026, 1038–39 (8th Cir. 1997); *Rudin v. Lincoln Land Cmty. Coll.*, 420 F.3d 712, 722 (7th Cir. 2005); *Mlynczak v. Bodman*, 442 F.3d 1050, 1050 (7th Cir. 2006).

[13] Pauline T. Kim, *Race-Aware Algorithms: Fairness, Nondiscrimination and Affirmative Action*, 110 Calif. L. Rev. 1539 (2022).

[14] Pauline T. Kim, *Manipulating Opportunity*, 106 Va. L. Rev. 69 (2020).

[15] Pauline T. Kim & Sharion Scott, *Discrimination in Online Employment Recruiting Symposium: Law, Technology, and the Organization of Work*, 63 St. Louis U. L.J. 93 (2018).

[16] See, e.g., Frank Dobbin & Alexandra Kalev, *The Promise and Peril of Sexual Harassment Programs*, 116 Proc. Nat'l Acad. Sci. 12255 (Jun. 2019); Frank Dobbin & Alexandra Kalev, *The Civil Rights Revolution at Work: What Went Wrong*, 47 Annual Review of Sociology 281 (Jul. 2021).

[17] EEOC Select Task Force on the Study of Harassment in the Workplace, Report of Co-Chairs Chai R. Feldblum & Victoria A. Lipnic (2016).

[18] Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 Wm. & Mary L. Rev. 857, 921 (2017).

[19] Civil Rights Principles for Hiring Assessment Technologies, The Leadership Conference Education Fund, available at: <https://civilrights.org/resource/civil-rights-principles-for-hiring-assessment-technologies/#>

[20] Emily Black, Manish Raghavan, and Solon Barocas, *Model Multiplicity: Opportunities, Concerns, and Solutions*, 2022 ACM Conference on Fairness, Accountability, and Transparency 850 (ACM Jun. 2022); Charles T. Marx, Flavio du Pin Calmon, and Berk Ustun, *Predictive Multiplicity in Classification*, arXiv:1909.06677 [cs, stat] (Sep. 2020).

[21] Pauline T. Kim, *Race-Aware Algorithms: Fairness, Nondiscrimination and Affirmative Action*, 110 Calif. L. Rev. 1539 (2022).

[22] Ass'n of Firefighters v. City of Cleveland, 478 U.S. 501 (1986) (“We have on numerous occasions recognized that Congress intended voluntary compliance to be the preferred means of achieving the objectives of Title VII”).



U.S. Equal Employment Opportunity Commission

Testimony of Jordan Crenshaw

Chair Burrows and distinguished members of the Commission, thank you for your invitation to testify. My name is Jordan Crenshaw, and I serve as the Vice President of the U.S. Chamber Technology Engagement Center (“C_TEC”). C_TEC is the technology policy hub of the U.S. Chamber, and its goal is to promote the benefits of technology in the economy and advocate for rational policy solutions that drive economic growth, spur innovation, and create jobs.

Today's hearing titled "Navigating Employment Discrimination in AI and Automated Systems: A New Civil Rights Frontier" is an important and timely discussion, and the Chamber appreciates the opportunity to participate.

The world has quickly entered its fourth industrial revolution, in which technology and artificial intelligence (or, “AI”) are helping propel humanity. Americans are witnessing the benefits of using AI daily, from its value in adapting vaccines to tailoring them to new variants to increasing patient safety during procedures like labor and delivery.¹ Artificial intelligence is also rapidly changing how businesses operate. From helping to assist employers in finding the perfect candidate who can help grow their business to new heights to the use of technology within organizations to help alleviate barriers for those with disabilities to participate fully within the workforce, the use of technology within organizations is a tremendous force for good in its ability to advance opportunities for all Americans.²

However, without public trust in the technology, the amazing benefits of this technology will never be fully realized. This is why the United States must lead globally in building trustworthy standards for artificial intelligence. These standards must be rooted in our unifying principles, such as individual liberties, privacy, and the rule of law. While the development and deployment of AI have become essential to facilitating innovation, this innovation will only reach its full potential and enable the United States to compete should the American public trust the technology and

the guardrails placed around its use are limited and well supported by facts and data. The business community understands that fostering this trust in AI technologies is essential to advance its responsible development, deployment, and use. This has been a core understanding of the U.S. Chamber, as it is the first principle within the 2019 “U.S. Chamber’s Artificial Intelligence Principles:

Trustworthy AI encompasses values such as transparency, explainability, fairness, and accountability. The speed and complexity of technological change, however, mean that governments alone cannot promote trustworthy AI. The Chamber believes that governments must partner with the private sector, academia, and civil society when addressing issues of public concern associated with AI. We recognize and commend existing partnerships that have formed in the AI community to address these challenges, including protecting against harmful biases, ensuring democratic values, and respecting human rights. Finally, any governance frameworks should be flexible and driven by a transparent, voluntary, and multi-stakeholder process.³”

AI also brings a unique set of challenges that should be addressed so that concerns over its risks do not dampen innovation and US trustworthy AI leadership. C_TEC shares the perspective with many of the leading government and industry voices, including the National Security Commission on Artificial Intelligence (NSCAI)⁴, the National Institute of Standards and Technology (NIST)⁵, that government policy to advance the ethical development of AI- based systems, sometimes called “responsible” or “trustworthy” AI, can enable future innovation and help the United States to be the global leader in AI.

Last year, the U.S. Chamber launched its Artificial Intelligence Commission on Competitiveness, Inclusion, and Innovation to advance U.S. leadership in using and regulating trustworthy AI technology.⁶ The Commission, led by co-chairs former Congressmen John Delaney and Mike Ferguson, is composed of representatives from industry, academia, and civil society to provide independent, bipartisan recommendations to aid policymakers with guidance on artificial intelligence policies as it relates to regulation, international research, development competitiveness, and future jobs.

Over a span of multiple months, the Commission heard oral testimony from 87 expert witnesses⁷ over five separate field hearings. The Commission heard from individuals such as Jacob Snow, Staff Attorney for the Technology & Civil Liberties Program at the ACLU of Northern California. In his testimony, he told the

Commission that the critical discussions on AI are “not narrow technical questions about how to design a product. They are social questions about what happens when a product is deployed to a society, and the consequences of that deployment on people’s lives.”⁸

Doug Bloch, then Political Director at Teamsters Joint Council 7, referenced his time serving on Governor Newsom’s Future of Work Commission: “I became convinced that all the talk of the robot apocalypse and robots coming to take workers’ jobs was a lot of hyperbole. I think the bigger threat to the workers I represent is the robots will come and supervise through algorithms and artificial intelligence.”⁹

Miriam Vogel, President and CEO of EqualAI and Chair of NAIAC, also addressed the Commission. She stated, “I would argue that it’s not that we need to be a leader, it’s that we need to maintain our leadership because our brand is trust.”

The Commission also received written feedback from stakeholders answering numerous questions that the Commission has posed in three separate requests for information (RFI), which asked questions about issues ranging from defining AI, balancing fairness and innovation,¹⁰ and AI’s impact on the workforce.¹¹ These requests for information outline many of the fundamental questions that the Commissioner looks to address in its final recommendations, which will help government officials, agencies, and the business community. The Commission is working on its recommendations and will look to release them this upcoming Spring, and we will make sure EEOC receives a copy.

While the Chamber is diligently taking a leading role within the business community to address many of the concerns which continue to be barriers to public trust and consumer confidence in the technology, this testimony will highlight how industry is using the technology, the importance of regulatory balance, and finally specific areas in which government can assist in providing the necessary incentives for the technology to be appropriately designed and deployed in a manner that helps all within society. Although these discussions do not explicitly address matters under EEOC’s purview, their broad applicability will make them relevant to furthering EEOC’s understanding of AI’s place in the workplace.

The following issues are considered in this testimony:

- Opportunities for the federal government and industry to work to together to develop Trustworthy AI

- How are Different Sectors Adopting Governance Models and Other Strategies to Mitigate Risks that Arise from AI Systems?
- Policy implications to consider while looking at regulating new technologies such as AI
- What Recommendations do you Have for how the Federal Government can Strengthen its Role for the Development and Responsible Deployment of Trustworthy AI Systems?

1. Opportunities for the Federal Government and Industry to Work Together to Develop Trustworthy AI

A. Support for Alternative Regulatory Pathways, Such As Voluntary Consensus Standards

New regulation is not always the answer for emerging or disruptive technologies. Non-regulatory approaches can often serve as tools to increase safety, build trust, and allow for flexibility and innovation. This is particularly true for emerging technologies such as artificial intelligence as the technology continues to evolve rapidly, while regulations are static and modifications are often obsolete upon issuance.

This is why the Chamber supports the National Institutes of Science and Technology's (NIST) work to draft the Artificial Intelligence Risk Management Framework (AIRMF). The AI RMF is meant to be a stakeholder-driven framework, which is "intended for voluntary use and to improve the ability to incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems."

The AI_RMFM also will look to develop "profiles" which enable organizations to "establish a roadmap for reducing the risk that is well aligned with organizational and sector goals, considers legal/regulatory requirements and industry best practices, and reflects risk management priorities." These profiles are beneficial in allowing sector-specific best practices to be developed openly and voluntarily.

Another example of non-regulation approach is the National Highway Traffic Safety Administration's ("NHTSA") Voluntary Safety Self-Assessments ("VSSA"). More than two dozen AV developers have submitted a VSSA to NHTSA, which has provided essential and valuable information to the public and NHTSA on how developers are addressing safety concerns arising from AVs. The flexibility provided by VSSAs, complemented by existing regulatory mechanisms, provides significant transparency into the activities of developers without compromising safety.

Voluntary tools provide significant opportunities for consumers, businesses, and the government to work together to address many of the underlying concerns with emerging technology while at the same time providing the necessary flexibility to allow the standards not to stifle innovation. These standards are pivotal in the United States' ability to maintain leadership in emerging technology as it is critical to ensuring our global economic competitiveness in this cutting-edge technology.

B. Stakeholder Driven Engagement

The U.S. Chamber of Commerce stands by and is ready to assist EEOC in any opportunity to improve consumer confidence and trust in AI systems for employment purposes. The business community has always viewed trust as a partnership, and only when government and industry work side by side can that trust be built. The opportunities to facilitate this work are great, but there are essential steps that industry and government can make today.

Last year the Chamber asked Americans about their perception of artificial intelligence. The polling results were very eye-opening, as there was a significant correlation between the trust and acceptance of AI and an individual's knowledge and understanding of the technology.¹² To build the necessary consumer confidence to allow artificial intelligence to grow for the betterment of all, all opportunities must be pursued of for industry and governments to work together in educating stakeholders about the technology.

Inclusive stakeholder engagement and transparency between government and industry are vital to building trust. C_TEC has continued highlighting NIST and its stakeholder engagement on the AI RMF as what agencies should strive to replicate. NIST has had three workshops during their process. The agency has also included three engagement opportunities for stakeholders to provide written feedback on

the development, direction, and critique of the AI RMF. This engagement by NIST has allowed for the development of trust between industry and the federal government. While extolling the virtues of the NIST process and their action on the RMF, it is prudent to highlight that NIST is only one entity within the federal government and that as other agencies, such as EEOC, look to receive crucial feedback from the business community, this open and transparent process should look to be modeled.

In contrast, policymakers should more skeptically view the Office of Science and Technology Policy (OSTP) and its development of the AI Bill of Rights, which needed a more transparent and open drafting process. Although OSTP in its “Blueprint,¹³” claims to highlight organizations for which OSTP met and received feedback, the process to obtain sufficient stakeholder input about these complex was substantively lacking. Furthermore, the only formal request for information from OSTP relied upon in the Blueprint focused on biometrics¹⁴ and not a comprehensive Bill of Rights”. OSTP failed to create a complete record of the use of the technology, which harmed trust in being a part of these critical conversations.

C. Awareness of the Benefits of Artificial Intelligence

Another excellent opportunity for industry and government to work together is highlighting the benefits and efficiencies of using technology within the government. The government’s utilization of AI can lead to medical breakthroughs¹⁵ to help to predict risk for housing and food insecurities.¹⁶ AI is helping government provide better assistance to the American public and is becoming a vital tool. Developing these resources does not occur in a vacuum, and most of these tools are brought about in partnership with the industry.

The Office of Personal Management (OPM) estimates that the Federal Workforce currently has 1.9 million¹⁷ employees. The number of government employees eligible for retirement is 14%, with the number jumping to 30% in 2023¹⁸. This significant reduction in the federal workforce and loss of knowledge could affect governments’ ability to operate. However, the federal government has an excellent opportunity to work alongside the private sector to help hire the future workforce through the use of Automated Employment Decision Tools (AEDT), which could

provide significant efficiencies to the process and allow for the benefits of the technology to be seen by all.

Highlighting these workstreams and the benefits such as AEDT and the efficiency they deliver for the American public can assist in fostering trust in technology and build overall consumer confidence in technology use outside of government. However, this would also require a foundational change in how government works, which includes addressing the “legacy culture” that has stifled the necessary investment and buildout of 21st-century technology solutions and harnessing data analytics.

1. How are Different Sectors Adopting Governance Models and Other Strategies to Mitigate Risks that Arise from AI Systems?

AI is a tool that does not exist in a legal vacuum. Policymakers should be mindful that activities performed and decisions aided by AI are often already subject to existing laws. Most notable for EEOC would be Title VII of the Civil Rights Act of 1964,¹⁹ which already protects employees and applicants against discrimination based on race, color, sex, national origin, and religion, as well as the Americans with Disabilities Act²⁰.

However, where new public policy considerations arise, governments should consider maintaining a sector-specific approach while removing or modifying those regulations that act as a barrier to AI’s development, deployment, and use. In addition, governments should avoid creating a patchwork of AI policies at the state and local levels and should coordinate across governments to advance sound and interoperable practices.

To begin with, companies are very risk-averse regarding potential legal liabilities associated with their use of technology. Further, companies have a market incentive to address associated risks with using artificial intelligence. This is why the Chamber applauds NIST’s development of a “Playbook,” which is “designed to inform AI actors and make the AI RMF more usable.”²¹ The playbook will provide a great resource for the business community and industry in helping them evaluate risk.

Every sector will have different risks associated with using AI, which is why it is important to maintain a sector-specific approach. However, policymakers must do necessary oversight to close current legal gaps. These critical assessments provide lawmakers and the industry with a comprehensive and baseline understanding of relevant regulations that are already in place and give more dialogue on where additional guidance may be necessary.

A. Standards and Best Practices have yet to be developed for audits.

As legislatures and agencies look to potentially regulate the use of AI, they must be aware that the technology is continuing to be developed and that the current processes to mitigate potential bias or concerns could soon be obsolete, which is why the Chamber discourages the use of one-size fits all solutions such as third-party audits. While outside assistance should never be discouraged, it should be noted that there is a well-documented risk²² of engaging third-party auditors. Given that there currently are no standards and certifications regarding third-party auditors, there is no guarantee that reviewers can deliver verifiable measurement methods that are valid, reliable, safe, secure, and accountable.

B. Perspective is key

Lawmakers need to be cognizant of the alternative to not using such technology and the implications that it can lead to. One of the critical benefits of AI is that it provides society with a tool that continues to help complement the workforce and provide efficiency and insights that have led to increased productivity and better outcomes. For this reason, it is essential that lawmakers also consider this when it's framing risks associated with the use of AI. This is why lawmakers should understand the importance of the "human-baseline approach," which asks that the outcomes of the use of the system be compared to the alternative of it being done by a human, not against vague AI-related risks without meaningful context. Furthermore, leaving out the human-baseline comparison could ultimately limit AI adoption, as organizations are only at the risk of the technology and not the totality of AI's benefit for the specific application.

Furthermore, Policymakers might be tempted to rush to regulate or ban the use of artificial intelligence in practices like hiring and employment but it is important to understand that AI can be used to provide opportunities to communities that have historically suffered from bias. Regulation should not hinder the very tools which promise to further equality of opportunity.

1. *Too much information on the system can be a bad thing:*

While transparency around creating automated decision systems and their outputs is critical for building public trust in AEDT, policymakers should avoid any federal mandate requiring the internal working of these systems to be fully divulged, as doing so could lead to the system being gamed and harming the overall trust in these systems. For this reason, should lawmakers look at how companies should provide transparency around the use of the systems, the Chamber would encourage them only to look to provide summaries or a set of takeaways that would provide the necessary transparency to create public confidence for the system while at the same time protecting intellectual property.

1. What Recommendations do you Have for how the Federal Government can Strengthen its Role for the Development and Responsible Deployment of Trustworthy AI Systems?

The federal government has the ability to take a leading role in strengthening the development and deployment of artificial intelligence. We believe that the following recommendations should be acted on now.

First, the federal government should conduct fundamental research in trustworthy AI: The federal government has played a significant role in building the foundation of emerging technologies through conducting fundamental research. AI is no different. A recent report that the U.S. Chamber Technology Center and the Deloitte AI Institute²³ released surveyed business leaders across the United States had found that 70% of respondents indicated support for government investment in fundamental AI research. The Chamber believes that enactment of the CHIPS and Science Act was a positive step as the legislation authorizes \$9 Billion for the

National Institutes of Standards Technology (NIST) for Research and Development and advancing standards for “industries of the future,” which includes artificial intelligence.

Furthermore, the Chamber has been a strong advocate for the National Artificial Intelligence Initiative Act, which was led by then-Chairwoman Eddie Bernice Johnson and Ranking Member Lucas, which developed the office of the National AI Initiative Office (NAIIO) to coordinate the Federal government’s activities, including AI research, development, demonstration, and education and workforce development.²⁴ The business community strongly advises Congress to appropriate these efforts fully.

Second, the Chamber encourages continued investment into Science, Technology, Engineering, and Math Education (STEM). The U.S. Chamber earlier this year polled the American public on their perception of artificial intelligence. The findings were clear; the more the public understands the technology, the more comfortable they become with its potential role in society. Education continues to be one of the keys to bolstering AI acceptance and enthusiasm as a lack of understanding of AI is the leading indicator for a push-back against AI adoption.²⁵

The Chamber strongly supported the CHIPS and Science Act, which made many of these critical investments, including \$200 million over five years to the National Science Foundation (NSF) for domestic workforce buildout to develop and manufacture chips, and also \$13 billion to the National Science Foundation for AI Scholarship-for-service. However, the authorization within the legislation is just the start; Congress should appropriate the funding for these important investments.

Third, the government should prioritize improving access to government data and models: High-quality data is the lifeblood of developing new AI applications and tools, and poor data quality can heighten risks. Governments at all levels possess a significant amount of data that could be used to improve the training of AI systems and create novel applications. When C_TEC asked leading industry experts about the importance of government data, 61% of respondents agree that access to government data and models is important. For this reason, the Chamber encourages EEOC to look at opening up government data which can assist with the training of AEDT’s.

Fourth, Increase widespread access to shared computing resources : In addition to high- quality data, the development of AI applications requires significant

computing capacity.

However, many small startups and academic institutions lack sufficient computing resources, which in turn prevents many stakeholders from fully accessing AI's potential. When we asked stakeholders within the business community about the importance of shared computing capacity, 42% of respondents supported encouraging shared computing resources to develop and train new AI models. Congress took a critical first step by enacting the National AI Research Resource Task Force Act of 2020. Now, the National Science Foundation and the White House's Office of Science and Technology Policy should fully implement the law and expeditiously develop a roadmap to unlock AI innovation across all stakeholders.

Fifth, Enable open source tools and frameworks : Ensuring the development of trustworthy AI will require significant collaboration between government, industry, academia, and other relevant stakeholders. One key method to facilitate collaboration is by encouraging the use of open source tools and frameworks to share best practices and approaches to trustworthy AI. An example of how this works in practice is the National Institute of Standards and Technology's (NIST) AI Risk Management Framework (RMF), which is intended to be a consensus-driven, cross-sector, and voluntary framework, akin to NIST's existing Cybersecurity Framework, whereby stakeholders can leverage as a best practice to mitigate risks posed by AI applications. Policymakers should recognize the importance of these types of approaches and continue to support their development and implementation

Conclusion

AI leadership is essential to global economic leadership in the 21st century. According to one study, AI will have a \$13 trillion impact on the global economy by 2030.²⁶ The federal government can play a critical role in incentivizing the adoption of trustworthy AI applications through the right policies. The United States has an enormous opportunity to transform our economy and society in positive ways through leading in AI innovation. As other economies around the world contemplate their approach to trustworthy AI it is imperative that U.S. policymakers pursue a wide range of options to advance trustworthy AI domestically and empower the United States to maintain global competitiveness in this critical technology sector. The United States must be the global leader in AI trustworthiness for the technology to develop in a balanced manner and takes into account fundamental values and

ethics. The United States can only be a global leader if the public and private sectors work together on a bipartisan basis.

1 [https://www.5newsonline.com/article/news/health/northwest-health-introducing-new-technology-to-
\(http://www.5newsonline.com/article/news/health/northwest-health-introducing-new-technology-to-\).enhance-maternal-and-fetal-safety/527-9c173d18-c56e-457b-8317-62ebaae93558](https://www.5newsonline.com/article/news/health/northwest-health-introducing-new-technology-to-
(http://www.5newsonline.com/article/news/health/northwest-health-introducing-new-technology-to-).enhance-maternal-and-fetal-safety/527-9c173d18-c56e-457b-8317-62ebaae93558)

2 <https://americaninnovators.com/research/data-for-good-promoting-safety-health-and-inclusion/>

3 <https://www.uschamber.com/technology/us-chamber-releases-artificial-intelligence-principles> (<http://www.uschamber.com/technology/us-chamber-releases-artificial-intelligence-principles>)

4 <https://www.nscai.gov/wp-content/uploads/2021/03/Full-Report-Digital-1.pdf> (<http://www.nscai.gov/wp-content/uploads/2021/03/Full-Report-Digital-1.pdf>)

5 <https://www.nist.gov/artificial-intelligence> (<http://www.nist.gov/artificial-intelligence>)

6 www.americaninnovators.com/aicommission (<http://www.americaninnovators.com/aicommission>)

7 <https://americaninnovators.com/aicommission/>

8 <https://americaninnovators.com/news/ai-for-all-experts-weigh-in-on-expanding-ais-shared-prosperity-and-reducing-potential-harms/>

9 <https://americaninnovators.com/news/ai-for-all-experts-weigh-in-on-expanding-ais-shared-prosperity-and-reducing-potential-harms/>

10 [https://americaninnovators.com/wp-content/uploads/2022/04/CTEC_RFI-Alcommission_2.pdf?](https://americaninnovators.com/wp-content/uploads/2022/04/CTEC_RFI-Alcommission_2.pdf?utm_source=sfmc&utm_medium=email&utm_campaign=&utm_term=RFI+3+)
[utm_source=sfmc&utm_medium=email&utm_campaign=&utm_term=RFI+3+](https://americaninnovators.com/wp-content/uploads/2022/04/CTEC_RFI-Alcommission_2.pdf?utm_source=sfmc&utm_medium=email&utm_campaign=&utm_term=RFI+3+)

+Workforce+-+20220518&utm_content=5/19/2022

11 https://uschambermx.iad1.qualtrics.com/jfe/form/SV_cMw5ieLrlsFwUP

12 [7] <https://americaninnovators.com/wp-content/uploads/2022/01/CTEC-US-Outlook-on-AI-Detailed-Analysis.pdf>

13 <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf> (<http://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>)_1414 Federal Register :: Notice of Request for Information (RFI) on Public and Private Sector Uses of Biometric Technologies

15 <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/deloitte-analytics/us-ai-institute-government-public-services-dossier.pdf>

16 <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/deloitte-analytics/us-ai-institute-government-public-services-dossier.pdf>

17 <https://www.opm.gov/policy-data-oversight/data-analysis-documentation/federal-employment-> (<http://www.opm.gov/policy-data-oversight/data-analysis-documentation/federal-employment->) reports/reports-publications/federal-civilian-employment/

18 <https://federalnewsnetwork.com/mike-causey-federal-report/2021/06/retirement-tsunami-this-time->

for-sure/

19 42 U.S.C. § 2000e-2(a).

20 42 U.S.C. § 12101.

21 <https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook-faqs> (<http://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook-faqs>)

Policy implications lawmakers should consider while looking at regulating new technologies such as AI

22 [https://www.gmfus.org/sites/default/files/2022-11/Goodman%20%26%20Trehu%20-\(http://www.gmfus.org/sites/default/files/2022-11/Goodman%20%26%20Trehu%20-\).](https://www.gmfus.org/sites/default/files/2022-11/Goodman%20%26%20Trehu%20-(http://www.gmfus.org/sites/default/files/2022-11/Goodman%20%26%20Trehu%20-).pdf)

%20Algorithmic%20Auditing%20-%20paper.pdf

23 <https://www.uschamber.com/technology/investing-trustworthy-ai>
(<http://www.uschamber.com/technology/investing-trustworthy-ai>)

24 <https://www.ai.gov/naiio/> (<http://www.ai.gov/naiio/>)

25 <https://americaninnovators.com/wp-content/uploads/2022/01/CTEC-US-Outlook-on-AI-Detailed-Analysis.pdf>

26 <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-> (<http://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier->) [modeling-the-impact-of-ai-on-the-world-economy](#)



U.S. Equal Employment Opportunity Commission

Testimony of ReNika Moore

Chair Burrows and Members of the Commission:

Thank you to Chair Burrows and the Commission for the invitation to testify at this public meeting. My remarks will be principally focused on the potential for employment discrimination when using algorithms, artificial intelligence (“AI”), and machine learning (“ML”) in automated decision-making (“ADM”) systems in hiring.

I am the Director of the Racial Justice Program at the American Civil Liberties Union (“ACLU”). In my role, I lead the ACLU’s racial justice litigation, advocacy, grassroots mobilization, and public education to dismantle barriers to equality for people of color. Prior to joining the ACLU, I served as Labor Bureau Chief of the New York Office of the Attorney General (“NYAG”). During my tenure, the Labor Bureau was nationally recognized for aggressively enforcing labor standards on behalf of low-wage workers who were disproportionately people of color and immigrants.

Before joining the NYAG, I supervised and coordinated the NAACP Legal Defense Fund’s economic justice litigation, public education, and public policy efforts. I litigated high-impact racial justice cases tackling a variety of civil rights issues, including major class actions challenging racial discrimination in employment. I also practiced at the plaintiff-side employment law firm Outten & Golden LLP, representing workers who had been unlawfully discriminated against or had been unlawfully denied their earned wages.

The ACLU is a nationwide, non-profit, non-partisan organization of nearly 2 million members dedicated to defending the principles of liberty and equality embodied in the U.S. Constitution and our nation’s civil rights laws. Founded more than 100 years ago, the ACLU has participated in numerous cases in state and federal court, including the U.S. Supreme Court, involving the scope and application of employment discrimination and other federal civil rights laws. The ACLU’s Racial Justice Program advocates in a range of issue areas including employment, education, housing, and the criminal legal system. We also work closely with our

ACLU colleagues who specialize in disability rights, women's rights, technology, data science, and analytics.

Thank you to my ACLU colleagues who provided guidance, suggestions, and feedback, and a special thank you to Olga Akselrod and Marissa Gerchick, who assisted with preparing this testimony.

In Part (I), I discuss the legacy and continuing reality of systemic discrimination in employment and the overarching ways in which bias and discrimination can infect technologically-driven ADM systems in employment. In Part (II), I detail the widespread use of tech-driven tools throughout the labor market and the specific tools used in employment, with a focus on hiring, and the ways these tools are vulnerable to bias based on protected characteristics. In Part (III), I offer recommendations to the Commission to improve employer compliance, transparency, and fairness for workers.

I. The legacy and continuing reality of systemic discrimination in employment

The COVID-19 pandemic led to sweeping changes in how huge numbers of jobs are filled. Technology led much of this massive change, with many employers dramatically expanding their use of technologically-driven ADM tools and products to recruit, hire, monitor, and evaluate workers. Yet, even as the use of employment-related technologies seems to become ubiquitous, the pandemic exposed that some of the oldest, most persistent dysfunctions of our labor markets and workplaces – discrimination, segregation, and exclusion based on race, ethnicity, gender, LGBTQ status, disability, and national origin – continue to limit opportunities for workers with marginalized identities. The history of discriminatory labor practices reaches far back and touches many different groups. The depth and breadth of this history demand that we prioritize equity and anti-discrimination protections for all workers. If we fail to acknowledge the pervasiveness of bias and discrimination in employment, we will fall short of taking the actions necessary, such as new guidance, research, and enforcement, to guarantee equal opportunity. We must have comprehensive public oversight, transparency, and accountability to guarantee that jobseekers and employees do not face the same old discrimination dressed up in new clothes.

1. A Deeply entrenched legacy of employment discrimination based on race, gender, and other protected characteristics persists.

Since the earliest days of the United States with its violent displacement of Indigenous people and dependence on the chattel slavery of Africans and their descendants, the most important aspects of work, such as who worked, in what job, under what conditions, and for what compensation, have been determined too often by a person's identity, *e.g.*, their race, ethnicity, or gender, rather than by what they were qualified to do. Examples of race and ethnicity limiting opportunity have been the rule rather than the exception. In the South, even after slavery was abolished, Jim Crow laws and customs limited the jobs Black people could hold. In the western United States, as Chinese immigration rose through the 1800s, Chinese immigrants were limited to dangerous, low-paying work building railroads and were denied job opportunities in most other sectors. In the West and the Southwest, Mexican-Americans and immigrants also faced violence, discrimination, and exploitation and were disproportionately restricted to low-wage farm labor. The lowest paid agricultural and domestic workers have been almost exclusively of color, including Black, Mexican, Filipino, and Central American. The New Deal established new protections for most workers, but agricultural and domestic workers were excluded from the federal minimum wage, overtime, collective bargaining, and other protections.

Title VII of the Civil Rights Act of 1964 outlawed employment discrimination explicitly based on race and gender, and other historically marginalized categories. [1] While employers began complying with the letter of the new law, almost immediately they began to undermine the spirit of Title VII. Employers began imposing educational and testing requirements to create new barriers for Black workers, barriers like those challenged in *Griggs v. Duke Power Co.*, [2] the seminal civil rights case first establishing the disparate impact theory of discrimination. The story of *Griggs* illustrates how new systems may appear at first glance to be unbiased or less-biased than the systems they replace, when in fact they may simply mask or worsen the same old discrimination. The *Griggs* example also highlights the critical, necessary role that the EEOC can play in protecting against evolving forms of discrimination.

Willie Boyd, a Black man, was one of the thirteen plaintiffs in *Griggs*. Mr. Boyd was the son of sharecroppers and he grew up toiling on the family's tobacco farm in North Carolina. When he began working at the Duke Power Company plant in the

mid-1950s, he saw the job as a significant improvement over the farm, but he found that his position was not so different from his sharecropper parents because there were no opportunities for Black workers to advance. The plant had four departments, but Black workers were only permitted to work in one, “the labor department,” doing the most menial jobs in the plant for the lowest pay – in fact, the Black workers referred to themselves as janitors. The highest-paid Black worker made less than the lowest-paid white worker. Prior to the passage of Title VII, the workforce was explicitly segregated by race: Black workers were forced to use segregated bathrooms, water fountains, and lockers. After Title VII was passed, Duke Power shed its explicitly racist practices and segregation.^[3] But it quickly adopted new requirements to work in every department except the labor department. The new requirements mandated that any employee who wanted to work in a department other than the labor department had to pass two general knowledge standardized tests. These new requirements effectively blocked all Black workers from transfers.

Mr. Boyd, who had become active in his local NAACP chapter, organized his Black coworkers and, with the help of the NAACP Legal Defense Fund, filed a charge with the then-newly established EEOC. The EEOC investigated and found that the tests were not job-related and discriminated against Black workers. The EEOC’s investigation laid the groundwork for the litigation that ultimately reached the U.S. Supreme Court. The combined efforts of the workers themselves, advocates, and the EEOC culminated in the Supreme Court ruling that the discriminatory tests were unlawful. Mr. Boyd went on to earn a promotion, becoming the first Black supervisor at Duke Power.

Since *Griggs*, the EEOC, advocates, and workers themselves have sought to identify and root out systemic barriers that discriminate based on historically marginalized characteristics. These efforts are only possible when the devices are known and can be investigated and evaluated.

There is also a long history of workers being denied opportunities because of their gender. Women have faced discrimination and segregation that cabined them into jobs in just a few sectors. Even when they have worked in male-dominated sectors, women have been paid less and had fewer opportunities for advancement. Employers, with the cooperation of newspapers, plainly advertised jobs to women and men separately.^[4] The jobs for women were for administrative support, domestic work, and other stereotypical “women’s work,” and the positions were

generally lower-paying, often part-time, and emphasized physical appearance as compared to jobs targeted to men.^[5] Hiring ads also reflected the occupational segregation based on race *and* gender with ads targeted, *e.g.*, to Black women for domestic work.^[6]

Through the late 20th century, women were disproportionately concentrated in teaching, administrative support, and domestic work.^[7] Black, Indigenous, Latina and other women of color fared even worse than white women and were consistently paid less than their white counterparts.^[8] Disproportionately high numbers of Black and Latina women continue to hold minimum and sub-minimum wage jobs as home health aides, childcare providers, waiters, and domestic and janitorial workers.^[9]

As the data on women of color demonstrate, race compounds the disadvantage of other characteristics too. For example, overall in 2021, people with disabilities were far less likely (by about half) to be employed than people without disabilities, but Black and Latino people with disabilities faced a higher unemployment rate than white people with disabilities.^[10] A survey of decades of data on people with disabilities found that people of color with disabilities were 40% less likely to be hired when unemployed than white jobseekers with disabilities.^[11] Among people who are LGBTQ, Black LGBTQ people, and especially Black trans people, experience higher rates of unemployment.^[12]

B. Tech-driven ADM tools, on their own, will not address systemic discrimination in employment.

It has been 50 years since Mr. Boyd successfully challenged Duke Power's hiring and promotion tests as unlawful. Despite this, anecdotal evidence and various data metrics show widespread employment discrimination based on race, gender, and other protected categories still exists. Throughout our labor markets, and most dramatically at the highest and lowest wage jobs, we still see disparities by race and gender in major employment indicators like unemployment rate, hiring, and pay. During the pandemic, Black and Latino workers experienced the highest rates of unemployment with Black and Latina women experiencing the highest rates within those groups.^[13] The most recent unemployment data from the U.S. Department of Labor show that the unemployment rate for all workers remains low but for 2022, the unemployment rate for Black workers was still at least 90% higher than – sometimes more than double – the rate for white workers.^[14] In hiring, a 2022 study

using over 80,000 fictitious applications to large employers found that otherwise similar applicants with traditionally Black names were less likely to advance than those with more traditionally white names.^[15]

At the same time, since the start of the pandemic in early 2020, the use of tech-driven ADM tools for recruiting and hiring has skyrocketed. As described in more detail in Part II, these new tools use algorithms or preset rules, AI, and ML to automate recruiting, sourcing, interviewing, and monitoring, among other employment processes. These tools are marketed as cheaper, more efficient, and non-discriminatory or less discriminatory than their predecessors. While these tools may theoretically be able to help employers identify and hire more diverse pools of candidates, these benefits are not proven. In fact, there is a dearth of controlled testing comparing human-driven hiring processes with AI-driven processes to evaluate for discrimination. To the contrary, there is research showing that AI-driven tools can lead to *more* discriminatory outcomes than human-driven processes. One recent study of human-driven hiring compared with typical AI-driven hiring found that the standard AI-driven tool selected 50% fewer Black applicants than humans did.^[16]

Furthermore, research has shown that there are various ways that bias and discrimination can creep in when employers rely on algorithms and AI in the hiring process and during employment:

- Overrepresentation in negative, undesirable data: Black and Latino people are over-represented in data sets that contain negative or undesirable information, such as records from **criminal legal proceedings, evictions, and credit history.**^[17] This is a consequence of many factors, including racial profiling of people of color by the police and harsher treatment within the criminal legal system that lead to longer and more serious consequences for Black, Indigenous, Latino and other people of color once arrested. Similarly, Black women are more likely to be targeted for eviction by landlords than other similarly situated groups.^[18] Data sets containing criminal records and eviction records are also notoriously poor quality; they contain incorrect or incomplete names, old and out of date entries, and non-uniform terms to describe charges, dispositions, and other information necessary to understand outcomes.^[19] Black people and many other people of color are similarly disadvantaged by credit history data. Though credit history data is not necessarily an undesirable source of data, employers generally only consider credit data to *disqualify* a candidate for a job

opportunity.^[20] A history of redlining, targeting people of color for predatory subprime loans more likely to be defaulted, and other barriers to accessing mainstream financial institutions has led to disproportionately low **credit scores** for people of color.^[21] 26 million people in the U.S. have no credit history, 19 million have insufficient credit history, and Black and Brown individuals are overrepresented in both categories.^[22] These realities result in people of color having relatively worse credit scores and histories than white people. As with **criminal** legal system and **eviction** records, this problem is compounded by data quality problems that have been documented in credit history data, including errors and misleading or incomplete information.^[23] These data sets are used for **background checks**.^[24] Thus, Black, Latino, and other people of color are more likely to be disadvantaged by and lose out on employment opportunities.

- Underrepresentation in the training data: Where data does not draw from a sufficiently diverse pool and significantly underrepresents groups in the data relative to the population for which the algorithm is used, the algorithm may be less accurate for people in the underrepresented group.^[25] For example, factors that may correlate negatively for white people may correlate positively for Black people, yet the algorithm may not have sufficient representation of data from Black people to accurately gauge the factor as it applies to them.^[26] People with disabilities^[27] and trans people^[28] are more likely to be missing from data altogether. There are many reasons these groups may be invisible. People with disabilities are more likely to have gaps in schooling and employment.^[29] Trans people and other LGBTQ people are more likely to use names and pronouns that do not match their government identification, thus obscuring their information in the data.^[30]
- Bias in the training data and target: Algorithms are trained with enormous amounts of data including past hiring decisions. In general, many algorithms are developed through analyses of correlations between a specified target outcome (e., some quantification of strong work performance) and patterns in the data. Selection of the target may itself introduce bias.^[31] For example, to the extent that the target is employees who will stay at the company for years and have good performance evaluations, those variables are the product of human decision-making and systems grounded in structural discrimination and subject to individual discrimination. Thus, inequities mar the outcomes in those systems, inequities such as prior discriminatory hiring decisions,

subjective performance evaluations, effects of a hostile workplace, or reduced access for people in protected categories to social network in a company.^[32] Similarly, data used to train the algorithm will reflect the outcomes of those same discriminatory decisions and systems, yet be treated by the algorithm as ground truth.^[33]

- Proxies in the training data and inputs: Even where race, gender, or other protected categories are withheld from the algorithm, many data points are proxies for those characteristics either in isolation or in combination, such as zip code, name, college attended, online browsing history, etc.^[34]
- Bias reinforcement through feedback loops: Many algorithms continue to learn after they are initially deployed, incorporating additional data as a kind of “feedback” through use of the algorithm.^[35] For example, targeting algorithms make predictions about who is likely to click on an ad – to the extent a user clicks on the ad as predicted, the algorithm often incorporates that successful click data into subsequent predictions.^[36] These feedback loops can reinforce discriminatory decisions, such as where an algorithm funnels predatory loan ads to Black users and their clicks on those ads lead to more such ads being delivered to those users.^[37]
- Impacts of the digital divide: Among the groups on the wrong side of the digital divide, Black, Indigenous, and Latino households are much less likely to have reliable high-speed internet access. Native Americans living on reservations have the lowest connectivity rates of any racial group.^[38] People with disabilities are also less likely to have high-speed internet access.^[39] Without internet service, people are less likely to engage digitally and/or online with many systems that produce data that is then used to train or otherwise develop tools. This lack of access may also create barriers to employment opportunities, including learning about job opportunities, submitting applications, or requesting accommodations or assistance.

II. Prevalence of algorithms, AI, and ML in employment and the potential for bias and discrimination.

Recent reports indicate at least seven out of ten employers are using ADM tools in their hiring process, including 99% of Fortune 500 companies.^[40] Media reports and employer announcements show increasing use of AI-driven hiring tools for lower

wage jobs in sectors like retail, logistics, and food services.^[41] Black and Latino workers are disproportionately concentrated in these sectors, and they may also interface with tech-driven ADM tools as they seek higher-paying managerial roles. At Amazon, the nation's second largest employer, Black and Latino workers are clustered in entry-level positions and have struggled to advance to the corporate levels, where they are consistently underrepresented. Amazon has faced lawsuits and reports of systemic discrimination.^[42] Against this backdrop, Amazon recently announced that it is moving to hire more employees through internally-developed AI-driven tools.^[43] Given the racial stratification of its workforce, reliance on such tools to select for employment opportunities raises questions about how fair these processes will be for Black and Latino workers – particularly given that Amazon's earlier attempt to use AI-driven tools for hiring is now one of the most frequently cited examples of algorithmic bias in employment because it discriminated against women applicants.^[44]

Employers are using automated tools in virtually every stage of the employment process, from recruiting and hiring to managing and surveilling employees.^[45] Often, workers may have little or no awareness that such tools are being used, let alone of how they work or that these tools may be making discriminatory decisions about them.^[46] While these tools may seem attractive to employers as a way to reduce the cost and time of otherwise resource-intensive employer processes^[47] and are marketed with claims that they are objective and less discriminatory,^[48] many of these tools instead pose an enormous danger of amplifying existing discrimination in the workplace and labor markets and exacerbating harmful barriers to employment based on race, gender, disability, and other protected characteristics.^[49]

This section discusses some of the tools that are being used, but this is by no means exhaustive. For a more detailed look at tools currently in use, please see the following sources:

- Upturn, Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias^[50]
- Upturn, Essential Work: Analyzing the Hiring Technologies of Large Hourly Employers^[51]
- Coworker, Little Tech is Coming for Workers^[52]
- Raghavan, et al., Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices^[53]

1. Recruitment and sourcing tools

In the sourcing stage, when employers seek to find and attract candidates, automated processes have come to play a pivotal role in determining who will and will not learn of a job opportunity. These processes can create major barriers to employment for people, especially people from groups that are already historically excluded from certain industries, and are invisible to most workers.^[54]

One example is the widespread use of targeted advertising for job opportunities, which funnels ads to individuals on job boards, social media, and other online sites based on data collected about their personal characteristics, online behaviors, interests, or location.^[55] Employers use various tools to select who will be shown a job ad. Some tools allow employers to select attributes from a dropdown menu of personal characteristics of people to whom the ad would be targeted. Other tools allow employers to use so-called “lookalike” tools to upload a list of people, which an algorithm then uses to curate an audience list of people with perceived similar attributes or interests.^[56] When ad targeting tools are used to show employment ads on the basis of people’s real or inferred personal characteristics and algorithmic predictions about their interests, others with different predicted characteristics or interests will never be shown the job opportunity.^[57]

Ad targeting tools have repeatedly been a vehicle of both intentional and unintentional discrimination in violation of civil rights laws. In 2018, for example, the ACLU filed a charge with the EEOC against Facebook and several employers that advertised on its platform for the use of trait selection menus and “lookalike” tools that included gender and other protected characteristics or close proxies.^[58] Employers were able to use the tools to directly exclude women and non-binary users from receiving their ads. Or, for example, employers could upload a list of current employees for use with a “lookalike” tool, and if that list was skewed towards white men due to historically biased hiring decisions, their ad would reach a primarily white male audience as the algorithm picked up on race and gender or proxies thereof in determining who would be similar to the list.^[59] While Facebook agreed in 2019 to changes^[60] to remove protected characteristics or close proxies from employers’ audience selection tools and to stop directly using them in Facebook’s determination of who would be “similar” to an audience the employer was seeking to reach, those changes were insufficient to remove discriminatory impact from the use of those tools – the algorithm continued to pick up on even distant proxies for protected characteristics.^[61] Moreover, even when employers

seek to reach a diverse audience, researchers have found that Facebook’s own ad-delivery algorithm and its predictions of what users “want” to see also continues to be biased and based in stereotypes. For example, a recent audit of Facebook’s ad-delivery system found that Facebook continues to withhold certain job ads from women in a way that perpetuates historical patterns of discrimination: ads for sales associates for cars were primarily shown to men, while ads for sales associates for jewelry were shown to women.^[62] While Facebook’s recent sweeping settlement with the Department of Justice (“DOJ”) and its agreement to expand the provisions in that settlement to employment ads will hopefully mean real progress in addressing discrimination on the platform,^[63] discriminatory ad targeting is not unique to Facebook.^[64]

Platforms such as LinkedIn,^[65] ZipRecruiter,^[66] Indeed,^[67] CareerBuilder,^[68] and Monster^[69] also play a crucial role in many employers’ recruitment and sourcing processes and in many job seekers’ search processes. These platforms perform a kind of matching: employers advertise open positions, job seekers upload or post information about their professional interests and backgrounds, and the platforms make recommendations, often in the form of ranked lists, to both candidates and employers about jobs they should apply for or candidates they should consider. These recommendations may be based on information provided by each kind of user – such as resumes provided by candidates or job descriptions provided by employers – as well as data about the user’s prior activity on the platform – like which job ads candidates have clicked on in the past or which candidates employers have reached out to for interviews.^[70] For employers, these platforms offer functionality that differs from the consumer-facing version with which job seekers interact. For example, LinkedIn’s offerings for employers include *LinkedIn Recruiter*, a tool that boasts usage by more than 1.6 million professionals and access to the more than 740 million users on LinkedIn.^[71]

Despite the pervasiveness of these platforms and their integral role in sourcing and recruitment for many employers, these ranking and recommendation systems are generally largely black boxes to candidates and the general public.^[72] What we do know about the candidate and job opportunity recommendations generated by these platforms raises serious concerns about the potential for these matching platforms to enable discrimination with little oversight or accountability, and demonstrates that there are multiple dangers with such recommender systems. For example, a predictive algorithm that assesses which jobseekers are similar to one another in making recommendations risks downplaying or even withholding job

opportunities based on protected characteristics or proxies thereof.^[73] In 2018, LinkedIn publicly shared that it had found that its recommendation system underpinning *LinkedIn Recruiter* generated results that unfairly ranked men over women, potentially enabling feedback loops in recruitment that perpetuated the gender bias.^[74] While LinkedIn has stated that it has taken steps to address this issue,^[75] it raises serious concerns that workers are wholly dependent on the employer/company to disclose and address algorithmic bias on its own. These kinds of biases are likely not limited to LinkedIn alone: researchers have found that recommender systems similar to those that comprise the core of job matching platforms can suffer from algorithmic bias in rankings and recommendations.^[76] We cannot rely solely on companies – which may have little incentive to share negative findings about their algorithms – to regularly self-evaluate for algorithmic bias. The Commission should examine not only the tools of vendors or employers, but also sourcing platforms like LinkedIn, Monster, ZipRecruiter, Indeed, and CareerBuilder, among others.^[77] Jobseekers need concrete protections that provide meaningful transparency and recourse, address algorithmic bias, and prevent discrimination enabled by these systems.

1. Screening and interviewing tools

ADM tools are also widely used at the screening stage, and applicants are now often rejected through algorithmic tools without any human review of their candidacy.^[78] An overwhelming number of employers – 99% of Fortune 500 companies and the vast majority of mid-size and large companies – use an Applicant Tracking System (“ATS”),^[79] many of which have built-in algorithmic tools that employers use to filter out or rank applicants with automated resume screening based on knockout questions, keyword requirements, or specific qualifications or characteristics.^[80] Many employers have also incorporated chatbots and text apps into their online hiring processes, which steer people through the application process, schedule interviews, or ask basic questions of jobseekers such as a jobseeker’s available days, hours, or work history.^[81] These chatbots (and indeed many screening and assessment tools) often do not have information about how to seek reasonable accommodations built into them or displayed in a way that is easy to find, creating additional barriers for persons with disabilities who want to ask for a reasonable accommodation.^[82] Some of these tools are designed to encourage or discourage applications based on answers to questions, and people interacting with these chatbots often will not know the impact their answers will have on their ability to apply for a role or advance in the interview process.^[83] Often, these automated

screening tools create rigid rules for highly specific certifications, credentials or particular descriptions of job experience, or screen for gaps in work history of more than 6 months, which can weed out qualified candidates that a human reviewer may have otherwise interviewed or hired and disproportionately create barriers for people with protected characteristics, such as pregnancy or a disability.^[84]

Employers also use various automated assessment tools to conduct personality testing. Some employers use online versions of multiple-choice personality tests that ask situational questions or questions about a person's outlook or approach to assess amorphous traits such as work style, dependability, whether they like to work in a team, communication style, emotion, enthusiasm, or attention. Other employers use gamified assessments that are video-game style tools that claim to assess similar traits through an automated analysis of how someone plays a game.^[85]

Employers also assess candidates through online video interviewing, whereby a candidate records an interview online in response to a set of standardized question prompts. Some employers solely use these tools as a means of conducting interviews without human labor, and humans later watch and evaluate the interview recording. Other employers use automated analysis tools so that a human never needs to watch the interview.^[86] Vendors of these tools often claim to be able to measure potentially vague and subjective personality traits similar to those in online tests and gamified assessments, sometimes using voice analysis that assesses content and audio factors such as tone, pitch, and word choice and/or video analysis that assesses visual factors such as facial expressions, eye contact, and posture.^[87] Some assessments are sold by vendors as standard applications for particular kinds of job functions.^[88] Others train their algorithms based on data obtained from the employer about its current staff, often having people identified by the employer as its best employees take the tests or undergo the interviews and then using their answers or performance as a baseline for candidate evaluation.^[89]

There are numerous concerns with these assessment tools and other automated screeners.

First, as discussed previously, any tools that rely on existing employee data to train the algorithm may exacerbate discrimination. Predictive hiring tools often rely on training data regarding who would be a successful employee that reflects existing institutional and systemic biases in employment.^[90] An employer's existing workforce may lack diversity, and employer decisions as to who to designate as a

successful employee to serve as the baseline for training is itself subjective and can reflect institutional and systemic biases in the workplace.^[91] The Amazon hiring algorithm that discriminated against women cited above, *supra* note 44, is an example of this.^[92]

Second, many ADM systems function by analyzing a large amount of data to uncover correlations and make predictions related to a target outcome, but the correlations that they uncover may not actually have a causal connection with being a successful employee, may not themselves be job-related, and may be proxies for protected characteristics.^[93] For example, one resume screening company found that its model identified being named Jared and playing lacrosse in high school as indicators of a successful employee, and another determined that there was a correlation between job tenure and residing within a certain distance of the office.^[94] Even when explicit consideration of race or other protected characteristics is removed from the model, the proxy-based correlations that an algorithm unearths to make its decisions can nevertheless lead to discriminatory decisions.^[95]

Moreover, as with traditional personality assessments, automated assessments are often designed to measure subjective and amorphous personality traits – characteristics such as optimism, positivity, ability to handle pressure, or extroversion – that are not clearly job related or necessary for the job, that may reflect standards and norms that are culturally specific, or that can screen out candidates with disabilities such as autism, depression, or attention deficit disorder.^[96] These problems are exacerbated even further with predictive tools that rely on facial and audio analysis or gamified assessments. Of course, there is cause for great skepticism that personality characteristics can be accurately measured through things such as how fast someone clicks a mouse, the tone of a person's voice, or facial expressions.^[97] But even if the tools are somehow generally able to make those measurements accurately, predictive tools that rely on analysis of facial, audio, or physical interaction with a computer raise even more risk that individuals will be automatically rejected or scored lower on the basis of protected characteristics.^[98] For example, there is a high risk that vocal assessments may perform more poorly on people with accents or with speech disabilities, and it has been established that video technology performs more poorly at recognizing the faces of women with darker skin.^[99] Likewise, tools can be inaccessible to people with disabilities when they rely on detection of color or reactions to visual images, measure physical reactions and speed, require verbal responses to question prompts, or are incompatible with screen readers.^[100]

The lack of transparency in the use of these tools only adds to the harm. Applicants know that they are being subjected to an online recorded interview or test assessment, but are rarely provided information on the standards that will be used to analyze them or what the interviews and tests are seeking to measure.^[101] As a result, applicants often do not have enough information about the process to know whether to seek an accommodation or alternative evaluation method.^[102] This dynamic is compounded by the fact that reasonable accommodation notices on online hiring sites are often difficult to find or unclear.^[103] Moreover, the lack of transparency makes it more difficult to detect discrimination, reducing the ability of individuals, the private bar, and government agencies to enforce civil rights laws.^[104]

1. Background checks

ATSs have made it easier than ever for employers to conduct background checks on applicants, allowing for easy integration of background check features for eviction and criminal legal records, finance records, and sometimes even social media searches, amongst others.^[105] As I discussed above, reliance on criminal legal system, eviction, and credit records can inject discrimination into the hiring process.^[106]

1. Post-hiring tools impacting workers

The ACLU's work on technologies used by employers has largely focused on the use of automated technologies for hiring, so my comments do not discuss in detail the tools employers use to evaluate and surveil their employees. But, I will briefly mention those tools and refer the EEOC to some of the useful resources that discuss the tools that are in use.

AI tools are increasingly used in worker evaluation and surveillance, especially in low wage jobs, and are being used by employers for key decisions such as setting hours, promotion, compensation, discipline, and termination.^[107] This includes tools that are used to monitor workers' movements, such as tools that monitor key strokes, time spent on particular tasks and breaks taken from those tasks, and GPS monitoring, tools that monitor worker communications both on and sometimes off the job, such as email and phone monitoring or social media monitoring, as well as

tools that algorithmically evaluate performance including analysis of recorded customer interactions for worker performance through vocal and sentiment analysis.^[108] Many of these tools raise similar concerns to the tools used for hiring, including discrimination based on disabilities and other protected characteristics, but raise additional concerns such as creating barriers to workers organizing, increased encroachments on worker privacy, and setting unreasonable pace and productivity expectations that can lead to increased injuries and harm workers' health. For a detailed discussion of these tools and the problems that they raise, I refer the Commission to the following resources:

- org, Little Tech is Coming for Workers^[109]
- Data and Society, The Constant Boss^[110]
- Data and Society, Algorithmic Management in the Workplace^[111]
- UC Berkeley Labor Center, Data and Algorithms at Work^[112]
- Center for Democracy & Technology, Warning: Bossware May Be Hazardous to Your Health^[113]
-

Discrimination in hiring and in the workplace is nothing new, and it has always been the EEOC's mission to prevent and remedy such discrimination. But the digital tools that are the focus of this hearing are the new frontier of discrimination, and they are more complex and less transparent than what workers have faced before, and threaten to exacerbate existing systemic inequities. In order to ensure that the protections of Title VII, the Americans with Disabilities Act ("ADA"), the Age Discrimination in Employment Act ("ADEA"), and other federal laws are enforced in this new automated landscape, the EEOC will need to meet the moment with robust regulation and enforcement using all of the tools in the EEOC's toolbox.

We applaud the EEOC for the work that it has undertaken to begin to address the harms of new technologies in the employment sphere. The EEOC's creation of the Initiative on AI and Algorithmic Fairness, as well as its collaboration with the DOJ to develop and issue guidance on the application of the ADA to new technologies, are critical first steps.^[114] This section lays out some recommendations for additional EEOC action that builds on that groundwork. I note that many of these recommendations are informed by the ACLU's work in coalition with numerous civil

rights and technology equity groups that have collaborated to advocate for federal government actors to center civil rights in their technology policies.**[115]**

1. The EEOC should issue additional guidance on the application of Title VII and ADEA to the use of tech-driven ADM systems in employment decisions.

The core guidance for employers and vendors on how to assess the fairness and validity of hiring and other selection procedures is the Uniform Guidelines on Employee Selection Procedures (“UGESP”), which was adopted 45 years ago, long before the advent of the kind of technological tools in use today. Many advocates and scholars have raised concerns that the UGESP is dated, including that the UGESP fails to address discrimination on the basis of disability, age, aspects of sex discrimination or intersectional discrimination, and that the UGESP do not clearly state whether employers can establish the validity of a procedure through evidence based on correlations between certain characteristics and job performance without showing such characteristics are necessary to perform the job.**[116]**

The EEOC should address the gaps in the application of the UGESP to new employment technologies. As a starting point, the EEOC should use its recent guidance on the application of the ADA to new AI and algorithmic technologies as a springboard for developing similar guidance on the application of Title VII and the ADEA, whether through technical assistance documents, Questions and Answers, or other guidance documents. Whatever the format, it is critical that the EEOC continue to educate employers – and software vendors – on how their use of these technologies can violate civil rights laws and advise on steps to take to come into compliance. The EEOC should also offer employers additional guidance under the ADA, Title VII and ADEA on the potential for discrimination in the use of technologies for monitoring worker performance and productivity, much of which directly impacts worker compensation, scheduling, benefits, termination, and other key employer decisions.

1. Any EEOC guidance should include more detailed and comprehensive best practice standards.

The EEOC’s recent guidance on the ADA contains some extremely important “promising practices” to help employers meet their obligations under the ADA, including providing reasonable accommodations and alternatives; using tools that

have been designed with accessibility in mind; providing plain language notice to applicants and employees regarding what traits are being assessed and how they are being measured; ensuring that the tools being used “only measure abilities or qualifications that are truly necessary for the job – even for people who are entitled to an on-the-job reasonable accommodation” and that “necessary abilities or qualifications are measured directly, rather than by way of characteristics or scores that are correlated with those abilities or qualifications”; and that employers inquire with vendors whether the tool asks questions of applicants or employees about disability information or are likely to lead to disclosure of such information.**[117]** The guidance also critically advises employers that they could be held liable for “the actions of their agents, which may include entities such as software vendors, if the employer has given them authority to act on the employer’s behalf.”**[118]**

These “promising practices” are indeed some of the critical steps needed to protect the rights of employees and applicants. Next, we recommend the EEOC further clarify *how* employers can ensure their tools conform with those principles – both for ADA compliance and with other civil rights laws. What kind of process will allow employers to determine whether their tools are following promising practices? What specifically should they ask of vendors? When does a tool pose too great of a risk of discrimination and, therefore, should not be used? Robust evaluation of algorithmic systems is crucial here, and because there are currently no industry standards for such evaluations or when mitigation or decommission measures should be employed, the EEOC can help to fill that void with research and detailed guidance about industry best practices for auditing and transparency measures, as well as guidance around what kinds of tools to avoid.

The EEOC can look to several existing sources for models on developing such standards.

First, the ACLU joined with the Center for Democracy & Technology (“CDT”) and a number of other civil society groups to draft the “Civil Rights Standards for 21st Century Employment Selection Procedures,” which were published in December. **[119]** The Civil Rights Standards provide a concrete, detailed road map for civil rights-focused guardrails for automated tools used in employment decisions, such as for pre- and post-deployment audits, short-form disclosures, procedures for requesting accommodations or opting out, record keeping, transparency and notice, and systems for oversight and accountability. The Standards also call for prohibition of “certain selection procedures that create an especially high risk of

discrimination. These include selection procedures that rely on analyzing candidates' facial features or movements, body language, emotional state, affect, personality, tone of voice, pace of speech, and other methods as determined by the enforcement agency.”^[120] One of the lead drafters of the Standards, Matt Scherer of CDT, is likewise testifying before this Commission and will provide further details on what the Civil Rights Standards contain.

Second, the EEOC should look to the White House's recently released Blueprint for an AI Bill of Rights,^[121] which contains comprehensive and robust measures that are very much in line with the growing consensus amongst civil society groups as to what is needed to address algorithmic discrimination and other harms from new technologies, including proactive measures throughout the entirety of an AI lifecycle, such as consultation with the communities directly impacted by system deployment, pre- and post-deployment testing and mitigation or decommissioning when necessary, independent auditing, transparent reporting, and notice and recourse measures for impacted individuals. The AI Bill of Rights framework includes useful discussions of five core principles: a right to safe and effective systems, protections from discriminatory or inequitable algorithmic systems, data privacy, notice and explanation, and human alternatives, consideration, and fallback.^[122]

Third, the EEOC can also look to the National Institute of Standards and Technology (“NIST”) proposal for “Managing Bias within Artificial Intelligence,” for an informative discussion of “technical characteristics needed to cultivate trust in AI systems: accuracy, explainability and interpretability, privacy, reliability, robustness, safety, and security (resilience) – and that harmful biases are mitigated.”^[123] The ACLU cautions that it has raised concerns to NIST that its proposal was too tech-determinist and did not sufficiently include non-technical sociological and ethical considerations, and it remains a work in progress.^[124] Nevertheless, NIST's work provides the EEOC with an opportunity for inter-agency collaboration around the development of clear standards for assessments.

1. Increased enforcement measures, including strategically selected targets.

While there has always been an information gulf between job applicants or workers and the ways that employment practices, especially hiring, may be discriminatory, the increased use of hiring technologies has widened that gulf. Many hiring technologies are invisible to workers, or workers are aware that a technology is being used but not how or the manner in which it is impacting them. This has made

it more challenging for individuals and the private bar to file complaints with the EEOC. It is therefore critical that the EEOC use the full force of its enforcement powers to proactively investigate discrimination in the use of hiring technologies. The EEOC can begin through research and information gathering to identify employers who are using the tools that are at greatest risk for discrimination, and where appropriate, use Commissioner charges under Title VII and the ADA,^[125] and direct investigations under the ADEA and the Equal Pay Act,^[126] to investigate systemic discrimination caused by these tools in the absence of individual complaints.

The ACLU is aware that the EEOC has published a draft of its Strategic Enforcement Plan and currently plans to separately submit comments on that draft during the open comment period.

1. The EEOC should take additional steps, including technical studies, to make hiring tech tools more transparent.

The auditing and notice standards mentioned above are critical to addressing transparency. But the EEOC can also use its additional authority to “make such technical studies as are appropriate to effectuate the purposes and policies of [Title VII] and to make the results of such studies available to the public[.]”^[127] We encourage the EEOC to use the full scope of its authority to conduct technical studies and examine other creative ways that it can encourage the private industry to share information about its practices.^[128] Public reporting on such studies is critical, but the EEOC could report such information in a summary or aggregated form where appropriate.

1. The EEOC should issue guidance on when digital platforms or software vendors can be held directly liable for their role in violations of civil rights laws.

While the EEOC’s recent guidance on the application of the ADA to automated tools discusses how employers can potentially be held liable for the actions of their vendors, more clarity is needed on when digital platforms or software vendors across the employment spectrum can themselves be liable under Title VII, the ADA, the ADEA, and other civil rights laws.^[129] The ACLU and others have argued that in targeting and delivering employment ads, Facebook could be held liable as an

employment agency.[130] In a recent complaint before the EEOC against Facebook, the complainants also argued that Facebook could be held liable for aiding and abetting employment discrimination, and could also be deemed an “employer” in their actions on behalf of an employer.[131] Similar arguments apply to other sourcing and recruiting platforms, and may likewise apply to vendors of other kinds of digital tools used in hiring and employment decisions. The EEOC should issue guidance that provides clarity in this area.

For additional recommendations, including adoption of the internet applicant rule, increased employer recordkeeping and reporting requirements, particularly for disability-related data, and others, please see the July 13, 2021 letter from ACLU and coalition partners to the EEOC.[132]

IV. Conclusion

Thank you to the Commission for convening this meeting to further explore and understand the challenges that these new technologies pose to equal employment opportunity and we look forward to working with the Commission to chart a course forward that protects the rights of all workers.

[1] 42 U.S.C. §§ 2000e, *et seq.*

[2] 401 U.S. 424 (1971).

[3] Robert Belton, *The Crusade for Equality in the Workplace: The Griggs v. Duke Power Story* (2014).

[4] See *Pittsburgh Press Co. v. Pittsburgh Comm’n on Hum. Rels.*, 413 U.S. 376 (1973) (upholding ordinance prohibiting segregated employment ads); Laura Tanenbaum & Mark Engler, *Help Wanted - Female*, *The New Republic* (Aug. 30, 2017), <https://newrepublic.com/article/144614/help-wantedfemale> (<https://newrepublic.com/article/144614/help-wantedfemale>).

[5] Tanenbaum & Engler, *supra* note 4.

[6] *Id.*

[7] Marina Zhavoronkova, Rose Khattar & Mathew Brady, *Occupational Segregation in America*, Ctr. for Am. Progress (Mar. 29, 2022)

<https://www.americanprogress.org/article/occupational-segregation-in-america/> (**<https://www.americanprogress.org/article/occupational-segregation-in-america/>**).

[8] *Id.*

[9] *Id.*

[10] *Disability Employment Statistics*, U.S. Dep't of Lab.: Off. of Disability Emp. Pol'y, **<https://www.dol.gov/agencies/odep/research-evaluation/statistics>** (**<https://www.dol.gov/agencies/odep/research-evaluation/statistics>**) (last visited Jan. 15, 2023) (white people with disabilities had a 9.2% unemployment rate, while Black people with disabilities had a 15.2% unemployment rate, and Latino people with disabilities had a rate of 13.9%).

[11] Edward Yelin & Laura Trupin, *Successful Labor Market Transitions for Persons with Disabilities: Factors Affecting the Probability of Entering and Maintaining Employment*, 1 *Rsch. in Soc. Sci. and Disability* 105–29 (2000).

[12] Movement Advancement Project, Center for American Progress, Human Rights Campaign, Freedom to Work & National Black Justice Coalition, *A Broken Bargain for LGBT Workers of Color* (Nov. 2013), at i, **<https://www.lgbtmap.org/file/a-broken-bargain-for-lgbt-workers-of-color.pdf>** (**<https://www.lgbtmap.org/file/a-broken-bargain-for-lgbt-workers-of-color.pdf>**).

[13] *Bearing the Cost: How Overrepresentation in Undervalued Jobs Disadvantaged Women During the Pandemic*, U.S. Dep't of Lab., 7 (Mar. 15, 2022).
<https://www.dol.gov/sites/dolgov/files/WB/media/BearingTheCostReport.pdf>
(**<https://www.dol.gov/sites/dolgov/files/WB/media/BearingTheCostReport.pdf>**)
.

[14] *Economic News Release: Table A-2, Employment Status of the Civilian Population by Race, Sex, and Age*, U.S. BUREAU OF LAB. STATS., (Jan, 6, 2023)
<https://www.bls.gov/news.release/empsit.t02.htm>
(**<https://www.bls.gov/news.release/empsit.t02.htm>**).

[15] Patrick Kline, Evan K. Rose & Christopher R. Walters, *Systemic Discrimination Among Large U.S. Employers*, 137 *Q. J. of Econ.* 1963, 1963 (2022),

**<https://academic.oup.com/qje/article/137/4/1963/6605934>
(<https://academic.oup.com/qje/article/137/4/1963/6605934>).**

[16] Learning Collider, *Hidden Bias in Hiring: Examining Applicant Screening Technologies*, 12 (2022),

**<https://static1.squarespace.com/static/60d0c05ace34212ef5a1131b/t/62ab8039e3a4642b49f2f730/1655406650864/Learning+Collider%27s+White+Paper+-+Hidden+Bias+in+Hiring+-+2022+Master.pdf>
(<https://static1.squarespace.com/static/60d0c05ace34212ef5a1131b/t/62ab8039e3a4642b49f2f730/1655406650864/Learning+Collider%27s+White+Paper+-+Hidden+Bias+in+Hiring+-+2022+Master.pdf>).**

[17] Valerie Schneider, *Locked Out by Big Data: How Big Data, Algorithms and Machine Learning May Undermine Housing Justice*, 52.1 Colum. Hum. Rts. L. Rev. 251, 270-74 (2020), **<https://hrlr.law.columbia.edu/hrlr/locked-out-by-big-data-how-big-data-algorithms-and-machine-learning-may-undermine-housing-justice/>** (**<https://hrlr.law.columbia.edu/hrlr/locked-out-by-big-data-how-big-data-algorithms-and-machine-learning-may-undermine-housing-justice/>**).

[18] Peter Hepburn, Renee Louis & Matthew Desmond, **Racial and Gender Disparities Among Evicted Americans**, 7 Socio. Sci. 649, 655 (2020), **https://sociologicalscience.com/download/vol-7/december/SocSci_v7_649to662.pdf
(https://sociologicalscience.com/download/vol-7/december/SocSci_v7_649to662.pdf).**

[19] Search Group Inc. & Bureau of Justice Statistics, *Data Quality of Criminal History Records*, Dep't of Just., 1-5 (Dec. 1, 1985) **<https://bjs.ojp.gov/library/publications/data-quality-criminal-history-records>
(<https://bjs.ojp.gov/library/publications/data-quality-criminal-history-records>)** ; Léon Digard & James Kang-Brown, *Yes, the New FBI Data is Poor Quality. But We've Always Needed Better*, Vera Inst. (Oct. 12, 2022), **<https://www.vera.org/news/yes-the-new-fbi-data-is-poor-quality-but-weve-always-needed-better>
(<https://www.vera.org/news/yes-the-new-fbi-data-is-poor-quality-but-weve-always-needed-better>).**

[20] Elizabeth Gravier, **Can Employers See Your Credit Score? How to Prepare for What They Actually See When They Run a Credit Check**, CNBC.com, (Nov. 18, 2022),

<https://www.cnbc.com/select/can-employers-see-your-credit-score/>
(<https://www.cnbc.com/select/can-employers-see-your-credit-score/>).

[21] Kristen Broady, Mac McComas & Amine Ouazad, *An Analysis of Financial Institutions in Black-Majority Communities: Black Borrowers and Depositors Face Considerable Challenges in Accessing Banking Services*, The Brookings Inst. (Nov. 2, 2021), **<https://www.brookings.edu/research/an-analysis-of-financial-institutions-in-black-majority-communities-black-borrowers-and-depositors-face-considerable-challenges-in-accessing-banking-services/>**
(<https://www.brookings.edu/research/an-analysis-of-financial-institutions-in-black-majority-communities-black-borrowers-and-depositors-face-considerable-challenges-in-accessing-banking-services/>).

[22] *CFPB Explores Impact of Alternative Data on Credit Access for Consumers Who Are Credit Invisible*, Consumer Fin. Prot. Bureau (Feb. 16, 2017), **<https://www.consumerfinance.gov/about-us/newsroom/cfpb-explores-impact-alternative-data-credit-access-consumers-who-are-credit-invisible/>**
(<https://www.consumerfinance.gov/about-us/newsroom/cfpb-explores-impact-alternative-data-credit-access-consumers-who-are-credit-invisible/>);
Explaining the Black-White Homeownership Gap, Urb. Inst. 8 (Oct. 2019), **https://www.urban.org/sites/default/files/publication/101160/explaining_the_black-white_homeownership_gap_2.pdf**
(https://www.urban.org/sites/default/files/publication/101160/explaining_the_black-white_homeownership_gap_2.pdf) (“For those with scores, more than half of white people have FICO scores above 700 compared with just 20.6 percent of black people. About one-third of black people do not have a FICO score.”);
Discriminatory Effects of Credit Scoring on Communities of Color, Nat’l Fair Hous. All. 15 (2012), **<https://nationalfairhousing.org/wp-content/uploads/2017/04/NFHA-credit-scoring-paper-for-Suffolk-NCLC-symposium-submitted-to-Suffolk-Law.pdf>** (**<https://nationalfairhousing.org/wp-content/uploads/2017/04/NFHA-credit-scoring-paper-for-Suffolk-NCLC-symposium-submitted-to-Suffolk-Law.pdf>**).

[23] Aaron Klein, *The real problem with credit reports is the astounding number of errors*, The Brookings Inst. (Sept. 28, 2017), **<https://www.brookings.edu/research/the-real-problem-with-credit-reports-is-the-astounding-number-of-errors/>** (**<https://www.brookings.edu/research/the-real-problem-with-credit-reports-is-the-astounding-number-of-errors/>**); Levi

Kaplan, Alan Mislove & Piotr Sapiezynski, *Measuring Biases in a Data Broker's Coverage*, Ne. Univ. (2022),

https://www.ftc.gov/system/files/ftc_gov/pdf/PrivacyCon-2022-Kaplan-Mislove-Sapiezynski-Measuring-Biases-in-a-Data-Brokers-Coverage.pdf (**https://www.ftc.gov/system/files/ftc_gov/pdf/PrivacyCon-2022-Kaplan-Mislove-Sapiezynski-Measuring-Biases-in-a-Data-Brokers-Coverage.pdf**); *Tenant Background Checks Market Report*, Consumer Fin. Prot. Bureau (Nov. 15, 2022), **<https://www.consumerfinance.gov/data-research/research-reports/tenant-background-checks-market-report/>** (**<https://www.consumerfinance.gov/data-research/research-reports/tenant-background-checks-market-report/>**) (examining problems with criminal background checks for tenants but applicable to employment-related criminal background checks as well).

[24] See Gravier, *supra* note 20.

[25] Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 Cal. L. Rev. 671, 680–81 (2016), **<https://www.californialawreview.org/wp-content/uploads/2016/06/2Barocas-Selbst.pdf>** (**<https://www.californialawreview.org/wp-content/uploads/2016/06/2Barocas-Selbst.pdf>**); see also Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, Conference on Fairness, Accountability and Transparency, 81 Proc. of Mach. Learning Rsch. 1 (2018), **<https://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf>** (**<https://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf>**); Allison Koenecke, et al., *Racial Disparities in Automated Speech Recognition*, 117 Proceedings of the National Academy of Sciences 7684, 7684 (2020), **<https://www.pnas.org/doi/epdf/10.1073/pnas.1915768117>** (**<https://www.pnas.org/doi/epdf/10.1073/pnas.1915768117>**).

[26] Pauline Kim, *Data-Driven Discrimination at Work*, 58 Wm. & Mary L. Rev. 857, 878–79 (2017), <https://scholarship.law.wm.edu/cgi/viewcontent.cgi?article=3680&context=wmlr>.

[27] Shari Trewin, *AI Fairness for People with Disabilities: Point of View*, IBM Accessibility Research (Nov. 26, 2018), **<https://arxiv.org/pdf/1811.10670.pdf>** (**<https://arxiv.org/pdf/1811.10670.pdf>**) (explaining challenges of “gathering a balanced set of training data [for disabilities], because there are so many forms and degrees of disability.”).

[28] AI Now Institute, *A New AI Lexicon: Gender: Transgender Erasure in AI; Binary Gender Data Redefining 'Gender' in Data Systems*, Medium (Dec. 15, 2021), <https://medium.com/a-new-ai-lexicon/a-new-ai-lexicon-gender-b36573e87bdc> (<https://medium.com/a-new-ai-lexicon/a-new-ai-lexicon-gender-b36573e87bdc>).

[29] Kevin Rockmael, *Higher Education Can Help Bridge the LD Employment Gap*, Nat'l Ctr. for Learning Disabilities (Oct. 13, 2021), <https://www.nclld.org/news/higher-education-can-help-bridge-the-ld-employment-gap/> (<https://www.nclld.org/news/higher-education-can-help-bridge-the-ld-employment-gap/>) (citing to lower high school and college graduation rates and employment rates for people with disabilities).

[30] See Sandy E. James, et al., *The Report of the 2015 U.S. Transgender Survey*, Nat'l Ctr. for Transgender Equal. 85–91 (2016), <https://transequality.org/sites/default/files/docs/usts/USTS-Full-Report-Dec17.pdf> (<https://transequality.org/sites/default/files/docs/usts/USTS-Full-Report-Dec17.pdf>) (out of over 27,000 respondents, “[m]ore than two-thirds (68%) of respondents did not have any ID or record that reflected both the name and gender they preferred.”); Silver Flight, *Name Changes: Do We Need Judicial Discretion?*, U. of Cincinnati L. Rev. (Oct. 1, 2021), <https://uclawreview.org/2021/10/01/name-changes-do-we-need-judicial-discretion/> (<https://uclawreview.org/2021/10/01/name-changes-do-we-need-judicial-discretion/>) (describing barriers to legal name changes for trans people).

[31] Barocas & Selbst, *supra* note 25, at 679–80.

[32] *Id.*

[33] *Id.* at 680–81.

[34] *Id.* at 690–92. See generally *Using Publicly Available Information to Proxy for Unidentified Race and Ethnicity*, Consumer Fin. Prot. BUREAU (2014), https://files.consumerfinance.gov/f/201409_cfpb_report_proxy-methodology.pdf (https://files.consumerfinance.gov/f/201409_cfpb_report_proxy-methodology.pdf).

[35] Aaron Rieke & Miranda Bogen, *Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias*, Upturn 44 (Dec. 10, 2018),

<https://www.upturn.org/work/help-wanted/>
(<https://www.upturn.org/work/help-wanted/>).

[36] See generally, *Good Questions, Real Answers: How Does Facebook Use Machine Learning to Deliver Ads?*, Facebook (June 11, 2020),

<https://www.facebook.com/business/news/good-questions-real-answers-how-does-facebook-use-machine-learning-to-deliver-ads>
(<https://www.facebook.com/business/news/good-questions-real-answers-how-does-facebook-use-machine-learning-to-deliver-ads>).

[37] See, e.g., Nicol Turner Lee, *Testimony Before the Task Force on Artificial Intelligence, House Committee on Financial Services: Hearing on “Perspectives on Artificial Intelligence: Where We Are and the Next Frontier in Financial Services,”* 7 (June 26, 2019),

<https://docs.house.gov/meetings/BA/BA00/20190626/109735/HHRG-116-BA00-Wstate-Turner-LeePhDN-20190626.pdf>
(<https://docs.house.gov/meetings/BA/BA00/20190626/109735/HHRG-116-BA00-Wstate-Turner-LeePhDN-20190626.pdf>).

[38] Hansi Lo Wang, *Native Americans On Tribal Land Are 'The Least Connected' To High-Speed Internet*, NPR (Dec. 6, 2018),

<https://www.npr.org/2018/12/06/673364305/native-americans-on-tribal-land-are-the-least-connected-to-high-speed-internet>
(<https://www.npr.org/2018/12/06/673364305/native-americans-on-tribal-land-are-the-least-connected-to-high-speed-internet>) (one out of three Native Americans generally lack access to high-speed internet but 47% of Native Americans living on reservations lack access).

[39] *Disability and the Digital Divide: Internet Subscriptions, Internet Use and Employment Outcomes*, Off. of Disability Emp. Pol’y & U.S. Dep’t of Lab. (June 2022),

<https://www.dol.gov/sites/dolgov/files/ODEP/pdf/disability-digital-divide-brief.pdf> (**<https://www.dol.gov/sites/dolgov/files/ODEP/pdf/disability-digital-divide-brief.pdf>**).

[40] Jennifer Alsever, *AI-Powered Speed Hiring Could Get You an Instant Job, but are Employers Moving Too Fast?*, Fast Co. (Jan. 6, 2023),

<https://www.fastcompany.com/90831648/ai-powered-speed-hiring-could-get-you-an-instant-job-but-are-employers-moving-too-fast>
(<https://www.fastcompany.com/90831648/ai-powered-speed-hiring-could-get-you-an-instant-job-but-are-employers-moving-too-fast>).

you-an-instant-job-but-are-employers-moving-too-fast) (“70% of companies now rely on automated tools for scoring candidates and conducting **background checks**, and AI-enabled tools are matching skills to jobs.”); Joseph B. Fuller, et al., *Hidden Workers: Untapped Talent*, Harv. Bus. Sch. & Accenture 20, (Oct. 4, 2021), <https://www.hbs.edu/managing-the-future-of-work/Documents/research/hiddenworkers09032021.pdf> (<https://www.hbs.edu/managing-the-future-of-work/Documents/research/hiddenworkers09032021.pdf>) (citing that 99% of Fortune 500 companies use an applicant tracking system).

[41] See, e.g., Patrick Thibodeau, *Food Industry Turns to AI Hiring Platform to Fill 1M Jobs*, TechTarget (Apr. 9, 2020), <https://www.techtarget.com/searchhrsoftware/news/252481461/Food-industry-turns-to-AI-hiring-platform-to-fill-1M-jobs> (<https://www.techtarget.com/searchhrsoftware/news/252481461/Food-industry-turns-to-AI-hiring-platform-to-fill-1M-jobs>).

[42] See Katherine Anne Long, *New Amazon Data Shows Black, Latino and Female Employees are Underrepresented in Best-Paid Jobs*, The Seattle Times (Apr. 14, 2021, 1:42 PM), <https://www.seattletimes.com/business/amazon/new-amazon-data-shows-black-latino-and-female-employees-are-underrepresented-in-best-paid-jobs/> (<https://www.seattletimes.com/business/amazon/new-amazon-data-shows-black-latino-and-female-employees-are-underrepresented-in-best-paid-jobs/>); Jason Del Rey, *Bias, Disrespect, and Demotions: Black Employees Say Amazon has a Race Problem*, Vox (Feb. 26, 2021, 8:00 AM), <https://www.vox.com/recode/2021/2/26/22297554/amazon-race-black-diversity-inclusion> (<https://www.vox.com/recode/2021/2/26/22297554/amazon-race-black-diversity-inclusion>).

[43] Jason Del Rey, *A Leaked Amazon Memo May Help Explain Why the Tech Giant is Pushing Out So Many Recruiters*, Vox (Nov. 23, 2022, 4:21 PM), <https://www.vox.com/recode/2022/11/23/23475697/amazon-layoffs-buyouts-recruiters-ai-hiring-software> (<https://www.vox.com/recode/2022/11/23/23475697/amazon-layoffs-buyouts-recruiters-ai-hiring-software>).

[44] See, e.g., Jessica Kim-Schmid & Roshni Raveendhran, *Where AI Can – and Can’t – Help Talent Management*, Harv. Bus. Rev (Oct. 13, 2022),

<https://hbr.org/2022/10/where-ai-can-and-cant-help-talent-management>
(<https://hbr.org/2022/10/where-ai-can-and-cant-help-talent-management>)
 (mentioning “the infamous Amazon AI tool that disadvantaged women applicants”).

[45] See Rieke & Bogen, *supra* note 35; Ifeoma Ajunwa, *An Auditing Imperative for Automated Hiring*, 34 Harv. J.L. & Tech. 2 (2021), **https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3437631**
(https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3437631); Mona Sloane, Emanuel Moss & Rumman Chowdhury, *A Silicon Valley Love Triangle: Hiring Algorithms, Pseudo-Science, and the Quest for Auditability*, 3 Patterns 1 (2022), **<https://www.sciencedirect.com/science/article/pii/S2666389921003081>**
(<https://www.sciencedirect.com/science/article/pii/S2666389921003081>).

[46] See *generally Online but Disconnected: Young Adults’ Experiences with Online Job Applications*, JobsFirstNYC 5 (Oct. 23, 2017), **https://jobsfirstnyc.org/wp-content/uploads/2019/11/Online_but_Disconnected.pdf**
(https://jobsfirstnyc.org/wp-content/uploads/2019/11/Online_but_Disconnected.pdf); Aaron Smith & Monica Anderson, *Americans’ Attitudes Toward Hiring Algorithms*, Pew Rsch. Ctr. (Oct. 4, 2017), **<https://www.pewresearch.org/internet/2017/10/04/americans-attitudes-toward-hiring-algorithms/>**
(<https://www.pewresearch.org/internet/2017/10/04/americans-attitudes-toward-hiring-algorithms/>).

[47] Fuller, et al., *supra* note 40, at 11.

[48] For examples of these types of claims, see, e.g., Table 2, Manish Raghavan, et al., *Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices*, Cornell Univ. & Microsoft Rsch. 8 (2020), **<https://arxiv.org/abs/1906.09208>**
(<https://arxiv.org/abs/1906.09208>).

[49] See, e.g., Rieke & Bogen, *supra* note 35, at 8; *Algorithm-Driven Hiring Tools: Innovative Recruitment or Expedited Disability Discrimination?*, Ctr. for Democracy & Tech. 7 (Dec. 2020), **<https://cdt.org/wp-content/uploads/2020/12/Full-Text-Algorithm-driven-Hiring-Tools-Innovative-Recruitment-or-Expedited-Disability-Discrimination.pdf>**
(<https://cdt.org/wp-content/uploads/2020/12/Full-Text-Algorithm-driven-Hiring-Tools-Innovative-Recruitment-or-Expedited-Disability-Discrimination.pdf>).

[50] Rieke & Bogen, *supra* note 35.

[51] Aaron Rieke, et al., *Essential Work: Analyzing the Hiring Technologies of Large Hourly Employers*, Upturn (May 2021), <https://www.upturn.org/work/essential-work/> (<https://www.upturn.org/work/essential-work/>).

[52] Wilneida Negrón, *Little Tech is Coming for Workers*, Coworker (2021), <https://home.coworker.org/wp-content/uploads/2021/11/Little-Tech-Is-Coming-for-Workers.pdf> (<https://home.coworker.org/wp-content/uploads/2021/11/Little-Tech-Is-Coming-for-Workers.pdf>).

[53] Raghavan, et al., *supra* note 48.

[54] Rieke & Bogen, *supra* note 35, at 14.

[55] *Id.* at 17.

[56] For more on lookalike audiences, see generally Piotr Sapiezynski, et al., *Algorithms that “Don’t See Color”: Measuring Biases in Lookalike and Special Ad Audiences*, Ne. Univ. (2019), <https://arxiv.org/pdf/1912.07579.pdf> (<https://arxiv.org/pdf/1912.07579.pdf>).

[57] See generally Linda Morris & Olga Akselrod, *Holding Facebook Accountable for Digital Redlining*, ACLU (Jan. 27, 2022), <https://www.aclu.org/news/privacy-technology/holding-facebook-accountable-for-digital-redlining> (<https://www.aclu.org/news/privacy-technology/holding-facebook-accountable-for-digital-redlining>); Ariana Tobin & Ava Kofman, *Facebook Finally Agrees to Eliminate Tool That Enabled Discriminatory Advertising*, ProPublica (June 22, 2022 4:30 PM), <https://www.propublica.org/article/facebook-doj-advertising-discrimination-settlement> (<https://www.propublica.org/article/facebook-doj-advertising-discrimination-settlement>); Muhammad Ali, et al., *Discrimination through Optimization: How Facebook’s Ad Delivery Can Lead to Biased Outcomes*, 3 Proc. of the ACM on Hum. Comput. Interaction 1 (2019), <https://dl.acm.org/doi/abs/10.1145/3359301> (<https://dl.acm.org/doi/abs/10.1145/3359301>); Till Speicher, et al., *Potential for Discrimination in Online Targeted Advertising: Conference on Fairness, Accountability, and Transparency*, Proc. of Mach. Learning Rsch. 2 (2018), <https://proceedings.mlr.press/v81/speicher18a/speicher18a.pdf> (<https://proceedings.mlr.press/v81/speicher18a/speicher18a.pdf>).

[58] See *Facebook EEOC Complaints*, ACLU, (Sept. 25, 2019), <https://www.aclu.org/cases/facebook-eec-complaints> (<https://www.aclu.org/cases/facebook-eec-complaints>).

[59] Ariana Tobin & Jeremy Merrill, *Facebook is Letting Job Advertisers Target Only Men*, ProPublica (Sept. 18, 2018, 6:39 PM), <https://www.propublica.org/article/facebook-is-letting-job-advertisers-target-only-men> (<https://www.propublica.org/article/facebook-is-letting-job-advertisers-target-only-men>); Ava Kofman & Ariana Tobin, *Facebook Ads Can Still Discriminate Against Women and Older Workers, Despite a Civil Rights Settlement*, ProPublica (Dec. 13, 2019, 5:00 AM), <https://www.propublica.org/article/facebook-ads-can-still-discriminate-against-women-and-older-workers-despite-a-civil-rights-settlement> (<https://www.propublica.org/article/facebook-ads-can-still-discriminate-against-women-and-older-workers-despite-a-civil-rights-settlement>).

[60] See, e.g., *Facebook Agrees to Sweeping Reforms to Curb Discriminatory Ad Targeting Practices*, ACLU (Mar. 19, 2019), <https://www.aclu.org/press-releases/facebook-agrees-sweeping-reforms-curb-discriminatory-ad-targeting-practices> (<https://www.aclu.org/press-releases/facebook-agrees-sweeping-reforms-curb-discriminatory-ad-targeting-practices>).

[61] See Jinyan Zang, *How Facebook's Advertising Algorithms Can Discriminate By Race and Ethnicity*, Tech. Sci. (Oct. 19, 2021), <https://techscience.org/a/2021101901/> (<https://techscience.org/a/2021101901/>).

[62] Karen Hao, *Facebook's Ad Algorithms are Still Excluding Women from Seeing Jobs*, MIT Tech. Rev. (Apr. 9, 2021), <https://www.technologyreview.com/2021/04/09/1022217/facebook-ad-algorithm-sex-discrimination/> (<https://www.technologyreview.com/2021/04/09/1022217/facebook-ad-algorithm-sex-discrimination/>); see also Sara Kingsley, et al., *Auditing Digital Platforms for Discrimination in Economic Opportunity Advertising*, Carnegie Mellon Univ. 1 (June 2020), <https://arxiv.org/ftp/arxiv/papers/2008/2008.09656.pdf> (<https://arxiv.org/ftp/arxiv/papers/2008/2008.09656.pdf>); Nicolas Kayser-Bril, *Automated Discrimination: Facebook Uses Gross Stereotypes to Optimize Ad Delivery*, AlgorithmWatch (Oct. 18, 2020), <https://algorithmwatch.org/en/automated->

[discrimination-facebook-google/ \(https://algorithmwatch.org/en/automated-discrimination-facebook-google/\)](https://algorithmwatch.org/en/automated-discrimination-facebook-google/).

[63] See Tobin & Kofman, *supra* note 57; *Justice Department and Meta Platforms Inc. Reach Key Agreement as They Implement Groundbreaking Resolution to Address Discriminatory Delivery of Housing Advertisements*, Dep't of Just. Off. of Pub. Affs. (Jan. 9, 2023), **<https://www.justice.gov/opa/pr/justice-department-and-meta-platforms-inc-reach-key-agreement-they-implement-groundbreaking>** (**<https://www.justice.gov/opa/pr/justice-department-and-meta-platforms-inc-reach-key-agreement-they-implement-groundbreaking>**).

[64] See, e.g., Pauline Kim & Sharion Scott, *Discrimination in Online Employment Recruiting*, 63 *St. Louis University Law Journal* 1, 8 (2018), **https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3214898** (**https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3214898**) (describing study that found that simulated users who were identified as male or female and who engaged in identical web browsing activities related to job searches on Google were shown very different ads, with an ad for coaching on higher paying executive jobs shown significantly more often to men).

[65] See LinkedIn, **<https://www.linkedin.com/>** (**<https://www.linkedin.com/>**) (last visited Jan. 17, 2023).

[66] See ZipRecruiter, **<https://www.ziprecruiter.com/>** (**<https://www.ziprecruiter.com/>**) (last visited Jan. 17, 2023).

[67] See Indeed, **<https://www.indeed.com/>** (**<https://www.indeed.com/>**) (last visited Jan. 17, 2023).

[68] See CareerBuilder, **<https://www.careerbuilder.com/>** (**<https://www.careerbuilder.com/>**) (last visited Jan. 17, 2023).

[69] See Monster, **<https://www.monster.com/>** (**<https://www.monster.com/>**) (last visited Jan. 17, 2023).

[70] Rieke & Bogen, *supra* note 35, at 19.

[71] See *LinkedIn Recruiter*, LinkedIn, **<https://business.linkedin.com/talent-solutions/recruiter>** (**<https://business.linkedin.com/talent-solutions/recruiter>**) (last visited Jan. 8, 2023).

[72] Rieke & Bogen, *supra* note 35.

[73] Rieke & Bogen, *supra* note 35, at 21.

[74] See Sheridan Wall & Hilke Schellmann, *LinkedIn's Job-Matching AI was Biased. The Company's Solution? More AI.*, MIT Tech. Rev. (June 23, 2021),

<https://www.technologyreview.com/2021/06/23/1026825/linkedin-ai-bias-ziprecruiter-monster-artificial-intelligence/>

(<https://www.technologyreview.com/2021/06/23/1026825/linkedin-ai-bias-ziprecruiter-monster-artificial-intelligence/>); Sahin Geyik, Stuart Ambler &

Krishnaram Kenthapadi, *Fairness-Aware Ranking in Search & Recommendation Systems with Application to LinkedIn Talent Search*, ACM KDD 2221–31 (July 2019),

<https://arxiv.org/pdf/1905.01989.pdf> (<https://arxiv.org/pdf/1905.01989.pdf>).

[75] Wall & Schellmann, *supra* note 74; Sahin Cem Geyik & Krishnaram Kenthapadi, *Building Representative Talent Search at LinkedIn*, LinkedIn (Oct. 10, 2018),

<https://engineering.linkedin.com/blog/2018/10/building-representative-talent-search-at-linkedin> (<https://engineering.linkedin.com/blog/2018/10/building-representative-talent-search-at-linkedin>).

[76] See generally Michael D. Ekstrand, et al., *All the Cool Kids, How Do They Fit In?: Popularity and Demographic Biases in Recommender Evaluation and Effectiveness: In Conference on Fairness, Accountability and Transparency*, 81 Proc. of Mach. Learning Rsch 1, 172–86 (2018), <https://proceedings.mlr.press/v81/ekstrand18b.html>

(<https://proceedings.mlr.press/v81/ekstrand18b.html>); Masoud Mansoury, et al., *Feedback Loop and Bias Amplification in Recommender Systems: In Proceedings of the 29th ACM International Conference on Information & Knowledge Management*,

Ass'n for Comput. Mach. 2145–48 (Oct. 2020), <https://arxiv.org/abs/2007.13019>

(<https://arxiv.org/abs/2007.13019>); Kyle Wiggers, *Researchers Find Evidence of Bias in Recommender Systems*, VentureBeat (July 29, 2020, 12:35 PM),

<https://venturebeat.com/ai/researchers-find-evidence-of-bias-in-recommender-systems/> (<https://venturebeat.com/ai/researchers-find-evidence-of-bias-in-recommender-systems/>).

[77] Alex Engler, *Auditing Employment Algorithms for Discrimination*, Brookings (Mar. 12, 2021), [https://www.brookings.edu/research/auditing-employment-](https://www.brookings.edu/research/auditing-employment-algorithms-for-discrimination/)

[algorithms-for-discrimination/](https://www.brookings.edu/research/auditing-employment-algorithms-for-discrimination/) (<https://www.brookings.edu/research/auditing-employment-algorithms-for-discrimination/>).

[78] Rieke & Bogen, *supra* note 35, at 13; see also Rieke, et al., *supra* note 51, at 20.

[79] Fuller, et al., *supra* note 40, at 20.

[80] Rieke & Bogen, *supra* note 35, at 26.

[81] *Id.*; see also Rieke, et al., *supra* note 51, at 13.

[82] *Cf.* Rieke, et al., *supra* note 51, at 24, 50.

[83] Rieke & Bogen, *supra* note 35, at 27.

[84] Fuller, et al., *supra* note 40, at 22.

[85] Rieke & Bogen, *supra* note 35, at 29; Rieke, et al., *supra* note 51, at 23; *Algorithm-Driven Hiring Tools*, *supra* note 49, at 6.

[86] Rieke & Bogen, *supra* note 35, at 36.

[87] *Algorithm-Driven Hiring Tools*, *supra* note 49, at 6.

[88] See Rieke & Bogen, *supra* note 35, at 29.

[89] Pymetrics is one example of a vendor of this kind of tool. See *id.* at 33.

[90] Kim, *supra* note 26, at 876; Barocas & Selbst, *supra* note 25, at 729–32; Rieke & Bogen, *supra* note 35, at 8.

[91] Kim, *supra* note 26.

[92] See, e.g., Jeffrey Dastin, *Amazon Scraps Secret AI Recruiting Tool that Showed Bias Against Women*, Reuters (Oct. 10, 2018, 7:04 PM), <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G> (<https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>).

[93] Kim, *supra* note 26; Barocas & Selbst, *supra* note 25, at 729–32; Rieke & Bogen, *supra* note 35, at 35.

[94] Dave Gershgorn, *Companies are on the Hook if their Hiring Algorithms are Biased*, Quartz (Oct. 22, 2018), <https://qz.com/1427621/companies-are-on-the->

[hook-if-their-hiring-algorithms-are-biased](https://qz.com/1427621/companies-are-on-the-hook-if-their-hiring-algorithms-are-biased)
(<https://qz.com/1427621/companies-are-on-the-hook-if-their-hiring-algorithms-are-biased>); Kim, *supra* note 26, at 863.

[95] Barocas & Selbst, *supra* note 25, at 729–32.

[96] See, e.g., Luke Stark & Jesse Hoey, *The Ethics of Emotion in Artificial Intelligence Systems: In Proceedings of ACM Conference on Fairness, Accountability, and Transparency (FAccT'21)*, ACM (Mar. 1, 2021), **<https://doi.org/10.1145/3442188.3445939>**
(<https://doi.org/10.1145/3442188.3445939>); *Algorithm-Driven Hiring Tools*, *supra* note 49, at 6; Rieke & Bogen, *supra* note 35; Lydia X. Z. Brown, *How Opaque Personality Tests Can Stop Disabled People from Getting Hired*, Ctr. for Democracy & Tech. (Jan. 6, 2021), **<https://cdt.org/insights/how-opaque-personality-tests-can-stop-disabled-people-from-getting-hired/>** (**<https://cdt.org/insights/how-opaque-personality-tests-can-stop-disabled-people-from-getting-hired/>**).

[97] See generally Luke Stark & Jevan Hutson, *Physiognomic Artificial Intelligence*, *Fordham Intell. Prop., Media & Ent. L. J.*, Forthcoming (Sept. 24, 2021), **https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3927300**
(https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3927300).

[98] *Id.*

[99] See, e.g., *Algorithm-Driven Hiring Tools*, *supra* note 49, at 9; Rieke & Bogen, *supra* note 35, at 36; Buolamwini & Gebru, *supra* note 25, at 77–91.

[100] See, e.g., *Algorithm-Driven Hiring Tools*, *supra* note 49, at 10; *Guidance on Web Accessibility and the ADA*, ADA.gov: U.S. Dep't of Just. Civ. Rts. Division (Mar. 18, 2022), **<https://www.ada.gov/resources/web-guidance/>**
(<https://www.ada.gov/resources/web-guidance/>).

[101] See, e.g., Rieke, et al., *supra* note 51, at 24.

[102] See, e.g., *Algorithm-Driven Hiring Tools*, *supra* note 49, at 10; Rieke, et al., *supra* note 51, at 24.

[103] *Id.*

[104] See generally Rieke & Bogen, *supra* note 35; Engler, *supra* note 77.

[105] Rieke, et al., *supra* note 51, at 21–22.

[106] This Commission has previously recognized some of the ways that background checks may lead to disparate impact based on race or other protected characteristics. *Background Checks: What Employers Need to Know*, EEOC (Mar. 11, 2014), <https://www.eeoc.gov/laws/guidance/background-checks-what-employers-need-know> (<https://www.eeoc.gov/laws/guidance/background-checks-what-employers-need-know>).

[107] Jodi Kantor & Aryan Sundaram, *The Rise of the Worker Productivity Score*, N.Y. Times (Aug. 14, 2022), <https://www.nytimes.com/interactive/2022/08/14/business/worker-productivity-tracking.html> (<https://www.nytimes.com/interactive/2022/08/14/business/worker-productivity-tracking.html>).

[108] *Id.*; see also Annette Bernhardt, Reem Suleiman & Lisa Kresge, *Data and Algorithms at Work: The Case for Worker Technology Rights*, U. Cal. Berkeley Lab. Ctr. (Nov. 3, 2021), <https://laborcenter.berkeley.edu/data-algorithms-at-work/> (<https://laborcenter.berkeley.edu/data-algorithms-at-work/>); Tom Simonite, *This Call May be Monitored for Tone and Emotion*, Wired (Mar. 19, 2018), <https://www.wired.com/story/this-call-may-be-monitored-for-tone-and-emotion/> (<https://www.wired.com/story/this-call-may-be-monitored-for-tone-and-emotion/>).

[109] Negrón, *supra* note 52; see also *Bossware and Employment Tech Database*, Coworker (Nov. 17, 2021), <https://home.coworker.org/worktech> (<https://home.coworker.org/worktech>).

[110] Aiha Nguyen, *The Constant Boss: Work Under Digital Surveillance*, Data & Soc’y (May 2021), https://datasociety.net/wp-content/uploads/2021/05/The_Constant_Boss.pdf (https://datasociety.net/wp-content/uploads/2021/05/The_Constant_Boss.pdf).

[111] Alexandra Mateescu & Aiha Nguyen, *Algorithmic Management in the Workplace*, Data & Soc’y (Feb. 2019), https://datasociety.net/wp-content/uploads/2019/02/DS_Algorithmic_Management_Explainer.pdf (https://datasociety.net/wp-content/uploads/2019/02/DS_Algorithmic_Management_Explainer.pdf).

[112] Bernhardt, Suleiman & Kresge, *supra* note 108.

[113] Center for Democracy & Technology, *Warning: Bossware May Be Hazardous to Your Health* (July 29, 2021), <https://cdt.org/insights/report-warning-bossware-may-be-hazardous-to-your-health/>.

[114] *The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees*, EEOC (May 12, 2022), <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence> (<https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence>).

[115] See, e.g., Letter from ACLU, American Association of People with Disabilities, Bazelon Center for Mental Health Law, Center for Democracy & Technology, Center on Privacy & Technology at Georgetown Law, Lawyers' Committee for Civil Rights Under Law, the Leadership Conference on Civil and Human Rights, and Upturn to EEOC, *Coalition Memo: Addressing Technology's Role in Hiring Discrimination* (July 13, 2021), <https://www.aclu.org/letter/coalition-memo-addressing-technologys-role-hiring-discrimination> (<https://www.aclu.org/letter/coalition-memo-addressing-technologys-role-hiring-discrimination>); Olga Akselrod, *How Artificial Intelligence Can Deepen Racial and Economic Inequities*, ACLU (July 13 2021), <https://www.aclu.org/news/privacy-technology/how-artificial-intelligence-can-deepen-racial-and-economic-inequities> (<https://www.aclu.org/news/privacy-technology/how-artificial-intelligence-can-deepen-racial-and-economic-inequities>) (discussing letter to federal administration signed by two dozen partner organizations asking the administration to take concrete action to address equity and civil rights concerns in AI and technology policy).

[116] See, e.g., Jenny Yang, *Testimony before the House Civil Rights and Human Services Subcommittee: The Future of Work: Protecting Workers' Civil Rights in the Digital Age*, Urb. Inst. 9–10 (Feb. 5, 2020), https://www.urban.org/sites/default/files/publication/101676/testimony_future_of_work_and_technology_-_jenny_yang_0_2.pdf (https://www.urban.org/sites/default/files/publication/101676/testimony_future_of_work_and_technology_-_jenny_yang_0_2.pdf); Raghavan, et al., *supra* note 48, at 17; Rieke & Bogen, *supra* note 35, at 11, 46; Rieke, et al., *supra* note 51, at 29–30.

[117] *The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees*, *supra* note 114, at Q14.

[118] *Id.* at Q3.

[119] Matt Scherer & Ridhi Shetty, *Civil Rights Standards for 21st Century Employment Selection Procedures*, Ctr. for Democracy & Tech. (Dec. 5, 2022), <https://cdt.org/insights/civil-rights-standards-for-21st-century-employment-selection-procedures/> (<https://cdt.org/insights/civil-rights-standards-for-21st-century-employment-selection-procedures/>).

[120] *Id.* at 8.

[121] *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People*, White House (Oct. 2022), <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf> (<https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>).

[122] *Id.*

[123] Reva Schwartz, et al., *A Proposal for Identifying and Managing Bias in Artificial Intelligence*, Nat'l Inst. Standards & Tech. Spec. Pub. 1270 i (June 2021), <https://doi.org/10.6028/NIST.SP.1270-draft> (<https://doi.org/10.6028/NIST.SP.1270-draft>).

[124] *ACLU Comment on NIST's Proposal for Managing Bias in AI*, ACLU (Sept. 10, 2021), <https://www.aclu.org/letter/aclu-comment-nists-proposal-managing-bias-ai> (<https://www.aclu.org/letter/aclu-comment-nists-proposal-managing-bias-ai>).

[125] 42 U.S.C. § 2000e-5(b).

[126] See 29 U.S.C. § 626; 29 U.S.C. § 211(a).

[127] 42 U.S.C. § 2000e-4(g)(5). Similar authority is granted by the ADA. 42 U.S.C. § 12117(a).

[128] See also Rieke, et al., *supra* note 51, at 37, 41.

[129] See, e.g., Rieke & Bogen, *supra* note 35, at 22 (discussing ambiguity in liability of recruiting platforms).

[130] *Facebook EEOC Complaint – Charge of Discrimination*, ACLU 6–7, 13 (Sept. 18, 2018), <https://www.aclu.org/legal-document/facebook-eec-complaint-charge-discrimination> (<https://www.aclu.org/legal-document/facebook-eec-complaint-charge-discrimination>); *Real Women in Trucking v. Meta Platforms, Inc. Charge* 41–42 (Dec. 1, 2022), <http://guptawessler.com/wp-content/uploads/2022/12/Real-Women-in-Trucking-Meta-Charge.pdf> (<http://guptawessler.com/wp-content/uploads/2022/12/Real-Women-in-Trucking-Meta-Charge.pdf>).

[131] *Meta Platforms, Inc. Charge*, *supra* note 130, at 42.

[132] See, e.g., Letter from ACLU, *supra* note 115.

Panel 2



U.S. Equal Employment Opportunity Commission

Testimony of Manish Raghavan

Thank you, Chair Burrows, Vice Chair Samuels, and Members of the Commission, for the opportunity to participate in today's hearing on employment discrimination in AI and automated systems.

My name is Manish Raghavan and I'm an assistant professor at the MIT Sloan School of Management and Department of Electrical Engineering and Computer Science. I hold a PhD in computer science from Cornell University. I research the impacts of algorithmic tools on society, and in particular, the use of machine learning in employment contexts. I've extensively studied the development of these tools and have had in-depth conversations with data scientists who build them. My testimony today will be on technical aspects of how the four-fifths rule of thumb has been applied to algorithmic systems.

Introduction

Predictive models

My testimony will focus on how the four-fifths rule of thumb has been applied to predictive models. By "predictive model" or simply "model," I mean a piece of software that takes as input data about an applicant (e.g., a resume) and outputs a score intended to measure the quality of the applicant. Developers typically create models based on historical data. For example, given a stack of resumes, each annotated with its quality (perhaps manually labeled by an expert, perhaps based on past interviewing decisions), a developer can build a model that can essentially extrapolate these quality labels to new resumes. This practice is commonly known as "machine learning." While my testimony today will not dwell on the technical details of how this is done, feel free to ask me questions on the subject.

Testing for algorithms for discrimination

Predictive models are frequently used in employment contexts to evaluate and score applicants and employees. As with other employment assessments, predictive models can be used either for binary reject/advance decisions or to give numeric scores to applicants. When binary decisions must be made, scores are often converted to decisions by thresholding: those with a score above the threshold get one outcome (e.g., an interview), while those with a score below the threshold get another (e.g., no interview).

Developers of these models can test them to see if they result in significantly different selection rates between different protected groups. Importantly, developers run these tests before the model is actually deployed. This requires that the developer collect a data set on which to measure selection rates. This data set must be representative – that is, they must resemble the actual population who will be evaluated. Using this data set, a developer can attempt to determine whether the model in question will satisfy the four-fifths rule of thumb (or a statistical test designed to look for selection rate disparities). If the model fails such a test, the developer can modify or re-build it to reduce selection rate disparities.

Statistical significance

When using quantitative tools to detect events like discrimination, we typically consider two properties: effect size and statistical significance. Effect size characterizes how salient an observation is – for example, when we observe that the selection rate of one group is 60% of the selection rate of another, we're looking at effect size. If it were instead 40% the selection rate of another group, we'd call that a bigger effect. Statistical significance considers how likely we would be to observe this outcome by random chance, as opposed to because of discrimination. Even if selection rates are very different, it would be hard to draw conclusions if an employer has only hired 3 people, as opposed to if they had hired 300. The more observations we have, the more statistically significant conclusions we can draw.

The four-fifths rule of thumb considers only effect size, not statistical significance. In practice, employers use a suite of formal statistical tests, as opposed to simply

relying on the four-fifths rule of thumb. The Uniform Guidelines recognize the importance of considering both effect size and statistical significance.^[1] Throughout this testimony, I will refer to the “four-fifths rule of thumb”; however, my remarks apply to this broader class of statistical techniques designed to consider both the effect size and statistical significance of differences in selection rates.

Limits of the four-fifths rule of thumb

The four-fifths rule of thumb and related statistical tests have technical and operational limitations, which have long been pointed out by psychologists. In what follows, I’ll lay out some of these shortcomings in the context of predictive models.

Retrospective and prospective uses

The four-fifths rule of thumb was initially designed to be used *retrospectively*: a selection rule would be deployed in practice, and an auditor could later analyze the selection rates of various groups. In contrast, the four-fifths rule of thumb is increasingly being used *prospectively* by employers or vendors of predictive models. In other words, before deploying or selling a model, a developer will attempt to determine whether this model will satisfy the four-fifths rule of thumb when deployed.

While prospective testing can be useful, it introduces an important limitation: the conclusions of any prospective test depend heavily on the data collected to perform that test. When applied retrospectively, this isn’t as much of a problem: the data have already been generated by past applicants. But for prospective uses, a developer must explicitly collect a dataset which they believe to be representative of the applicant pool. In effect, they must try to guess what the applicant pool will look like. If the data collected differ significantly from the applicant pool, then it is impossible to draw valid conclusions from this dataset. Moreover, because applicant pools differ by region and occupation, a developer must find some way to maintain representative data for each context in which a model will be deployed.

When a firm collects a dataset on which to evaluate a model, there is no guarantee that it will do so in good faith. In fact, regulations that require or encourage prospective auditing can create incentives to curate datasets that make it “easier” for a model to pass a statistical test. If a firm is worried that their model under-selects applicants from a particular demographic group, they can simply add more qualified applicants from that demographic group to their dataset, thereby increasing the group’s measured selection rate on that dataset. This doesn’t make the model itself any more or less discriminatory; it simply affects whether it passes the test. Thus, a prospective audit is only as reliable as the data collector. It is well within a data collector’s power to alter a dataset such that a model appears to satisfy the four-fifths rule of thumb even if it will not do so when deployed in practice. For statistical tests that measure statistical significance in addition to effect size, firms may try to collect smaller datasets in general, since these will make selection rate disparities harder to detect statistically.

Auditing with centralized data

One tempting response to the problems introduced by data collection is to attempt to centralize collection. If a third party (e.g., a regulator) collects and maintains data, firms will lose their ability to manipulate datasets used for statistical tests. This approach faces a major hurdle: datasets used to evaluate a predictive model must contain exactly the information required as input to that model. A model that makes predictions based on recorded video interviews requires a dataset containing such interviews. A model that makes predictions based on questionnaires requires a dataset of responses to questionnaires. Thus, the dataset used for a firm’s model must be specific to the firm in question; a regulator cannot simply collect a common dataset to be used by all firms. Centralized data collection would require the regulator to collect a new dataset for each firm or model to be evaluated, which may be prohibitively expensive or simply infeasible.

Thresholding a model’s outputs

The four-fifths rule of thumb is designed for binary decisions (i.e., yes/no decisions). In contrast, predictive models are often continuous, meaning a model may output a number instead of a yes/no decision. For example, many models are designed to produce a score between 0 and 1, reflecting the predicted likelihood that (say) an applicant is qualified. But statistical tests are typically defined with respect to binary

labels; an applicant was either selected or not. To produce binary labels from continuous model outputs, practitioners often use a threshold: scores above the threshold are treated as “selections,” while scores below the threshold are not. Importantly, the choice of threshold can affect whether a model passes or fails a statistical test on a given dataset. A model may pass a test when we use one threshold, but fail the test when we use another.

In some cases, thresholding scores is a reasonable approximation to how employers use models in practice. Some employers simply set thresholds and interview all applicants who score above the threshold. But in other cases, model predictions are used in far more complex ways. An employer might simply rank applicants and interview them sequentially until they make an offer. Or a human evaluator may take the scores into account as one of many factors in their decision-making process. In such cases, running a statistical test on thresholded scores does not reflect the conditions under which the scores are deployed.

Concretely, consider a model that outputs scores between 0 and 1. When analyzed before deployment, suppose an employer tests the model using a threshold of 0.5. When they do, suppose the selection rates of different demographic groups do not differ significantly. However, when the model is used in practice, the employer gets far more applicants than expected and must raise the threshold that they use in order to reduce the number of applicants they interview. They decide to only interview candidates with a score above 0.75 instead of 0.5. Now, it’s possible for the model to produce significantly different selection rates when used with this new threshold. Even though selection rates were similar at a threshold of 0.5, there’s no guarantee that they will be with a threshold of 0.75. Thus, when analyzing selection rates using thresholded scores, it’s important to choose a threshold that reflects real-world conditions. If there’s uncertainty about what real-world usage will look like, one approach employers can take is to test a model across a range of possible thresholds instead of just picking one.

Validity

A final point of concern with developers’ and auditors’ focus on the four-fifths rule of thumb is that it detracts from questions regarding a model’s overall validity. Models

that satisfy the four-fifths rule of thumb do not necessarily have much predictive value – for example, a model that selects a completely random subset of the population will in general satisfy the four-fifths rule of thumb. Thus, satisfying the four-fifths rule of thumb provides no guarantee that a model actually does a good job of predicting the thing it claims to predict.

When developers focus on the four-fifths rule of thumb without concern for validity, this can have negative downstream consequences for applicants from marginalized groups. For example, the four-fifths rule of thumb does not rule out *differential validity*, where a model makes more accurate predictions for one demographic group than another.

Concretely, a model suppose that when a developer builds their model, they ensure that selection rates between different groups are roughly equal.^[2] They evaluate the model's validity and find that it has reasonably good predictive power. However, if they were to disaggregate the model's validity and specifically look at the validity for white vs. Black applicants, they may find significant differences. We call such differences differential validity, and they often arise in predictive applications when the developer has more historical data on one group than on another.^[3] There is a key difference between selection rate disparities and differential validity: selection rate disparities cause fewer applicants from one group than another to be selected, which a regulator can observe after the fact. Differential validity can cause fewer *qualified* applicants from one group to be selected. This is much harder for a regulator to detect. Were there simply fewer qualified applicants from that group to begin with, or was this a consequence of differential validity? The four-fifths rule of thumb is designed to detect selection rate disparities. It does nothing to prevent differential validity.

Desirable properties

Despite its limitations, the four-fifths rule of thumb has some desirable properties. For one, it does not depend on labeled outcomes (i.e., past decisions or evaluations of quality). Unlike other measures (including differential validity), the four-fifths rule of thumb is unaffected by so-called ground truth. This has the advantage that it is unchanged by inaccurate or biased labels. Suppose a firm builds a model based on

past hiring decisions. If those past decisions were discriminatory, a model can replicate those discriminatory decisions without appearing to have differential validity, simply because it is accurately reflecting those discriminatory decisions. In contrast, an employer or auditor using the four-fifths rule of thumb in this hypothetical case will notice that the model in question produces selection rate disparities, regardless of whether the model is “accurate” according to historical data. As a result, the four-fifths rule of thumb can serve as a check against poor or biased measures of outcomes.

In this sense, the four-fifths rule of thumb can be viewed as aspirational in nature.^[4] Instead of purely assessing the world as it is, it incentivizes the reduction of significant differences between demographic groups. While this may not accurately reflect the present state of affairs, it can provide incentives to push for expanded opportunities for those historically underrepresented. For example, in order to reduce disparities in selection rates, a firm may increase its outreach to encourage qualified individuals from underrepresented backgrounds to apply.

Finally, the four-fifths rule of thumb can create some benefits on the margin by pressuring firms to search for equally accurate models with minimal selection rate disparities.^[5] While research in the past has found strong trade-offs between selection rate disparities and validity,^[6] the introduction of more modern machine learning techniques has made this trade-off less stark; models with very similar accuracy can vary dramatically in their subgroup-specific selection rates, and the four-fifths rule of thumb can encourage firms to seek to minimize adverse impact across models with similar performance.^[7] The cost of this search for alternatives is dropping, and as it does, it should become standard practice for model developers.

^[1] In particular, they state: “Smaller differences in selection rate may nevertheless constitute adverse impact, where they are significant in both statistical and practical terms”

^[2] There are a variety of techniques developers can use in practice to do this. One common approach is to remove any model inputs that appear to create selection rate disparities until those disparities fall to an acceptable range. See, e.g.,

Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020, January). Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 conference on fairness, accountability, and transparency* (pp. 469-481).

[3] Buolamwini, J., & Gebru, T. (2018, January). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency* (pp. 77-91). PMLR.; Koenecke, A., Nam, A., Lake, E., Nudell, J., Quartey, M., Mengesha, Z., ... & Goel, S. (2020). Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences*, 117(14), 7684-7689.

[4] Friedler, S. A., Scheidegger, C., & Venkatasubramanian, S. (2021). The (im) possibility of fairness: Different value systems require different mechanisms for fair decision making. *Communications of the ACM*, 64(4), 136-143.

[5] Raghavan, M., & Barocas, S. (2019). Challenges for mitigating bias in algorithmic hiring. Brookings.

[6] Mcdaniel, M. A., Kepes, S., & Banks, G. C. (2011). The Uniform Guidelines are a detriment to the field of personnel selection. *Industrial and Organizational Psychology*, 4(4), 494-514.

[7] Black, E., Raghavan, M., & Barocas, S. (2022). Model Multiplicity: Opportunities, Concerns, and Solutions. In *Conference on Fairness, Accountability, and Transparency* (pp. 850-863)



U.S. Equal Employment Opportunity Commission

Testimony of Nancy T. Tippins

Thank you, Chair Burrows, Vice Chair Samuels, and members of the Commission, for the opportunity to participate in today's hearing on artificial intelligence and algorithmic fairness.

My name is Nancy Tippins. I am an industrial and organizational (I/O) psychologist, and I am here today, representing the Society for Industrial and Organizational Psychology (SIOP), which is the professional organization for I/O psychologists, who study human behavior in the context of work. Many of our 10,000+ members are rigorously trained in the development, validation, and implementation of employee selection procedures.[1] For over a century, I/O psychologists have worked with employers to develop a wide range of tests (e.g., multiple choice test, open ended responses tests, interviews, work samples, etc.) that measure a variety of skills and abilities and demonstrate the relationship of test scores to future behavior such as job performance, absenteeism, accidents, etc.

The Society for Industrial and Organizational Psychology

SIOP is vitally interested in AI-based assessments, their development, their statistical and psychometric characteristics, and their operational use.[2] To that end, SIOP has engaged in multiple efforts to share scientific knowledge regarding tests and assessments used for employment decisions. SIOP sets professional standards and guidelines for tests used for hiring and promotion by publishing and regularly updating its *Principles for the Validation and Use of Personnel Selection Procedures* (*Principles*, 2018), which reflect current scientific research and best practices in testing for hiring and promotion.[3] In addition, a SIOP task force (SIOP, 2022) studied artificial intelligence-based (AI-based) assessments and established five requirements that supplement the *Principles*:

- The content and scoring of AI-based assessments should be clearly related to the job(s) for which the assessment is used.
- AI-based assessments should produce scores that are fair and unbiased.
- AI-based assessments should produce consistent scores (e.g., upon re-assessment) of job-related characteristics.
- AI-based assessments should produce scores that accurately predict future job performance (or other relevant outcomes).
- All steps and decisions relating to the development, validation, scoring, and interpretation of AI-based assessments should be documented for verification and auditing.

The SIOP task force has developed a more detailed supplement to the *Principles* that explains how various professional requirements for employment testing also apply to AI-based assessments and was just released (SIOP, 2023). In addition, SIOP is working with the Society for Human Resource Management (SHRM) to provide workshops to Human Resource professionals on AI-based assessments. We believe that these collective efforts provide guidance on the development, validation, and use of AI-based assessments, reflecting contemporary science and practice related to employment tests in general and AI-based assessments specifically.

From our perspective, much of the *Principles* is aligned with the *Uniform Guidelines on Employee Selection Procedures* (UGESP, 1978) and applies to AI-based assessments. However, I would like to highlight five key areas in which the 2018 *Principles* go beyond the 1978 UGESP with implications for AI-based assessments.

Job Analysis

The UGESP requires some form of job analysis (UGESP, 1978, Section 14A) to determine measures of work behaviors or performance relevant to the job. Although a review of job requirements is expected, the appropriate method of job analysis is not specified in the UGESP. The information from the job analysis should be used to determine the appropriateness of the criterion used in validation research. When AI-based assessments are developed, the job analysis information should be used to identify appropriate criteria (e.g., job performance, turnover) against which supervised machine learning algorithms will be trained.

The *Principles* require job analysis not only to justify the criteria in most cases but also to determine what knowledge, skills, abilities, or other characteristics (KSAOs) to measure. The theoretical foundation of personnel selection asserts that jobs are composed of sets of tasks and those tasks require KSAOs to perform them.^[4] Measures of the KSAOs that are important and needed at entry are included in a test battery used for employment decisions. Thus, a job analysis determines what KSAOs are important and needed at entry, provides the foundation for evaluating the extent to which the critical KSAOs are covered, and helps to ensure the criterion measure is appropriate.

The job analysis provides evidence to support the job relevance of a selection procedure. A predictor-criterion relationship (e.g., test-job performance relationship) is an important source of evidence that supports job relevance; however, a correlation alone is not sufficient to indicate relevance. A job analysis facilitates our understanding of how a predictor relates to the requirements of the job and the criterion.

The amount of rigor needed in the job analysis depends in part on the type of test and its purpose. For example, a job knowledge test usually requires a detailed specification of the knowledge domain required for a specific job so that it may be sampled appropriately using test items. In contrast, universally relevant criteria such as absenteeism or turnover typically need less rigor to justify their importance. When the job analysis results are used to identify the KSAOs to measure, more rigor is usually needed. With AI-based assessments, this can be challenging because these assessments often use hundreds (or even thousands) of predictors that are not designed to measure specific KSAOs.

Validity

The UGESP identifies three approaches to validation: criterion-related validity, content validity, and construct validity. The *Principles* define validity as “the degree to which evidence and theory support the interpretations of test scores for proposed uses of tests” (*Principles*, 2018, p. 96) and describes similar approaches to accumulating evidence of validity. In the *Principles*, validity is a unitary concept. There are not different types of validity; instead, there are different sources of validity evidence. Importantly, the validity evidence should support the interpretation of the test score to be made. For example, if an employer wants to predict job performance, the validity evidence might come from a test-job performance relationship. Scores from machine learning algorithms that use quasi-

criteria and identify applicants like or unlike a group of “good” employees should not be interpreted as predicting job performance without additional evidence.

Due to the nature of most AI-based assessments that rely on an algorithm derived from machine learning, the focus has been on criterion-related validity evidence, which is based on a statistical predictor-criterion relationship, instead of content or construct validity evidence. Although content validation strategies are theoretically possible to execute, in many situations, there are simply too many variables (or features) for subject matter experts to make judgments regarding their relationship to the job requirements. Similarly, construct validation studies are possible when the KSAOs being measured are defined, but evidence from a construct validation study relating, for example, one measure of customer service to another does not support a prediction about future job performance. Such evidence only indicates the two measures are related to each other.

Evidence of validity is necessary for multiple reasons. From a legal perspective, validation is required of selection procedures for which there is adverse impact in most situations (UGESP, 1978, Section 3A). From a professional perspective, validation is necessary to demonstrate the accuracy and value of the selection procedure regardless of whether or not adverse impact exists. In addition, validation evidence is necessary for employers to evaluate alternative selection procedures and identify those that have greater or substantially equal validity and less adverse impact.

There are a number of challenges to establishing evidence of validity for any assessment, including those based on artificial intelligence. Sufficient sample sizes can pose problems for traditional approaches to validation based on correlations as well as newer ones based on machine learning models. In either case, the researcher should determine *a priori* the sample size that is required for the statistics used and ensure an adequate sample is achieved.

Statistical requirements alone do not drive sample size; the need for proper representation of the applicant pool and incumbents also affects sampling. To generalize the validity of a test, regardless of its nature, the validation sample should represent the applicant population. Ideally, the applicant population and the incumbent population are similar. If the AI-based assessment is trained on data from an incumbent population that is not similar to the applicant population, the employer runs the risks of using an algorithm that is only applicable to a limited segment of the applicant population.

Traditionally, assessments have been updated when there are substantial changes in the job and its requirements or in the applicant pool or when there is evidence that the assessment has been compromised and is no longer useful. The platforms on which many AI-based assessments are administered have the capability of updating algorithms whenever new data are available. Dynamic updating of this nature poses significant challenges in the documentation of the validity of each version of the algorithm and comparisons of applicants whose scores depend on different algorithms.

The traditional metric for criterion-related validity is some form of correlation, for example, r or R^2 . AI-based assessments derived from machine learning models often use other metrics, such as mean absolute error or mean squared error. To compare the validity of different selection procedures, processes for equating different metrics need to be identified, agreed upon, and reported in technical documentation. Alternatively, scores resulting from the application of algorithms could be validated using appropriate criteria in traditional ways (e.g., correlations, regression).

Fairness

The UGESP calls for studies of fairness when technically feasible and suggests the user “review the A.P.A. Standards regarding investigation of possible bias in testing” (UGESP, 1978, Section 14B(8)).^[5] Psychologists take a broad perspective on fairness. In the context of employment testing, fairness is a multi-dimensional term, each with different meanings for psychologists (*Principles*, 2018, pp. 38-42).

- The term *equal outcomes* refers to equal pass rates or mean scores across groups. Although relevant to assessing disparate impact, this definition of fairness has been rejected by testing professionals. However, the *Principles* suggest that a lack of equal outcomes should serve as an impetus for further investigation as to the source of those differences. Many AI-based assessments incorporate routines to eliminate or minimize group differences that may not be appropriate if those procedures adjust scores on the basis of group membership.
- *Equitable treatment* refers to equitable testing conditions, access to practice materials, performance feedback, opportunities for retesting, and

opportunities for reasonable accommodations. The *Principles* recommend that employers audit their selection systems to ensure equal treatment for all applicants.

The proliferation of computer-based testing often results in applicants taking tests on different devices with different internet connections in situations that vary in the level of distractions. The effect of the device and internet connection depends in part on the type of test. For example, scores on an untimed, multiple-choice measure of personality may be unaffected by device and internet connection, but scores on an AI-based gamified assessment may partially depend on the speed with which the test taker can respond. Tests like video-based interviews may require a stable, high-speed internet connection to capture responses accurately. Employers should inform applicants of the ideal conditions for taking an assessment and provide alternatives to applicants who lack appropriate conditions or access to equipment that meets the test's technical requirements. Equitable treatment also incorporates the opportunity for reasonable accommodations.

- *Equal access* to constructs refers to the opportunity for all test takers to show their level of ability on the job-relevant KSAOs being measured without being unduly advantaged or disadvantaged by job-irrelevant personal characteristics, such as race, ethnicity, gender, age, and disability. Thus, a video-based interview that evaluates response content, facial features, and voice characteristics should not limit an individual with a disability from demonstrating relevant skills unless facial features and voice characteristics can be demonstrated to be job-related.
- *Freedom from bias* refers to a lack of systematic errors that result in subgroup differences. *Measurement bias* refers to systematic errors in test scores or criterion measures that are not related to the KSAOs being measured. For example, items regarding leadership experiences on sports team might disadvantage women. One way to examine measurement bias is through a sensitivity review conducted by subject matter experts who examine items and instructions to determine if a predictor is differentially understood by

demographic, cultural, or linguistic groups. However, when hundreds of variables are used in an algorithm, demonstrating freedom from measurement bias may be difficult because evaluating each item may not be feasible.

Predictive bias refers to systematic errors that result in subgroup differences in the predictor-criterion relationship. In traditional forms of employment testing, predictive bias is usually evaluated by comparing the slopes and intercepts of the regression lines of each group. The methods for evaluating bias when complex algorithms are used have not been widely researched or tested in court decisions.

Although a finding of unequal outcomes is not sufficient evidence of unfairness, all of these forms of fairness are important. Employers should take steps to ensure tests are unbiased and administered appropriately.

Documentation

Documentation of the development and validation of an employment test should encompass all the information in Section 15B of the UGESP, including the underlying data on which the computations were made. In addition, employment testing professionals recommend that documentation of AI-based assessments should include details that are specific to such assessments, e.g., information on how the algorithm was selected, how the model was developed, and how the algorithmic model is translated into an AI-based assessment. Documentation should be sufficient for computational reproducibility.

Adverse Impact

The *Guidelines* are clear on the requirements for documenting adverse impact of the overall selection process (UGESP, 1978, 3A). In addition, adverse impact of the components should be documented if the overall score has adverse impact (UGESP, 1978, 4C). Subsequent court decisions (*Connecticut v. Teal*) also require analysis of the adverse impact of each step of a multiple hurdle selection process.

Although the UGESP describes the “four-fifths rule” as an appropriate measure to determine if evidence of adverse impact exists (UGESP, 1978, 4D), it may not be

sufficient. Because the *Principles* are a document focused on professional standards for employment tests, it does not discuss adverse impact. However, because legal compliance is critically important to employers, the *Principles* encourage testing professionals to take legal considerations into account (*Principles*, pp. 43-44). In practice, most I/O psychologists recognize the complexity of evaluating adverse impact and assess it in a variety of ways, including the four-fifths rule, the binomial distribution, chi-square, Fisher's exact test, etc. (Morris & Dunleavy, 2017; Outtz, 2010).

Conclusion

AI-based assessments hold the promise of being effective tools for predicting future behavior in systematic, unbiased ways. SIOP has carefully developed standards for employment tests that represent the consensus of opinion among I/O psychologists and are aligned with the requirements of the UGESP. We believe that the *Principles* should apply to all tests used for employment decisions. The challenge before us is to determine how best to apply these existing standards to AI-based selection procedures.

Again, thank you for the opportunity to participate today. SIOP members have important and unique expertise in personnel selection. We are willing to engage at any time to address further questions and concerns around these matters.

References

American Educational Research Association, American Psychological Association, National Council on Measurement in Education. (2014). *Standards for educational and psychological testing*. Washington, DC: American Psychological Association.

Connecticut v. Teal, 457 U.S. 440, 102 S. Ct. 2525 (1982).

Guion, R. M. (1998). *Assessment, measurement, and prediction for personnel decisions*. Lawrence Erlbaum Associates Publishers.

Morris, S.B. & Dunleavy, E.M. (2017). *Adverse impact analysis: Understanding data, statistics, and risk*. Routledge.

Ottz, J. L. (Ed.). (2010). *Adverse impact: Implications for organizational staffing and high stakes selection*. Routledge/Taylor & Francis Group.

Society for Industrial and Organizational Psychology. (2018). *Principles for the validation and use of personnel selection procedures* (5th ed.). Retrieved from <https://www.apa.org/ed/accreditation/about/policies/personnel-selection-procedures.pdf> (available here:

<https://www.apa.org/ed/accreditation/about/policies/personnel-selection-procedures.pdf>

(<https://www.apa.org/ed/accreditation/about/policies/personnel-selection-procedures.pdf>)

Society for Industrial and Organizational Psychology. (January 29, 2022). SIOP statement on the use of artificial intelligence (AI) for hiring: Guidance on the effective use of AI-based assessments.

https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligence.pdf?ver=mSGVRY-z_wR5iluE2NWQPQ%3d%3d

Society for Industrial and Organizational Psychology. (January 23, 2023).

Considerations and Recommendations for the Validation and Use of AI-Based Assessments for Employee Selection. **<https://www.siop.org/Portals/84/SIOP-AI%20Guidelines-Final-010323.pdf?ver=iwuP4slt7y21h66ELuiPzQ%3d%3d>**

(<https://www.siop.org/Portals/84/SIOP-AI%20Guidelines-Final-010323.pdf?ver=iwuP4slt7y21h66ELuiPzQ%3d%3d>)

Uniform Guidelines on Employee Selection Procedure (1978); 43 FR __ (August 25, 1978).

[1] Selection procedure refers to any tool used for employee selection, including traditional forms of tests, assessments, interviews, work samples, and simulations as well as more recent forms of tests and assessments that rely on artificial intelligence such as games, recorded interviews, and algorithms that use data from resumes, applications, social media.

[2] Artificial intelligence refers to a broad range of technologies and statistical techniques that are applied to candidate information and used to predict future job performance or other criteria (e.g., turnover, safety behavior, accidents).

[3] The current edition of the SIOP *Principles* (2018) is aligned with the latest version of another set of testing guidelines, *Standards for Educational and Psychological Testing* (2014) (Standards). The *Standards* is written for all types of educational and psychological testing; the *Principles* are specific to employment testing.

[4] See Guion (1998).

[5] “A.P.A. Standards” refers to the *Standards for Educational and Psychological Testing*. The latest edition was published in 2014.



U.S. Equal Employment Opportunity Commission

Testimony of Gary D. Friedman

I. Introduction

Chair Burrows and Commissioners, thank you for inviting me to testify before the Equal Employment Opportunity Commission (“EEOC” or “Commission”) on this important emerging topic. I am a senior partner at Weil, Gotshal & Manges LLP in its Employment Litigation Practice Group, and I represent employers in a wide range of employment-related matters, including discrimination and other complex employment class and collective actions, trade secrets and restrictive covenant litigations, and internal investigations. I have handled scores of matters on behalf of employers before the EEOC, including matters involving Commissioner’s charges and Commission-initiated investigations. I have also testified before public bodies on behalf of management on myriad topics, including proposed changes to the discrimination and harassment laws. I have also written and spoken on artificial intelligence (“AI”) issues in the employment context and have been advising global businesses across sectors on the use of AI in the workplace.

A 2022 study by the *Society for Human Resource Management* (“SHRM”) found that “nearly 1 in 4 organizations report using automation or artificial intelligence (AI)¹ to support HR-related activities,” including recruitment and hiring.² The proliferation of AI shows no signs of abating. In 2021, one market research firm valued the AI market at \$59.67 billion, and estimated that number would grow to \$422.37 billion by 2028.³ In fact, just last week, Microsoft announced that it was making a multi-billion dollar investment in OpenAI, the start-up behind the viral ChatGPT chatbot. Because of AI’s enormous growth potential, and the potential issues it raises with respect to bias and other workplace concerns, I applaud the Commission for taking proactive steps to help ensure employers’ use of AI tools complies with our existing federal employment discrimination laws.

In my testimony, I hope to bring forward the perspective of employers which use, or may in the future use, automated decision-making tools. I do not intend to speak on behalf of the management bar per se, Weil, or any particular client, but I can synthesize what I see as the important considerations for employers on this topic. I think what you will find is that, as a general matter, the objectives of employers, employees and applicants, and the Commission on AI in the workplace are aligned in many ways.

II. Companies want to use AI responsibly.

1. Companies are increasingly using

We've seen that more and more companies are using AI in their recruitment and hiring practices, performance evaluations, and general decision-making regarding employment. Survey statistics support this anecdotal evidence. A recent study showed that up to 83% of large employers surveyed are using some form of AI in employment decision-making.⁴ According to a February 2022 survey from SHRM, 79% of employers use AI and/or automation for recruitment and hiring.⁵ Modern Hire, a vendor of AI hiring technology, advertises that its clients include more than half of the Fortune 100 companies, including FedEx, LG, Macy's, Pepsico, Delta, Starbucks, Sysco, Volvo, and Roche.⁶ In my experience, corporate employers are not only using AI more frequently, but are also focused on using it responsibly, recognizing the benefits of AI in, among other things, reducing unconscious biases that are often present in human decision-making.

Companies' responsible use of AI to reduce unconscious bias in employment decision-making should not come as a surprise. In the past few years, increasingly more public and private companies have prioritized diversity, equity, and inclusion initiatives. After the death of George Floyd and the subsequent protests around the country, businesses pledged \$200 billion to increase efforts toward racial justice.⁷ In 2022, a survey by the American Productivity & Quality Center showed that, in the previous year, 36% of respondents increased staff dedicated to Diversity, Equity, and Inclusion ("DEI"), 32% increased their DEI budgets, and 30% disclosed DEI metrics publicly and invested more in employee resource and affinity groups.⁸ Moreover, pay equity audits are on the rise, with 58% of organizations reporting in 2021 that they have reviewed their pay structures and decisions.⁹ Businesses are incentivized now, more than ever, to take action to improve diversity and reduce bias in the workplace.

B. Companies use AI in the hiring and employment context.

AI is here, and it's not going anywhere. A survey of 7,300 human resources managers worldwide found that the proportion who said their department uses predictive analytics increased nearly 400%, from 10% in 2016 to 39% in 2020.¹⁰ That is not surprising because there are many obvious incentives for companies to use AI. AI can reduce costs through efficiencies in hiring and decision-making processes. And, more importantly, AI can help foster a more diverse workforce and minimize unconscious human biases, which are key goals for 21st century employers.¹¹

Companies today are using AI to assist in a variety of contexts, including anonymizing resumes and interviewees, performing structured interviews, and using neuroscience games to identify traits, skills, and behaviors.¹² During the interviewing stage, some companies will conduct video interviews of applicants and then use AI to analyze factors including facial expression, eye contact, and word choice.

Companies also use AI to monitor employee productivity and performance, and to make decisions regarding promotion and salary increases.¹³ For example, UPS uses AI to monitor and report on driver safety and productivity by tracking driver movement and when drivers put their trucks in reverse.¹⁴ Other companies may use AI to track employee login times, and monitor whether employees are paying attention to their computer screens using webcams and eye-tracking software.¹⁵

C. Companies are trying to use AI technology to mitigate bias and improve diversity.

The use of AI technology can help avoid decisions that treat similarly situated applicants and employees differently based on entrenched bias or even just the whims of individual decision makers. For example, if the criteria for hiring or promotion are set in advance, using an algorithm to assess employees can help reduce the bias of individual managers by applying the criteria uniformly. A Yale study showed that when evaluating candidates for police chief, human evaluators justified choosing men without college degrees over women with college degrees because "street smarts" purportedly was the most important criteria. However, when the names on the applications were reversed, evaluators chose men with college degrees over women without college degrees, claiming that degrees were

the more important criteria. If the criteria had been set in advance, unconscious biases against women could have been mitigated because evaluators would not have been able to justify their decisions post hoc. Importantly, AI can be trained on certain criteria and, unlike humans, AI tools won't deviate from pre-selected criteria to rationalize a biased decision.¹⁶ Shortly, I will discuss some illustrative examples of how AI has been shown to reduce bias in real- world hiring decisions.

The McKinsey Global Institute has reported that AI can reduce the effect of humans' subjective interpretation of data because machine-learning algorithms learn to consider only variables that improve predictive accuracy.¹⁷ For example, algorithms can consider various characteristics on a resume—including a candidate's name, prior experience, education, and hobbies. AI algorithms can be trained to consider *only* those characteristics or traits that predict a desired outcome, such as whether a candidate will perform well once on the job.¹⁸

AI can also be instrumental in detecting existing workplace discrimination. Professors Kleinberg, Ludwig, Mullainathan, and Sunstein provide a useful hypothetical on this issue. Consider a firm that is trying to decide which of its sales professionals it will steer toward its most lucrative clients based on two predictive inputs: (1) past sales levels; and (2) manager ratings. Suppose that for men, the firm's managers provide meaningful assessments that provide useful information about employee performance that is not fully captured in the past sales data, but for women, the same managers provide meaningless assessments infused with negative bias towards women, which assigns them lower performance scores. An algorithm that is designed to be *cognizant* of gender would be able to identify the discriminatory manager ratings. If the algorithm is tasked with the function of determining whether manager ratings are predictive of future sales proficiency, it will identify the gender discriminatory manager assessments, as they would *not* be predictive of women's future sales proficiency based on the more objective data inputs of past sales levels.¹⁹

Despite its advantages, AI technology can also perpetuate discrimination depending on the data sets used to train the AI tool. A well-publicized cautionary tale involves Amazon. Amazon began working on an AI tool to screen job applicants in 2014, and in 2018 news broke that Amazon scrapped the tool because it determined that it resulted in certain bias against female applicants.²⁰ As reported, Amazon fed the tool resumes submitted to the company over the course of the prior 10 years as "training data." The tool then recognized patterns among the resumes, constructed

an image for itself of the “ideal candidate,” and then searched an applicant pool and scored applicants on a scale of 1 to 5. Most of the resumes in the training data were those of men, reflecting the disproportionate number of men in the technology sector. The AI tool taught itself that men were preferable candidates because of patterns in the training data. The tool then attributed a lower score to resumes of people who attended “women’s” colleges or who played on the “women’s” chess team. Importantly, Amazon scrapped the tool when it realized the adverse consequences of the algorithm.

Recognizing the risks associated with AI, some companies have collaborated to develop polices to mitigate its potential discriminatory effects. Data & Trust Alliance is a corporate group that has developed “Algorithmic Bias Safeguards for Workforce” with the goal of detecting, mitigating, and monitoring algorithmic bias in workforce decisions.²¹ It has signed up major employers such as American Express, CVS Health, Deloitte, General Motors, Humana, IBM, MasterCard, Meta, Nike, and Walmart.²² According to a recent *New York Times* article reporting on the group, “[c]orporate America is pushing programs for a more diverse work force.”²³

Data and Trust Alliance’s proposed safeguards include 55 questions for companies to ask an AI vendor, education and assessment for evaluating vendor responses, a scorecard to compare vendors, and guidance for integrating the safeguards.²⁴ For example, some questions ask about the use of “proxy” data in AI, including cellphone type (which could be indicative of class or age), sports affiliations, and social club memberships, in otherwise seemingly neutral datasets.²⁵ Other questions ask how bias is minimized during training models, what steps are used to remediate bias, and what practices are used to mitigate bias during deployment.

D. Studies show that AI can help mitigate bias.

As referenced above, there is growing evidence that AI can be used to mitigate unconscious bias. In a forthcoming paper, Bo Cowgill at Columbia Business School studied the performance of a job-screening algorithm in hiring software engineers. A large company trained an algorithm to predict which candidates would pass its interview. Cowgill found that a candidate picked by the machine (and not by a human) is 14% more likely to pass an interview and receive a job offer and 18% more likely to accept a job offer when extended. He found the algorithm also increases hiring of what he calls “non-traditional” candidates—*i.e.*, women, racial minorities, candidates without a job referral, candidates from non-elite colleges,

candidates with no prior work experience, and candidates who did not work for competitors. He concluded that while completely eliminating bias may be extremely difficult, reducing bias is more feasible.²⁶

In an example outside of the employment context, Professors Kleinberg, Lakkaraju, Leskovec, Ludwig, and Mullainathan studied the use of AI in predicting the risk of a criminal defendant's failure to appear in court in the future. The study compared judges' decisions to those made by an algorithm based on three factors: (1) age, (2) current offense, and (3) **prior criminal record**. The professors found that, as compared to the algorithm, judges detain many low-risk people and release many high-risk people. The judges were overweighting the severity of the charge, but the machine learned that a person's prior record matters far more in terms of future risk, which judges were not considering. The professors concluded that using the algorithm's release recommendations would reduce jail population by up to 42% without any increase in crime.²⁷ The authors did not use race as an input in their prediction, but recognized that other variables could correlate with race and result in an algorithm that aggravated racial disparities. However, they found that, in this case, the algorithm could reduce crime and jail populations while simultaneously reducing racial disparities in detention rates.²⁸

A study at the Fisher College of Business analyzed the use of machine learning in selecting board directors by comparing human-selected boards with the predictions of machine learning algorithms.²⁹ The main measure of performance was based on shareholder support that directors receive in annual director re-elections. The study found that in comparison to algorithm-selected directors, management-selected directors were more likely to be male, had larger networks, and had many past and current directorships. By contrast, the machine algorithm found that directors who were not friends of management, had smaller networks, and had different backgrounds than those of management were more likely to be effective directors, including by monitoring management more rigorously and offering potentially more useful opinions about policy, suggesting that directors from diverse backgrounds would be more effective.³⁰

Moreover, AI vendors are self-reporting about their products' ability to improve diversity. Pymetrics, an AI vendor, has conducted several case studies on the effectiveness of its product. In one study, Pymetrics worked with a top food production company that was looking to more effectively review 6,000 job applications for 40 job openings. The recruiting team had been using indicators

such as GPA, past experience, and/or extracurricular activities to screen applications. Pymetrics selected a group of top-performers at the company and built a predicative model based on game play. Pymetrics then developed a candidate evaluation process whereby candidates completed the same set of Pymetrics core exercises, a numerical and logical reasoning assessment, and a digital interview process with standardized questions. The company was able to review 16% of the applications that were pre-screened by Pymetrics (as opposed to 40% when it was screening manually using CVs). Importantly, for the first time, the gender split of the finalists was 50:50.³¹

In another case study, Pymetrics evaluated the use of AI in hiring at a large investment firm. When the firm first started working with Pymetrics, it was receiving 20,000 applications for its job vacancies. Pymetrics again developed an assessment based on data collected from gameplay of current employees. The Pymetrics-recommended candidates were 97% more likely to ultimately receive a job offer. The Pymetrics evaluation process also expanded the diversity of universities represented (from 9 universities to 66 different schools), and increased female representation among recommended candidates by 44% and minority representation by 9%.³²

Of course, AI is not perfect, and it will likely be unable to completely eradicate bias and discrimination in the hiring and employment context. But the studies evaluating AI so far are promising, and suggest that AI can be developed to improve diversity in the workplace. My experience in advising corporate clients has been that companies are making good faith efforts to use AI responsibly, thus contributing to the development of more equitable and efficient uses of AI.

III. The use of AI brushes up against a number of concerns in the employment context.

Most employers are aware that federal antidiscrimination laws prohibit them from making employment decisions based on race, color, religion, sex, national origin, age, disability, or genetic information (*i.e.*, a protected class). And although outside of the EEOC's purview, employers are also forbidden from making employment decisions based on military veteran and union membership status.

The use of AI and automation tools implicate federal antidiscrimination laws in a variety of ways.

First, employers are not permitted to disfavor individuals based on their membership in a protected class. Aggrieved individuals assert these claims under either a disparate treatment or disparate impact theory.

A disparate treatment claim may arise as the result of an employer's use of a tool if the employee can show that an employer intentionally programmed or used a tool to disadvantage individuals in a protected class. A tool programmed, for instance, to filter out candidates above a certain age would fall into this category. This issue was at the heart of the EEOC's complaint filed last year against iTutorGroup, Inc. 33—a provider of online English language tutoring services to students in China—where it was alleged that iTutorGroup programmed its application software to automatically reject female applicants over the age of 55 and male applicants over the age of 60. Although the technology allegedly used by iTutorGroup was not technically artificial intelligence but rather a form of automated screening, the allegations in the complaint illustrate the perils of using impermissible data inputs to develop a hiring algorithm that would clearly discriminate against members of a protected class.

Second, a disparate impact claim arising out of an employer's use of an AI tool may be a high-tech version of neutral factors that have adverse consequences, but it is certainly not novel. In a disparate impact claim, an employee would show that an employer's practice—such as the use of an AI tool to screen applicants—results in members of a protected class being disfavored at a higher rate than members of another class. This would be the case even where an employer had no intention to discriminate. To avoid disparate impact liability, an employer using such a policy or practice must show that the practice is “job related for the position in question and consistent with business necessity” and that no alternative requirement would suffice.³⁴ A company that uses an AI tool that, for instance, disproportionately disfavors women may be subject to liability on a disparate impact theory.

The ADA poses a host of other issues, as the EEOC pointed out in its May 2022 guidance. The major issue here is that employers must, under the ADA, provide employees and job seekers with reasonable accommodations that allow them to perform the essential functions of their positions so long as the accommodations are not an undue hardship on employers. Individuals with disabilities, for instance, might have limited dexterity that results in programs assigning them lower scores on computer “games” that some employers use to build personality profiles of candidates. If an employer can provide an alternative testing format, and doing so would not constitute an undue hardship, the employer must offer this

accommodation. Employers may also be required to offer more time to complete such assessments or might be required to provide assessments compatible with accessible technology such as screen readers.

Given these potential legal concerns, the EEOC has an interest in keeping companies on the right side of the line. And it should go without saying that most employers want to avoid discrimination, and its associated liability, as well. It should be noted that although Amazon's tool disproportionately disfavored certain protected classes, Amazon properly stress-tested the tool and scrapped it because it realized it was producing biased results. This illustrates the important point that errors in designing an AI program used to assist with workplace decision-making can be easily identified and corrected, as compared to human decision-making, where rooting out conscious or unconscious bias can be far more challenging.

IV. So far, regulations are a mixed bag from the employer perspective.

Although some recent proposed and actual state and local regulations have been a step in the right direction, there is still ambiguity that complicates compliance for employers. I am hopeful that states and municipalities can act as laboratories for regulating AI, but it is premature to draw any conclusions from these regulations with respect to their impact in actually reducing workplace bias.

Two states, Maryland and Illinois, have enacted statutes regulating the use of AI. Illinois' Artificial Intelligence Video Interview Act requires employers using AI analysis of applicant-submitted video interviews to (1) notify applicants that such AI will be used to analyze an applicant's video interview, (2) provide applicants with "information before the interview explaining how the artificial intelligence works and what general types of characteristics it uses to evaluate applicants," and (3) obtain the applicant's consent to be evaluated by the AI program.³⁵ The Illinois law was amended on January 1, 2022 to require employers who "[rely] solely upon an artificial intelligence analysis of a video interview to determine whether an applicant will be selected for an in-person interview" to report annually to the Illinois Department of Commerce and Economic Opportunity ("DCEO") the race and ethnicity of applicants who are not offered in-person interviews after the use of AI analysis and the race and ethnicity of applicants who are actually hired.³⁶ The DCEO will then analyze the reported data and create a report discussing whether the data discloses racial bias in the use of AI. The first report has not yet been

issued, but is due by July 1, 2023. Maryland's law requires only that applicants consent to the use of facial recognition technology during an interview.³⁷

On the plus side, these laws notify applicants that an AI tool will be used. This gives applicants the opportunity to conduct research to better understand how AI functions in the interview process and request ADA accommodations, if needed, and hopefully this transparency will foster trust in these tools. Also, since applicants know the tools will be used, applicants can challenge the use of these tools as discriminatory. This incentivizes employers to ensure the tools do not disproportionately disadvantage members of protected classes. On the other hand, these laws are somewhat vague and provide employers minimal guidance. By way of example, the Illinois statute fails to even define "AI."

California's Fair Employment and Housing Council has proposed regulations with regard to the use of "Automated Decision Systems."³⁸ These proposed regulations would apply to any computational process that "screens, evaluates, categorizes, recommends, or otherwise makes a decision or facilitates human decision making that impacts employees or applicants."³⁹ At bottom, the proposed regulations clarify that California's antidiscrimination laws apply equally to decisions made by AI tools. For instance, "[t]he use of and reliance upon automated-decision systems that limit or screen out, or tend to limit or screen out, applicants based on protected characteristic(s)...may constitute unlawful disparate treatment or disparate impact."⁴⁰ This clarification that antidiscrimination laws apply equally to employment decisions made in reliance on AI tools should motivate employers to evaluate tools for any discriminatory impact.

Significantly, California's proposed regulation also takes a page from the European Union's General Data Protection Regulation and would impose liability on *vendors* of AI tools, as the proposed regulation applies to "agents" of employers, defined to include those that "provide[] services related to recruiting, applicant screening...or the administration of automated-decision systems for an employer's use in recruitment, hiring, performance evaluation, or other assessments..."⁴¹ This type of regulation imposing liability on AI tool vendors would likely be outside of the EEOC's scope, but serious consideration should be given to ensuring accountability of the software developers in this space.

The New York City law places regulations on AI tools that are somewhat burdensome and may dissuade employers from using them. Regulators must be careful to strike the right balance between encouraging the correct use of AI tools—

reaping the previously discussed benefits of eliminating unconscious bias and fostering diversity in the workplace—and incentivizing employers to ensure the tools do not disadvantage members of protected classes. The New York City law, set to take effect April 15, 2023, covers “Automated Employment Decision Tools” (“AEDTs”), defined as those tools that use a computational process that issues a “simplified output, including a score, classification, or recommendation, that is used to substantially assist or replace discretionary decision making” in employment decisions.⁴² It requires that an independent party conduct bias audits of these tools, which is a major positive of the New York City law and could help employers ensure that their tools do not unintentionally discriminate. I’ll touch on the potential limits of such independent auditing in a bit.

However, the law also requires summaries of the bias audits to be published, which may discourage employers from adopting these tools. Employers already have plenty of incentive to ensure their tools do not screen out protected classes because they want to avoid liability under federal and state antidiscrimination laws. Under the New York City law, employers must also inform applicants that an AEDT tool will be used to evaluate them “no less than 10 business days before” the tool is used.⁴³ This requirement is burdensome, and perhaps prohibitive, for employers who want to have an expeditious hiring process. If employers have to wait 10 days before using AI tools, employers may lose candidates as they find employment elsewhere and will be unable to fill open roles quickly. The danger with these sorts of laws is that companies may decide to scrap AI tools altogether despite their potential to reduce unconscious bias and subjectivity in the hiring process.

V. Some recommendations for approaching AI in the employment context.

It is unlikely that there is going to be a one-size-fits-all approach to using AI effectively and responsibly. Guidelines will need to be tailored to different sectors. For example, the types of considerations relevant for AI tracking the productivity of truck drivers will be different from those relevant for AI tracking the performance of sales representatives. Of course, sectors will certainly be able to learn from each other.

Regardless of the industry, there are some key guideposts that can help companies use AI responsibly and help mitigate the risk of violating antidiscrimination laws. First, transparency—companies should be upfront about the use of AI, as required

by some of the state and city laws we've discussed today. At this time, there is no federal requirement for employers to disclose the AI technology they use.

Nevertheless, applicants and employees should know when they are being evaluated by a machine algorithm as opposed to a human reviewer. Companies should not need to provide excruciating detail of how they are using AI, but general notice will give applicants the opportunity to request more information and help identify instances of potential discrimination.

The second guidepost is auditing—whether it is self-auditing or third-party auditing, it is important that companies are proactive in mitigating potential biases of AI. As mentioned above, New York City's AI law will require independent parties to conduct bias audits of AI tools, and will require employers to post summaries of the bias audit findings on their websites. Although well-intentioned, independent auditing may be difficult to implement effectively in practice. Recently, Pymetrics paid a team of computer scientists from Northeastern University to audit its AI hiring algorithm.⁴⁴ The audit evaluated the “fairness” of Pymetrics' algorithm under the EEOC's four-fifths rule stating that hiring procedures should select roughly the same proportion of men and women, and people from different racial groups. That is, if 100% of men are passing a test, at least 80% of women must pass it.

Northeastern University's audit showed that Pymetric's algorithm satisfies the four-fifths rule, but it did not show that the tool was bias-free or that it chose the most qualified candidates. The audit compared men versus women, and one racial group against another, but did not address disparities between people who belong to more than one protected class. The tool could not determine if the algorithm was biased against Asian men or Black women, for example. Moreover, the audit was funded by Pymetrics, which creates a risk of the auditor being influenced by the client. As independent auditing companies pop up in response to the New York City law or client demand, companies should be cautious not to take their assessments at face value. They should look at which metrics independent auditing companies are using to evaluate AI technology, and consider how the auditing companies are compensated.

To date, there is a lack of consensus of which metrics and data auditors should use to audit AI technology. There are no clear standards for which biases to test for in AI, and it could be difficult to define which data points are the most useful in detecting bias. IBM has suggested that it become standard practice for auditing companies to disclose the assumptions used for determining relevant protected characteristics

used in an audit bias.⁴⁵ As companies conduct audits to comply with NYC's recent AI legislation (and potentially future legislation), more standards may develop as to what constitutes a valid and effective bias audit of AI technology.

Relatedly, vendor vetting can also help companies decrease the likelihood that the AI tools they are using do not perpetuate bias. First, companies can ask vendors questions such as those proposed by Data and Trust Alliance: (1) what measures are taken to detect and mitigate bias; (2) what approaches are used to remediate bias; (3) what measures have been taken to demonstrate that the system performs as intended, and as claimed; and (4) what are the vendor's commitments to ethical practices. Companies can also ask the vendors about the types of data used to train models: (1) how is the data collected, (2) where does that data come from, (3) how often is the data updated, and (4) how often is the data audited. In short, companies should have a sense of how vendors are developing AI algorithms and what steps they are taking to regularly mitigate potential biases.

Finally, companies should develop their own internal policies to regulate and mitigate biases in AI technology. Just as companies have had to recently assess and develop social media policies, they will have to work with consultants and counsel to draft and implement best practices. Some major companies have already created such guidelines for mitigating bias in the use of AI technology. For example, Google has a set of guidelines to ensure that machine learning is "fair." It encourages developers to (1) design models using concrete goals for fairness and inclusion (*i.e.*, making tools accessible in different languages or for different age groups); (2) use representative datasets to train and test models; (3) check the system for unfair biases, including by organizing a pool of diverse testers to identify who may experience adverse impacts; and (4) analyze the performance of the machine learning model.

Similarly, IBM has developed five "pillars of trustworthy AI": Explainability, Fairness, Robustness, Transparency, and Privacy.⁴⁶ It encourages companies and developers to (1) take accountability for the outcomes of their AI systems in the real world; (2) be sensitive to a wide range of cultural norms and values;

(3) work to address biases and promote inclusive representation; (4) ensure humans can understand an AI decision process; and (5) preserve and fortify users' power over their own data.

Companies can learn from each other and develop standardized industry regulations for using AI responsibly in hiring practices. As technology and auditing systems develop, we will get a sense of what works, and what doesn't. Self-regulation may also emerge at the corporate board level, as boards become aware of how AI can be used effectively and in a non-discriminatory manner.

VI. Government's Role

In terms of regulation, I think it is important to keep in mind a few over-arching concepts. First, employers do not want to use AI tools that discriminate. Second, these tools are new, and not going away, so regulators should leave room for experimentation and take into account employers' efforts to get the use of these tools right. Third, these tools may be better than the alternative, which is human subjectivity. In contrast to the human mind, AI tools can at least be audited. Unconscious bias in humans is extremely difficult to audit.

The Commission could and should put employers on notice that AI tools are subject to all the same rules and regulations as other processes and procedures used to make employment decisions. The Commission need not look far to provide guidance on how to audit employment processes and procedures for disparate impact. The Commission's 1978 Uniform Guidelines on Employee Selection Procedures apply equally to intelligence/aptitude tests—the original problem disparate impact theories of discrimination were meant to solve—and AI tools.⁴⁷ If an AI tool has a disparate impact on members of a protected class, employers must show that the selection procedure “is predictive of or significantly correlated with important elements of job performance.”⁴⁸ The Commission also uses statistical analysis in other areas such as in determining whether differences in pay between a protected class and a comparator group are statistically significant.⁴⁹

An approach to regulating this space that allows employers leeway to regularly audit, revise, and (if needed) scrap AI tools that result in discriminatory outcomes would allow for needed experimentation with AI. The concept of allowing employers that audit themselves and self-correct is not new in the law. For instance, in Massachusetts, employers have a defense to claims under the Massachusetts Equal Pay Act if, in the past three years, the employer “has both completed a self-evaluation of its pay practices in good faith and can demonstrate that reasonable progress has been made towards eliminating wage differentials based on gender for comparable work, if any, in accordance with that evaluation.”⁵⁰ Such measures

encouraging employer self-correction recognize that it is not possible for employers to make completely unbiased decisions all the time, and has the salutary effect of encouraging regular auditing and self-correcting. In suggesting this, I am in no way suggesting that an audit should shield *intentional* discrimination.

At bottom, due to the emergence of powerful social movements over the past half-decade and the expansive federal, state and local regulation of issues that intersect with Title VII, the ADEA, the ADA and the Equal Pay Act, employers' awareness of and focus on matters of importance to legally protected groups is at an all-time high, and as they embark on deploying artificial intelligence tools in the workplace, they do so with these principles top of mind. This is an inflection point at which the EEOC can partner with the business community to issue guidance that will allow employers to continue to increase their focus on diversity, equity and inclusion while giving them room to use these tools to enhance workplace culture and performance.

I want to thank the Commission for giving me this opportunity to share my perspectives, and I look forward to working with all of you on this important initiative.

1 “Automation” and “AI” are two concepts that should not be conflated. As a general matter, “automation is a broad category describing an entire class of technologies,” which has been around for many decades, and is designed primarily to relieve humans of repetitive, monotonous and/or mundane tasks. See Michael Gaynor, *Automation and AI Sounds Similar, But May Have Vastly Different Impacts on the Future of Work*, BROOKINGS (Jan. 29, 2020), <https://www.brookings.edu/blog/the-avenue/2020/01/29/automation-and-artificial-intelligence-sound-similar-but-may-have-vastly-different-impacts-on-the-future-of-work/>. “AI,” however, is designed to simulate human thinking. It “refers to how computer systems can use huge amounts of data to imitate human intelligence and reasoning, allowing the system to learn, predict and recommend what to do next.” Or Shani, *AI Automation: What You Need To Know*, MARKETING ARTIFICIAL INTELLIGENCE INSTITUTE (Nov. 9, 2021),

<https://www.marketingaiinstitute.com/blog/automation-and-ai-what-you-need-to-know>.

(<http://www.marketingaiinstitute.com/blog/automation-and-ai-what-you-need-to-know>)

2 SOCIETY FOR HUMAN RESOURCE MANAGEMENT, *Automation & AI in HR* (Apr. 2022), at 3, https://advocacy.shrm.org/SHRM-2022-Automation-AI-Research.pdf?_ga=2.112869508.1029738808.1666019592-61357574.1655121608.

3 ZION MARKET RESEARCH, *Global Artificial Intelligence (AI) Market to Register an Annual Growth of 39.4% During Forecast Period*

(June 22, 2022), <https://www.zionmarketresearch.com/news/global-artificial-intelligence-market>. (**(<http://www.zionmarketresearch.com/news/global-artificial-intelligence-market>)**)

4 Keith E. Sonderling, *Op-ed: Artificial Intelligence Is Changing How HR Is Handled at Companies. But Do Robots Care About Your Civil Rights?*, CHI. TRIB. (Sept. 20, 2021), [https://www.c \(http://www.chicagotribune.com/opinion/commentary/ct-opinion-robots-ai-civil-rights-\)](https://www.chicagotribune.com/opinion/commentary/ct-opinion-robots-ai-civil-rights-) <http://www.chicagotribune.com/opinion/commentary/ct-opinion-robots-ai-civil-rights-> (**(<http://www.chicagotribune.com/opinion/commentary/ct-opinion-robots-ai-civil-rights->)**)_amazon-20210920-tef7m7az3rgjtacauazvw3u224-story.html.

5 SOCIETY FOR HUMAN RESOURCE MANAGEMENT, *Automation & AI in HR* (Apr. 2022), at 4, https://advocacy.shrm.org/SHRM-2022-Automation-AI-Research.pdf?_ga=2.112869508.1029738808.1666019592-61357574.1655121608.

6 See MODERN HIRE, <https://modernhire.com/> (last visited January 12, 2023).

7 Earl Fitzhugh, JP Julien, Nick Noel & Shelley Stewart, *It's Time For A New Approach to Racial Equity*, MCKINSEY & COMPANY, at 2 (May 25, 2021).

8 Dale Buss, *12 Ways Companies Are Boosting Their DEI*, SOCIETY FOR HUMAN RESOURCE MANAGEMENT (Mar. 9, 2022), [https://www.s \(<http://www.shrm.org/resourcesandtools/hr-topics/behavioral-competencies/global-and-cultural-effectiveness/pages/12-ways->\)](https://www.shrm.org/resourcesandtools/hr-topics/behavioral-competencies/global-and-cultural-effectiveness/pages/12-ways-) [hrm.org/resourcesandtools/hr-topics/behavioral-competencies/global-and-cultural-effectiveness/pages/12-ways-](http://www.shrm.org/resourcesandtools/hr-topics/behavioral-competencies/global-and-cultural-effectiveness/pages/12-ways-) (**([http://www.shrm.org/resourcesandtools/hr-topics/behavioral-](http://www.shrm.org/resourcesandtools/hr-topics/behavioral-competencies/global-and-cultural-effectiveness/pages/12-ways-)**)

competencies/global-and-cultural-effectiveness/pages/12-ways-)_companies-are-boosting-their-dei.aspx.

9 Dr. Brian Marentette, *Pay Equity Audits: On the Rise and Becoming Essential*, HR DAILY ADVISOR (Jan. 13, 2023), <https://hrdailyadvisor.blr.com/2022/11/17/pay-equity-audits-on-the-rise/#:~:text=More%20companies%20are%20adding%20at,level%20businesses%20taking%20the%20lead.>

10 Hilke Schellmann, *Auditors Are Testing Hiring Algorithms for Bias but There's No Easy Fix*, MIT TECH. REV. (Feb. 11, 2021), <https://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/>.
(<http://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/>)
<http://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/>
(<http://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/>)

11 See Kimberly A. Houser, *Can AI Solve the Diversity Problem in the Tech Industry? Mitigating Noise and Bias in Employment Decision-Making*, 22 STAN. TECH. L. REV. 290, 294 (2019).

12 See *Id.* at 324.

13 Keith E. Sonderling, Bradford J. Kelley & Lance Casimir, *The Promise and The Peril: Artificial Intelligence and Employment Discrimination*, 77 U. MIA. L. REV. 1, 33 (2022), <https://repository.law.miami.edu/umlr/vol77/iss1/3>.

14 *Id.*

15 *Id.* at 15.

16 See Kimberly A. Houser, *Can AI Solve the Diversity Problem in the Tech Industry? Mitigating Noise and Bias in Employment Decision-Making*, 22 STAN. TECH. L. REV. 290, 326 (2019). (citing Eric Luis Uhlmann & Geoffrey L. Cohen, *Constructed Criteria Redefining Merit to Justify Discrimination*, 16 PSYCH-OL. SCI. 474, 474 (2005)).

17 Jake Silberg & James Manyika, *Notes From the AI Frontier: Tackling Bias in AI (And In Humans)*, MCKINSEY GLOBAL INSTITUTE, at 2 (June 2019),

<https://www.mckinsey.com/~media/mckinsey/featured%20insights/artificial%20intelligence/tackling%20bias%20in%20artificial%20intelligence%20and%20in%20humans/mgi-tackling-bias-in-ai-june-2019.pdf>
(<https://www.mckinsey.com/~media/mckinsey/featured%20insights/artificial%20intelligence/tackling%20bias%20in%20artificial%20intelligence%20and%20in%20humans/mgi-tackling-bias-in-ai-june-2019.pdf>)
(<https://www.mckinsey.com/~media/mckinsey/featured%20insights/artificial%20intelligence/tackling%20bias%20in%20artificial%20intelligence%20and%20in%20humans/mgi-tackling-bias-in-ai-june-2019.pdf>)

18 Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan & Cass R. Sunstein, *Discrimination in the Age of Algorithms*, at 33 (Feb. 5, 2019), <https://ssrn.com/abstract=3329669>.

19 *Id.* at 33-34.

20 See Jeffrey Dastin, *Amazon Scraps Secret AI Recruiting Tool that Showed Bias against Women*, REUTERS (Oct. 10, 2018), <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>.
(<http://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>)

21 DATA AND TRUST ALLIANCE, *Algorithmic Safety: Mitigating Bias in Workforce Decisions*, <https://dataandtrustalliance.org/our-initiatives/algorithmic-safety-mitigating-bias-in-workforce-decisions> (last visited Jan. 13, 2023).

22 DATA AND TRUST ALLIANCE, <https://dataandtrustalliance.org/> (last visited Jan. 13, 2023).

23 Steve Lohr, *Group Backed by Top Companies Moves to Combat A.I. Bias in Hiring*, N.Y. TIMES (Dec. 8, 2021), <https://www.nytimes.com/2021/12/08/technology/data-trust-alliance-ai-hiring->
(<http://www.nytimes.com/2021/12/08/technology/data-trust-alliance-ai-hiring->
(<http://www.nytimes.com/2021/12/08/technology/data-trust-alliance-ai-hiring->

hiring-)

bias.html#:~:text=The%20Data%20%26%20Trust%20Alliance%2C%20announced,o
rganization%20or%20a%20think%20tank.

24 DATA AND TRUST ALLIANCE, *Algorithmic Safety: Mitigating Bias in Workforce Decisions*, <https://dataandtrustalliance.org/our-initiatives/algorithmic-safety-mitigating-bias-in-workforce-decisions> (last visited Jan. 13, 2023).

25 Steve Lohr, *Group Backed by Top Companies Moves to Combat A.I. Bias in Hiring*, N.Y. TIMES (Dec. 8, 2021), <https://www.n>

**(<http://www.nytimes.com/2021/12/08/technology/data-trust-alliance-ai-hiring->
[ytimes.com/2021/12/08/technology/data-trust-alliance-ai-hiring-](http://www.nytimes.com/2021/12/08/technology/data-trust-alliance-ai-hiring-)
<http://www.nytimes.com/2021/12/08/technology/data-trust-alliance-ai-hiring->
[hiring-](http://www.nytimes.com/2021/12/08/technology/data-trust-alliance-ai-hiring-))**

bias.html#:~:text=The%20Data%20%26%20Trust%20Alliance%2C%20announced,o
rganization%20or%20a%20think%20tank.

26 Bo Cowgill, *Bias and Productivity in Humans and Machines*, COLUM. UNIV. (Mar. 21, 2020),

https://conference.iza.org/conference_files/MacroEcon_2017/cowgill_b8981.pdf.

27 Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig & Sendhil Mullainathan, *Human Decisions and Machine Predictions*, Q. J. ECON. 237 (2018).

28 Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig & Sendhil Mullainathan, *Human Decisions and Machine Predictions*, Q. J. ECON. 237 (2018); Kleinberg, Jon and Ludwig, Jens and Mullainathan, Sendhil and Sunstein, Cass R., *Discrimination in the Age of Algorithms* (Feb. 5, 2019), available at <https://ssrn.com/abstract=3329669>.

29 Isil Erel, Lea H. Stern, Chenhao Tan & Michael S. Weisbach, *Selecting Directors Using Machine Learning*, Fisher College of Bus. Working Paper No. 2018-03-005, Charles A. Dice Ctr. Working Paper No. 2018-05, Eur. Corp. Governance Inst. (ECGI) - Finance Working Paper No. 605/2019 at 23 (Dec. 13, 2020), <https://ssrn.com/abstract=3144080>.

30 *Id.*

31 PYMETRICS, *Leading Food Production Company*, <https://www.pymetrics.ai/case-studies/highest-potential-talent-identified> (last visited Jan. 16, 2023).

32 PYMETRICS, *Leading Global Investment Firm*, <https://www.pymetrics.ai/case-studies/leading-global-investment-firm> (last visited Jan. 17, 2023).

33 See Amended Complaint, *EEOC v. iTutorGroup, Inc.*, No. 1:22-cv-02565-PKC-PK (E.D.N.Y. Aug. 3, 2022).

34 See 42 U.S.C. § 2000e-2(k)(1)(A).

35 Artificial Intelligence Video Interview Act, 820 ILL. COMP. STAT. ANN. 42/5 (West 2022).

36 *Id.* at § 42/20.

37 MD. CODE ANN. LAB. & EMPL. § 3-717 (West 2020).

38 Fair Employment & Housing Council Draft Modification to Employment Regulations Regarding Automated-Decision Systems, CAL. CODE REGS. tit2, § 11008 *et seq.* (proposed Mar. 15, 2022), <https://calcivilrights.ca.gov/wp-content/uploads/sites/32/2022/03/AttachB-ModtoEmployRegAutomated-DecisionSystems.pdf>.

39 *Id.* at § 11008(e).

40 *Id.* at § 11016(c)(5).

41 *Id.* at § 11008(a).

42 N.Y.C. ADMIN. CODE § 20-870.

43 *Id.* at § 20-871(b).

44 Hilke Schellmann, *Auditors Are Testing Hiring Algorithms for Bias but There's No Easy Fix*, MIT TECH. REV. (Feb. 11, 2021), <https://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/>.
[\(https://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/\)](https://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/)
[echnologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/](https://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/).
[\(https://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/\)](https://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/)

45 *Standards for Protecting At-Risk Groups in AI Bias Auditing*, IBM (Nov. 2022), <https://www.ibm.com/downloads/cas/DV4YNKZL>.
[bm.com/downloads/cas/DV4YNKZL](https://www.ibm.com/downloads/cas/DV4YNKZL).
[\(https://www.ibm.com/downloads/cas/DV4YNKZL\)](https://www.ibm.com/downloads/cas/DV4YNKZL)

46 *Everyday Ethics for Artificial Intelligence*, IBM, <https://www.ibm.com/watson/assets/duo/pdf/everydayethics.pdf>.
[\(https://www.ibm.com/watson/assets/duo/pdf/everydayethics.pdf\)](https://www.ibm.com/watson/assets/duo/pdf/everydayethics.pdf)
[bm.com/watson/assets/duo/pdf/everydayethics.pdf](https://www.ibm.com/watson/assets/duo/pdf/everydayethics.pdf)
[\(https://www.ibm.com/watson/assets/duo/pdf/everydayethics.pdf\)](https://www.ibm.com/watson/assets/duo/pdf/everydayethics.pdf) (last visited Jan. 17, 2023).

47 See 29 C.F.R. § 1607 (2022).

48 *Id.* at § 1607.5(B) (2022).

49 See U.S. EQUAL EMP. OPPORTUNITY COMM'N, EEOC-915.003, EEOC COMPL. MAN. § 10-III.

50 MASS. GEN. LAWS ch. 149, § 105A(d) (2018).



U.S. Equal Employment Opportunity Commission

Testimony of Adam T. Klein

Introduction

Good morning, Chair Burrows and members of the Commission. Thank you for this opportunity to share my insights with you about the use of AI and automated systems relating to sourcing, recruiting, and applicant selection, and ways that the EEOC can provide additional guidance in this area to protect against technology-sponsored discrimination. My name is Adam Klein and I serve as Managing Partner at Outten & Golden LLP. In that capacity, I represent classes of workers in civil rights and workplace equity litigation, including cases focusing on discriminatory hiring selection procedures relating to **criminal background records** and social media/online job advertising.

As a starting point, there are numerous types of AI and automated systems used by employers situated on a spectrum of complexity and utility – from simple applicant tracking to complex gamified psychometrics and unsupervised machine learning deployed to source and recruit applicants on social media platforms. In this context, AI/algorithms are specifically designed to process large data sets and efficiently differentiate applicants with limited or no human participation. For the remainder of this discussion, I will focus on the more complex use case for AI/algorithms in the workplace.

The advantages of using AI/algorithms by employers are clear and obvious: a computer algorithm can easily and cheaply source, recruit, and select applicants for employment. The disadvantages are equally obvious: there is a fundamental and profound lack of a theoretical or practical nexus between the key competencies or requirements of a target position – using a job analysis and competency model – with the actual selection criteria used by AI systems.

Moreover, there is a serious concern that these AI sourcing/hiring selection systems will essentially automate ingrained biases that tend to perpetuate disturbing and longstanding patterns of hiring discrimination based on protected characteristics. I urge the EEOC to take additional proactive measures to address these emerging trends in the workplace.

Recommendations

First, as noted, there is a disturbing lack of scientific evidence supporting claims that machine-learning, automated hiring processes provide any practical utility other than user convenience. Predictive algorithms claim to identify the “best” or preferred candidates but may instead perpetuate biased representation rates and identified traits and interests of “favored” incumbent employees that are not job relevant. The EEOC should issue guidance requiring employers to document the use of these emerging technologies and provide a sound scientific basis for its use for sourcing, recruiting, and hiring selection.

Second, employers must have the ability, and be incentivized, to audit data from AI systems and isolate each discrete selection step so they can monitor for adverse impact. This is important because the algorithmic “tests” used in selection are constantly changing (or “learning”), and typically proceed with no underlying conceptual framework. Consequently, evidence of adverse impact is extremely problematic and should be eliminated to the extent possible. Moreover, many of the AI systems are maintained by outside vendors with no real accountability.

Third, applicants exposed to AI hiring selection systems should be informed of its use and be provided with disclosures sufficient to understand whether a potential violation of federal anti-discrimination statutes may have occurred.

Fourth, the federal government has a unique role to play to address the use of these emerging technologies. I recommend a coordinated government response – including drawing resources from federal agencies with particular subject-matter expertise in the use of AI and machine learning systems.

Growing Technologies in Pre-Employment

Legal scholars and practitioners have written at length about the discriminatory impact on workers resulting from the increase in employers’ use of AI or

algorithmic/automated decision-making in pre-employment recruitment, selection/screening, and assessment practices.¹

First, **Online Recruitment/Job Advertising**. Online advertising is big business, reaching ever-widening audiences. Social media platforms mine user data, and their algorithms employ that information to target particular audiences for job advertisements. Social media platforms, such as Meta (Facebook), previously required job advertisers to select filters—location (as a proxy for race), age, and gender—to target their ads until this practice was challenged in several class action lawsuits brought by my law firm, co-counsel, and nonprofit organizations. In a landmark settlement, Facebook agreed to settle these legal challenges and discontinued the underlying practices. In addition, employers who published the discriminatory job ads settled as well with agreements to discontinue job ad targeting based on protected characteristics.²

Notwithstanding such changes, researchers have commented that Facebook's AI in job advertisements may still target audiences in ways consistent with gender-, race-, or age-based stereotypes (for example, male users disproportionately receiving jobs for lumberjacks or truck drivers). In response to newly filed legal challenges and DOJ investigations, earlier this month, Facebook announced its plan to create a "Variance Reduction System" to advance equitable distribution of ads, including employment ads, with the goal of reducing the variances between the eligible and actual audiences along perceived sex and race/ethnicity identifiers in the delivery of ads. Time will tell if this newest intervention is successful in diminishing unlawful targeting.³

It appears that Facebook is just one of several job advertising platforms that utilize AI in delivering ads to prospective applicants—others appear to include LinkedIn, ZipRecruiter, CareerBuilder, and Monster. While it remains largely opaque how the AI operates, journalists report that these sites utilize AI "matching engines," which are optimized to produce applications based on categories of user-provided information; data assigned to the user based on other users' skill sets, experiences, and interests; and behavioral data based on how a user interacts or responds to job postings. Another example is the rise of TikTok and its foray into the job recruitment space through the pilot "TikTok Resumes" program, which invites applicants to submit a TikTok video resume for employer review. The clear concern is that this type of video resume technology may perpetuate age, appearance-based

(implicating gender discrimination), and race-based discriminatory hiring practices.⁴

Second, **Applicant Screening Tools**. Employers also use AI-driven automated screening tools to sort, rank, and select applicants for employer review. Job applicants often interact with automated hiring platforms to submit their employment application, including providing personal information, agreeing to **background checks**, and completing personality/skills assessments.

These automated hiring platforms then utilize algorithms to sort the applications; they are ubiquitous in certain industries, including retail. One example of bias is when AI-based hiring programs screen out applicants with gaps in employment history, disparately impacting groups who have taken time off for caregiving responsibilities. Holistically, we should be asking if the predictive algorithms are designed to select the perceived “best” candidates or qualified candidates based on actual job-related criteria. It should be the latter: because if the applicant is qualified based on the ability to perform the job, the applicant should have the opportunity to compete and not get washed out early in the process.

Further, third-party vendors harvest online information creating datasets of attributes and behaviors, then create automated decision-making programs that analyze the datasets to find statistical relationship between variables. What these vendors are doing is predicting who is a good match for an employer by identifying patterns through inferring characteristics from a dataset of information about candidates. These algorithms make predictions based upon statistical correlations or observed patterns but are not based on causal factors and further lack any connection with job performance, making them prone to error and biases. Such screening and selection procedures should be validated at a minimum by SIOP Principles, which note that “[v]ariables chosen as predictors [for employment] should have theoretical, logical, or empirical foundation.”⁵ When selection procedures are challenged as having a disparate impact, employers bear the burden of demonstrating the selection procedure is job-related and consistent with we **announced (<https://about.fb.com/news/2022/06/expanding-our-work-on-ads-fairness/>)** our plan to create the Variance Reduction System (VRS) to help advance the equitable distribution of ads on Meta technologies. After more than a year of collaboration with the DOJ, we have now launched the VRS in the United States for housing ads. Over the coming year, we will extend its use to US **employment (<https://www.facebook.com/business/help/153775906681893?>**

ref=search_new_0) and credit ads.

([https://www.facebook.com/business/help/1157846251802527?](https://www.facebook.com/business/help/1157846251802527?ref%3Dsearch_new_2&sa=D&source=docs&ust=1671573500422142&usg=AOvVaw3TvyY_m1isUT4IAJa3jK56)

ref%3Dsearch_new_2&sa=D&source=docs&ust=1671573500422142&usg=AOvVaw3TvyY_m1isUT4IAJa3jK56).

Additionally, we discontinued **the use of Special Ad Audiences,** **(<https://www.facebook.com/business/help/2408667629202904>)** an additional commitment in the settlement.”), last visited on January 28, 2023. business necessity.⁶ In this way, employer bear the burden of demonstrating that the model is statistically valid and substantively meaningful, as opposed to merely job related.

Lastly, **Psychometric Assessments.** Newer psychometric assessments include personality tests, video interviews, and gamified assessments. Of course, psychometric testing has been around for decades but is currently making a comeback with the help of AI. In the 1950s, psychometric tests began to be used in the workplace by companies outside of the armed services. In the 1960s and 1970s, IO psychologists began reintroducing personality tests based on new behavioral and social science research and techniques.

Vendors selling these services promote the idea that certain tests can accurately assess candidates for certain job competencies, values, and intelligence. But personality tests have a long history of legal challenges including privacy concerns, accessibility and disability discrimination, as well as disparate impact concerns. The shortcomings of facial recognition programming are now well documented. Before discontinuing the practice in 2021, it was publicly reported that HireVue’s virtual interview program would sort and grade video job applicants and uses AI algorithms to evaluate their performance, analyze the interview, and predict their performance based on the interview.⁷

“Gamification” in psychometric assessment goes beyond personality questions by “add[ing] features such as rules; competition; scores; medals, badges, or trinkets won; levels of progress; and comparisons of performance against other ‘players,’ typically in work-related scenarios.”⁸ For example, it is reported that one widely-used vendor provides “an online technology platform that enables hiring managers to hold blind audition challenges,” in which “job applicants are given mini assignments that are designed to assess the applicant for the specific skills required for the open position.”⁹

To address this new wave of assessment, the Society of Industrial Organization Psychologists (SIOP) published a white paper that recognized that gamification testing for hiring has not developed enough to be scientifically studied and needs “further empirical testing in accuracy of job performance predictivity and accuracy in general.”¹

Vendors of this type of testing have offered voluntary audits of its AI assessments, which beg new questions regarding the AI-auditing industry. These voluntary audits have been criticized for being self-funded, creating “a risk of the auditor being influenced by the fact that this is a client,” failing to account for intersectionality, and questioning whether auditing reveals if AI products assist employers with making better hiring choices.¹¹

EEOC Regulatory Enforcement and Proposed Actions

The EEOC’s October 2021 launch of the Artificial Intelligence and Algorithmic Fairness Initiative and its May 2022 technical assistance document¹² about AI and disability discrimination have been important steps in engaging stakeholders and the public to update its Uniform Guidelines on Employee Selection Procedures. The DOL’s Office Federal Contract Compliance Program also issued guidance in 2019, stating that AI-based pre-employment screening and selection programs would be subject to the Uniform Guidelines if an adverse impact was found, and that contractors would be required to validate the selection procedure.¹³ Generally, as attorneys, we may not be the best equipped with the expertise to propose technical solutions and guidance, and should first be informed by IO psychologists, like SIOP (as the Uniform Guidelines had been), mathematicians, computer scientists, social scientists, and others.

Commissioner Sonderling summarized proposed solutions and approaches to addressing employment discrimination in AI-based pre-employment tools include:¹⁴

- The Algorithmic Accountability Act, granting FTC authority to promulgate regulations to require large companies to assess AI tools for potential bias.

- State-level proposals to expand liability for employers and third-party vendors using, selling, or administering AI tools used in employment decision-making.
- Model risk management (MRM), or self-
- Improved data collection
- Looking to the European Union’s Artificial Intelligence Act’s risk-based approach to regulation, which also focuses on vendor liability rather than solely employer liability— which will impact companies doing business in both the US and the EU.

clearly related to the job; (3) AI-based assessments should produce scores that predict future job performance (or other relevant outcomes) accurately; (4) AI-based assessments should produce consistent scores that measure job- related characteristics (e.g., upon re-assessment); and (5) All steps and decisions relating to the development and scoring of AI-based assessments should be documented for verification and auditing.

I agree with Commissioner Sonderling in his proposal that “the EEOC should consider using Commissioner charges and directed investigations to address AI-related employment discrimination” because they can “facilitate and may expedite the initiation of targeted bias probes.”¹⁵ For example, “the EEOC can initiate [directed] investigations without an underlying charge from an identifiable victim” and “Commissioner charges are useful for identifying and remedying possible systemic or pattern-or-practice discrimination rather than single plaintiff discrimination because they are initiated from a broader enforcement perspective.”¹⁶ These proposals contemplate the difficulty that potential plaintiffs may face when the source of the discrimination—AI in the various stages of recruitment and hiring—is largely unknown as the reason for the employment decision.

Other preventative measures that can be taken include voluntary compliance by employers (facilitated by an EEOC-created voluntary compliance program), along with attorneys’ adherence to professional responsibility duties in technology competence to advise clients in using AI-based technologies responsibly, ethically, and legally. These approaches are tied to suggestions of auditing (whether self-auditing or third-party auditing), and as the Director of OFCCP Jenny R. Yang

recognized: “the EEOC could be empowered to establish standards for auditors concerning qualifications and independence” and that the “government could establish an auditing framework and set core requirements for retention and documentation of technical details, including what training data must be disclosed for review during an investigation.”¹⁷

Workplace advocates agree that the EEOC should provide more frequent and consistent guidance to clarify the law and help encourage technology vendors and employers to be proactive in preventing discriminatory effects through issuing more opinion letters on the topic, and to work with state and local agencies where new laws directed at AI are becoming more prevalent. The EEOC could work collectively with localities that are out in front protecting workers, including New York City’s Local Law Int. No. 144, which just took effect on January 1, 2023, and will be enforced starting April 15, 2023, as the NYC Department of Consumer and Worker Protection considers proposed rules to implement the law.¹⁸ This new law regulates the use of “automated employment decision tools” in hiring and promotion decisions within NYC. The law, which applies to employers and employment agencies alike, requires that: any AI tool undergo an annual, independent “bias audit,” with a publicly available summary; employers provide each candidate (internal or external) with 10 business days’ notice prior to being subject to the tool; the notice list the “job qualifications and characteristics” used by the tool to make its assessment; the sources and types of data used by the tool, as well as the applicable data-retention policy, be made available publicly (or upon written request from the candidate); and candidates be able to opt out and request an alternative selection process or accommodation.

AI Promoting Worker Empowerment

As a workplace fairness advocate, I’m particularly attuned to how marginalized workers are most disadvantaged from these new technologies. As mathematician Cathy O’Neil recognized in her book, *Weapons of Math Destruction*, algorithms have a destructive disparate impact on poor candidates because wealthier individuals are more likely to benefit from personal input. “A white-shoe law firm or an exclusive prep school will lean far more on recommendations and face-to-face interviews than will a fast-food chain or cash-strapped urban school district. The privileged . . . are processed more by people, the masses by machines.”¹⁹

Conclusion

AI/algorithm technologies that are deployed for sourcing, recruitment, and hiring selection are designed to discriminate, that is, differentiate and select prospective job applicants and candidates based on complex statistical analyses. These tools function primarily as time-saving and cost-effective ways to sort and hire workers and have become ubiquitous in low-wage industries. For example, Workstream, a hiring and onboarding platform, states on its website: “Workstream is the mobile-first hiring and onboarding platform for the deskless workforce.

Powered by automation and two-way texting, our platform enables businesses to source, screen and onboard hourly workers faster. More than 24,000 businesses trust Workstream to hire - and save up to 70% of time on hiring.”²⁰

However, cost-effectiveness cannot drive employment decision-making if it runs afoul of anti-discrimination laws. We need to construct the means for testing the validity and reliability of these models. For any algorithmic decision making, the algorithm should be tested, by experts, as valid *in advance* for the type of job at issue before it is applied. “Congress or state legislatures could codify, with stiff penalties, the Uniform Guidelines approach that before using a selection tool for hiring, an employer should perform a job analysis to determine which measures of work behaviors or performance are relevant to the job or group of jobs in question. Then the employer must assess whether there is ‘empirical data demonstrating that the selection procedure is predictive of or significantly correlated with important elements of job performance.’”²¹

Moreover, people who are exposed to this technology should be given adequate notice of its use and sufficient information to assess whether their civil rights have been implicated or violated. Finally, a coordinated federal and state/local inter-agency government response is clearly warranted to develop the technical expertise required to evaluate and regulate the

¹ See, e.g., Pauline Kim and Matthew T. Bodie, *Artificial Intelligence and the Challenges of Workplace Discrimination and Privacy*, 35 ABA Journal of Labor and Employment Law 2, 289 (2021); Nantiya Ruan, *Attorney Competence in the Algorithm Age*, 35 ABA Journal of Labor & Employment Law 317 (2021); Jenny R. Yang, *Adapting Our Anti-Discrimination Laws to Protect Workers’ Rights in the Age of*

Algorithmic Employment Assessments and Evolving Workplace Technology, 35 ABA Journal of Labor & Employment Law 207 (2021).

² See <https://www.aarp.org/work/age-discrimination/facebook-settles-discrimination-lawsuits/> (<https://www.aarp.org/work/age-discrimination/facebook-settles-discrimination-lawsuits/>) and <https://www.propublica.org/article/facebook-ads-discrimination-settlement-housing-employment-credit>. (<https://www.propublica.org/article/facebook-ads-discrimination-settlement-housing-employment-credit>)

³ See <https://about.fb.com/news/2023/01/an-update-on-our-ads-fairness-efforts/> (<https://about.fb.com/news/2023/01/an-update-on-our-ads-fairness-efforts/>). (“As a part of our settlement with the Department of Justice (DOJ), representing the US Department of Housing and Urban Development (HUD), we are using these new technologies in the workplace to protect against systemic violations of our nation’s civil rights statutes.”)

⁴ See <https://newsroom.tiktok.com/en-us/find-a-job-with-tiktok-resumes> (<https://newsroom.tiktok.com/en-us/find-a-job-with-tiktok-resumes>). (“Interested candidates are encouraged to creatively and authentically showcase their skillsets and experiences, and use #TikTokResumes in their caption when publishing their video resume to TikTok.”), last visited on January 28, 2023.

⁵ <https://www.apa.org/ed/accreditation/about/policies/personnel-selection-procedures.pdf> (<https://www.apa.org/ed/accreditation/about/policies/personnel-selection-procedures.pdf>) at 12.

⁶ *Griggs v. Duke Power Co.*, 401 U.S. 424 (1971).

⁷ See Joe Avella & Richard Feloni, *We Tried the AI Software Companies Like Goldman Sachs and Unilever Use to Analyze Job Applicants*, Bus. Insider (Aug. 29, 2017), <https://www.businessinsider.com/hirevue-uses-ai-for-job-interview-applicants-goldman-sachs-unilever-2017-8> [<https://perma.cc/6YJL-ZNXM>].

8 Jessica M. Walker & Don Moretti, *Visibilit Comm., Soc’y for Indus. & Org. Psych., Recent Trends in Psychometric Assessment 4* (2018), [**http://www.siop.org**](http://www.siop.org) ([**\(http://www.siop.org/Portals/84/docs/White%20Papers/PreAssess.pdf**](http://www.siop.org/Portals/84/docs/White%20Papers/PreAssess.pdf)) (“The intent is to provide a more captivating candidate experience that assesses specific skills while keeping the applicant engaged.”).

9 Stephanie Bornstein, *Reckless Discrimination*, 105 Cal. L. Rev. 1055, 1102 (2017) (citing Marianne Cooper, *The False Promise of Meritocracy*, Atlantic (Dec. 1, 2015), [**http://www.theatlantic.com/business/archive/2015/12/meritocracy/418074;**](http://www.theatlantic.com/business/archive/2015/12/meritocracy/418074;) [**\(http://www.theatlantic.com/business/archive/2015/12/meritocracy/418074%3B**](http://www.theatlantic.com/business/archive/2015/12/meritocracy/418074%3B) [**B\) Discover Great Talent “The Voice” Way, GapJumpers,**](https://www.gapjumpers.me) [**https://www.gapjumpers.me\).**](https://www.gapjumpers.me)

10 SIOP Statement on the use of Artificial Intelligence (AI) for Hiring: Guidance on the Effective use of AI- Based Assessments, Society for Industrial and Organizational Psychology, (January 29, 2022), [**https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligenc**](https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligenc) ([**\(https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligence.pdf?ver=mSGVRY-z_wR5iluE2NWQPQ%3d%3d**](https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligence.pdf?ver=mSGVRY-z_wR5iluE2NWQPQ%3d%3d) [**\(https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligence.pdf?ver=mSGVRY-z_wR5iluE2NWQPQ%3d%3d\)**](https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligence.pdf?ver=mSGVRY-z_wR5iluE2NWQPQ%3d%3d) Such advice includes: (1) AI-based assessments should produce scores that are considered fair and unbiased; (2) The content and scoring of AI-based assessments should be

11 Hilke Schellman, *Auditors are testing hiring algorithms for bias, but there’s no easy fix*, MIT Technology Review (February 11, 2021), [**https://questionstechnologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-**](https://questionstechnologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-) ([**\(https://questionstechnologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/\)**](https://questionstechnologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/) [**algorithms-bias-big-questions-remain/**](https://questionstechnologyreview.com/2021/02/11/1017955/auditors-testing-) [**\(https://questionstechnologyreview.com/2021/02/11/1017955/auditors-testing-**](https://questionstechnologyreview.com/2021/02/11/1017955/auditors-testing-)

ai-hiring-algorithms-bias-big-questions-remain/ (quoting Professor Pauline Kim).

¹² Press Release, U.S. Equal Emp. Opportunity Comm'n, U.S. EEOC and U.S. Department of Justice Warn Against Disability Discrimination (May 12, 2022), [https://www.eeoc.gov/newsroom/us-eeoc-and-us-department-w.eeoc.gov/newsroom/us-eeoc-and-us-department-\(http://www.eeoc.gov/newsroom/us-eeoc-and-us-department-\)justice-warn-against-disability-discrimination](https://www.eeoc.gov/newsroom/us-eeoc-and-us-department-w.eeoc.gov/newsroom/us-eeoc-and-us-department-(http://www.eeoc.gov/newsroom/us-eeoc-and-us-department-)justice-warn-against-disability-discrimination).

¹³ See Off. of Fed. Cont. Compliance Programs, *Validation of Employee Selection Procedures*, U.S. DEP'T.

OF LAB (2019), <https://www.dol.gov/agencies/ofccp/faqs/employee-selection-procedures> (<https://www.dol.gov/agencies/ofccp/faqs/employee-selection-procedures>).

¹⁴ Keith E. Sonderling, Bradford J. Kelley, and Lance Casimir, *The Promise and The Peril: Artificial Intelligence and Employment Discrimination*, 77 U. Miami L. Rev. 1, 53-61 (2022). Available at: <https://repository.law.miami.edu/umlr/vol77/iss1/3>.

¹⁵ *Id.* at 66.

¹⁶ *Id.* at 67.

¹⁷ Jenny R. Yang, *Adapting Our Anti-Discrimination Laws to Protect Workers' Rights in the Age of Algorithmic Employment Assessments and Evolving Workplace Technology*, 35 ABA Journal of Labor & Employment Law 207, 227 (2021).

¹⁸ NYC Admin. Code § 20-1201.

¹⁹ Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatened Democracy*, 8 (2016).

²⁰ <https://www.workstream.us/>, (<https://www.workstream.us/>) last visited January 23, 2023.

²¹ Lori Andrews and Hannah Bucher, *Automating Discrimination: AI Hiring Practices and Gender Inequality*, 44 *Cardozo L. Rev.* 145, 200–01 (2022).

Panel 3



U.S. Equal Employment Opportunity Commission

Testimony of Matthew Scherer

Chair Burrows, Vice Chair Samuels, and Commissioners,

Thank you for the opportunity to testify on employment discrimination in A.I. and automated systems. My name is Matt Scherer, and I am Senior Policy Counsel for Workers' Rights and Technology at the Center for Democracy & Technology. CDT is a nonprofit, nonpartisan 501(c)(3) organization based in Washington, D.C. that advocates for stronger civil rights protections in the digital age. CDT's work includes a project focused on how algorithmic tools that are used in employment decisions can interfere with workers' access to employment and limit their advancement opportunities.¹

CDT has worked with a broad coalition of civil rights and civil society organizations over the past several years to help develop principles and standards regarding these technologies that center and advance the interests of workers, particularly those from historically marginalized and disadvantaged groups. In particular, over the past two years, we worked with several of these organizations to create the *Civil Rights Standards for 21st Century Employment Selection Procedures*.² We were proud to partner with the ACLU, American Association of People with Disabilities, Upturn, the Leadership Conference for Civil and Human Rights, and the National Women's Law Center in drafting the principles, and to receive endorsements from the Bazelon Center for Mental Health Law, Color of Change, the National Employment Law Project, the Autistic Self-Advocacy Network, and other groups.

As the Commission is aware, more and more employers are using artificial intelligence and other automated systems to make employment decisions that determine the course of workers' careers and lives. Automated employment decision tools (AEDTs) come in many forms, including tools that analyze the words candidates use in resumes and programs that use computer games or quizzes to estimate a candidate's personality traits.

But these tools rarely, if ever, make an effort to directly measure a worker's ability to perform the essential duties and tasks that will be expected of whomever the employer hires for the position. They also often pose a risk of discrimination against already-disadvantaged groups of workers, who are often underrepresented in the data used to train employment decision tools and whose relevant skills and abilities may not be as obvious to an automated system. My testimony will discuss how the current legal framework fails to adequately account for the unique risks of discrimination that AEDTs present and discuss how the Civil Rights Standards are a key resource that the Commission should use to inform future guidance, technical assistance, and regulatory efforts.

The Current Legal Framework Does Not Adequately Address the Heightened Discrimination Risks That AEDTs Pose

From a civil rights perspective, the current legal landscape governing AEDTs needs clarification and refinement. While the Uniform Guidelines for Employee Selection Procedures (UGESPs)³ remain in effect, they do not adequately reflect the many changes in law and social science that have occurred in the five decades since they were drafted.

The Commission and its sister agencies adopted the UGESPs in 1978, more than a decade before Congress passed the Americans with Disabilities Act (1990). By their own terms, the UGESPs do not address discrimination against people with disabilities or age discrimination, nor do they address the full scope of sex discrimination. The UGESPs also have not been updated to expressly incorporate modern scientific standards regarding validation and fairness.⁴ This makes further action by the EEOC urgent, to clarify how the EEOC will interpret and apply the statutes and regulations it enforces to meet the unique risks posed by automated tools.

To begin, Title VII (and the ADA) state that where an employment practice has a disparate impact, it constitutes unlawful discrimination unless the employer demonstrates “that the practice is job related *for the position in question* and consistent with business necessity.”⁵ The phrase “for the position in question” means that for a test that has a disparate impact to be valid, an employer must link it to the duties *of the specific job for which it is being used*. This echoes the Supreme Court's admonition in *Griggs v. Duke Power Co.* that any test or screening

mechanism for job applicants “must measure the person for the job and not the person in the abstract” to survive a Title VII challenge.⁶

Many of the AEDTs being marketed today fail to meet this job-specific validity requirement because they have one (or both) of two characteristics: (1) they measure abstract or amorphous characteristics not tailored to the job in question; and/or (2) they rely on machine-learning techniques that use correlation alone—rather than a logical or causal relationship with job functions—to establish a link between test results and job performance.

Tools That Measure Abstract Candidate Characteristics

Tools offered by some of the more prominent AEDT vendors claim to rate candidates not on specific knowledge or abilities, but on highly abstract and subjective qualities like “empathy,” “influence,” and “personality.”⁷ CDT’s 2020 report, *Algorithm-driven Hiring Tools: Innovative Recruitment or Expedited Disability Discrimination?*, describes in detail how these tools can discriminate against workers based on attributes including race, sex, national origin, and disability status. When such tests result in disparate impacts or tend to screen out disabled workers, federal law requires employers to establish job-relatedness in order to survive a discrimination claim. That showing is not tenable with tools that purport to measure generic personality traits or other characteristics untethered from the specific duties or essential functions of the jobs for which candidates are being evaluated. That is precisely the sort of measurement of “the person in the abstract,” rather than for a specific job, that the Supreme Court warned against and that the text of Title VII and the ADA expressly prohibits.⁸

The guidance document that the EEOC published last year on AEDTs and the ADA is a great first step in pushing back against the use of such tests. It recognizes that to minimize the risk of unlawful disability discrimination, employers should ensure that “tools only measure abilities or qualifications that are truly necessary for the job—even for people who are entitled to an on-the-job reasonable accommodation.”⁹ I encourage the full Commission to take the next step by elevating this from a best-practice recommendation to a rule that the EEOC will enforce. I respectfully submit that the plain terms of our antidiscrimination laws require nothing less.

Tools That Rely Exclusively on Correlation

Another problem with many AEDTs stems from the fact that machine-learning algorithms do not examine whether the attributes that a model uses to predict job performance are logically or causally related to the essential functions of a job, nor do they analyze whether the attributes in the training data include a set of variables that are representative of the skills needed to perform a particular job.¹⁰ Instead, AEDTs that rely on such algorithms depend on correlation—which, in the employment selection processes, means the degree to which differences in “predictor” attributes that could be used to assess candidates (such as years of experience or schools attended) are associated with differences in some target “criterion” connected to job performance (such as supervisory ratings or sales figures). Using correlation alone to select which “predictor” variables an AEDT will use can lead to both invalidity and discrimination for two reasons.

First, a tool built on correlation-based techniques alone is highly unlikely to capture all (or a representative set) of the essential functions of a specific job. Say a company wants to screen marketing specialists using a resume scanning algorithm that uses machine learning to decide which words in resumes are predictive of job performance. Even if thousands of qualified marketing specialists’ resumes are included in the algorithm’s training data, there are aspects to marketing jobs (such as interpersonal skills) that cannot be reliably extracted from a resume alone—and the fact that certain terms often show up in the resumes of successful workers does not necessarily mean those terms are the best indicators of a person’s ability to do the functions of a job (as discussed further below).

Few, if any, data sets are rich enough to cover all of the essential knowledge and abilities needed for a given job, much less the nuances of how such knowledge and abilities will be needed for a role at a specific company. This means that any tool that operates solely by searching for correlations in historical data sets will create an incomplete picture of a candidate’s ability to perform the job in question.¹¹ If a company then over-relies on such a tool when making employment decisions, its decisions will not be based on an adequate representation of essential job functions, as both the ADA and the UGESPs require.

Second, the use of correlation-driven statistical methods increases the risk that a tool will capitalize on correlations that are due to chance rather than due to a logical, causal, or organic relationship with job performance. As a result, AEDTs may discover and use attributes that are *construct-irrelevant*—that is, attributes that are

tied not to job-performance factors (the “construct” that employment tests are supposed to capture), but to irrelevant characteristics.¹² This can lead to differences in scores or selection decisions that are due to factors unrelated to a candidate’s ability to perform essential job functions.¹³ This can happen, for example, when a test measures something more or different than the relevant aspects of job performance (e.g., if a test of oral communication skills is affected by a test-taker’s proficiency in written English); or when outcomes reflect cultural differences rather than (or in addition to) differences in job-related competencies.

A hypothetical example from an article I co-authored illustrates how this can lead to unlawful discrimination:

[S]ay that a company was training an algorithmic tool to recognize good software engineers using training data that reflects the demographics of their best current network engineers, who are predominantly white males. If these employees share, as is likely, construct-irrelevant characteristics that are reflected in the training data, the tool will learn to associate those characteristics with good job performance. This could have two related adverse impacts on qualified candidates who are not white males. First, if the ablest female and nonwhite candidates have attributes (whether construct-relevant or not) that differ from those of the white males who dominate the current sample, the tool’s accuracy will be lower when scoring those candidates, just as the gender classification programs in the MIT [Gender Shades] study^[14] were less accurate when attempting to classify individuals with darker skin. Second, the individuals that the tool identifies as the best candidates from the underrepresented groups may have scored highly not because of characteristics that affect their actual competence, but because of the construct-irrelevant characteristics they share with the current software engineers.

Both of those factors may drive down the number of qualified female and minority candidates that the tool selects. In addition, the candidates who the tool does recommend from the disadvantaged group are less likely to be the most competent candidates from that group, which may reduce the likelihood that they are ultimately hired and retained. Through these mechanisms, an employer’s adoption of an algorithmic tool could inadvertently reinforce existing demographics.¹⁵

When adverse impacts arise because members of a group perform differently on improperly included or excluded aspects of job performance, the resulting discriminatory impacts would constitute Title VII and ADA violations.¹⁶

The UGESPs, having been written long before the rise of machine-learning algorithms that can comb through hundreds or thousands of potential predictors and build a model based solely on correlation, do not adequately address this source of discrimination and invalidity in AEDTs.

Again, the ADA guidance that the EEOC published last year is encouraging in this regard. That guidance suggests that employers ensure that “necessary abilities or qualifications are measured directly, rather than by way of characteristics or scores that are correlated with those abilities or qualifications.”¹⁷

Here too, I encourage the full Commission to take further formal action to stop the proliferation of discriminatory tools that rely on aimless correlations. The Commission should issue additional guidance explaining that, consistent with the proper interpretation of “job-relatedness” under our antidiscrimination laws and *Griggs*, correlation alone does not suffice to establish

job-relatedness absent other evidence or explanation addressing *why* the attributes measured by an automated tool should be expected to predict a candidate’s ability to perform essential functions. Additionally, the EEOC should issue guidance and, if practicable, engage in rulemaking to address how impact and validity analyses should be conducted in light of the unique requirements of the ADA—a statute to which the UGESPs do not apply and that thus presents an important area for agency action.

Civil Rights Standards for 21st Century Employment Selection Procedures

Despite the threats to validity and the risk of discrimination that AEDTs pose, some vendors and allied special interest groups have actively sought policy changes that would weaken or undercut existing protections or confuse employers and consumers about what current law requires. They often do so under the pretense that their technologies are less biased than longer-established employment decision processes, and that their proposed policy changes thus represent a pro-civil-rights position. The evidence and arguments used to support these efforts are generally incomplete at best, and highly misleading at worst.¹⁸

Faced with this combination of (i) the risks of wide-scale discrimination posed by AEDTs and (ii) intensifying efforts to insulate AEDTs from discrimination accountability, CDT partnered with many of the nation’s leading civil rights

organizations—including the ACLU, which is here today—to create the *Civil Rights Standards*.¹⁹ While the rise of automated tools was the impetus for the *Standards*, the *Standards* themselves apply to all formalized selection procedures and thus lay out a path to updating existing rules and guidance regarding employee assessments.

It is our hope that the Commission considers the *Standards* as it completes its Strategic Enforcement Plan, and more generally as it moves forward in its regulatory, educational, and compliance efforts regarding automated tools.

The *Standards* advance the five *Civil Rights Principles for Hiring Assessment Technologies*²⁰ that were first developed in 2020 by a broad coalition of civil rights groups, including CDT. Those five principles are: Nondiscrimination, job-relatedness, auditing, notice and explanation, and oversight and accountability. The *Standards* expand on and operationalize these five Principles, providing a concrete alternative to proposals that would set very weak notice, audit, and fairness standards for automated tools. They are designed for inclusion in regulatory guidance, for adoption by vendors and companies, and for workers who deserve to know their rights.

The *Standards*' key provisions include:

Nondiscrimination

Targeting and reducing the risk of all forms of unlawful discrimination by:

- Requiring companies to take a proactive approach to eliminating potential sources of discrimination
- Mandating that employers use the least discriminatory selection procedure (SP) available
- Banning certain SPs that pose an especially high risk of discrimination

Job-Relatedness

Ensuring that SPs only measure traits and skills that are important to job performance by:

- Requiring SPs to measure only the essential functions of the job(s) for which they are used
- Requiring audits to include a description of the essential functions for which the SP is being used and an explanation of why those functions are essential to the position
- Requiring correlation-based evidence of validity to be supported by theoretical, logical, or causal reasoning sufficient to explain *why* the SP's predictors should be expected to predict the ability to perform essential functions
- Prohibiting the use of SPs whose validity cannot be assessed

Auditing

Requiring both *pre-deployment* and *ongoing audits* by an independent auditor.

Audits must:

- Include a thorough job analysis for each position for which the SP would be used
- Analyze the SP's validity and risk of various forms of discrimination
- Determine whether valid and less discriminatory assessment methods are available

Notice and Explanation

Creating three layers of disclosure requirements, each tailored to a different intended audience:

- A short-form disclosure describing for candidates how the SP works and how they can raise concerns and access accommodations
- Detailed audit summaries, intended for sophisticated stakeholders like regulators and workers' advocates
- Comprehensive recordkeeping obligations so all information is preserved in case of investigation or litigation

Oversight and Accountability

Giving all stakeholders a role in ensuring that selection procedures do not violate the law, by:

- Making it unlawful to use or sell discriminatory SPs
- Giving candidates the right to communicate concerns about SPs prior to their use, and the right to request human review of automated SPs' recommendations
- Providing for agency enforcement, as well as a private right of action for certain unlawful practices
- Making employers and vendors jointly responsible for audits, and jointly and severally liable for discrimination

The *Civil Rights Standards* provide a roadmap to managing the risks associated with modern selection tools while centering the rights and dignity of workers, particularly those most vulnerable to marginalization and discrimination. They contain provisions that would address the unique threats of discrimination discussed above. They are designed to be modular; each standard reinforces and strengthens the others, but each also stands on its own. Again, we hope they can be a resource to the Commission as it completes its Strategic Enforcement Plan and continues its important regulatory, educational, and compliance efforts regarding automated tools.

Conclusion

We have seen many ways in which new technologies have made the workplace and labor market fairer and more efficient. The rise of the Internet, for example, enhanced workers' ability to search and apply for jobs and career paths. But not all new technologies represent progress. History is replete with examples of supposed innovations that, despite the hype and assurances of the companies promoting

them, failed to live up to their potential or were rushed to market before the technology was ready for prime time. Where, as here, the careers and livelihoods of so many workers are at stake, there is a risk of great harm if ineffective and potentially harmful technologies are allowed to proliferate without proper scrutiny.

As the Commission completes its Strategic Enforcement Plan over the coming weeks and months, we urge it to use its platform and authority to ensure that workers, not machines, remain at the center of the future labor market. The rights of workers, particularly vulnerable and marginalized workers, must not be trampled or glossed over for the sake of convenience or efficiency. Thank you.

¹<https://cdt.org/area-of-focus/privacy-data/workers-rights/>
(<https://cdt.org/area-of-focus/privacy-data/workers-rights/>).

² Ctr. for Democracy & Tech., et al., *Civil Rights Standards for 21st Century Employment Selection Procedures*, [https://cdt.org/insights/civil-rights-standards-for-21st-century-employment-selection- \(https://cdt.org/insights/civil-rights-standards-for-21st-century-employment-selection-procedures/\) procedures/ \(https://cdt.org/insights/civil-rights-standards-for-21st-century-employment-selection-procedures/\)](https://cdt.org/insights/civil-rights-standards-for-21st-century-employment-selection- (https://cdt.org/insights/civil-rights-standards-for-21st-century-employment-selection-procedures/) procedures/ (https://cdt.org/insights/civil-rights-standards-for-21st-century-employment-selection-procedures/).).

³ 29 C.F.R. § 1607.1 et seq.

⁴ See generally Am. Educ. Research Ass'n, et al., *Standards for Educational and Psychological Testing*, (4th ed. 2014); Soc'y for Indus. & Organizational Psychology, *Principles for the Validation and Use of Personnel Selection Procedures* (5th ed. 2018). See also Matthew U. Scherer, et al., *Applying Old Rules to New Tools: Employment Discrimination Law in the Age of Algorithms*, 71 So. Car. L. Rev. 449 (2019), available at https://ssrn.com/abstract_id=3472805
(https://ssrn.com/abstract_id%3D3472805).

⁵ 42 U.S.C. § 2000e-2(k)(1)(A)(i) (Title VII) and 42 U.S.C. § 12112(b)(6) (ADA) (emphasis added).

⁶ *Griggs v. Duke Power Co.*, 401 U.S. 424, 436 (1971).

⁷ See, e.g., Matthew Scherer, HireVue “AI Explainability Statement” Mostly Fails to Explain What it Does (2022), <https://cdt.org/insights/hirevue-ai-explainability->

statement-mostly-fails-to-explain-what-it-does/?

(<https://cdt.org/insights/hirevue-ai-explainability-statement-mostly-fails-to-explain-what-it-does/>

utm_source=rss&utm_medium=rss&utm_campaign=hirevue-ai-explainability-statement-mostly-fails-to-explain-what-it-does)

utm_source=rss&utm_medium=rss&utm_campaign=hirevue-ai-explainability-statement-mostly-fails-to-ex (<https://cdt.org/insights/hirevue-ai-explainability-statement-mostly-fails-to-explain-what-it-does/>

utm_source=rss&utm_medium=rss&utm_campaign=hirevue-ai-explainability-statement-mostly-fails-to-explain-what-it-does) plain-what-it-does

(<https://cdt.org/insights/hirevue-ai-explainability-statement-mostly-fails-to-explain-what-it-does/>

utm_source=rss&utm_medium=rss&utm_campaign=hirevue-ai-explainability-statement-mostly-fails-to-explain-what-it-does)

(noting how the competencies that one vendors' assessments claim to measure "are not moored to the actual responsibilities and functions of specific jobs").

⁸ 42 U.S.C. § 2000e-2(k)(1)(A)(i); 42 U.S.C. § 12112(b)(6).

⁹ EEOC, The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees, Q14 (Promising Practices for Employers), <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence#Q14>.

¹⁰ See Scherer et al., *supra* note 4, at 487.

¹¹ This shortcoming is referred to as *construct deficiency* in the literature on test validity. That is, the test does not capture all of the relevant aspects of the construct (in this case, the ability to perform essential job functions) that the tool is supposed to be measuring.

¹² See Keith E. Sonderling, et al., *The Promise and The Peril: Artificial Intelligence and Employment Discrimination*, 77 U. Miami L. Rev. 1, 24 ("In analyzing a large quantity of data, an algorithm might identify a statistical correlation between a specific characteristic of a job applicant and future job success that nevertheless lacks a causal relationship.").

¹³ The technical term for this phenomenon is *construct-irrelevant variance* or *contamination*.

¹⁴ In that study, facial recognition technology was found to be less accurate in correctly identifying the gender of darker-skinned individuals than lighter-skinned ones--and the darker the individual's skin, the less accurate the tool was.

¹⁵ Scherer, *supra* note 4, at 488.

¹⁶ See 42 U.S.C. § 2000e-2(k)(1)(A)(i) (demonstration of job-relatedness required to overcome showing of adverse impact).

¹⁷ EEOC ADA Guidance, *supra* note 11.

¹⁸ See, e.g., Scherer, *supra* note 7; Hilke Schellmann, *Auditors are testing hiring algorithms for bias, but there's no easy fix*, MIT Technology Review, Feb. 11, 2021, <https://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/> (<https://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/>); Mona Sloane, *The Algorithmic Auditing Trap*, OneZero (Medium), Mar. 16, 2021, <https://onezero.medium.com/the-algorithmic-auditing-trap-9a6f2d4d461d> (<https://onezero.medium.com/the-algorithmic-auditing-trap-9a6f2d4d461d>).

¹⁹ As of January 13, 2023, the following organizations have endorsed the *Standards*: Center for Democracy & Technology (CDT), American Association for People with Disabilities (AAPD), American Civil Liberties Union (ACLU), Autistic People of Color Fund, Autistic Self Advocacy Network (ASAN), Autistic Women & Nonbinary Network (AWN), Bazelon Center for Mental Health Law, Color Of Change, The Leadership Conference on Civil and Human Rights, National Employment Law Project (NELP), National Women's Law Center (NWLC), TechEquity Collaborative, Upturn.

²⁰ *Civil Rights Principles for Hiring Assessment Technologies* (2020), <https://civilrights.org/resource/civil-rights-principles-for-hiring-assessment-technologies/>

hiring-assessment-technologies/ (https://civilrights.org/resource/civil-rights-principles-for-hiring-assessment-technologies/).



U.S. Equal Employment Opportunity Commission

Testimony of Heather Tinsley-Fix

Chair Burrows and distinguished Commissioners, on behalf of our 38 million members and all older Americans nationwide, thank you for the opportunity to speak to you today regarding the intersection of AI-enabled employment decisions and the potential for age discrimination. I am honored to be here. My name is Heather Tinsley-Fix and I am Senior Advisor, Employer Engagement at AARP. AARP believes that any type of discrimination in the workplace is unacceptable. Too often, when discussing discrimination, age is not included, although ageism continues to be a widespread problem.

My remarks today will focus on (a) ways in which the current use of AI in hiring and other workforce decisions can affect older workers and (b) what employers and the EEOC can do to mitigate the risks of unintended age discrimination.

Current Use of AI in Hiring and Other HR Technologies

Advances in technology over the past two decades have drastically changed the way companies recruit, hire, and manage talent. The prevalence and reach of platforms designed to connect job seekers to the right jobs means that companies can no longer manually process the flood of resumes they may receive for any one job opening. In addition, the competition for skilled workers coupled with historically low unemployment rates have intensified the demand for automated solutions that help organizations find, hire, train, and promote the best candidates for the job. Furthermore, the tantalizing promise of outsourcing the analysis and selection of job candidates to a bloodless algorithm which will curb or even eliminate human biases is difficult to resist.

Before I dive into ways in which AI and automation have the potential for discouraging or discriminating against older candidates, I want to make two points. The first is that, across the AI-enabled hiring process, the inputs used to define and

then train the algorithms build iteratively on each other toward the ultimate goal of predicting which candidate should be hired. This makes unpicking the source of bias extremely challenging because the algorithm not only spots patterns based on what it's been told to look for, it also learns from the decisions introduced to the process by human actors. Throughout the creation and implementation of such systems, human definitions, decisions, and inputs mingle with the data stream to the extent that the "A" in "AI" is more of an augmentation of existing human intelligence rather than an artificial replacement of it. And the second is that not all companies use AI across all aspects of the hiring process – some may simply use an Applicant Tracking System that scans resumes while others might leverage matching and ranking functionality, or chatbots, or online games – which makes analysis of what works and what doesn't challenging.

In terms of age bias and discrimination, the potential pitfalls associated with the use of AI in hiring and workforce management platforms are, at the root, the same for older candidates as they are for other protected classes – namely, the quality or relevance of available data used to train algorithms, and the normative judgments baked into the process about what "good" looks like. However, the way those pitfalls affect older workers can look a little different or come from unexpected places. Here are some examples:

- Type and amount of data collected – to the extent that algorithms scrape and use data from social media posts and activity, professional digital profiles, internet browsing history, mobile device use, etc. to power their predictive rankings, older adults may be left out of the consideration set due to either a lack of those types of data in their digital footprint or the fact that fewer older job candidates are considered when building "ideal candidate" profiles. Furthermore, any data point collected that explicitly reveals or serves as a proxy for age – such as date of birth, years of experience, or date of graduation – can be noticed by the algorithm as part of a pattern denoting undesirable candidates and signal the algorithm to lower their ranking or screen them out entirely.
- Cultural norms – there are a host of unconscious assumptions baked into our culture that associate age with slowing, cognitive decline, an inability to learn new things, and resistance to change. These norms inform the way job descriptions are worded, target variables are defined, interviews are conducted, and assessments are designed and scored. For example, if reaction

time is a variable on which candidates are scored, older workers may be at a disadvantage. Research shows that older brains exhibit slower processing speeds but greater contextual knowledge.[1] However, if skills assessments or the analysis of interview footage are optimized toward younger brains by the data scientists working on them, older workers could receive arbitrarily lower scores. Additionally, older workers could be excluded at the start of the hiring process because they never see the job ads to begin with. In 2017, ProPublica revealed that Facebook was allowing organizations to age-target their employment ads, in some cases excluding workers over the age of 35, but in most cases excluding workers over 50. This can also include the way job descriptions are worded – phrases like “recent college grad” and “digital native” are explicitly ageist, but even subtle references such as “fast-paced,” “high-energy,” and “super fun,” have been shown to deter older workers from applying.[2]

- The feedback loop of decisions taken by recruiters or hiring managers – AARP research on the experiences of older workers shows that age discrimination remains stubbornly with us. Our most recent survey reveals that 64% of workers aged 40+ face age discrimination at work. To the extent that algorithms learn from the preferences and decisions made against older candidates during the recruiting process, they will spot the patterns in the data that indicate an older candidate, and subsequently promote those candidates less frequently and less far through the automated process. For example, if an older candidate makes it past the resume screening process but gets confused by or interacts poorly with the chatbot and ultimately gets rejected by the recruiter, that data could teach the algorithm that candidates who have similar dates of graduation and hesitate when chatting with a bot should be ranked lower. This applies to performance data as well; research shows that performance reviews tend to level out or even decline with age despite weak to no correlation between increased age and a drop in productivity. To the that extent performance evaluations, or indeed a wider host of employment-related decisions such as who is selected for training, innovative projects, high-performing teams, etc. are influenced by ageism and that data is fed into ranking algorithms as proof points, older workers could be disadvantaged.

What Employers and the EEOC Can Do to Mitigate the Risks of AI-Enabled Age Discrimination

Employers use AI-enabled hiring technologies and platforms shorten the time it takes to fill open positions, to find the best match between available job seekers and available jobs, and ideally to continue in the quest to remove as much human bias from the process as possible. The task at hand is not to convince them not to use such technologies but to provide them with the best information and guidance to make informed decisions, and to shore up that awareness with regulatory guardrails.

In the realm of practical guidance, the non-profit organization Upturn, which seeks to advance equity and justice in the design, governance, and use of technology, has a comprehensive set of recommendations that serve as a template and starting point for the goal of using such technologies wisely.^[3] In addition, there are many steps employers can take to specifically address the risk of unintended age discrimination and bias. They are as follows:

1. Stop asking for age-related data in applications, such as dates of birth or graduation, unless there is a proven business reason to do so. If employers must know the age of a candidate, they should not limit the age a candidate can be, such as limiting the years listed in a drop-down menu. Alternatively, platforms could simply verify that candidates are at least 18 years old if that is a business requirement.
2. Pay close attention to the words used in job descriptions, and don't cap the years of experience required. Replacing "2 – 5 years' experience" with "at least 2 years' experience" signals that candidate of all ages are welcome to apply. AARP has a guide to age-inclusive job posting language which can be found at [**aarp.org/employers \(http://www.aarp.org/employers\)**](http://www.aarp.org/employers).
3. Don't age-target employment ads on platforms that allow such targeting, even if that includes filters that approximate age such as job seniority or years of experience.
4. Look for vendors who work with certified Industrial/Organizational Psychologists, who are trained in the development and evaluation of tests, assessments, and other selection procedures. In particular, any use of non-employment-related data should be vigorously scrutinized for its potential to rely on correlation rather than causation. The Society for Industrial Organizational Psychologists have recently published guidelines for evaluating

AI-enabled selection technologies which can be used when evaluating vendors.

[4]

5. Request (or conduct) regular and independent audits of algorithmic performance to see whether adverse impact is occurring, and in what part of the hiring funnel.
6. Include age as an element of diversity, equity, and inclusion initiatives. Driving awareness of the value of age diversity at work will help shift a culture of unconscious ageism.
7. And finally, empower recruiters to challenge hiring managers who exhibit conscious or unconscious preferences for candidates based on age. There is a strong business case for the inclusion of older workers as part of an age-diverse workforce. Visit [aarp.org/employers](http://www.aarp.org/employers) (<http://www.aarp.org/employers>) to learn more.

On the legislative front, AARP is supporting federal and state legislative initiatives to ban age-related questions during the application process that have the effect of screening out and deterring older applicants. Such information is simply irrelevant to a candidate's qualifications and skills. Connecticut and Delaware have both enacted such bans and AARP is working with legislators in New York and Oregon to enact similar bans.

At the federal level, AARP supports the *Protect Older Job Applicants Act* (POJA), which would clarify that job applicants are allowed to bring disparate impact claims under the Federal Age Discrimination in Employment Act. Two appellate cases, *Villareal v RJ Reynolds* (11th Cir. 2016) and *Kleber v PriceWaterhouse* (7th Cir. 2019), interpreted the law to mean that applicants may not bring disparate impact claims under Section 4(a)(2) of the ADEA, only employees. POJA closes this inadvertent gap in the ADEA to ensure the legal rights of applicants for jobs are protected as well.

AARP continues to advocate for passage of the bipartisan *Protecting Older Workers Against Discrimination Act* (POWADA) to overturn *Gross v. FBL Financial Services, Inc.* and amend the Age Discrimination in Employment Act (ADEA), Title VII's retaliation provision, the Americans with Disabilities Act, and the Rehabilitation Act of 1973, to clarify that the same standards of proof apply to all claims under all of these laws.

Again, thank you for providing AARP the opportunity to testify today. I look forward to answering any questions.

[1] <https://news.brown.edu/articles/2014/11/age>

[2] <https://www.nber.org/papers/w30287>

[3] <https://www.upturn.org/work/help-wanted/>
(<https://www.upturn.org/work/help-wanted/>)

[4]

[https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligence.pdf?ver=mSGVRY-](https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligence.pdf?ver=mSGVRY-z_wR5iluE2NWQPQ%3d%3d)

[z_wR5iluE2NWQPQ%3d%3d](https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligence.pdf?ver=mSGVRY-z_wR5iluE2NWQPQ%3d%3d)

([https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligence.pdf?ver=mSGVRY-](https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligence.pdf?ver=mSGVRY-z_wR5iluE2NWQPQ%3d%3d)

[z_wR5iluE2NWQPQ%3d%3d](https://www.siop.org/Portals/84/docs/SIOP%20Statement%20on%20the%20Use%20of%20Artificial%20Intelligence.pdf?ver=mSGVRY-z_wR5iluE2NWQPQ%3d%3d))



U.S. Equal Employment Opportunity Commission

Testimony of Dr. Ifeoma Ajunwa

1. Introduction

Greetings to all. Chair Burrows and Commissioners, thank you for inviting me to testify at this important public meeting on employment discrimination in AI and automated systems. My name is Dr. Ifeoma Ajunwa. I am a law professor at the University of North Carolina School of Law where I am also the founding director of the AI Decision-Making Research Program. I am also a founding board member of the Labor Tech Research Network which is an international group of scholars dedicated to research on ethical and legal issues associated with AI in the workplace. I have published several law review articles on automated hiring systems and I have a forthcoming book, *The Quantified Worker*, which discusses pressing legal issues arising from the use of AI and automated decision-making technologies in the workplace. I have previously testified before this commission in 2016, and in 2020, I also testified before the Congressional Education and Labor Committee on the issue of workers' rights in the digital age.

1. Automated Hiring's Potential for Discrimination

In several writings, I have documented the capability of automated hiring programs to both replicate unlawful employment discrimination and obfuscate it. [1] First, I make note of the business trend towards the use of automated hiring programs. In an informal survey that a co-author and I conducted, we found that the top twenty private employers in the Fortune 500 list all made use of automated hiring systems. [2] Notably also, this list comprised of mostly retail companies with large numbers of workers in low-wage or entry level work. It is true that many businesses turn to

automated hiring in an attempt to diversify their workplace and eliminate the human bias that might stand in the way of that, yet, there is evidence that such algorithmic decision-making processes may stymie the progress made by antidiscrimination laws and may also serve to magnify inequality in the workplace. [3] In a study of 135 archival texts tracing the development of hiring platforms from 1990-2006, my co-author and I found that although one presumed reason for automated hiring was to reduce hirer bias, in actuality, the automated hiring platforms were initially advertised as a way to “clone your best worker,” a slogan that, in effect, replicated bias. [4] This is precisely because automated hiring systems take the nebulous concept of “cultural fit” for a given corporation or organization and concretize it into select algorithmic variables that are stand-ins for protected categories. [5] The creation of these proxy variables can have racial, gender, and age discrimination effects in contravention of antidiscrimination laws for which this commission has regulatory power such as Title VII of the Civil Rights Act and the Age Discrimination in Employment Act (ADEA). Furthermore, the designs of automated hiring systems, both in their user interface and data retention protocols can further enable unlawful employment discrimination, in a manner that is both discreet and difficult to document. [6]

I’ll illustrate these problems first with an example shared by an employment and labor law attorney who had been hired to audit an automated hiring system. As reported by the lawyer, when the automated system was presented with training data (presumably the resumes of top performers) and then queried as to which two variables were found to be the most relevant – the system reported back those variables as the name “Jared” and whether an applicant had played “high school lacrosse.” [7] As I detail in my forthcoming book, the significance of this result is that these are variables that may be considered proxy variables. As confirmed by social security records, the name “Jared” is highly correlated to individuals who are both white and male. [8] Furthermore, lacrosse is a sport that isn’t found in all high schools, rather it is an expensive sport that is found in well-funded high schools located in affluent neighborhoods that are more likely to be predominantly white, given the history of racial segregation in the United States. [9] Thus, in this insidious manner, proxy variables as part of automated hiring systems can enact unlawful racial and gender employment discrimination.

I share two more stories as examples of how automated hiring systems can enable unlawful discrimination through the platform authoritarianism of their user interface and through their data retention rules. The first story is that of a man in

Massachusetts who found that he was unable to complete an application because he could not register his college graduation year in the drop-down menu which limited choices to more recent graduation years.[10] In this case, the year of college graduation which is highly correlated to age is masquerading as a proxy variable for age discrimination. The second story is that my co-author and I conducted a field experiment in which we found that we could not complete the application for a major retail employer. Although we had selected a part-time position, the automated hiring platform still required the applicant to indicate unlimited availability to submit the application.[11] What was stealthy about this was that the design features of most automated hiring systems would not retain a record of the failed attempt to complete the application.[12] In this scenario, requiring unlimited availability disallows applicants with caregiving obligations from applying. This would have a disproportionate impact on women as women are more likely to be caregivers.

Finally, I point to the use of automated video interviewing with facial analysis and emotion detection as an inherently flawed and discriminatory automated hiring practice.[13] In 2018, 60% of organizations were using video interviews, but the use of automated video interviews sharply increased to 86% in 2020 due to the Covid-19 pandemic.[14] This automated hiring practice is akin to the discredited pseudoscience of phrenology and thus should be banned. Automated video interviews that are scored by speech algorithms provide opportunity for accent discrimination[15] and the ones that claim to detect emotion from facial analysis further enable racial and gender discrimination given cultural and gender differences in how emotions are expressed.[16]

- **Proposals for Governance of Automated Hiring**

Given the demonstrated capability of automated systems to be used in the service of unlawful discrimination, I offer four proposals that the EEOC should consider as part of its enforcement of employment antidiscrimination laws.

1. Discrimination Per Se

First, the EEOC should consider the addition of a third cause of action for Title VII of the Civil Rights Act. Currently, there are two causes of action: disparate treatment

and disparate impact. Both causes of action present high burdens of proof for the applicant. In the first, the applicant must in essence find the “smoking gun” that proves direct and intentional discrimination. In the second, the applicant must provide the requisite statistics to show that a specific hiring practice that had a disparate impact on a protected category. However, courts have lacked consistency in deciding appropriate statistics to consider for a disparate impact cause of action, with the effect that the plaintiff’s burden of proof has become unreasonably high.

[17] A third cause of action, which I call discrimination per se, would shift the burden of proof from applicant to employer, so long as an applicant is able to point to a feature or requirement of an automated hiring systems that seems egregiously discriminatory for a particular protected class.**[18]** The employer would then bear the burden of proving with the statistical results of the audits of its automated hiring system, that the identified automated hiring feature does not in fact have a disparate impact on a protected class.**[19]**

2. Mandated Audits of Automated Hiring Systems

This brings me to the second proposal. I propose that the EEOC should mandate employer audits of any automated hiring systems in use.**[20]** There is some question as to whether this should be an internal audit or whether the mandate should require an external audit with an independent third-party auditor. I argue that the EEOC should require external audits as internal audits are not enough. The EEOC could chose to take on these external audits or it could certify third party vendors that would provide those audits.

3. The EEOC Should Create Its Own Audit Tools

A third proposal follows from the second, as I argue that the EEOC (perhaps in cooperation with the FTC) should develop its own automated governance tools in the form of AI or automated systems that could be used to audit automated hiring systems.**[21]** The EEOC could then provide audit services to corporations deploying automated hiring systems or it could use such an audit tool in its investigation and enforcement actions.**[22]**

4. The Uniform Guidelines on Employee Selection Procedures Should Govern

Finally, my fourth proposal relates to automated video interviewing and other automated hiring practices based on shaky science and which have no probative value for hiring. I propose that the EEOC should release an advisory notice that the Uniform Guidelines on Employee Selection Procedures govern the use of variables in algorithmic hiring and the design of automated hiring systems to retain full records of both completed application and application attempts.^[23] All variables selected for automated hiring systems should then meet criterion, content, and construct validity for the specified job position:

“Evidence of the validity of a test or other selection procedure by a criterion-related validity study should consist of empirical data demonstrating that the selection procedure is predictive of or significantly correlated with important elements of job performance. Evidence of the validity of a test or other selection procedure by a content validity study should consist of data showing that the content of the selection procedure is representative of important aspects of performance on the job for which the candidates are to be evaluated. Evidence of the validity of a test or other selection procedure through a construct validity study should consist of data showing that the procedure measures the degree to which candidates have identifiable characteristics which have been determined to be important in successful performance in the job for which the candidates are to be evaluated.”^[24]

This will lessen opportunities for proxy variables to be deployed for unlawful employment discrimination against protected categories of workers and would dissuade the use of facial analysis for emotion recognition.

Once again, my gratitude to Chair Burrows and her fellow commissioners for the opportunity to present these remarks and proposals in support of the EEOC mission of equal employment opportunity.

[1] See, Ifeoma Ajunwa, *Automated Video Interviewing as the New Phrenology*, 36 Berkeley Tech. L.J. 1173 (2022);

Ifeoma Ajunwa, *The Auditing Imperative for Automated Hiring*, 34 Harv. J.L. & Tech. 621 (2021); Ifeoma Ajunwa, *The Paradox of Automation as Anti-Bias Intervention*, 41 Cardozo. L. Rev. 1671 (2020); Ifeoma Ajunwa,

Protecting Workers' Civil Rights in the Digital Age, 21 N.C.J.L & Tech. 1 (2020), Ifeoma Ajunwa,

Age Discrimination by Platforms, 40 Berkeley J. Emp. & Lab. L.1 (2019); Ifeoma Ajunwa and Daniel Greene, "Platforms at Work: Automated Hiring Platforms and Other New Intermediaries in the Organization of the Workplace" In *Work and Labor in the Digital Age. Research in the Sociology of Work*. Published online: 14 Jun 2019; 61-91.

[2] "Platforms at Work: Automated Hiring Platforms and Other New Intermediaries in the Organization of the Workplace" In *Work and Labor in the Digital Age. Research in the Sociology of Work*. Published online: 14 Jun 2019; 61-91.

[3] "Advocates applaud the removal of human beings and their flaws from the assessment process." Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 4 (2014). Algorithms or automated systems are often seen as fair because they are "claimed to rate all individuals in the same way, thus averting discrimination." *Id.*

[4] Ifeoma Ajunwa and Daniel Greene, "Platforms at Work: Automated Hiring Platforms and Other New Intermediaries in the Organization of the Workplace" In *Work and Labor in the Digital Age. Research in the Sociology of Work*. Published online: 14 Jun 2019; 61-91, page 77.

[5] Ifeoma Ajunwa, *The Paradox of Automation as Anti-Bias Intervention*, 41 Cardozo. L. Rev. 1671, 1712-1715 (2020).

[6] Ifeoma Ajunwa, *The Auditing Imperative for Automated Hiring*, 34 Harv. J.L. & Tech. 621, 622 (2021).

[7] Ifeoma Ajunwa, *The Paradox of Automation as Anti-Bias Intervention*, 41 Cardozo. L. Rev. 1671, 1690 (2020).

[8] Ifeoma Ajunwa, *The Quantified Worker* (forthcoming from Cambridge University Press, April 2023).

[9] Ifeoma Ajunwa, *The Quantified Worker* (forthcoming from Cambridge University Press, April 2023).

[10] Patricia G. Barnes, *Behind the Scenes, Discrimination by Job Search Engines*, AGE DISCRIMINATION EMP. (Mar. 29, 2017), <https://www.agediscriminationinemployment.com/behind-the-scenes-discrimination-by-job-search-engines>

[11] Ifeoma Ajunwa & Daniel Greene, *Platforms at Work: Data Intermediaries in the Organization of the Workplace*, in *WORK AND LABOR IN THE DIGITAL AGE* (2019)

[12] See generally Cathy O’Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (2016).

[13] “Akin to how phrenology sought to quantify human character through observable, physical traits, video interview technologies also seek to quantify and objectively understand human behavior as it relates to job success. An inherent underlying assumption of these technologies is that there exist observable, physical manifestations that give insight into the character and behavioral traits that define a successful individual. Video interviewing technology is purportedly motivated by objectivity, yet it ranks candidates based on judgments rooted in normative comparisons. The automated video interviewing algorithms are trained to search for certain traits deemed to be valuable, but these normative conclusions are based on samples of existing employees, and these samples are not random and may not be representative.” Ifeoma Ajunwa, *Automated Video Interviewing as the New Phrenology*, 36 *Berkeley Tech. L.J.* 1173, 1190 (2022). See also, Evan Selinger & Woodrow Hartzog, *The Inconsistency of Facial Surveillance*, 66

LOY. L. REV. 101, 105 (2019).

[14] Nilam Oswal, *The Latest Recruitment Technology Trends and How to Really Use Them*, PC WORLD (Feb. 9, 2018, 4:56 PM), <https://www.pcworld.idg.com.au/article/633219/latest-recruitment-technology-trends-how-really-use-them/>; *Gartner HR Survey Shows 86% of Organizations Are Conducting Virtual Interviews to Hire Candidates During Coronavirus Pandemic*, gartner: newsroom (Apr. 30, 2020), <https://www.gartner.com/en/newsroom/press-releases/2020-04-30-gartner-hr-survey-shows-86-of-organizations-are-cond>

[15] “Accent discrimination is a credible threat of automated video systems given that many video interview algorithms employ vocal analysis. In fact, a recent audit of HireVue’s algorithms suggest[ed] that accent discrimination may already be present in the company’s assessment outcomes.” Ifeoma Ajunwa, *Automated Video Interviewing as the New Phrenology*, 36 Berkeley Tech. L.J. 1173, 1199 (2022).

[16] In *Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices*, Raghavan and his coauthors point to “[a] wave of studies [which have] shown that several commercially available facial analysis techniques suffer from disparities in error rates across gender and racial lines. Ifeoma Ajunwa, *Automated Video Interviewing as the New Phrenology*, 36 Berkeley Tech. L.J. 1173, 1190 (2022); See also, Manish Raghavan, Solon Barocas, Jon Kleinberg & Karen Levy, *Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices*, 2020 conf. on fairness, accountability, and transparency (fat*) 469, 475 (2020); Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, 81 proc. mach. learning rsch. 1, 2 (2018); Lauren Rhue, *Racial Influence on Automated Perceptions of Emotions, race, ai, & emotions* 1, 1 (2018).

[17] Charles A. Sullivan, *Disparate Impact: Looking Past the Desert Palace Mirage*, 47 WM. & MARY L. REV. 911, 989 (2005).

[18] Ifeoma Ajunwa, *The Paradox of Automation as Anti-Bias Intervention*, 41 Cardozo. L. Rev. 1671, 1726 -1728 (2020).

[19] *Id.*

[20] Ifeoma Ajunwa, *The Auditing Imperative for Automated Hiring*, 34 Harv. J.L. & Tech. 621, 622 (2021).

[21] Ifeoma Ajunwa, *Automated Governance*, 101 N.C. L. Rev 355 (2023).

[22] I offer several mechanisms to serve as guardrails to automated governance. These mechanisms include: 1) Standing Advisory Council of Technologists and Social Scientists, 2) Stakeholder and Constituency Engagement, 3) Congressional Overview and Review. Ifeoma Ajunwa, *Automated Governance*, 101 N.C. L. Rev 355, 398-402 (2023).

[23] Ifeoma Ajunwa, *The Auditing Imperative for Automated Hiring*, 34 Harv. J.L. & Tech. 621, 675-677 (2021).

[24] Charles Sullivan, *Employing AI*, 63 *vill. l. rev.* 395, 423 (2018). *See also*, 29 C.F.R. § 1607.5B (2018).



U.S. Equal Employment Opportunity Commission

Testimony of Alex C. Engler

My name is Alex Engler, I am fellow at the Brookings Institution, an associate fellow at the Center for European Policy Studies, and an adjunct professor at Georgetown University. In these roles, I primarily study the interaction between algorithms and social policy. This research is informed by a decade of experience working as a data scientist in government, think tanks, and academia.

First, I would like to commend the EEOC on last year's technical assistance, detailing how AI hiring tools can be discriminatory against people with disabilities, and how they might comply with the Americans with Disabilities Act. The EEOC guidance is reasoned and well attuned in its underlying goal to meaningfully improve the market of AI hiring software for people with disabilities. I applaud this work, and will continue to hold it up as an example to other federal agencies, especially for how it considers the entire socio-technical process of hiring, and not just the algorithms alone.¹ Further, I commend the EEOC on providing guidance not just for non-discrimination, but also implementing principles of disclosure—that people with disabilities deserve to know when they are being evaluated by an algorithm—as well as the availability of a reasonable accommodation and/or alternative, non-algorithmic, process.

This work from EEOC is especially encouraging, because the story of AI hiring is not unique. Almost all critical decisions in the employment are experiencing 'algorithmization'—meaning the steady expansion of algorithms to more and more tasks. This now includes AI's application to targeted job ads, recruitment, background-checks, task allocation, evaluation of employee performance, wage-setting, promotion, at times even termination, and others.

Unfortunately, while many of these AI systems have significant value when used responsibly, they have too often been deployed with inflated promises and insufficient testing or worker protections. Much like AI hiring, this can lead to

discriminatory outcomes, worker disenfranchisement through to black-box AI decisions, and unjust decisions resulting from algorithmic mistakes.

The most comprehensive U.S. federal document on AI harms, the Blueprint for an AI Bill of Rights, states that these AI applications pose meaningful risks to equal opportunity, and warrant government scrutiny. The European Union's draft AI Act also recognizes this, and when it passes, it will categorize nearly all of these AI applications as "high-risk," and will create significant new regulatory requirements, as well as government enforcement capacity.

While AI hiring is perhaps the most visible in the media and best analyzed by academics and civil society, these other AI employment systems are used by thousands of businesses and affect millions of Americans. It is difficult to precisely interpret the limited survey evidence about the market penetration of AI employment tools. Still, the prevailing evidence suggests that, for medium- and large-businesses, algorithmic systems contribute significantly to, or perform outright, the majority of all employment decisions in the categories mentioned above.

That most employment decisions will be assisted by, or made by, an AI system is a sea change in the employer-employee relationship, and in turn, requires profound change at the EEOC. Continuing the work of the Artificial Intelligence and Algorithmic Fairness Initiative, the EEOC should systematically consider these AI applications, develop tailored guidance for each under all of the EEOC's legal authorities, and build necessary enforcement capacity.

I understand that this is an enormous undertaking, and that it will take time and resources. I also expect that, over time, it will affect the structure and core functions of the EEOC. While a great challenge, this is the appropriate response to the new algorithmic paradigm in employment.

Beyond new policy, the EEOC must also develop new capacity. An important takeaway from my research is that the transition to AI employment systems represents a possibility for a more fair and just labor market, these better outcomes are absolutely not guaranteed. In a Brookings Institution paper, I argue that the market incentives around AI hiring, are not sufficient to produce fair outcomes on their own. Further, an effective and independent auditing market that might self-regulate AI hiring systems will not emerge on its own, without any government enforcement.²

The European Union recognizes this challenge, and the EU AI Act will enable significant government oversight, notably requiring that developers to make available data and documentation to regulators, which will enable algorithmic audits, to ensure conformity with the EU AI requirements. Notably, the EU AI Act will also require registration of all covered AI employment systems in a public database, potentially leading to an informative resource for U.S. policymakers.

I was encouraged to see “Technology-related employment discrimination” mentioned in the EEOC’s Draft Strategic Enforcement Plan. In order to provide meaningful enforcement, the EEOC should actively build capacity, such as by hiring data scientists who specialize in regulatory compliance, as well as algorithmic auditors, who will be essential in the investigation and litigation of AI employment systems. Even before any specific enforcement actions, the EEOC should look to acquire and evaluate AI employment systems in order to improve public knowledge. This effort might be modeled after the National Institute for Standards and Technology’s Face Recognition Vendor Testing Program, which evaluates facial recognition software and publishes results from this testing. In total, this development of new EEOC capacity for algorithmic oversight will be as critical as the development of policy guidance and technical assistance.

To summarize, I urge the EEOC to:

- 1. Consider a wide range of AI employment systems, not just in hiring, but also targeted job ads, recruitment, task allocation, evaluation of employee performance, wage setting, promotion, and termination.*
- 2. Encourage and enforce the whole range of AI principles on these AI employment systems, as advocated in exemplar policy documents, such as the Blueprint for an AI Bill of Rights or the EU AI Act, to the extent possible under EEOC’s legal authorities.*
- 3. Develop the capacity to provide oversight, such as by using investigations to audit these critical AI systems and ensure their compliance with federal law, as well as to use information gathering authorities to inform the EEOC and the public on their proliferation and impact.*

¹ Alex C. Engler, “The EEOC wants to make AI hiring fairer for people with disabilities.” May 26th, 2022.

<https://www.brookings.edu/blog/techtank/2022/05/26/the-eeoc-wants-to-make-ai-hiring-fairer-for-people-with-disabilities/>
[\(https://www.brookings.edu/blog/techtank/2022/05/26/the-eeoc-wants-to-make-ai-hiring-fairer-for-people-with-disabilities/\)](https://www.brookings.edu/blog/techtank/2022/05/26/the-eeoc-wants-to-make-ai-hiring-fairer-for-people-with-disabilities/)
[disabilities/](https://www.brookings.edu/blog/techtank/2022/05/26/the-eeoc-wants-to-make-ai-hiring-fairer-for-people-with-disabilities/)
[\(https://www.brookings.edu/blog/techtank/2022/05/26/the-eeoc-wants-to-make-ai-hiring-fairer-for-people-with-disabilities/\)](https://www.brookings.edu/blog/techtank/2022/05/26/the-eeoc-wants-to-make-ai-hiring-fairer-for-people-with-disabilities/)

² Alex C. Engler, “Auditing employment algorithms for discrimination.” March 12th, 2021. <https://www.brookings.edu/research/auditing-employment-algorithms-for-discrimination/> (<http://www.brookings.edu/research/auditing-employment-algorithms-for-discrimination/>)