



Artificial Intelligence (AI)— Emerging Issues and Developments with the Use of AI in the Workplace and Labor Negotiations

**1 General CLE credit
ID 108028 (part of 105093)**

**From the Oregon State Bar CLE seminar *Labor and Employment
Annual Conference—Finding a New Balance: Navigating the
Intersection of Business and Individual Workplace Needs*,
presented on October 13, 2023**

Chapter 6

Employment Law Issues with Artificial Intelligence

MATTHEW SCHERER
Senior Policy Counsel
Workers' Rights and Technology Policy
Center for Democracy & Technology
Portland, Oregon

Contents

Presentation Slides: Employment Law Issues with Artificial Intelligence—What Practitioners Should Know About the Legal Landscape Surrounding New Technologies	6–1
S. 2419—No Robot Bosses Act	6–13
Vermont H.114	6–61
“NY City Council Rams Through Once-Promising but Deeply Flawed Bill on AI Hiring Tools”. . .	6–75
“Civil Rights Standards for 21st Century Employment Selection Procedures”.	6–85
Center for Industry Self-Regulation “Principles for Trustworthy AI in Recruiting & Hiring”. . . .	6–125
Center for Industry Self-Regulation “Independent Certification Protocols for AI-Enabled Hiring & Recruiting Technologies”.	6–133
<i>Mobley v. Workday</i> Complaint	6–143
<i>Mobley v. Workday</i> Motion to Dismiss	6–159
EEOC Press Release: iTutor Group to Pay \$365,000 to Settle EEOC Discriminatory Hiring Suit	6–191
EEOC Strategic Enforcement Plan Excerpt	6–193
“Applying Old Rules to New Tools: Employment Discrimination Law in the Age of Algorithms” .	6–197



Oregon State Bar Labor & Employment Annual Conference

Employment Law Issues With Artificial Intelligence

What practitioners should know about the legal landscape surrounding new technologies

Matthew Scherer, Senior Policy Counsel for Workers' Rights and Technology

October 13, 2023

What is AI?

Machine Learning

- machines using training data to make better predictions

Predictive AI

- Analyzes training data to predict future outcomes

Generative AI

- Creates new content based on training data



Why are employers using AI?



Increase Efficiency

Saving time to review applications, resumes, and cover letters



Monitor/analyze workforce

Using data analysis tools to track and manage workers



Diversify Candidates

Identifying passive candidates



Examples of Automated Employment Decision Tools (AEDTs)

- ★ Personality and aptitude tests
 - Candidates respond to questions/images or play a video game
 - Responses/performance used to infer personality traits
- ★ Resume screening tools
 - Search for specific terms or combinations of terms in resumes to identify candidates who may have job-relevant characteristics
- ★ Video interview analysis
 - Candidates recorded while answering questions
 - Responses (and sometimes their body language, facial expressions, etc) are then analyzed



What is ESAM / Bossware?

- ★ **Refers to employers' use of two different types of tech:**
 - *Electronic surveillance*: Track workers' activities, physical movements, and pace of work continuously and in unprecedented detail
 - *Algorithmic management*: Automate tasks traditionally employed by human managers, including assessing workers' productivity and performance and making disciplinary decisions, often leaving workers with little recourse.
- ★ **Concerning trends in employers' use of bossware**
 - Increase pace of work
 - Eliminate breaks and downtime
 - Automate discipline and other personnel management functions



Risks when using AI in the workplace

How bossware can threaten workers' health and safety

- ★ **Discouraging and penalizing lawful, health-enhancing employee conduct**
 - Namely, taking breaks to rest or use toilet facilities
- ★ **Enforcing faster work pace**
 - Increases risk of physical injuries, particularly those stemming from fatigue and repetitive motion
- ★ **Increasing job strain**
 - Job strain occurs when workers face high job demands but low job control
 - Decades of research shows negative mental and physical health consequences



Risks when using AI in the workplace

How AEDTs can lead to discrimination

- ★ “Algorithmic bias” is the umbrella term for systemic errors that create unfair outcomes for certain individual or groups
- ★ Bias can stem from a number of sources, each of which can overlap and amplify each other to generate discriminatory outputs or reinforce patterns of underrepresentation:
 - Historical patterns of discrimination or underrepresentation may be baked into the data on which a machine-learning tool is trained
 - Developers may fail to properly clean or weight the data, resulting in an algorithm that discriminatory hiring/decision practices
 - The tool may be deployed/operationalized in a way that leads to biased outcomes--e.g., a tool that is inaccessible to disabled workers or for which need for accommodation is not considered



Risks when using AI in the workplace

How AEDTs can lead to discrimination

- ★ Instead of identifying ability to perform essential functions, the tools may measure:
 - Personality traits and aptitudes typical but not necessary for position
 - Attributes that appear most frequently in resumes of successful workers
 - Personality traits and aptitudes based on movements, vocal intonation, speech patterns
- ★ Results can be affected by gender and cultural norms and disability



Risks when using AI in the workplace

Accessibility and accommodation concerns

- ★ Tools may not be provided in an accessible format, and may not be able to communicate with a person about format's inaccessibility
- ★ Lack of insight into how tools work and will influence employment decisions
 - Hard to detect whether tools may have made an adverse decision discriminatory
 - Hard to anticipate if accommodations are needed



Risks when using AI in the workplace

Privacy

- ★ Increasing number of laws impose restrictions or disclosure requirements on companies that process personal data
- ★ The more types of personal information employers collect for AEDT uses, the more sources of risk they take on
 - Esp when acquired from third parties, but many types of data trigger privacy concerns regardless of how it's obtained
- ★ High-risk sources include:
 - Health / biometric data
 - Social media data
 - Financial / credit data
 - Criminal history



The Policy Landscape

Federal Developments: Administrative

- ★ Interagency
 - Multiple interagency statements involving DOJ, DOL, EEOC, and FTC stating that they plan to coordinate future enforcement efforts
- ★ EEOC
 - May 2022: Guidance on disability discrimination in relation to the use of AI-based selection tools
 - January 2023: Hearing focused on civil rights implications of AEDTs
 - May 2023: Guidance on Title VII in the same vein
 - Sept 2023: Strategic Enforcement Plan released that repeatedly highlighted algorithmic discrimination as an area of regulatory focus in coming years



The Policy Landscape

Federal Developments: Administrative

- ★ DOJ
 - May 2022: Issued guidance on AEDTs and disability discrimination, mirroring EEOC's
- ★ OFCCP
 - Updated scheduling letter asks contractors to: “Identify and provide information and documentation of policies, practices, or systems used to recruit, screen, and hire, including the use of artificial intelligence, algorithms, automated systems or other technology-based selection procedures.”



The Policy Landscape

Federal Developments: Legislative

- ★ Stop Spying Bosses Act (S. 262)
 - Introduced by Sen. Casey (D-PA)
 - Would prohibit employers from:
 - Surveilling workers engaged in union or protected labor activity;
 - Collecting health or disability data;
 - Using an automated system to predict non-job-related behavior;
 - Monitoring workers while they are off-duty;
 - Using surveillance in a manner that threatens workers' health; or
 - Selling/licensing worker data to third parties
 - Would require companies to disclose key aspects of surveillance practices



The Policy Landscape

Federal Developments: Legislative

- ★ No Robot Bosses Act (S. 2419)
 - Would require employers to:
 - Disclose use of AEDTs to candidates, including descriptions of what AEDT purports to measure, how it measures them, and how the characteristics relate to job functions;
 - Engage in robust auditing and discrimination testing; and
 - Give candidates right to appeal AEDT decisions
- ★ Other developments
 - Sens. Blumenthal and Hawley introduced bipartisan framework on AI
 - Majority Leader Schumer conducting listening sessions on AI regulation
 - HELP Chair Sen. Sanders also quietly looking to shape legislation



The Policy Landscape

State/Local Developments: Existing privacy legislation

- ★ Current legislation
 - California Consumer Privacy Act (CCPA)
 - Already in partial effect; full effect by March 2024
 - Imposes notice, access, and right-to-delete requirements on many types of worker data that companies collect
 - Other states with data privacy laws (mostly NOT applicable workers)
 - In effect: CA, CO, CT, VA
 - Going into effect by end of year: UT
 - Passed in 2023: DE, FL, IA, IN, MT, **OR**, TN, TX

- ★ And if your practice touches Europe, the GDPR (and its UK analogue) are even more intense



The Policy Landscape

State/Local Developments: Existing AEDT legislation

- ★ Illinois AI Video Interview Act
 - Technically exists and requires disclosure when a company uses AI to analyze a recorded video interview, but it has no enforcement mechanism, so it can be safely ignored
- ★ NYC Local Law 144
 - Hyped up bill that is, in fact, very narrow and imposes almost no new requirements
 - Applies only to tools that effectively replace human decision-making
 - Requires minimal disclosure to candidates before using AEDT
 - Technically requires performing a disparate impact test--but one that federal regulations already effectively require employers to do



The Policy Landscape

State/Local Developments: Proposed legislation and regulations

- ★ California
 - AB 331 would have required companies to perform impact assessments on AEDTs and workplace surveillance technologies, but bill appears dead
 - Civil Rights Council (formerly DFEH) published draft regulations that would have notably clarified and tightened rules surrounding AEDTs, but they appear in limbo
- ★ Other states (*though none of these bills appear close to passage*)
 - NY/NJ both considering bills modeled on (stronger version of) NYC LL 144
 - DC: Stop Discrimination by Algorithms Act would impose significant notice and audit requirements
 - MA / VT: Bills targeting both AEDTs and workplace surveillance



The Policy Landscape

Civil Society Proposals: Civil Rights Standards

Full title: *Civil Rights Standards for 21st Century Employment Selection Procedures*

- ★ Published last year and endorsed by more than a dozen civil rights and workers' rights organizations, incl. CDT, AAPD, ACLU, Color of Change, Leadership Conference, NWLC, and Upturn
 - Expands on 2020 *Civil Rights Principles for Hiring Assessment Technologies*
- ★ Provide a strong, civil-rights-focused roadmap to fairness in hiring:
 - Detecting and preventing discrimination before a tool is used on actual candidates
 - Transparency requirements for use of the tools
 - Accountability for when tools go awry



The Policy Landscape

Civil Society Proposals: BBB-National Programs *Principles and Protocols*

- ★ Two documents
 - *Principles for Trustworthy AI in Recruiting and Hiring*
 - Provides “practical baseline standards for trusted business practices” surrounding development and deployment of AI in hiring
 - *Independent Certification Protocols for AI-Enabled Hiring and Recruiting Technologies*
 - Framework for (self-)certification regime for compliance with *Principles*
- ★ Goal was to establish a set of voluntary standards for auditing and disclosure for companies to comply with
- ★ Focus is more on technical auditing than on candidate disclosure and fairness/bias concerns.
- ★ It would require companies to be sure they are complying with all antidiscrimination laws



The Policy Landscape

Civil Society Proposals: Proposals focused on technology and privacy rights

- ★ Worker Privacy Act (Georgetown Center on Privacy & Technology)
 - Aims to “prevent invasive worker data collection and increase worker control over workplace data” by...
 - Increasing employee input in decisions around data and surveillance
 - Establishing clear prohibitions on the use of data
 - Creating a dedicated office to worker privacy
- ★ Berkeley Labor Center framework on workplace technology rights
 - Strict “data minimization requirements”--data collection/monitoring can only be done for specific purposes, and must be narrowly tailored
 - Strong notice/disclosure requirements
 - Bars uses of surveillance that harm workers’ health, safety, rights



Employer recommendations

- ★ Keep your ear close to the ground on legal developments
 - AI regulation is a hot topic, states are getting increasingly active in privacy legislation, and workers' rights is an increasingly sympathetic issue for both parties
- ★ Do your due diligence when acquiring AI or bossware tools—especially AEDTs
 - Vendors are massively overhyping tools' capabilities and underplaying their legal risks
- ★ Don't just adopt new technologies because you can. Be wary of vendors...
 - Promising productivity gains--but only by using tools that will stress workers' minds and bodies
 - Claiming that they have miraculously automated decision processes that humans still struggle with



THANK YOU | CDT.ORG



118TH CONGRESS
1ST SESSION

S. 2419

To prohibit certain uses of automated decision systems by employers, and
for other purposes.

IN THE SENATE OF THE UNITED STATES

JULY 20, 2023

Mr. CASEY (for himself, Mr. SCHATZ, Mr. FETTERMAN, and Mr. SANDERS)
introduced the following bill; which was read twice and referred to the
Committee on Health, Education, Labor, and Pensions

A BILL

To prohibit certain uses of automated decision systems by
employers, and for other purposes.

1 *Be it enacted by the Senate and House of Representa-*
2 *tives of the United States of America in Congress assembled,*

3 **SECTION 1. SHORT TITLE.**

4 This Act may be cited as the “No Robot Bosses Act”.

5 **SEC. 2. DEFINITIONS.**

6 For purposes of this Act:

7 (1) **AUTOMATED DECISION SYSTEM.**—

8 (A) **IN GENERAL.**—The term “automated
9 decision system” means a system, software, or
10 process that—

(i) uses computation, in whole or in part, to determine outcomes, make or aid decisions (including through evaluations, metrics, or scoring), inform policy implementation, or collect data or observations, including such a system, software, or process derived from machine learning, statistics, or other data processing or artificial intelligence techniques; and

(ii) is not passive computing infrastructure.

(B) PASSIVE COMPUTING INFRASTRUCTURE.—For purposes of this paragraph, the term “passive computing infrastructure” means any intermediary technology that does not influence or determine the outcome of a decision, make or aid in a decision (including through evaluations, metrics, or scoring), inform policy implementation, or collect data or observations, including web hosting, domain registration, networking, caching, data storage, or cybersecurity.

(2) AUTOMATED DECISION SYSTEM OUTPUT.—The term “automated decision system output” means any information, assumption, prediction,

1 score, recommendation, decision, evaluation, metric,
2 conclusion, inference, or profile generated by an
3 automated decision system.

4 (3) CANDIDATE.—The term “candidate”, with
5 respect to an employer, means an individual who ap-
6 plies, or applied, to be employed by, or otherwise
7 perform work for remuneration for, the employer.

8 (4) COVERED INDIVIDUAL.—The term “covered
9 individual”, with respect to an employer, means an
10 individual—

11 (A) who is employed by, or otherwise per-
12 forming work for remuneration for, the em-
13 ployer, including such an individual who is—

14 (i) any individual performing work for
15 remuneration for an employer described in
16 clauses (i)(I) and (ii) of paragraph (6)(A);

17 (ii) any individual performing work
18 for remuneration for an entity described in
19 paragraph (6)(A)(i)(II);

20 (iii) any individual performing work
21 for remuneration for an employing office
22 described in paragraph (6)(A)(i)(III);

23 (iv) any individual performing work
24 for remuneration for an employing office
25 described in paragraph (6)(A)(i)(IV); or

1 (v) any individual performing work for
2 remuneration for an employing agency de-
3 scribed in paragraph (6)(A)(i)(V); or
4 (B) who is a candidate with respect to the
5 employer.

6 (5) EMPLOY.—The term “employ” has the
7 meaning given such term in section 3 of the Fair
8 Labor Standards Act of 1938 (29 U.S.C. 203).

9 (6) EMPLOYER.—

10 (A) IN GENERAL.—The term “employer”
11 means any person who is—

12 (i)(I) a covered employer who is not
13 described in any other subclause of this
14 clause;

15 (II) an entity employing a State em-
16 ployee described in section 304(a) of the
17 Government Employee Rights Act of 1991
18 (42 U.S.C. 2000e–16c(a));

19 (III) an employing office, as defined
20 in section 101 of the Congressional Ac-
21 countability Act of 1995 (2 U.S.C. 1301);

22 (IV) an employing office, as defined in
23 section 411(c) of title 3, United States
24 Code; or

1 (V) an employing agency covered
2 under subchapter V of chapter 63 of title
3 5, United States Code; and

4 (ii) engaged in commerce (including
5 government), or an industry or activity af-
6 fecting commerce (including government).

7 (B) COVERED EMPLOYER.—In subpara-
8 graph (A), the term “covered employer”—

9 (i) means any person engaged in com-
10 merce or in any industry or activity affect-
11 ing commerce who employs, or otherwise
12 engages for the performance of work for
13 remuneration, 11 or more covered individ-
14 uals;

15 (ii) includes—

16 (I) any person who acts, directly
17 or indirectly, in the interest of a cov-
18 ered employer in relation to any indi-
19 vidual performing work for remunera-
20 tion for such covered employer;

21 (II) any successor in interest of a
22 covered employer;

23 (III) any public agency; and

1 (IV) the Government Account-
2 ability Office and the Library of Con-
3 gress; and

4 (iii) does not include any labor organi-
5 zation (other than when acting as an em-
6 ployer) or anyone acting in the capacity of
7 officer or agent of such labor organization.

8 (C) PUBLIC AGENCY.—For purposes of
9 this paragraph, a public agency shall be consid-
10 ered to be a person engaged in commerce or in
11 an industry or activity affecting commerce.

12 (D) DEFINITIONS.—For purposes of this
13 paragraph, the terms “commerce”, “person”,
14 and “public agency” have the meanings given
15 the terms in section 3 of the Fair Labor Stand-
16 ards Act of 1938 (29 U.S.C. 203).

17 (7) EMPLOYMENT-RELATED DECISION.—The
18 term “employment-related decision” includes a deci-
19 sion by an employer with regard to—

20 (A) hiring a covered individual (including
21 any decision with regard to recruiting, screen-
22 ing, interviewing, or selecting a candidate);

23 (B) firing, retaining, taking a disciplinary
24 action against, demoting, or reassigning duties
25 of a covered individual; or

1 (C) any other term, condition, or privilege
2 of employment or work of the covered indi-
3 vidual, such as relating to pay, scheduling,
4 health care or long-term care coverage, benefits,
5 or hours worked or promoting a covered indi-
6 vidual.

7 (8) GOVERNMENT ENTITY.—The term “govern-
8 ment entity” means—

9 (A) a Federal agency (as such term is de-
10 fined in section 3371 of title 5, United States
11 Code);

12 (B) a State or political subdivision thereof;

13 (C) any agency, authority, or instrumen-
14 tality of a State or political subdivision thereof;
15 or

16 (D) a Tribal government or political sub-
17 division thereof.

18 (9) INDIAN TRIBE.—The term “Indian Tribe”
19 means any Indian or Alaska Native tribe, band, na-
20 tion, pueblo, village, community, component band, or
21 component reservation individually identified (includ-
22 ing parenthetically) in the list published most re-
23 cently as of the date of enactment of this Act pursu-
24 ant to section 104 of the Federally Recognized In-
25 dian Tribe List Act of 1994 (25 U.S.C. 5131).

1 (10) LABOR ORGANIZATION.—The term “labor
2 organization” has the meaning given the term in
3 section 2(5) of the National Labor Relations Act (29
4 U.S.C. 152(5)), except that such term shall also in-
5 clude—

6 (A) any organization composed of labor or-
7 ganizations, such as a labor union federation or
8 a State or municipal labor body; and

9 (B) any organization which would be in-
10 cluded in the definition for such term under
11 such section 2(5) but for the fact that the orga-
12 nization represents—

13 (i) individuals employed by the United
14 States, any wholly owned Government cor-
15 poration, any Federal Reserve Bank, or
16 any State or political subdivision thereof;

17 (ii) individuals employed by persons
18 subject to the Railway Labor Act (45
19 U.S.C. 151 et seq.); or

20 (iii) individuals employed as agricul-
21 tural laborers.

22 (11) PREDISPUTE ARBITRATION AGREEMENT.—
23 The term “predispute arbitration agreement” means
24 any agreement to arbitrate a dispute that has not

1 yet arisen at the time of the making of the agree-
2 ment.

3 (12) PREDISPUTE JOINT-ACTION WAIVER.—The
4 term “predispute joint-action waiver” means an
5 agreement, whether or not part of a predispute arbi-
6 tration agreement, that would prohibit, or waive the
7 right of, one of the parties to the agreement to par-
8 ticipate in a joint, class, or collective action in a ju-
9 dicial, arbitral, administrative, or other forum, con-
10 cerning a dispute that has not yet arisen at the time
11 of the making of the agreement.

12 (13) SECRETARY.—The term “Secretary”
13 means the Secretary of Labor.

14 (14) STATE.—The term “State” means each of
15 the several States of the United States, the District
16 of Columbia, or any territory or possession of the
17 United States.

18 (15) STATE ATTORNEY GENERAL.—The term
19 “State attorney general” means—

20 (A) with respect to a State, the attorney
21 general or chief law enforcement officer of the
22 State, or another official or agency designated
23 by the State to bring civil actions on behalf of
24 the State or the residents of the State; and

1 (B) with respect to a Tribal government,
2 the attorney general or chief law enforcement
3 officer of the Tribal government, or another of-
4 ficial or agency designated by the Tribal gov-
5 ernment to bring civil actions on behalf of the
6 Tribal government or the Indian Tribe of the
7 Tribal government.

8 (16) STATE PRIVACY REGULATOR.—The term
9 “State privacy regulator” means—

10 (A) the chief consumer protection officer of
11 a State; or

12 (B) a State consumer protection agency
13 with expertise in data protection, including the
14 California Privacy Protection Agency.

15 (17) TRIBAL GOVERNMENT.—The term “Tribal
16 government” means the recognized governing body
17 of an Indian Tribe.

18 **SEC. 3. USE OF AN AUTOMATED DECISION SYSTEM BY AN**
19 **EMPLOYER.**

20 (a) EMPLOYMENT-RELATED DECISIONS.—

21 (1) IN GENERAL.—An employer—

22 (A) may not rely exclusively on an auto-
23 mated decision system in making an employ-
24 ment-related decision with respect to a covered
25 individual; and

1 (B) may not use an automated decision
2 system output in making an employment-re-
3 lated decision with respect to a covered indi-
4 vidual unless—

5 (i) the automated decision system
6 used to generate such automated decision
7 system output has had pre-deployment
8 testing and validation with respect to—

9 (I) the efficacy of the system;

10 (II) the compliance of the system
11 with applicable employment discrimi-
12 nation laws, including—

13 (aa) title VII of the Civil
14 Rights Act of 1964 (42 U.S.C.
15 2000e et seq.);

16 (bb) the Age Discrimination
17 in Employment Act of 1967 (29
18 U.S.C. 621 et seq.);

19 (cc) title I of the Americans
20 with Disabilities Act of 1990 (42
21 U.S.C. 12111 et seq.);

22 (dd) title II of the Genetic
23 Information Nondiscrimination
24 Act of 2008 (42 U.S.C. 2000ff et
25 seq.);

12

1 (ee) section 6(d) of the Fair
2 Labor Standards Act of 1938 (29
3 U.S.C. 206(d));
4 (ff) sections 501 and 505 of
5 the Rehabilitation Act of 1973
6 (29 U.S.C. 791; 793); and
7 (gg) the Pregnant Workers
8 Fairness Act (division II of the
9 Consolidated Appropriations Act,
10 2023 (Public Law 117–328));
11 (III) the lack of any potential
12 discriminatory impact of the system,
13 including discriminatory impact based
14 on race, color, religion, sex (including
15 pregnancy, sexual orientation, or gen-
16 der identity), national origin, age, or
17 disability and genetic information (in-
18 cluding family medical history); and
19 (IV) the compliance of the sys-
20 tem with the Artificial Intelligence
21 Risk Management Framework re-
22 leased by the National Institute of
23 Standards and Technology on Janu-
24 ary 26, 2023, or successor framework;

1 (ii) such automated decision system
2 is, not less than annually, independently
3 tested for discriminatory impact described
4 in clause (i)(III) or potential biases and
5 the results of such test are made publicly
6 available;

7 (iii) the employer has provided the
8 disclosure required under paragraph (2)
9 with respect to such use of an automated
10 decision system output;

11 (iv) such use is designed for purposes
12 of making such an employment-related de-
13 cision;

14 (v) the employer independently cor-
15 roborates, via meaningful oversight by a
16 human with appropriate and relevant expe-
17 rience, such automated decision system
18 output;

19 (vi) not later than 7 days after mak-
20 ing such an employment-related decision,
21 the employer provides full, accessible, and
22 meaningful documentation in plain lan-
23 guage to such covered individual (at no
24 cost to such covered individual) on the

1 automated decision system output, includ-
2 ing—

3 (I) a description of the auto-
4 mated decision system used to gen-
5 erate such automated decision system
6 output;

7 (II) a description and expla-
8 nation, in plain language, of the input
9 data to such automated decision sys-
10 tem used to generate such automated
11 decision system output and a ma-
12 chine-readable copy of such data;

13 (III) a description and expla-
14 nation of how such automated deci-
15 sion system output was used in mak-
16 ing such employment-related decision;
17 and

18 (IV) the reasoning for the use of
19 such automated decision system out-
20 put in such employment-related deci-
21 sion; and

22 (vii) the employer enables the covered
23 individual to, after receiving such docu-
24 mentation—

1 (I) dispute (in a manner that is
2 accessible and equitable and does not
3 pose an unreasonable burden on the
4 covered individual) such automated
5 decision system output to a human
6 with appropriate and relevant experi-
7 ence; and

8 (II) appeal such employment-re-
9 lated decision to a human with appro-
10 priate and relevant experience who is
11 not the human for purposes of the
12 corroboration under clause (v).

13 (2) DISCLOSURE.—

14 (A) IN GENERAL.—An employer that uses
15 or intends to use an automated decision system
16 output in making an employment-related deci-
17 sion with respect to a covered individual shall,
18 in accordance with subparagraph (B), disclose
19 to such covered individual—

20 (i) that the employer uses or intends
21 to use an automated decision system out-
22 put in making such an employment-related
23 decision;

24 (ii) a description and explanation of
25 the automated decision system used or in-

- 1 tended to be used to generate such auto-
- 2 mated decision system output, including—
- 3 (I) the types of data collected or
- 4 intended to be collected as inputs to
- 5 the automated decision system and
- 6 the circumstances of such collection;
- 7 (II) the characteristics that the
- 8 automated decision system measures
- 9 or is intended to measure, such as the
- 10 knowledge, skills, or abilities of the
- 11 covered individual;
- 12 (III) how such characteristics re-
- 13 late or would relate to any function
- 14 required for the work or potential
- 15 work of the covered individual;
- 16 (IV) how the system measures or
- 17 is intended to measure such charac-
- 18 teristics; and
- 19 (V) how the covered individual
- 20 can interpret the automated decision
- 21 system output in plain language;
- 22 (iii) the identity of the individual or
- 23 entity that operates the automated decision
- 24 system that provides such an automated
- 25 decision system output;

1 (iv) how the employer uses or intends
2 to use such an automated decision system
3 output in making such an employment-re-
4 lated decision; and

5 (v) how the covered individual may
6 dispute or appeal an employment-related
7 decision made with respect to the covered
8 individual using an automated decision sys-
9 tem output.

10 (B) TIMING OF NOTICE.—

11 (i) INITIAL DISCLOSURE.—An em-
12 ployer shall provide the disclosure required
13 under subparagraph (A)—

14 (I) in the case of a covered indi-
15 vidual for whom an employment-re-
16 lated decision with regard to the hir-
17 ing of the covered individual—

18 (aa) was made before the
19 date of enactment of this Act, to
20 the covered individual not later
21 than 30 days after such date of
22 enactment; or

23 (bb) is made on or after the
24 date of enactment of this Act, to
25 the covered individual, except as

1 provided in subclause (II), prior
2 to making such employment-re-
3 lated decision; and

4 (II) in the case of a candidate
5 who applies to the employer on or
6 after the date of enactment of this
7 Act, prior to accepting an application
8 by the candidate to be employed by,
9 or otherwise perform work for remu-
10 nation for, the employer.

11 (ii) SUBSEQUENT DISCLOSURES.—Not
12 later than 30 days after any information
13 provided by an employer to a covered indi-
14 vidual through a disclosure required under
15 clause (ii) or (iv) of subparagraph (A) sig-
16 nificantly changes or after any significant
17 new information required to be provided in
18 such a disclosure becomes available, the
19 employer shall provide the covered indi-
20 vidual with an updated disclosure.

21 (3) TRAINING.—An employer that uses or in-
22 tends to use an automated decision system output in
23 making an employment-related decision with respect
24 to a covered individual shall train any individual or
25 entity that operates the automated decision system

1 that provides such an automated decision system
2 output or uses such automated decision system out-
3 put on the use of such system, including on—

4 (A) the input information used by such
5 automated decision system;

6 (B) the appeals process for such an auto-
7 mated decision system output;

8 (C) potential biases in automated decision
9 systems;

10 (D) any limitations of the automated deci-
11 sion system;

12 (E) any potential adverse effects to covered
13 individuals due to the automated decision sys-
14 tem;

15 (F) any potential errors or problems re-
16 lated to the automated decision system; and

17 (G) examples of inappropriate uses of the
18 automated decision system.

19 (b) MANAGEMENT BY AN AUTOMATED DECISION
20 SYSTEM.—An employer that manages a covered individual
21 through an automated decision system shall enable the
22 covered individual to opt out of such management and be
23 managed through a human manager who is able to make
24 employment-related decisions with respect to the covered
25 individual.

1 **SEC. 4. ESTABLISHMENT OF TECHNOLOGY AND WORKER**
2 **PROTECTION DIVISION.**

3 (a) IN GENERAL.—There is established in the De-
4 partment of Labor the Technology and Worker Protection
5 Division.

6 (b) ADMINISTRATOR OF THE TECHNOLOGY AND
7 WORKER PROTECTION DIVISION.—The President shall
8 appoint an Administrator of the Technology and Worker
9 Protection Division to head the Technology and Worker
10 Protection Division.

11 (c) EMPLOYEES AND ADVISORY BOARDS OF THE DI-
12 VISION.—

13 (1) IN GENERAL.—The Administrator—

14 (A) may select, appoint, and employ, with-
15 out regard to the provisions of sections 3309
16 through 3318 of title 5, United States Code, in-
17 dividuals, including technologists, directly to po-
18 sitions in the competitive service, as defined in
19 section 2102 of such title, to carry out the du-
20 ties of the Administrator under this Act; and

21 (B) may fix the compensation of the indi-
22 viduals described in subparagraph (A) without
23 regard to chapter 51 and subchapter III of
24 chapter 53 of title 5, United States Code, relat-
25 ing to classification of positions and General
26 Schedule pay rates, except that the rate of pay

1 for such individuals may not exceed the rate
2 payable for level V of the Executive Schedule
3 under section 5316 of that title.

4 (2) ADVISORY BOARDS.—

5 (A) ESTABLISHMENT.—The Administrator
6 shall establish the following advisory boards to
7 advise and consult with in the exercise of the
8 functions of the Administrator under this Act
9 and to provide information on emerging prac-
10 tices relating to the treatment of data by em-
11 ployers:

12 (i) The User Advisory Board, which
13 shall be comprised of experts in consumer
14 protection, privacy, civil rights, disability
15 rights, labor organizations, and ethics.

16 (ii) The Research Advisory Board,
17 which shall be comprised of individuals
18 with academic and research expertise in
19 privacy, cybersecurity, computer science,
20 innovation, design, ethics, economics, civil
21 rights law, disability law, labor organiza-
22 tions and public policy and representatives
23 of labor organizations.

24 (iii) The Product Advisory Board,
25 which shall be comprised of technologists,

1 computer scientists, designers, product
2 managers, attorneys, representatives of
3 labor organizations, workplace technology
4 experts, accessibility experts, and other
5 representatives of employers and employ-
6 ees.

7 (iv) The Labor Advisory Board, which
8 shall be comprised of representatives of
9 labor organizations and representatives of
10 workers.

11 (B) APPOINTMENTS.—The Administrator
12 shall appoint members to the advisory boards
13 established under subparagraph (A) without re-
14 gard to party affiliation.

15 (C) MEETINGS.—Each advisory board es-
16 tablished under subparagraph (A) shall meet—

17 (i) at the call of the Administrator;
18 and

19 (ii) not less than 2 times annually.

20 (D) COMPENSATION AND TRAVEL EX-
21 PENSES.—A member of an advisory board es-
22 tablished under subparagraph (A) who is not an
23 officer or employee of the Federal Government
24 shall—

1 (i) be entitled to receive compensation
2 at a rate fixed by the Administrator while
3 attending meetings of the advisory board,
4 including travel time; and

5 (ii) receive travel expenses, including
6 per diem in lieu of subsistence, in accord-
7 ance with applicable provisions under sub-
8 chapter I of chapter 57 of title 5, United
9 States Code.

10 (E) EXEMPTION FROM THE FEDERAL AD-
11 VISORY COMMITTEE ACT.—Each advisory board
12 established under subparagraph (A) shall be ex-
13 empt from chapter 10 of title 5, United States
14 Code.

15 (3) USE OF VOLUNTARY SERVICES.—The Ad-
16 ministrator may, as may from time to time be need-
17 ed, use any voluntary or uncompensated services.

18 (4) ATTORNEYS.—Attorneys appointed under
19 this subsection may appear for and represent the
20 Administrator in any litigation.

21 (d) OFFICES.—

22 (1) IN GENERAL.—The principal office of the
23 Technology and Worker Protection Division shall be
24 in the District of Columbia.

1 (2) REGIONAL, LOCAL, AND OTHER OFFICES.—
2 The Administrator may establish regional, local, or
3 other offices, including an office in the city of San
4 Francisco, California or the San Francisco Bay area
5 in California.

6 **SEC. 5. REGULATIONS.**

7 (a) IN GENERAL.—

8 (1) AUTHORITY.—

9 (A) IN GENERAL.—Except as provided in
10 paragraph (2), the Secretary, acting through
11 the Administrator, may prescribe such regula-
12 tions as may be necessary to carry out this Act
13 with respect to covered individuals described in
14 section 2(4)(A) (other than covered individuals
15 described in clauses (iii) through (v) of such
16 section) and other individuals affected by em-
17 ployers described in subclause (I) or (II) of sec-
18 tion 2(6)(A)(i), including individuals who are
19 covered individuals described in section 2(4)(B)
20 with respect to such employers.

21 (B) CONSULTATION.—In prescribing any
22 regulations authorized under this paragraph,
23 the Secretary, acting through the Adminis-
24 trator, may consult with—

1 (i) Federal agencies that have juris-
2 diction over Federal privacy laws or exper-
3 tise in privacy, including the Federal
4 Trade Commission;

5 (ii) Federal agencies that have juris-
6 diction over labor and employment issues,
7 including the Equal Employment Oppor-
8 tunity Commission, the National Science
9 Foundation, and the National Labor Rela-
10 tions Board; and

11 (iii) any other Federal agencies with
12 relevant expertise, including the United
13 States Access Board and the Office of
14 Science and Technology Policy.

15 (2) GOVERNMENT ACCOUNTABILITY OFFICE; LI-
16 BRARY OF CONGRESS.—The Comptroller General of
17 the United States and the Librarian of Congress
18 shall prescribe any regulations described in para-
19 graph (1)(A) with respect to covered individuals of
20 the Government Accountability Office and the Li-
21 brary of Congress, respectively, and other individuals
22 affected by the Comptroller General of the United
23 States and the Librarian of Congress, respectively.

24 (b) EMPLOYEES COVERED BY CONGRESSIONAL AC-
25 COUNTABILITY ACT OF 1995.—

1 (1) AUTHORITY.—Not later than 45 days after
2 the Secretary prescribes any regulation under sub-
3 section (a)(1)(A), the Board of Directors of the Of-
4 fice of Compliance shall prescribe (in accordance
5 with section 304 of the Congressional Accountability
6 Act of 1995 (2 U.S.C. 1384)) such regulations as
7 may be necessary to carry out this Act with respect
8 to covered individuals described in section
9 2(4)(A)(iii) and other individuals affected by em-
10 ployers described in section 2(6)(A)(i)(III), including
11 individuals who are covered individuals described in
12 section 2(4)(B) with respect to such employers.

13 (2) AGENCY REGULATIONS.—The regulations
14 prescribed under paragraph (1) shall be the same as
15 substantive regulations promulgated by the Sec-
16 retary under subsection (a)(1)(A) except insofar as
17 the Board may determine, for good cause shown and
18 stated together with the regulations prescribed
19 under paragraph (1), that a modification of such
20 regulations would be more effective for the imple-
21 mentation of the rights and protections involved
22 under this section.

23 (c) EMPLOYEES COVERED BY CHAPTER 5 OF TITLE
24 3, UNITED STATES CODE.—

1 (1) AUTHORITY.—Not later than 45 days after
2 the Secretary prescribes any regulation under sub-
3 section (a)(1)(A), the President (or the designee of
4 the President) shall prescribe such regulations as
5 may be necessary to carry out this Act with respect
6 to covered individuals described in section
7 2(4)(A)(iv) and other individuals affected by employ-
8 ers described in section 2(6)(A)(i)(IV), including in-
9 dividuals who are covered individuals described in
10 section 2(4)(B) with respect to such employers.

11 (2) AGENCY REGULATIONS.—The regulations
12 prescribed under paragraph (1) shall be the same as
13 substantive regulations promulgated by the Sec-
14 retary under subsection (a)(1)(A) except insofar as
15 the President (or designee) may determine, for good
16 cause shown and stated together with the regula-
17 tions prescribed under paragraph (1), that a modi-
18 fication of such regulations would be more effective
19 for the implementation of the rights and protections
20 involved under this section.

21 (d) EMPLOYEES COVERED BY CHAPTER 63 OF TITLE
22 5, UNITED STATES CODE.—

23 (1) AUTHORITY.—Not later than 45 days after
24 the Secretary prescribes any regulation under sub-
25 section (a)(1)(A), the Director of the Office of Per-

1 sonnel Management shall prescribe such regulations
2 as may be necessary to carry out this Act with re-
3 spect to covered individuals described in section
4 2(4)(A)(v) and other individuals affected by employ-
5 ers described in section 2(6)(A)(i)(V), including indi-
6 viduals who are covered individuals described in sec-
7 tion 2(4)(B) with respect to such employers.

8 (2) AGENCY REGULATIONS.—The regulations
9 prescribed under paragraph (1) shall be the same as
10 substantive regulations promulgated by the Sec-
11 retary under subsection (a)(1)(A) except insofar as
12 the Director may determine, for good cause shown
13 and stated together with the regulations prescribed
14 under paragraph (1), that a modification of such
15 regulations would be more effective for the imple-
16 mentation of the rights and protections involved
17 under this section.

18 **SEC. 6. WHISTLEBLOWER PROTECTIONS.**

19 No employer shall discriminate or retaliate (including
20 through intimidation, threats, coercion, or harassment)
21 against any covered individual of the employer—

22 (1) for exercising, or attempting to exercise,
23 any right provided under this Act; or

1 (2) because the covered individual (or another
2 individual acting at the request of the covered indi-
3 vidual) has—

4 (A) filed a written or oral complaint to the
5 employer or a Federal, State, or local govern-
6 ment entity of a violation of section 3;

7 (B) sought assistance or intervention with
8 respect to a worker privacy-related concern
9 from the employer, a Federal, State, or local
10 government, or a worker representative;

11 (C) instituted, caused to be instituted, or
12 otherwise participated in any inquiry or pro-
13 ceeding under or related to this Act;

14 (D) given, or is about to give, any informa-
15 tion in connection with any inquiry or pro-
16 ceeding relating to any right provided under
17 this Act; or

18 (E) testified, or is about to testify, in any
19 inquiry or proceeding relating to any right pro-
20 vided under this Act.

21 **SEC. 7. ENFORCEMENT.**

22 (a) IN GENERAL.—

23 (1) DEFINITION.—For purposes of this sub-
24 section:

1 (A) COVERED INDIVIDUAL.—The term
2 “covered individual” means a covered indi-
3 vidual—

4 (i) described in section 2(4)(A) (other
5 than covered individuals described in
6 clauses (iii) through (v) of such section); or

7 (ii) described in section 2(4)(B) with
8 respect to an employer.

9 (B) EMPLOYER.—The term “employer”
10 means an employer described in subclause (I)
11 or (II) of section 2(6)(A)(i).

12 (2) ENFORCEMENT BY THE TECHNOLOGY AND
13 WORKER PROTECTION DIVISION.—

14 (A) INVESTIGATION.—

15 (i) IN GENERAL.—To ensure compli-
16 ance with the provisions of this Act, or any
17 regulation or order issued under this Act,
18 the Secretary, acting through the Adminis-
19 trator—

20 (I) may investigate and gather
21 data regarding the wages, hours, and
22 other conditions and practices of em-
23 ployment in any industry subject to
24 this Act, and may enter and inspect
25 any place or record (and make such

1 transcriptions thereof), question any
2 covered individual, and investigate any
3 facts, conditions, practices, or matters
4 as the Secretary may deem necessary
5 or appropriate to determine whether
6 an employer has violated any provi-
7 sion of this Act, or which may aid in
8 the enforcement of the provisions of
9 this Act; and

10 (II) may require, by general or
11 special orders, an employer, to file
12 with the Secretary, in such form as
13 the Secretary may prescribe, annual
14 or special reports or answers in writ-
15 ing to specific questions, furnishing to
16 the Secretary such information or
17 records as the Secretary may require
18 as to the organization, business, con-
19 duct, practices, management, and re-
20 lation to other corporations, partner-
21 ships, and individuals, of the em-
22 ployer.

23 (ii) REPORTS AND ANSWERS.—An em-
24 ployer shall file the reports and answers
25 (including information and records) re-

1 quired under clause (i)(II) in such manner,
2 including under oath or otherwise, and
3 within such reasonable time period as the
4 Secretary may require.

5 (iii) JOINT INVESTIGATIONS.—The
6 Secretary, acting through the Adminis-
7 trator, may conduct investigations and
8 make requests for information, as author-
9 ized under this Act, on a joint basis with
10 another Federal agency, a State attorney
11 general, or a State agency.

12 (iv) OBLIGATION TO KEEP, PRESERVE,
13 AND MAKE AVAILABLE RECORDS.—An em-
14 ployer shall make, keep, preserve, and
15 make available to the Secretary records
16 pertaining to compliance with this Act in
17 accordance with section 11(c) of the Fair
18 Labor Standards Act of 1938 (29 U.S.C.
19 211(c)) and in accordance with any regula-
20 tion or order issued by the Secretary.

21 (B) ENFORCEMENT.—With respect to em-
22 ployers and covered individuals, the Secretary,
23 acting through the Administrator, shall receive,
24 investigate, and attempt to resolve complaints
25 of violations of section 3 or 6 in the same man-

ner that the Secretary receives, investigates, and attempts to resolve complaints of violations of sections 6 and 7 of the Fair Labor Standards Act of 1938 (29 U.S.C. 206 and 207).

(C) REFERRAL FOR CRIMINAL PROCEEDINGS.—If the Secretary, in the course of the performance of any act or duty under this Act, obtains evidence that any employer has engaged in conduct that may constitute a violation of Federal criminal law, the Secretary shall refer the matter to the Attorney General for prosecution under any applicable law. Nothing in this paragraph shall affect any other authority of the Secretary to disclose information.

(D) LITIGATION.—The Solicitor of Labor may appear for and represent the Secretary on any litigation brought under this subsection.

(3) PRIVATE RIGHT OF ACTION.—

(A) IN GENERAL.—

(i) COVERED INDIVIDUAL.—Notwithstanding any action by the Secretary under paragraph (2)(B), any covered individual adversely affected by an alleged violation of section 3 or 6, may commence a civil action against any person that violates such

1 section in any Federal court of competent
2 jurisdiction.

3 (ii) LABOR ORGANIZATION.—Notwith-
4 standing any action by the Secretary under
5 paragraph (2)(B), any labor organization
6 adversely affected by an alleged violation of
7 3 or 6 may commence a civil action against
8 any person that violates such section in
9 any Federal court of competent jurisdic-
10 tion.

11 (B) RELIEF.—

12 (i) IN GENERAL.—In a civil action
13 brought under subparagraph (A) in which
14 the covered individual or labor organization
15 prevails, the court may award the covered
16 individual or labor organization—

17 (I) damages of—

18 (aa) an amount equal to the
19 sum of any actual damages sus-
20 tained by the covered individual;
21 or

22 (bb) not more than treble
23 damages;

24 (II) statutory damages described
25 in clause (iv);

1 (III) injunctive relief; and

2 (IV) equitable relief.

3 (ii) ATTORNEY'S FEES.—In a civil ac-
4 tion brought under subparagraph (A) in
5 which the covered individual or labor orga-
6 nization prevails, the court shall award the
7 covered individual or labor organization
8 reasonable attorney's fees and litigation
9 costs.

10 (iii) TEMPORARY RELIEF FOR WHIS-
11 TLEBLOWERS.—In a civil action brought
12 under subparagraph (A) regarding a viola-
13 tion of section 6, the court may award the
14 covered individual or labor organization
15 temporary relief while the case is pending,
16 including reinstatement.

17 (iv) STATUTORY DAMAGES.—The
18 court may, in accordance with clause (v),
19 award statutory damages under clause
20 (i)(II) against a person in the following
21 amounts:

22 (I) USING AN AUTOMATED DECI-
23 SION SYSTEM FOR PROHIBITED AC-
24 TIVITIES.—For each violation of sec-
25 tion 3 by an employer with respect to

1 a covered individual, the court may,
2 subject to clause (vi), award—

3 (aa) damages of an amount
4 not less than \$5,000 and not
5 more than \$20,000; or

6 (bb) for any willful or re-
7 peated violation by the employer,
8 damages of an amount not less
9 than \$10,000 and not more than
10 \$40,000.

11 (II) RETALIATION ON WHISTLE-
12 BLOWERS.—For each violation of sec-
13 tion 6, the court may, subject to
14 clause (vi), award—

15 (aa) damages of an amount
16 not less than \$5,000 and not
17 more than \$50,000; or

18 (bb) for any willful or re-
19 peated violation, damages of an
20 amount not less than \$10,000
21 and not more than \$100,000.

22 (v) CONSIDERATIONS FOR STATUTORY
23 DAMAGES.—In determining the amount of
24 statutory damages assessed under clause
25 (iv), the court shall consider any relevant

1 circumstances presented by the parties to
2 the action, including—

3 (I) the nature and seriousness of
4 the violation;

5 (II) the number of violations;

6 (III) the persistence of the mis-
7 conduct;

8 (IV) the length of time over
9 which the misconduct occurred;

10 (V) the willfulness of the mis-
11 conduct of the employer; and

12 (VI) the assets, liabilities, and
13 net worth of the employer.

14 (vi) ADJUSTMENT FOR INFLATION.—

15 The dollar amounts referred to subclauses
16 (I) and (II) of clause (iv) shall be in-
17 creased annually, for fiscal year 2025 and
18 every fiscal year thereafter, by the percent
19 increase, if any, in the consumer price
20 index for all urban consumers for the most
21 recent 12-month period for which applica-
22 ble data is available.

23 (C) RIGHTS OF THE SECRETARY AND A
24 STATE ATTORNEY GENERAL.—Prior to an cov-
25 ered individual or labor organization bringing a

1 civil action under subparagraph (A), such cov-
2 ered individual or labor organization shall, in
3 writing, notify the Secretary and any relevant
4 State attorney general of the intent to com-
5 mence such civil action. Upon receiving such
6 notice, the Secretary and State attorney general
7 shall each, not later than 60 days after receiv-
8 ing such notice—

9 (i) determine whether to intervene in
10 such action and, upon intervening—

11 (I) be heard on all matters aris-
12 ing in such action; and

13 (II) file petitions for appeal of a
14 decision in such action; and

15 (ii) notify such covered individual or
16 labor organization.

17 (D) REMEDIES FOR STATE EMPLOYEES.—

18 (i) WAIVER OF SOVEREIGN IMMU-
19 NITY.—A State’s receipt or use of Federal
20 financial assistance for any program or ac-
21 tivity of a State shall constitute a waiver
22 of sovereign immunity, under the 11th
23 Amendment to the Constitution of the
24 United States or otherwise, to a suit
25 brought by a covered individual of that

1 program or activity, or a labor organiza-
2 tion representing such a covered individual,
3 under this paragraph for equitable, legal,
4 or other relief authorized under this para-
5 graph.

6 (ii) OFFICIAL CAPACITY.—An official
7 of a State may be sued in the official ca-
8 pacity of the official by any covered indi-
9 vidual or such a labor organization who
10 has complied with the procedures under
11 this paragraph, for injunctive relief that is
12 authorized under this paragraph. In such a
13 suit the court may award to the prevailing
14 party those costs authorized by section 722
15 of the Revised Statutes (42 U.S.C. 1988).

16 (iii) APPLICABILITY.—With respect to
17 a particular program or activity, clause (i)
18 applies to conduct that occurs—

19 (I) after the date of enactment of
20 this Act; and

21 (II) on or after the day on which
22 a State first receives or uses Federal
23 financial assistance for that program
24 or activity.

1 (iv) DEFINITION OF PROGRAM OR AC-
2 TIVITY.—In this subparagraph, the term
3 “program or activity” has the meaning
4 given the term in section 606 of the Civil
5 Rights Act of 1964 (42 U.S.C. 2000d–4a).

6 (E) REMEDIES FOR TRIBAL GOVERNMENT
7 EMPLOYEES.—

8 (i) WAIVER OF SOVEREIGN IMMU-
9 NITY.—A Tribal government’s receipt or
10 use of Federal financial assistance for any
11 program or activity of the Tribal govern-
12 ment shall constitute a waiver of sovereign
13 immunity to a suit brought by a covered
14 individual of that program or activity, or a
15 labor organization representing such a cov-
16 ered individual, under this paragraph for
17 equitable, legal, or other relief authorized
18 under this paragraph.

19 (ii) OFFICIAL CAPACITY.—An official
20 of a Tribal government may be sued in the
21 official capacity of the official by any cov-
22 ered individual or such a labor organiza-
23 tion who has complied with the procedures
24 under this paragraph for injunctive relief
25 that is authorized under this paragraph. In

1 such a suit the court may award to the
2 prevailing party those costs authorized by
3 section 722 of the Revised Statutes (42
4 U.S.C. 1988).

5 (iii) APPLICABILITY.—With respect to
6 a particular program or activity, clause (i)
7 applies to conduct that occurs—

8 (I) after the date of enactment of
9 this Act; and

10 (II) on or after the day on which
11 a Tribal government first receives or
12 uses Federal financial assistance for
13 that program or activity.

14 (iv) DEFINITION OF PROGRAM OR AC-
15 TIVITY.—In this subparagraph, the term
16 “program or activity” has the meaning
17 given the term in section 606 of the Civil
18 Rights Act of 1964 (42 U.S.C. 2000d–4a).

19 (4) ENFORCEMENT BY THE GOVERNMENT AC-
20 COUNTABILITY OFFICE AND LIBRARY OF CON-
21 GRESS.—Notwithstanding any other provision of this
22 subsection, in the case of the Government Account-
23 ability Office and the Library of Congress, the au-
24 thority of the Secretary under this subsection shall
25 be exercised respectively by the Comptroller General

1 of the United States and the Librarian of Congress,
2 respectively.

3 (b) EMPLOYEES COVERED BY CONGRESSIONAL AC-
4 COUNTABILITY ACT OF 1995.—The powers, remedies, and
5 procedures provided in the Congressional Accountability
6 Act of 1995 (2 U.S.C. 1301 et seq.) to the Board (as de-
7 fined in section 101 of that Act (2 U.S.C. 1301)), or any
8 person, alleging a violation of section 202(a)(1) of that
9 Act (2 U.S.C. 1312(a)(1)) shall be the powers, remedies,
10 and procedures this Act provides to that Board, or any
11 person, with regard to an allegation of a violation of sec-
12 tion 3 or 6 against a covered individual described in sec-
13 tion 2(4)(A)(iii) or described in section 2(4)(B) with re-
14 spect to an employer described in section 2(6)(A)(i)(III).

15 (c) EMPLOYEES COVERED BY CHAPTER 5 OF TITLE
16 3, UNITED STATES CODE.—The powers, remedies, and
17 procedures provided in chapter 5 of title 3, United States
18 Code, to the President, the Merit Systems Protection
19 Board, or any person, alleging a violation of section
20 412(a)(1) of that title, shall be the powers, remedies, and
21 procedures this Act provides to the President, that Board,
22 or any person, respectively, with regard to an allegation
23 of a violation of section 3 or 6 against a covered individual
24 described in section 2(4)(A)(iv) or described in section

1 2(4)(B) with respect to an employer described in section
2 2(6)(A)(i)(IV).

3 (d) EMPLOYEES COVERED BY CHAPTER 63 OF TITLE
4 5, UNITED STATES CODE.—The powers, remedies, and
5 procedures provided in title 5, United States Code, to an
6 employing agency, provided in chapter 12 of that title to
7 the Merit Systems Protection Board, or provided in that
8 title to any person, alleging a violation of chapter 63 of
9 that title, shall be the powers, remedies, and procedures
10 this Act provides to that agency, that Board, or any per-
11 son, respectively, with regard to an allegation of a viola-
12 tion of section 3 or 6 against a covered individual de-
13 scribed in section 2(4)(A)(v) or described in section
14 2(4)(B) with respect to an employer described in section
15 2(6)(A)(i)(V).

16 (e) ENFORCEMENT BY STATES.—

17 (1) IN GENERAL.—In any case in which a State
18 attorney general or a State privacy regulator has
19 reason to believe that an interest of the residents of
20 a State has been or is adversely affected by any per-
21 son who violates any provision of section 3 or 6, in-
22 cluding a regulation or order prescribed under this
23 Act, the State attorney general or State privacy reg-
24 ulator, as *parens patriae*, may bring a civil action on
25 behalf of the residents of the State in an appropriate

1 State court or an appropriate district court of the
2 United States to—

3 (A) enjoin further violation of such provi-
4 sion by the person;

5 (B) compel compliance with such provision;

6 (C) obtain damages (including statutory
7 damages described in paragraph (4)), civil pen-
8 alties, restitution, or other compensation on be-
9 half of the residents of the State; or

10 (D) obtain reasonable attorney's fees and
11 other litigation costs reasonably incurred.

12 (2) RIGHTS OF AGENCY.—Before initiating a
13 civil action under paragraph (1), the State attorney
14 general or State privacy regulator, as the case may
15 be, shall notify the Secretary in writing of such civil
16 action. Upon receiving such notice, the Secretary
17 may—

18 (A) intervene in such action; and

19 (B) upon intervening—

20 (i) be heard on all matters arising in
21 such civil action; and

22 (ii) file petitions for appeal of a deci-
23 sion in such action.

24 (3) PREEMPTIVE ACTION BY AGENCY.—In any
25 case in which a civil action is instituted by or on be-

1 half of the Secretary for violation of this Act or a
2 regulation promulgated under this Act, a State at-
3 torney general or State privacy regulator may not,
4 during the pendency of such action, institute a civil
5 action against any defendant named in the com-
6 plaint in the action instituted by or on behalf of the
7 Secretary for a violation that is alleged in such com-
8 plaint. In a case brought by the Secretary that af-
9 fects the interests of a State, the State attorney gen-
10 eral or State privacy regulator may intervene as of
11 right pursuant to the Federal Rules of Civil Proce-
12 dure.

13 (4) STATUTORY DAMAGES.—In a civil action in-
14 stituted under paragraph (1), a court may award
15 statutory damages under paragraph (1)(C) against a
16 person for a violation of any provision of section 3
17 or 6—

18 (A) in an amount not more than \$50,000
19 for each such violation; or

20 (B) in the case of such a violation that re-
21 sults in the discharge of an employee or other
22 serious economic harm to an employee by such
23 a person who has, within the preceding 5 years,
24 committed another such violation resulting in
25 such a discharge or other serious economic

1 harm, not more than \$100,000 for each such
2 violation.

3 (5) PRESERVATION OF STATE POWERS.—Ex-
4 cept as provided in paragraph (3), no provision of
5 this subsection shall be construed as altering, lim-
6 iting, or affecting the authority of a State attorney
7 general or State privacy regulator to—

8 (A) bring an action or other regulatory
9 proceeding arising solely under the laws in ef-
10 fect in that State; or

11 (B) exercise the powers conferred on the
12 State attorney general or State privacy regu-
13 lator by the laws of the State, including the
14 ability to conduct investigations, administer
15 oaths or affirmations, or compel the attendance
16 of witnesses or the production of documentary
17 or other evidence.

18 (f) ARBITRATION AND CLASS ACTION.—Notwith-
19 standing any other provision of law, no predispute arbitra-
20 tion agreement or predispute joint-action waiver shall be
21 valid or enforceable with respect to any alleged violation
22 of section 3 or 6.

23 **SEC. 8. COORDINATION.**

24 In carrying out this Act, the Secretary, acting
25 through the Administrator, shall coordinate with any ap-

1 appropriate Federal agency or State regulator to promote
2 consistent regulatory treatment of automated decision sys-
3 tems.

4 **SEC. 9. RELATION TO OTHER LAWS.**

5 Except as explicitly provided otherwise in this Act,
6 nothing in this Act shall be construed to preempt, modify,
7 limit, or supersede—

8 (1) any provision of Federal or State law; or

9 (2) the authority of the Federal Trade Commis-
10 sion, Equal Employment Opportunity Commission,
11 National Labor Relations Board, or any other Fed-
12 eral agency.

13 **SEC. 10. SEVERABILITY.**

14 If any provision of this Act or the application of such
15 provision to any person or circumstance is held to be un-
16 constitutional, the remainder of this Act and the applica-
17 tion of the provisions of such to all other persons or cir-
18 cumstances shall not be affected thereby.

○

BILL AS INTRODUCED
2023

H.114
Page 1 of 13

1 H.114

2 Introduced by Representatives Priestley of Bradford, Anthony of Barre City,
3 Graning of Jericho, Headrick of Burlington, Mulvaney-Stanak
4 of Burlington, Sims of Craftsbury, and Small of Winooski

5 Referred to Committee on

6 Date:

7 Subject: Labor; employment practices; fair employment; electronic
8 monitoring; automated decision systems

9 Statement of purpose of bill as introduced: This bill proposes to restrict the
10 use of electronic monitoring of employees and the use of automated decision
11 systems for employment-related decisions.

12 An act relating to restricting electronic monitoring of employees and
13 employment-related automated decision systems

14 It is hereby enacted by the General Assembly of the State of Vermont:

15 Sec. 1. 21 V.S.A. § 495o is added to read:

16 § 495o. ELECTRONIC MONITORING OF EMPLOYEES; AUTOMATED
17 DECISION SYSTEMS; RESTRICTIONS; EXCEPTIONS

18 (a) Definitions. As used in this section:

19 (1) “Algorithm” means a computerized procedure consisting of a set of
20 steps used to accomplish a determined task.

BILL AS INTRODUCED
2023

H.114
Page 2 of 13

1 (2) “Automated decision system” means an algorithm or computational
2 process that is used to make or assist in making employment-related decisions,
3 judgments, or conclusions. The term “automated decision system” includes
4 algorithms and computational processes that are derived from machine
5 learning, statistics, data processing, or artificial intelligence.

6 (3) “Automated decision system output” means information, data,
7 assumptions, predictions, scoring, recommendations, decisions, or conclusions
8 generated by an automated decision system.

9 (4) “Data” means information obtained by any means that, directly or
10 indirectly, identifies, relates to, describes, may reasonably be associated with,
11 or could reasonably be linked to an employee, including:

12 (A) personal identifying information;

13 (B) biometric information;

14 (C) health, medical, lifestyle, and wellness information;

15 (D) information related to workplace activities;

16 (E) human resources information and personnel files; and

17 (F) information related to the employee’s Internet and digital
18 activities, including electronic communications, private social media activity,
19 and personal Internet Protocol address.

20 (5) “Electronic monitoring” means the collection of information
21 concerning employee activities or communication by any means other than

BILL AS INTRODUCED
2023

H.114
Page 3 of 13

1 direct, in-person observation, including through the use of a digital device,
2 computer, telephone, wire, radio, camera, or electromagnetic, photoelectronic,
3 or photo-optical system.

4 (6) “Employee” means an individual who, in consideration of direct or
5 indirect gain or profit, is employed by an employer.

6 (7) “Employer” means a person who employs one or more individuals in
7 Vermont, and includes any agent or contractor acting on the employer’s behalf.

8 (8) “Employment-related decision” means any decision by an employer
9 that:

10 (A) affects an employee’s compensation, benefits, or terms and
11 conditions of employment;

12 (B) relates to the discipline, evaluation, promotion, or termination of
13 an employee; or

14 (C) relates to the hiring of an individual or employee for a position or
15 job.

16 (9) “Essential job function” means a fundamental duty of a job or
17 position that an employee with that job or position must be able to perform.

18 (b) Employee monitoring restricted. An employer shall not engage in
19 electronic monitoring of an employee unless all of the following requirements
20 are met:

21 (1) The employer’s purpose in utilizing the electronic monitoring is to:

BILL AS INTRODUCED
2023

H.114
Page 4 of 13

1 (A) assist or allow the employee to accomplish an essential job
2 function;

3 (B) monitor production processes or quality;

4 (C) ensure compliance with applicable employment or labor laws;

5 (D) protect the health, safety, or security of the employee; or

6 (E) track time worked or production output for purposes of
7 determining the employee's compensation.

8 (2) Electronic monitoring is necessary to accomplish the purpose
9 identified pursuant to subdivision (1) of this subsection (b).

10 (3) The specific form of electronic monitoring is the least invasive
11 means, with respect to the employee, of accomplishing the purpose identified
12 pursuant to subdivision (1) of this subsection (b).

13 (4) The specific form of electronic monitoring is used with the smallest
14 number of employees and collects the smallest amount of data necessary to
15 accomplish the purpose identified pursuant to subdivision (1) of this subsection
16 (b).

17 (5) The employer ensures that only authorized persons have access to
18 any data produced through the electronic monitoring and that the data is only
19 used for the purpose and duration that the employee has been notified of
20 pursuant to subsection (c) of this section.

21 (c) Required notice for employee monitoring.

1 (1) At least 15 calendar days prior to commencing any form of
2 electronic monitoring, an employer shall provide notice of the electronic
3 monitoring to each employee who will be subject to it. The notice shall, at a
4 minimum, include the following information:

5 (A) the specific form of electronic monitoring;

6 (B) a description of the intended purpose of the electronic monitoring
7 and why the electronic monitoring is necessary to accomplish that purpose;

8 (C) a description of how any data generated by the electronic
9 monitoring will be used;

10 (D) a description of the technologies that will be used to conduct the
11 electronic monitoring;

12 (E) whether and, if so, how the data generated by the electronic
13 monitoring will be used to inform employment-related decisions;

14 (F) the name of any person conducting electronic monitoring on the
15 employer's behalf and any associated contract language related to the
16 monitoring;

17 (G) the name of any person, apart from the employer, who will have
18 access to any data generated by the electronic monitoring and the reason why
19 the person will have access to the data;

20 (H) the positions within the employer that will have access to any
21 data generated by the electronic monitoring;

BILL AS INTRODUCED
2023

H.114
Page 6 of 13

1 (I) when, where, and how frequently monitoring will occur;

2 (J) the period of time for which any data generated by the electronic
3 monitoring will be retained by the employer or another person and when that
4 data will be destroyed;

5 (K) notice of how the employee may access the data generated by the
6 electronic monitoring and the process to correct any errors in the data; and

7 (l) notice of the employee's rights pursuant to this section.

8 (2)(A) Notice of electronic monitoring provided pursuant to this section
9 shall be written in plain, clear, and concise language, and provided to each
10 employee in the employee's primary language.

11 (B) An employer shall provide a new, updated notice to employees if
12 it makes any significant changes to the manner of electronic monitoring or to
13 the way that the employer utilizes the electronic monitoring or any data
14 generated by it.

15 (3) Notwithstanding subdivision (1) of this subsection (c), prior notice
16 of electronic monitoring shall not be required if:

17 (A) the employer has reasonable grounds to believe that the
18 employee is engaged in conduct that:

19 (i) is illegal;

20 (ii) violates the legal rights of the employer or another employee;

21 or

BILL AS INTRODUCED
2023

H.114
Page 7 of 13

1 (iii) creates a hostile work environment; and

2 (B) the electronic monitoring is reasonably likely to produce
3 evidence of the conduct.

4 (4)(A) An employer that utilizes electronic monitoring shall annually
5 provide each of its employees with a list of all electronic monitoring systems
6 currently in use by the employer in relation to that employee. The list shall be
7 provided in the primary language of the employee.

8 (B) As used in this subdivision (4), “currently in use” means that the
9 employer:

10 (i) is currently using the system in relation to the employee;

11 (ii) used the electronic monitoring system in relation to the
12 employee within the past 90 days; or

13 (iii) intends to use the electronic monitoring system in relation to
14 the employee within the next 30 days.

15 (d) Prohibitions on employee monitoring. Notwithstanding the purposes
16 for electronic monitoring set forth in subdivision (b)(1) of this section,
17 electronic monitoring shall not be used:

18 (1) in any manner that violates State or federal labor or employment
19 laws;

20 (2) in relation to workers who are off-duty and not performing work-
21 related tasks;

1 (3) to identify employees exercising legal rights;

2 (4) for audio-visual monitoring of bathrooms, locker rooms, changing
3 areas, breakrooms, smoking areas, employee cafeterias, lounges, or other
4 similarly private areas;

5 (5) to determine the frequency with which employees visit or use
6 bathrooms, locker rooms, changing areas, breakrooms, smoking areas,
7 employee cafeterias, lounges, or other similarly private areas; or

8 (6) for audio-visual monitoring of any space within an employee's
9 residence or personal vehicle, or a property owned or rented by the employee,
10 unless the monitoring is necessary to ensure the employee's health and safety
11 or to verify the security of employer or client data.

12 (e) Restriction of employee monitoring through personal devices.

13 (1) An employer shall not require an employee to install an application
14 on a personal device for purposes of electronic monitoring or to wear a device
15 or attach, embed, or physically implant a device on the employee's clothing
16 that can be used for electronic monitoring, unless the electronic monitoring is:

17 (A) necessary to accomplish the employee's essential job functions;
18 and

19 (B) limited to only the times and activities necessary to accomplish
20 the essential job functions.

1 (2) Any location tracking function of an application or device shall be
2 disabled outside of the times when the employee is engaged in activities
3 necessary to accomplish essential job functions.

4 (3) An employer shall not require an employee to physically implant a
5 device on the employee's body for purposes of employee monitoring.

6 (f) Restrictions on use of automated decision systems.

7 (1) An employer shall not use an automated decision system in a manner
8 that:

9 (A) violates or results in a violation of State or federal law;

10 (B) makes predictions about an employee's behavior that are
11 unrelated to the employee's essential job functions;

12 (C) identifies, profiles, or predicts the likelihood that an employee
13 will exercise the employee's legal rights;

14 (D) makes predictions about an employee's emotions, personality, or
15 other sentiments; or

16 (E) uses customer or client data, including customer or client reviews
17 and feedback, as an input of the automated decision system.

18 (2)(A) An employer shall not solely rely on outputs from an automated
19 decision system when making employment-related decisions.

20 (B) An employer may utilize an automated decision system in
21 making employment-related decisions if:

BILL AS INTRODUCED
2023

H.114
Page 10 of 13

1 (i) the automated decision system outputs considered in making
2 the employment-related decision are corroborated by human oversight of the
3 employee, including supervisory or managerial observations and
4 documentation of the employee’s work, personnel records, and consultations
5 with the employee’s coworkers;

6 (ii) the employer has conducted an impact assessment of the
7 automated decision system pursuant to subsection (g) of this section; and

8 (iii) the employer is in compliance with the notice requirements of
9 subsection (h) of this section.

10 (3) An employer shall not use any automated decision system outputs
11 regarding an employee’s physical or mental health in relation to an
12 employment-related decision.

13 (g) Impact assessment of automated decision systems.

14 (1) Prior to utilizing an automated decision system, an employer shall
15 create a written impact assessment of the system that includes, at a minimum:

16 (A) a detailed description of the automated decision system and its
17 purpose;

18 (B) a description of the data utilized by the system;

19 (C) a description of the outputs produced by the system and the types
20 of employment-related decisions in which those outputs may be utilized;

BILL AS INTRODUCED
2023

H.114
Page 11 of 13

1 (D) an assessment of the necessity for the system, including reasons
2 for utilizing the system to supplement nonautomated means of decision
3 making;

4 (E) a detailed assessment of the potential risks of utilizing the system,
5 including the risk of:

6 (i) errors;

7 (ii) discrimination against employees on the basis of race, color,
8 religion, national origin, sex, sexual orientation, gender identity, ancestry,
9 place of birth, age, crime victim status, or physical or mental condition;

10 (iii) violating employees' legal rights or chilling employees'
11 exercise of legal rights;

12 (iv) directly or indirectly harming employees' physical health,
13 mental health, safety, sense of well-being, dignity, or autonomy;

14 (v) harm to employee privacy, including through potential security
15 breaches or inadvertent disclosure of information; and

16 (vi) negative economic and material impacts to employees,
17 including potential effects on compensation, benefits, work conditions,
18 evaluations, advancement, and work opportunities;

19 (F) a detailed summary of measures taken by the employer to address
20 or mitigate the risks identified pursuant to subdivision (E) of this subdivision
21 (g)(1); and

BILL AS INTRODUCED
2023

H.114
Page 12 of 13

1 (G) a description of any methodology used in preparing the
2 assessment.

3 (2) An employer shall provide a copy of the assessment prepared
4 pursuant to subdivision (1) of this subsection (g) to an employee upon request.

5 (3) An employer shall update the assessment required pursuant to this
6 subsection any time a significant change or update is made to the automated
7 decision system.

8 (h) Prohibition on facial, gait, and emotion recognition technology.
9 Electronic monitoring and automated decision systems shall not incorporate
10 any form of facial, gait, or emotion recognition technology.

11 (i) Protection of employee privacy.

12 (1) An employer; any person that develops, operates, or maintains
13 electronic monitoring or an automated decision system on behalf of an
14 employer; and any person who collects, stores, analyzes, interprets,
15 disseminates, or otherwise uses data produced or utilized by electronic
16 monitoring or an automated decision system shall implement reasonable
17 security procedures and practices appropriate to the nature of the data to
18 protect employees' personal information from unauthorized or illegal access,
19 destruction, use, modification, or disclosure.

20 (2) Any person that develops, operates, or maintains electronic
21 monitoring or an automated decision system on behalf of an employer and any

BILL AS INTRODUCED
2023

H.114
Page 13 of 13

1 person who collects, stores, analyzes, interprets, disseminates, or otherwise
2 uses data produced or utilized by electronic monitoring or an automated
3 decision system shall, upon termination of the contract with the employer:

4 (A) return all data and automated decision system outputs to the
5 employer; and

6 (B) destroy all data and automated decision system outputs in the
7 person's possession.

8 (j) Employee right to access data. An employer shall, upon request,
9 provide an employee with any data that relates to the employee that was
10 produced or utilized by electronic monitoring or an automated decision system
11 used by the employer.

12 (k) Retaliation prohibited. An employer shall not discharge or in any other
13 manner retaliate against an employee who exercises or attempts to exercise the
14 employee's rights under this section. The provisions against retaliation set
15 forth in subdivision 495(a)(8) of this subchapter shall apply to this section.

16 (l) Enforcement. The provisions of section 495b of this subchapter shall
17 apply to this section.

18 Sec. 2. EFFECTIVE DATE

19 This act shall take effect on July 1, 2023.

Equity in Civic Technology, Privacy & Data

NY City Council Rams Through Once-Promising but Deeply Flawed Bill on AI Hiring Tools

November 12, 2021 / [Matt Scherer](#), [Ridhi Shetty](#)

New York City became the first place in the U.S. this week to pass significant legislation aimed at mitigating the risks of discrimination associated with the use of automated employment decision tools. Unfortunately, [Introduction 1894](#) (I.1894), the once-promising bill the City Council adopted, is deeply flawed. Many of its most harmful provisions were introduced and rammed through in a rushed process that excluded workers, civil rights groups, and other stakeholders from providing any input.

The Council's decision to pass this weakened bill without consulting impacted stakeholders should worry workers, advocates, and the broader public in New York City and the rest of the country alike.

In recent years, employers have increasingly used [automated tools when making crucial employment decisions](#). This trend has generated a wide-ranging conversation about how these technologies can discriminate against candidates, especially people with disabilities, capturing the interest of employers and workers, [civil society organizations](#), and recently prompting new initiatives from [White House](#) and [U.S. Equal Employment Opportunity Commission](#). The New York City bill could have been a model for jurisdictions around the country to follow. Instead, it is a missed opportunity that fails to hold companies accountable, and leaves important forms of discrimination unaddressed.

Below, we delve into the concerns about the New York City legislation with the hope that the Council will amend the bill—or, at the very least, that other policymakers avoid the bill’s harmful flaws.

Background on I.1894: How a Bill Should Not Become a Law

The original version of I.1894 was introduced in February of 2020. The Council’s Committee on Technology held a hearing on the bill last

“... ..”



[original bill](#) noted the bill’s failure to address the high potential for automated tools to worsen discrimination against disabled workers. CDT and a coalition of disability advocacy and other civil society organizations [laid out for the Council](#) the ways in which the draft bill failed to protect disabled workers.

Given those concerns, we were initially hopeful when we saw that a revised version of the bill had landed on the Council’s website on Tuesday. Perhaps the Council had heeded our recommendations to [strengthen the bill’s auditing provisions](#), give candidates a meaningful opportunity to raise concerns about the bill, and shore up the bill’s regulatory and enforcement provisions? To our dismay, the revised bill was, if anything, *more* flawed than the original, with a narrower scope and weaker audit requirements and enforcement provisions, as we discuss below. But CDT and other interested parties never had an opportunity to raise the alarm about these weakening changes because of how rapidly the revised bill was adopted and passed.

On the Council’s website, [the revised bill \(download\)](#) is dated November 9th, accompanied by Mayor Bill de Blasio’s [“message of necessity” \(download\)](#) (also dated November 9th) certifying the need for the Council to pass the draft ordinance “immediately.” The Committee held a “hearing” on the revised bill the very next morning. The scare quotes around “hearing” are necessary – the [Committee Report \(download\)](#) does not indicate that the Committee actually solicited or heard any testimony on the revised bill at this hearing, instead referring readers back to the materials from the November 2020 hearing on the original bill. As far as we can tell, there was no public notice of yesterday’s Committee hearing. The lack of advance notice effectively precluded interested parties from reviewing and responding to the revised bill.

The Committee approved the bill 4-1 immediately following the hearing, sending it to the full Council. Less than two hours later, the Council passed the bill, sending it to the Mayor’s desk for signature. The time from when the Council formally announced the revised bill (which was dated November

9th) to when the Council approved it (November 10th) was approximately 24 hours, if not less. Given the message of necessity that the Mayor had issued the previous day, it seems likely that he will sign it, though we hold out hope that strong pressure from public interest organizations may dissuade him from doing so.

Where the Revised Bill Takes Steps Backwards

There are a number of areas in which the revised bill, which is now set to become law, is not merely deficient, but weaker than the original bill.

Narrowing of scope and application

At the highest level, the revised bill applies to a narrower class of employment decisions than the original. Where the original bill defined an “employment decision” as including all decisions on hiring, compensation, and all other “terms, conditions or privileges of employment,” the revised bill applies only to hiring and promotion. This leaves many substantial employment decisions that dramatically impact workers’ lives – including those relating to compensation, scheduling, and working conditions – outside the bill’s scope.

The revised bill also applies only to workers who are residents of New York City, rather than to all employees of NYC-based employers. Given the sheer volume of non-NYC residents employed by NYC employers – particularly in the dawning post-COVID age – this represents a significant narrowing of the bill’s applicability.

Weaker bias audit requirements

I.1894 requires employers to conduct a “bias audit” on an automated tool before using it to evaluate candidates. The original bill required bias audits to assess a tool’s “predicted compliance” with New York City’s employment discrimination laws “and any other applicable law relating to discrimination in employment.” This definition of bias audit, while somewhat vague, implied that employers must assess a tool’s compliance with *all* anti-discrimination laws. This meant that the bill would cover “a wide range of protected traits that are too often neglected when examining bias,” such as age and disability.

The revised bill replaced this broad definition of bias audits with language that, in essence, merely reaffirms employers’ *existing requirements under federal regulations* to check whether the tool would have a disparate impact based on race, ethnicity, or sex. This narrow requirement is accompanied by language that gives the illusion of greater inclusivity, saying that a bias audit “shall include but not be limited to” checking for disparate impact on race, ethnicity, and sex. But in practice, the revised bill gives employers and vendors no reason to do more than this. After all, why would an employer choose to conduct analyses that might reveal (and thus put them on notice of) violations of law that might otherwise remain hidden?

The main effect of the revisions is, therefore, to relieve employers of any incentive to check for other forms of discrimination, such as:

- discrimination against disabled, older, or LGBTQ+ workers;
- intersectional discrimination where a multiply marginalized candidate is disadvantaged based on a combination of protected characteristics; and

- forms of disparate treatment that can arise when automated systems “learn” from biased training data to recognize (and discriminate against) protected characteristics without being explicitly programmed to do so.

The bill also requires employers to post a “summary” of the bias audit on their website, but it says nothing about what this summary must contain. This leaves open the possibility that employers will omit information about potential sources of discrimination even with respect to the limited analysis that the bias audit requires.

Inadequate notice requirements

The revised bill appears, at first blush, to improve on the original in its notice and disclosure requirements. The original bill required candidates to be notified that an automated tool had assessed them, as well as the “job qualifications or characteristics” that were the subject of the assessment. But where the original bill did not require such notice until *after* the evaluation occurred, the revised bill requires pre-assessment notice. The revised bill also requires the notice to “allow a candidate to request an alternative selection process or accommodation.” Lastly, upon a written request from a worker, the employer must provide additional information regarding the data that the automated tool uses, if the information is not already on their website.

A closer examination reveals these notices still fail to give candidates who may otherwise face discrimination a meaningful chance to voice concerns or request accommodation. The notice does not have to include any information on *how* candidates’ qualifications and characteristics will be

assessed. Workers, particularly disabled workers, need this information to know whether they might require accommodations or an alternative selection procedure to ensure they get a fair assessment.

Moreover, “job qualifications and characteristics” are never defined, and employers could easily interpret them in a narrow way that further limits the information candidates receive. This too is particularly problematic for disabled workers, since the bill has been stripped of any other provisions that require employers to consider impacts on disabled workers.

Because nothing else in the bill requires employers or tool designers to audit for disability discrimination or consider potential accommodations when designing or deploying a tool, this provision is likely to have little practical effect. Consequently, employers might only provide accommodations that are ill-suited for candidates’ needs, if they even offer accommodations at all.

Lastly, the disclosure regarding the data the tool uses need only be provided if a candidate requests it in writing. Even if such a request is made, the bill is once again vague on exactly what information about “data collection” and “sources” must be provided, giving employers leeway to provide only cursory information on the data undergirding a tool.

Removal of rulemaking authority, and weakened enforcement provisions

The revised bill makes it harder to mitigate the bill’s flaws by stripping the [New York City Commission on Human Rights](#) of the authority to clarify and expand on the bill’s provisions and by transferring enforcement authority away from the Commission. The original bill gave “[t]he department [of

Consumer and Worker Protection], the commission on human rights and any other agency designated by the mayor” authority to “promulgate such rules as it deems necessary to implement and enforce the provisions of this subchapter.” This broad authority would have allowed the Commission, which has specific and deep expertise in civil rights laws, to flesh out the bill’s audit and notice requirements. By deleting those provisions, the revised bill creates ambiguity as to which entity, if any, has that authority. Certainly, it is clear that the Commission, which is the only City agency with the necessary expertise, will not have that authority.

The revisions also transferred enforcement authority from the Commission to New York City’s Corporation Counsel, a generalist office tasked with a wide range of legal responsibilities. As a result, the odds of aggressive enforcement of the law’s already weakened provisions are greatly reduced.

The Committee Report Accompanying the Revised Bill Misleads

The Committee Report further demonstrates how ill-advised it was for the Council to pass the revised bill without seeking additional public input. The report is highly selective in its characterization of data-driven human resources (HR) tools, focusing on certain types of problematic automated tools and assessments while overlooking others.

The report speaks positively of technology that is “used to quickly evaluate resumes,” not acknowledging the [serious potential for discrimination](#) that resume screening technologies [have demonstrated](#). It refers only fleetingly to risks for “ethnic or minority workers,” and only with respect to

“automated searches” that exclude these workers due to housing segregation when looking for workers who live near worksites.

The report does allude to potential issues with “questionable assessments of observable physical factors,” but makes no mention of the [serious concerns that also arise when tools assess candidates based on personality traits or other characteristics associated with mental health or cognitive disabilities](#). Moreover, while the report mentions the risk that some tools that measure physical factors “might systematically disadvantage applicants with disabilities,” it neglects to mention that the revised bill actually *stripped* the language that would have required employers to check for such disability discrimination.

Overall, the report paints a picture of automated employment decision tools that fails to truly recognize how such tools can harm marginalized workers; in fact, the main risk that the report cites is simply that automated tools may “potentially recommend[] an unqualified candidate.” The report thus suggests that the Council acted on far-from-complete information about the risks HR technologies pose.

Conclusion

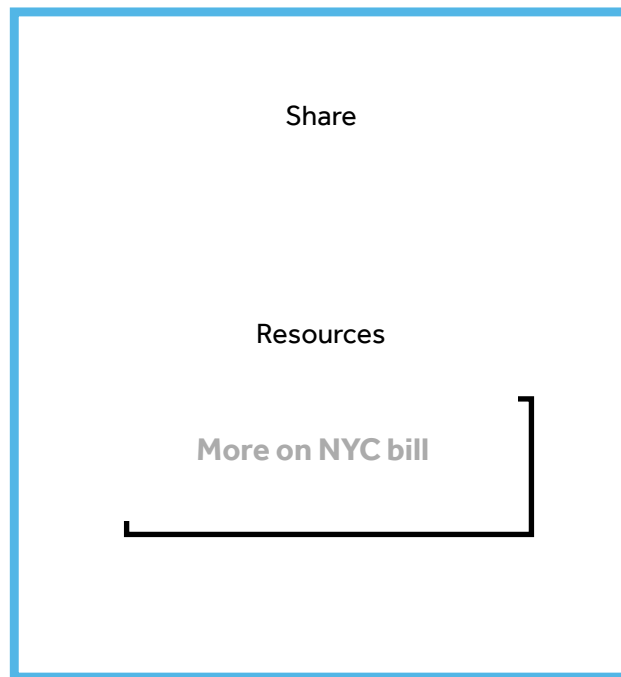
I.1894 was already the most significant proposed legislation on algorithmic, data-driven, or AI-powered HR tools in the U.S. to date. Now that it is about to become law, other jurisdictions might view it as a potential model for regulating HR technologies. That would be a serious mistake.

While we appreciate the Council’s willingness to tackle these issues, we fear that the ordinance will create a situation worse than the *status quo*

ante. The intent of the ordinance, and indeed its ultimate effect, seems to be to make it easy for vendors and employers to comply. By marketing their compliance with such weak protections, vendors and employers would create the false impression that their tools are valid and nondiscriminatory. That could well create additional pressure for employers to adopt such tools, creating a boon for vendors but threatening further harm to vulnerable workers.

Civil society organizations must be vocal in pointing out the many ways in which this ordinance falls short, and fight hard to ensure that future legislation does not suffer from the same serious flaws.


Updated on November 15, 2021.



Related Reading

Civil Rights Standards

for 21st Century Employment Selection Procedures



December 2022

Civil Rights Standards for 21st Century Employment Selection Procedures

Authors

Center for Democracy & Technology (CDT)
American Association for People with Disabilities (AAPD)
American Civil Liberties Union (ACLU)
The Leadership Conference on Civil and Human Rights
National Women’s Law Center
Upturn

Acknowledgements

These Standards are made possible by the support of the Ford Foundation.

Special thanks to Chai Feldblum, Sarah David Heydemann, Aaron Rieke, Clarence Okoh, Amalea Smirniotopoulos, and Katurah Topps for their invaluable input during the creation of the Standards.

December 2022



Endorsements



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


Table of Contents

Introduction	5
Executive Summary	7
Scope	7
Nondiscrimination	8
Job-Relatedness	9
Auditing	9
Notice and Explanation	10
Oversight and Accountability	11
Civil Rights Standards	13
Standard 1: Statement of Purpose	13
Standard 2: Definitions	14
Standard 3: Pre-deployment audits	20
Standard 4: Short-form disclosures	23
Standard 5: Procedure to request accommodation, communicate concerns, or opt out; right to explanation for adverse actions	25
Standard 6: Ongoing audits and adjustments	27
Standard 7: Audit summaries	29
Standard 8: Recordkeeping	30
Standard 9: Unlawful employment practices	32
Standard 10: Enforcement, remedies, and liability	34
Standard 11: Enforcement agency authority and responsibilities	36
Standard 12: Construction	38

Introduction

More than ever, employers are implementing new selection methods for virtually every stage of the employment process, from candidate sourcing and recruitment to employee evaluation and termination. Workers that face these tools are at an extreme information disadvantage, with little insight into how they are assessed or whether they face a risk of an unfair or discriminatory decision. In 2020, a broad coalition of civil rights and technology policy organizations published the [Civil Rights Principles for Hiring Assessment Technologies](#) (the “Principles”) in an effort “to guide the development, use, auditing, and oversight of hiring assessment technologies, with the goals of preventing discrimination and advancing equity in hiring.” In the two years that have followed, an increasing number of cities and states have considered legislation or regulations on hiring technologies that fail to implement – or even actively undermine – the Principles.

The Civil Rights Standards for 21st Century Employment Selection Procedures (the “Standards”) were drafted to operationalize and expand on the Principles. The Standards provide a concrete alternative to recent proposals that would set very weak notice, audit, and fairness standards for automated tools. They also map out a more rigorous and rights-focused approach as compared to the outdated rules that currently govern how employers assess whether their selection procedures are discriminatory and whether they actually measure the worker characteristics they claim to measure. The Standards have been drafted so that policymakers, industry groups, and employers alike can reference them when determining what information candidates should receive, how selection procedures should be audited, and how to ensure accountability when selection procedures threaten workers’ civil rights.



The Standards are the culmination of a year-long collaboration among a number of civil society groups. The coalition behind the Standards includes organizations that focus on racial justice, disability rights, digital rights, workers' rights, and a number of other civil rights and technology policy issues.

Workers that face these [new employment selection] tools are at an extreme information disadvantage, with little insight into how they are assessed or whether they face a risk of an unfair or discriminatory decision.

Executive Summary

Scope

The Civil Rights Standards cover all *workers* defined by Labor Department regulations as part of the labor force, regardless of whether they are classified as employees or contractors. Standard 2(ac). Consistent with existing antidiscrimination laws, they apply to all employers and employment agencies. The Standards cover the developers and sellers of selection procedures by explicitly classifying them as employment agencies. Standard 2(m). This reflects the fact that such vendors are increasingly filling roles traditionally performed by employment agencies.

Under the Standards, the definition of *selection procedure* (Standard 2(y)) includes every worker assessment that meets the following criteria:

- a. It is sold by a vendor or other employment agency, or is used to assess at least 100 workers per year;
- b. The score, recommendation, or other output it generates is primarily the result of automated, algorithmic, or deterministic processes; and
- c. Its output is used as a basis, factor, or recommendation in connection with *employment decisions*.

While the rapid rise of automated selection procedures is a key motivator for the Civil Rights Standards, the Standards also cover other worker assessments so long as their scores, recommendations, or other outputs are assigned through algorithms, standardized rubrics, or similar processes. The Standards cover traditional multiple-choice Scantron tests, for instance, because the output of such a

test is computed using a (non-computerized) algorithm, with the candidates' responses as inputs and the score as the output.

The Civil Rights Standards' definition of *employment decision* (Standard 2(n)) largely tracks with the definition in the [Uniform Guidelines for Employee Selection Procedures \(UGESP\)](#), which cover most major personnel decisions. The Standards expand the UGESp definition to include decisions setting terms or conditions of employment and selecting workers for targeted recruitment or advertising.

Nondiscrimination

The auditing standards (Standards 3 and 6) would require companies to take a proactive approach to mitigating discrimination risk by:

- Identifying and anticipating discriminatory barriers throughout a selection procedure's lifecycle;
- Exploring alterations, accommodation, and alternative selection procedures that might reduce or eliminate potential sources of discrimination; and
- Requiring companies to choose the least-discriminatory valid method for measuring candidates' essential job functions.

The Standards would further extend civil rights laws' prohibitions against selection procedures that constitute or contribute to employment discrimination. Specifically, Standard 9 would prohibit companies from:

- Using or marketing discriminatory selection procedures (including selection procedures that result in disparate treatment or disparate impact); failing to provide candidates with reasonable accommodation; and failing to choose the least discriminatory valid method of candidate assessment.
- Failing to alter a selection procedure or provide reasonable accommodation, where alteration or accommodation is necessary to ensure that the selection procedure validly measures candidates' ability to perform essential job functions.
- Retaliating against workers who request reasonable accommodation or otherwise exercise their rights under the Civil Rights Standards.

The Standards also call for a ban on 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. Standard 9(a)(11).

Job-Relatedness

Before using a selection procedure that might adversely impact members of a protected class, the Standards would require employers to demonstrate that the selection procedure is a *valid* method of measuring candidates' ability to perform the *essential functions* of each position for which it is used. See Standards 3(c)(3); 6(b)(3); 9(a)(3). An employer establishes the essential functions of a position through objective evidence, such as:

- Workers' past/present experiences and performance
- Time workers spend on each function
- Consequences of workers not performing the function

Regarding *validity* – that is, the extent to which a selection procedure is an accurate and effective means of measuring the essential job functions that it purports to measure – the Civil Rights Standards look to contemporary standards of social science, rather than the outdated UGESP. Standard 2(ab). They would preclude employers from establishing validity simply by showing that the selection procedure's output correlates with existing measures of job performance. Such exclusive reliance on correlational evidence has become increasingly common with the advent of automated selection procedures, but blind reliance on correlation can result in selection procedures that incorporate systemic biases or cultural norms that disadvantage vulnerable groups of workers, in addition to (or even to the exclusion of) characteristics that have a causal link to workers' ability to perform essential job functions.¹

Auditing

The Civil Rights Standards provide for both a pre-deployment audit before a selection procedure is first used to assess candidates, as well as ongoing audits conducted at regular intervals for as long as the selection procedure is in use. Both types of audits would be conducted by an independent auditor certified by the enforcement agency.

Standard 3 spells out a thorough *pre-deployment audit* that would have to be completed before a company could use a selection procedure. The Standards place this responsibility jointly on employers and vendors to encourage all parties involved in the development and use of a selection procedure to work together to ensure it is audited for the specific context(s) in which it will be used. The audit itself is conducted by an independent auditor, who would:

¹ In one instance, a vendor developed a resume screener that determined that the two factors most indicative of job performance were whether the candidate's name was Jared and whether the candidate had played high school lacrosse. Dave Gershorn, 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/>.

- Identify existing and potential future sources of discrimination;
- Evaluate the selection procedure's job-relatedness/validity for each position for which the selection procedure will be used;
- Determine what alterations or accommodation might be required to ensure the selection procedure fairly assesses all candidates; and
- Explore potential alternative approaches to candidate assessment to determine if a valid and less-discriminatory alternative is available.

The auditor would also examine the company's existing employment decision practices to identify existing sources and patterns of discrimination.

After a selection procedure is deployed, Standard 6 would require companies to conduct ongoing audits at least annually. The components of ongoing audits largely mirror those of the pre-deployment audit, except that the auditor would review the selection procedure's impact and validity in light of the real-world performance data and candidate feedback received since the last audit. The ongoing audit also analyzes any changes made to the selection procedure or to the essential functions of the position(s) for which it is being used.

Notice and Explanation

The Civil Rights Standards contemplate three levels of disclosure and transparency:

- Short-form disclosures (Standard 4);
- Detailed summaries of all audits (Standard 7); and
- Comprehensive recordkeeping (Standard 8).

The short-form disclosure is a candidate-facing document that is designed to help alleviate the severe information disadvantage that workers currently face when employers subject them to selection procedures. This disclosure would provide workers with key information about both the selection procedure and how workers can exercise their rights under the Civil Rights Standards and antidiscrimination laws. The proposed short-form disclosure would include:

- The position(s) for which the selection procedure is used, the characteristics the selection procedure is supposed to measure, and how those characteristics relate to the position's essential functions;
- Potential sources of discrimination identified during prior audits of the selection procedure;
- How candidates can raise concerns about the selection procedure and/or request accommodation; and
- A link where candidates can review the detailed audit summaries required by Standard 7.

Standard 7 would require employers to prepare and publish detailed summaries of each pre-deployment (Standard 3) and ongoing audit (Standard 6). These audit summaries are designed to provide enforcement agencies and workers' advocates with enough information to determine whether an investigation or complaint regarding a selection procedure is warranted. To that end, the audit summaries would:

- Describe the audit's methodology, findings, results, and conclusions for each element specified in the audit;
- Explain any changes made to the selection procedure during the course of the audit; and
- Be posted on the employer's website, filed with the relevant enforcement agency, and linked in the short-form disclosure that Standard 4 requires.

Standard 8's comprehensive recordkeeping requirements provide the final layer of disclosure and transparency. Employers would retain "all data, code, and other information necessary to allow for subsequent independent audits and investigations regarding the lawfulness and validity of the selection procedure," and provide those records to the enforcement agency or an assigned auditor upon request. Covered information would be retained for five years or for as long as the selection procedure is used, whichever period of time is greater.

Oversight and Accountability

The Standards call for strengthening oversight and accountability of selection procedures by giving candidates greater input and a right to an explanation in the selection process (Standard 5), and by providing a strong regime for enforcement and remedies (Standards 9-11).

Under Standard 5, employers and employment agencies would have to:

- Allow candidates to request accommodation or raise concerns regarding the selection procedure, and provide human review of such communications.
- Provide a post-assessment explanation that identifies the information that led to an adverse employment decision, and allow candidates to submit corrections or supplementary information in response.
- Provide accommodation, an alternative selection method, or reevaluation if necessary to avoid unlawful discrimination.
- Give candidates facing assessment by automated selection procedures the right to opt out and be assessed instead through human review, a non-automated selection procedure, or other alternative means.

In addition to barring discrimination and retaliation (see *Nondiscrimination* section), Standard 9 calls for policymakers to make it unlawful for employers and employment agencies to:

- Fail to comply with the Standards' requirements pertaining to audits, notice, disclosure, and recordkeeping.
- Provide knowingly false or materially misleading or incomplete information in the Standards' required documentation.
- Extract a candidate's biometric data available through administration of a selection procedure, or disseminate such data without the candidate's consent for commercial third-party use.

The Standards call for enforcement and remedies to be made available through both civil and administrative actions. Enforcement agencies would be able to pursue civil penalties, injunctive relief, actual damages, and other remedies against employers and vendors, as well as against auditors where applicable under Standard 9, which would also create a private right of action for certain violations of the Standards.

Standard 10(e) calls for imposing joint and several liability on all employers and vendors involved in the development and use of a discriminatory selection procedure. This approach incentivizes both employers and vendors to proactively prevent discrimination resulting from their selection procedures, and also ensures that workers have a remedy in cases where a party is judgment-proof or a jury has difficulty allocating responsibility for discrimination.

Civil Rights Standards

Standard 1: Statement of Purpose

The policies specified in these Standards would:

- a. Reinforce and elaborate upon state and federal civil laws' prohibitions against the use of employment selection procedures that have the intent or effect of discriminating against members of any protected class on the basis of a protected attribute;
- b. Require designers, developers, vendors, and employers to ensure that selection procedures are audited, both prior to deployment and regularly during the course of their use, for all forms of unlawful discrimination, including but not limited to disparate treatment discrimination, disparate impact discrimination, failure to provide reasonable accommodation, and discrimination in advertising employment opportunities;
- c. Establish notice, disclosure, and recordkeeping requirements for the users of selection procedures;
- d. Ensure that selection procedures assess candidates solely on the basis of valid measurements of essential job functions using the least discriminatory method available;
- e. Establish mechanisms for investigation and enforcement that vindicate the rights of candidates affected by the use of discriminatory selection procedures while encouraging transparency and cooperation by the users of selection procedures; and
- f. Provide adequate remedies for members of protected classes who experience unlawful discrimination, or any other violation of their legal rights, as a result of a selection procedure.

Standard 2: Definitions

For the purposes of these Standards, the following terms have the following meanings:

- a. *Accessibility*. The term “accessibility” means the degree to which workers with disabilities are able to access the functionality of, and benefits associated with, a device, good, service, or program, in a manner equally as effective as the access that others are able to utilize. In the context of selection procedures, this includes, but is not limited to, the degree to which:
 1. Workers with disabilities are able to acquire the same information, engage in the same interactions, and be assessed in a manner comparable to workers without disabilities, with substantially equivalent ease of use; and
 2. Potential access barriers for workers with disabilities have been avoided or eliminated in the design or administration of the selection procedure, or mitigated by making appropriate and effective accommodation available.
- b. *Accommodation*. The term “accommodation” means, with respect to a specific selection procedure, the provision of tools or changes to the environment or the way in which the selection procedure is usually administered that a worker can request at or before the time the selection procedure is administered. An effective accommodation is an accommodation that allows a worker with a protected attribute to access the selection procedure and be measured by it in an equally effective manner as, and on equal footing with, other workers.
- c. *Adverse impact*. The term “adverse impact” means a substantial difference in scores, selection rates, or other outputs or effects of a selection procedure that disadvantages members of any protected class in an employment decision.
- d. *Alteration*. The term “alteration” means a change made to the design, structure, functioning, or content of a selection procedure.
- e. *Alternative selection procedure*. The term “alternative selection procedure” means, with respect to a given selection procedure, a selection procedure that:
 1. Was developed or validated separately from the selection procedure at issue;
 2. Measures significantly different knowledge, skills, abilities or other characteristics; or
 3. Significantly differs from the selection procedure at issue in the method or process by which it measures the knowledge, skills, abilities, or other characteristics.

This includes a selection procedure that has undergone such a significant alteration that it satisfies subparagraph (2) or (3) of this paragraph.

- f. *Applicant*. An “applicant” is a candidate who meets the following criteria:
1. The candidate submitted an expression of interest in employment with an employer or employment agency;
 2. The employer or employment agency considered the candidate for a particular position, for multiple positions, or for employment opportunities that may arise in the future; and
 3. The candidate had not removed themselves from further consideration or otherwise indicated that they were no longer interested in working for the employer or employment agency.
- g. *Auditor*. An “auditor” is a person licensed by the enforcement agency pursuant to Standard 12(a)(3) to conduct the audits described in Standard 3 and Standard 6; who is independent of all employers, employment agencies, and other persons and entities that designed, developed, or used the selection procedure being audited; and whose methodologies for conducting such audits have been approved by the enforcement agency.
- h. *Automated selection procedure*. An “automated selection procedure” means a selection procedure that is based in whole or in significant part on machine learning, artificial intelligence, computerized algorithms, automated statistical or probabilistic modeling, or similar techniques.
- i. *Candidate*. The term “candidate” means any worker who is the subject of an employment decision made by a selection procedure, regardless of whether that worker applied for, expressed an interest in, or removed themselves from consideration for the position(s) for which the selection procedure is used.
- j. *Candidate pool*. The term “candidate pool” means the population of candidates on which the selection procedure is being used or, for a selection procedure that has not yet been deployed, the population of candidates on which the selection procedure will be used.
- k. *Deployed*. The term “deployed” means, with respect to a particular selection procedure, the period starting with the first time the selection procedure is used to make employment decisions for a position. “Pre-deployment” means the period before the first such use of a selection procedure, including but not limited to periods where the selection procedure is first being designed, developed, trained, tested, and validated for use for a position or multiple positions.
- l. *Employer*. The term “employer” means a person who retains or accepts labor or services from a worker, including an agent of such a person, which is of sufficient size to be deemed an employer pursuant to relevant employment discrimination laws.

- m. *Employment agency*: An “employment agency” is any person who procures workers for an employer; procures for workers opportunities to work for an employer; knowingly sells, offers for sale, or distributes assessments, software, or technology that is used to make or inform employment decisions; engages in a contract with an employer or another employment agency to provide services, software, or technology that collects, stores, analyzes, or interprets candidate information; or operates an online job board, platform, or other service that employers or employment agencies use to assist in the making of employment decisions. “Employment agency” also includes an agent of such a person, but does not include the developer or distributor of software or other technology if that person:
1. Was not aware that the software or other technology would be used to make employment decisions; and
 2. Placed the software or other technology in the public domain without any license or reservation of rights, or made it available under a copyleft, GPL, BSD, or similar license that allows any member of the public to copy, distribute, or modify the source code or other technology without payment, royalties, or fees.
- n. *Employment decision*. The term “employment decision” includes but is not limited to hiring, promotion, demotion, referral, retention, termination, compensation; setting the terms, conditions, or privileges of employment; selecting workers for recruitment, interviewing, or targeted job or career advertising; and licensing and certification, to the extent that licensing and certification may be covered by applicable federal, state, or local laws against employment discrimination. Other decisions, such as training or transfer, may also be considered employment decisions if they alter a worker’s terms or conditions of employment or lead to any of the decisions listed in the preceding sentence.
- o. *Enforcement agency*. The term “enforcement agency” refers to the office responsible for investigation and enforcement within the agency, department, or division of government responsible for interpreting and enforcing laws against employment discrimination.
- p. *Essential functions*. The term “essential functions” means the fundamental job duties of a position and does not include the marginal functions of the position. Essential functions are to be determined based on objective evidence such as the amount of time workers spend performing each function, the direct consequences of not requiring workers in the position to perform the function, the direct consequences of a worker failing to perform or inaccurately performing the function, the terms of any applicable collective bargaining agreement, and workers’ past and present work experiences and performance in the position in question. Past and current written job descriptions and the employer’s reasonable, non-discriminatory judgment as to which functions are essential may be evidence as to which functions are essential for achieving the purpose of the job, but may not be the sole basis

for this determination absent the objective evidence described above. “Essential functions” does not include prerequisites that the employer establishes that do not relate to the work activities of the job itself, such as being able to work all shifts, to work overtime, or to arrive at work at a specified time.

- q. *High-risk selection procedure.* The term “high-risk selection procedure” means a selection procedure that relies on analysis of a candidate’s affect or emotional state; personality; facial features or movements, body language, gait, tone of voice, vocal pitch, or pace of speech; heart rate, respiration, or other bodily functions regulated by the autonomic nervous system; or any other technique or methodology identified by the enforcement agency as creating an especially high risk of unlawful discrimination.
- r. *Interactive process.* The term “interactive process” means an informal communication or series of communications with a candidate with a disability to clarify whether the candidate requires accommodation or an alternative selection procedure under applicable law and to identify appropriate accommodation(s) or alternative selection procedure(s).
- s. *Opt-out.* The term “opt-out” means a candidate’s decision not to be assessed by an automated selection procedure, and to instead be assessed through human review or an alternative, non-automated selection procedure.
- t. *Person.* The term “person” includes any natural person, entity, public body, trust, or unincorporated organization.
- u. *Position.* The term “position” means a particular job or role at a particular employer.
- v. *Protected attribute.* The term “protected attribute” means a personal characteristic or trait that is protected from employment discrimination under any applicable federal, state, or local law.
- w. *Protected class.* The term “protected class” means a group or class of persons sharing one or more protected attributes in common.
- x. *Proxy.* The term “proxy” in the phrase “proxy for a protected attribute” means a facially neutral attribute or set of attributes that are so closely associated with or predictive of a protected attribute that the selection procedure’s use of the facially neutral attribute(s) has substantially the same practical effect on workers with the attribute(s) as use of the protected attribute(s) themselves.
- y. *Selection procedure.* The term “selection procedure” means any measure, combination of measures, test, method, or process to assess workers that meets the following criteria:

1. It is:
 - A. Sold or distributed by an employment agency; or
 - B. Used to assess at least 100 workers per year;
2. It outputs a score, ranking, recommendation, evaluation, or other judgment that is primarily the result of:
 - A. Automated processes, including processes that are based in whole or in significant part on machine learning, artificial intelligence, computerized algorithms, automated statistical or probabilistic modeling, or similar techniques; and/or
 - B. Standardized processes, whether automated or non-automated, where outputs are generated algorithmically or deterministically; and
3. The output described in paragraph (2) is used as a basis for any employment decision, as a factor in any employment decision, to provide a recommendation with respect to any employment decision, or to assist, influence, or inform human decision-makers or automated systems in the making of any employment decision.

A group of two or more purported selection procedures, each of which satisfies both paragraph (2) and paragraph (3), is the same purported selection procedure for purposes of paragraph (1)(B) if they derive from the same development process; were tested, trained, or validated together; rely on common data; or are marketed under common trade or product names.

- z. *Small employer.* The term “small employer” means an employer with fewer than 15 full-time equivalent workers.
- aa. *Use.* To “use” a selection procedure means to utilize the selection procedure to make an employment decision about a candidate.
- ab. *Validity.* The term “validity” means the extent to which a selection procedure is an accurate and effective means of measuring the essential job functions that it purports to measure, using the principles of test validation under contemporary standards of social science at the time the selection procedure is used, but a selection procedure is not valid for purposes of these Standards if the evidence for validity is based solely on correlation between the output of the selection procedure and measures of job performance, unless the employer or employment agency using the selection procedure supports the correlational evidence with theoretical, logical, or causal reasoning sufficient to explain why the specific attributes measured by the selection procedure should be predictive of the ability to perform essential job functions.
- ac. *Worker.* The term “worker” means an employee, contractor, paid or unpaid intern, applicant, or any other person who offers or provides labor or services in exchange for compensation or other benefits. “Worker” also includes any individual who is

considered part of the labor force under the applicable standards and guidance issued by the United States Department of Labor's Bureau of Labor Statistics, regardless of whether the individual is currently working. In any proceeding involving the terms of these Standards or brought under any provision of these Standards, a plaintiff or complainant who claims to be a worker should be presumed to be a worker, and the employer or employment agency answering the complaint or other action should bear the burden of demonstrating that the plaintiff or complainant is not a worker.

Standard 3: Pre-deployment audits

An employer or employment agency should not use a selection procedure unless:

- a. Prior to procurement or use of the selection procedure, an auditor has examined the employer's existing employment decision practices to identify disparities between protected classes, the use of proxies for protected attributes, and other potentially discriminatory patterns of disparate treatment and disparate impact on protected classes with respect to employment decisions. This requirement should not apply to small employers.
- b. An auditor has conducted a pre-deployment audit on the selection procedure for each position for which the selection procedure is to be used. Each employer and employment agency that uses, sells, distributes, or develops the selection procedure should have a joint and non-delegable responsibility for ensuring that an audit compliant with this Standard is performed before the selection procedure is deployed. Such employers and employment agencies may enter into contracts assigning obligations, duties, and indemnification responsibilities relating to the conduct of a pre-deployment audit, but such contracts should not abrogate any party's duty to ensure that a proper audit is conducted or liability under these Standards in the event of non-compliance. The audit should:
 1. Identify and describe essential functions for each position for which the selection procedure will be used to evaluate candidates, explain why these functions are in fact essential, and demonstrate that the selection procedure is scientifically valid in measuring candidates' ability to perform these essential functions;
 2. Identify and describe the methods and techniques used to design the selection procedure, the attributes and criteria on which the selection procedure relies, and any other input or aspect of the design, development, validation, or testing of the selection procedure that the enforcement agency determines necessary;
 3. For any automated selection procedure, describe the sources of the training/modeling data, and the steps taken to ensure that the training data and samples are accurate and representative in light of the position's candidate pool;
 4. Determine whether the decisions, recommendations, scores, or other outputs of the selection procedure have an adverse impact on members of any protected class using at least one reasonable and appropriate test of statistical significance and one reasonable and appropriate test of effect size, describe the tests of statistical significance and effect size used to test for the presence and extent of such adverse impacts, and describe the nature and extent of any adverse impacts detected;
 5. Determine whether the administration of the selection procedure or its results limits accessibility for persons with disabilities, or for persons with any specific disability;

6. Explain whether and how the selection procedure reduces or otherwise addresses discriminatory outcomes identified under paragraph (a), or risks of such outcomes, in the employer's employment practices;
 7. Consider and describe potential sources of adverse impact against protected classes that may arise after the selection procedure is deployed;
 8. Identify and describe any attributes on which the selection procedure relies and determine whether the selection procedure engages in disparate treatment by relying on any protected attribute or any proxy for a protected attribute to make an employment decision;
 9. Determine, for any adverse impacts or limitations on accessibility detected during the audit, whether alterations to the selection procedure can be made, whether effective accommodation can be provided, and whether less discriminatory alternative selection procedures or other assessment methods are available, that would mitigate the adverse impact or limitation on accessibility while retaining validity in measuring candidates' ability to perform essential functions;
 10. Identify any other reasonable alterations needed to ensure that the selection procedure validly and effectively assesses the abilities of candidates from each protected class to perform each position's essential functions; and
 11. Include any other information or sections required under the rules and regulations of the enforcement agency.
- c. If the pre-deployment audit described in paragraph (b) of this Standard identifies any reliance on any protected attribute or proxy for a protected attribute, adverse impact, or limitation on accessibility:
1. The employer alters the selection procedure, and the auditor reasonably determines that the alteration eliminated the potential discrimination or inaccessibility;
 2. In the case of reliance on a protected attribute or proxy for a protected attribute, the auditor reasonably determines that the use of the attribute or proxy is lawful pursuant to a valid affirmative action plan;
 3. In the case of adverse impact or limitations on accessibility, the auditor reasonably determines that the selection procedure is both valid and the least discriminatory method of assessing the candidate's ability to perform the essential job function(s); or
 4. In the case of limitations on accessibility, if the limitation cannot be eliminated by alteration, the employer or employment agency makes an effective accommodation available.
- d. The employer incorporates and implements any reasonable alterations identified in paragraph (b)(10) and/or paragraph (c) of this Standard; and

- e. The conditions and manner in which the employer uses the selection procedure, and purpose for which the employer uses the procedure, comport with the specifications of the selection procedure as implemented after the incorporation of alterations described in paragraph (d) of this Standard.

Standard 4: Short-form disclosures

- a. Any employer or employment agency that uses a selection procedure should prepare a short-form disclosure for each such selection procedure that:
 1. States the positions for which the selection procedure is or will be used and what types of employment decisions will be made or informed by the selection procedure;
 2. Describes, for each position:
 - A. The knowledge, skills, abilities, and other characteristics that the selection procedure measures;
 - B. How those characteristics relate to the position's essential function(s);
 - C. How the selection procedure measures those characteristics; and
 - D. How to interpret the results or other outputs of the selection procedure;
 3. Identifies any reasonably foreseeable accommodation that candidates may require;
 4. States that candidates may, and provides up-to-date information on how to, request and access any accommodation or alternative selection procedures, communicate concerns regarding the selection procedure, file a civil or administrative complaint with the enforcement agency, or submit supplementary information, as set forth in Standard 5;
 5. For an automated selection procedure, specifies when and how the candidate may opt out of the selection procedure, as described in Standard 5(a)(4); and
 6. Clearly identifies a functional URL that links to the detailed summaries, described in Standard 7, of all prior audits and adverse impact assessments.
- b. The short-form disclosure should be:
 1. Provided in English, in any non-English language spoken by at least one percent (1%) of the population of the relevant jurisdiction, and in any other language that the employer or employment agency regularly uses to communicate with workers or candidates;
 2. Written in clear and plain language;
 3. Made available in formats that are accessible to people who are blind or have other disabilities; and
 4. Otherwise presented in a manner that ensures the disclosure clearly and effectively communicates the required information to candidates.
- c. An employer or employment agency should attach or conspicuously link an up-to-date short-form disclosure within each posting, advertisement, or recruitment communication regarding a position for which it may use the selection procedure.
- d. Except as provided in paragraph (e) of this Standard, each employer or employment agency that uses a selection procedure should:

1. Provide the short-form disclosure to each candidate within a reasonable time prior to the use of the selection procedure;
 2. Publish short-form disclosures for all selection procedures the employer uses on the employer's website and on any platform used to receive or process applications for a position or positions. The form should be posted in PDF or HTML format, or in another digital format if that format has been authorized by the enforcement agency.
- e. When an employment agency uses a selection procedure on behalf of a small employer, the employment agency should be responsible for fulfilling the obligations set forth in paragraph (d) of this Standard.

Standard 5: Procedure to request accommodation, communicate concerns, or opt out; right to explanation for adverse actions

- a. Prior to using a selection procedure on an applicant, an employer or employment agency should:
 1. Provide each applicant who may be subjected to the selection procedure a copy of the short-form disclosure described in Standard 4, and in accordance with the format and accessibility standards specified in that Standard;
 2. Provide the applicant with a meaningful opportunity to request accommodation or an alternative selection procedure or other assessment method, or to otherwise communicate concerns to the employer or employment agency regarding the selection procedure's ability to validly evaluate the applicant's ability to perform the position's essential functions;
 3. Engage in an interactive process with candidates with disabilities if the candidate requests accommodation or if the employer or employment agency knows of the candidate's need for accommodation; and
 4. If the selection procedure is an automated selection procedure, allow the applicant to opt out of using the selection procedure and assess the applicant through human review, a non-automated selection procedure, or other means of assessment, on equal footing with applicants who are assessed through the automated selection procedure.
- b. After subjecting a candidate to a selection procedure, an employer or employment agency should:
 1. Provide an explanation that identifies the factors, candidate characteristics, and other information that led the selection procedure to render an adverse employment decision with respect to each position for which the selection procedure assessed the candidate; and
 2. Provide the candidate with a meaningful opportunity to submit corrections or otherwise provide supplementary information challenging factors identified under Standard 5(b)(1) and/or the selection procedure's overall ability to validly measure the candidate's ability to perform the position's essential functions.
- c. Employers and employment agencies who receive requests, corrections, or other information from candidates pursuant to Standard 5(a)(2) or (b)(2) should be deemed to have knowledge of the information and requests for accommodation included therein and do each of the following within a reasonable amount of time after receiving the information from the candidate:
 1. Assign a natural person to review the information before the employment decision is finalized and determine whether provision of an accommodation, assessment of the candidate by alternative means, or reevaluation is warranted or needed to ensure compliance with applicable antidiscrimination laws;

2. Provide the requested accommodation, reevaluation, or alternative selection procedure or other assessment method if failing to do so would create a substantial risk of unlawful discrimination;
 3. If the candidate made a specific request, provide the candidate with a decision on whether the candidate's request will be granted or denied, and the reasons for that decision;
 4. Retain all information and documentation relating to the candidate communication in accordance with Standard 8; and
 5. Review the information and requests as part of the next ongoing audit, as described in Standard 6.
- d. A candidate's failure to invoke or utilize any of the rights or procedures described in this Standard should not:
1. Waive or affect the availability of any other rights, procedures, or remedies under these Standards or any other applicable law; or
 2. Be used as evidence regarding the candidate's need for accommodation if the candidate is ultimately selected or hired.

Standard 6: Ongoing audits and adjustments

- a. After a selection procedure has been deployed, the selection procedure should undergo ongoing audits at standardized intervals that ensure the selection procedure is audited at least once per year for each position for which the selection procedure is used. Each employer and employment agency that used, sold, distributed, or developed the selection procedure should have a joint and non-delegable responsibility for ensuring that ongoing audits compliant with this Standard are performed. Each ongoing audit should be conducted by an auditor who analyzes and documents in detail:
 1. Whether and how each position's essential functions, the characteristics of the candidate pool, or other features relevant to the validity of the selection procedure have changed since the pre-deployment audit and, if applicable, the last ongoing audit;
 2. If the parameters, training data, or other input components of the selection procedure have changed, whether and how the changes have affected the pre-deployment audit determinations identified in Standard 3(b);
 3. Whether and how the decisions, recommendations, scores, or other outputs of the selection procedure have had an adverse impact on members of any protected class, using the adverse impact testing standards and procedures specified in Standard 3(b)(4);
 4. Any new sources of adverse impact that may arise if the employer continues to use the selection procedure;
 5. The effectiveness of efforts to mitigate any potential adverse impacts identified during the pre-deployment audit;
 6. What new or additional reasonable alterations to the selection procedures or individual accommodation, if any, would improve the selection procedure's accessibility or its ability to fairly, validly, and effectively assess the abilities of candidates from each protected class to perform each position's essential functions without unlawful discrimination;
 7. The clarity and completeness of the adverse action notices and explanations described in Standard 5(b)(1);
 8. Whether less discriminatory alternative selection procedures or other methods of assessment are now available, using the standards specified in Standard 3(b)(9); and
 9. Any other information or issues required under the rules and regulations of the enforcement agency.
- b. An employer or employment agency should cease use of a selection procedure for a position if the ongoing audit reveals:
 1. That the essential functions of the position have changed since the pre-deployment audit, in which case the employer should conduct a new pre-deployment audit;

2. Any reliance on any protected attribute or proxy for a protected attribute, unless it has been eliminated or the auditor reasonably determines that the use of the attribute or proxy is lawful pursuant to a valid affirmative action plan;
 3. Any adverse impact or limitation on accessibility, unless it has been eliminated or:
 - A. The employer demonstrates that the selection procedure is the least-discriminatory valid method for assessing candidates' ability to perform essential job functions; or
 - B. In the case of a limitation on accessibility, if the limitation cannot be eliminated by altering the selection procedure, the employer or employment agency makes effective accommodation available; or
 4. A need for a reasonable alteration as described in paragraph (a)(6) of this Standard, until the employer has implemented the reasonable alteration.
- c. An ongoing audit that complies with this Standard should be conducted regardless of whether a selection procedure has been changed since the pre-deployment audit.
- d. The enforcement agency should have discretion to require certain selection procedures to be audited more frequently than once per year, but nothing in these Standards should be construed as:
1. Suggesting that the enforcement agency allow any selection procedure to be audited less frequently than once per year; or
 2. Preventing employers from auditing selection procedures more frequently than once per year, or more frequently than required under applicable rules and standards issued by the enforcement agency.

Standard 7: Audit summaries

- a. For each pre-deployment or ongoing audit, the auditor should prepare a detailed summary of the results of each audit that:
 1. Includes the plain-language definitions of “audit” and “adverse impact” published by the enforcement agency;
 2. Explain that the selection procedure was audited and tested for adverse impacts and why;
 3. Describes the audit’s methodology, findings, results, and conclusions for each element described in Standard 3(b) or Standard 6(a), as applicable;
 4. If the selection procedure has been changed pursuant to Standard 3(e), Standard 6(a)(6), or Standard 6(b)(2)-(4), or ceased pursuant to Standard 6(b), describes the nature of any changes and the reasons for any changes or cessation of use of the selection procedure; and
 5. Includes all other information required by the rules and regulations of the enforcement agency.
- b. Each employer and employment agency should:
 1. If it has a website, post each audit summary on a dedicated and easily searchable page on its website, and keep the audit summary posted throughout the period of time that the employer or employment agency uses the selection procedure or any derivative of the selection procedure and one year thereafter; and
 2. File each audit summary with the enforcement agency, which should catalog audit summaries for each employer and employment agency and post them on the enforcement agency’s website within seven days of filing.

Standard 8: Recordkeeping

- a. For each selection procedure that the employer or employment agency uses, sells, or distributes, the employer and employment agency should retain all data, code, and other information in their possession or control necessary to allow for subsequent independent audits and investigations regarding the lawfulness and validity of the selection procedure, including:
 1. All documentation of impact and validity evidence for the selection procedure required under Section 15 of the Uniform Guidelines on Employee Selection Procedures, 41 CFR § 60-3.15;
 - A. Subject to subparagraph (1)(B), an employer or employment agency with fewer than 100 full-time equivalent workers should be permitted to retain simplified records relating to a selection procedure, as specified in paragraph (A)(1) of 41 CFR § 60-3.15, except the records should include the simplified information about all protected classes, and should include information on all protected attributes shared by more than one percent (1%) of the labor force for which the selection procedure is used.
 - B. An employment agency that sells selection procedures to employers; distributes selection procedures to employers; or administers, distributes, or uses selection procedures on employers' behalf, should retain full records of the impact and validity of the selection procedures, regardless of the size of the employment agency.
 2. Copies of every version of each short-form disclosure described in Standard 4 and every audit summary described in Standard 7;
 3. Records of each request for accommodation, request to be assessed by alternative means, or other communication received pursuant to Standard 5; the employer or employment agency's response to each such communication or request; and the reasons for its response;
 4. Records of each notice and explanation of adverse action described in Standard 5(b)(2);
 5. Any other information that the employer or employment agency relied upon when it decided whether to develop, procure, or use the selection procedure;
 6. Any other information that an auditor relied upon during a pre-deployment or ongoing audit; and
 7. All other data, code, records, or other information regarding the selection procedure required under the rules and regulations of the enforcement agency.
- b. All documentation, data, results, and other records and information described in this Standard should:
 1. Be retained for as long as the selection procedure is used, or five years, whichever period is greater;
 2. Be secured during the retention period in a manner conforming to

- contemporary cybersecurity industry standards, applicable health and personal information privacy laws, and any rules or regulations issued by the enforcement agency; and
3. Upon request, be securely provided to the agency or agencies responsible for enforcing laws relating to employment discrimination, as well as to any auditor duly retained to conduct an audit on a selection procedure to which the information is relevant.
- c. To the extent that the records and information retained under this Standard include a candidate's personal information, the employer or employment agency should not sell, release, transfer, provide access to, or divulge in any manner the records or information to any third party, other than to an auditor in connection with an audit that is in progress, unless:
1. The third party is an employer or employment agency that wishes to make an employment decision regarding the candidate;
 2. The employer or employment agency provides the candidate with a notice that discloses:
 - A. The name of the third party to which the employer or employment agency wishes to disclose the candidate's personal information;
 - B. The position(s) for which the third party wishes to make an employment decision;
 - C. The type(s) of employment decision the third party wishes to make regarding the candidate; and
 - D. The specific personal information that will be disclosed to the third party and the purpose that this information will fulfill in the third party's decision; and
 3. The employer or employment agency obtains separate, specific, and affirmative written consent from each such candidate with respect to each position with each third party.
- d. Candidates should be protected from retaliation if they refuse to consent to the sharing of their data under paragraph (c) of this Standard.

Standard 9: Unlawful employment practices

Laws and regulations should be updated, amended, enacted, or promulgated to the extent necessary to ensure that the following are deemed unlawful employment practices under the laws of the jurisdiction in question:

- a. For an employer or employment agency to use, sell, or distribute a selection procedure:
 1. That contains technical elements, methods, or features that individually, or in concert, result in adverse impact for protected classes, unless the employer or employment agency establishes that the selection procedure was the least discriminatory valid method of measuring candidates' ability to perform the essential functions of the relevant position(s) at the time the employer or employment agency used the selection procedure;
 2. That contains elements, methods or features that individually, or in concert, use protected attributes, or proxies of such attributes, to limit, segregate, classify, or deprioritize candidates for employment opportunities, including but not limited to selection procedures that use such attributes to make less available an advertisement for an employment opportunity on a physical or digital media platform or webpage, except pursuant to a valid affirmative action plan;
 3. Without altering the selection procedure or providing reasonable accommodation, where such alteration or accommodation is needed to ensure that the selection procedure validly and effectively assesses the abilities of candidates from a protected class to perform each position's essential functions;
 4. That, if used, would otherwise violate any prohibition on discrimination or lead to any unlawful employment practice under the laws of the United States or this jurisdiction;
 5. That contains elements, methods, or features that cannot be adequately assessed for scientific validity, accuracy, or compliance with the provisions of these Standards;
 6. That has not been subjected to the audits described in Standards 3 and 6 or that unreasonably continued to use a selection procedure without adequately addressing potential sources of discrimination or invalidity identified during such audits;
 7. Without complying with the notice, disclosure, candidate communication, and other provisions of Standards 4 and 5;
 8. For which compliant summaries have not been prepared, submitted, or published, as described in Standard 7;
 9. Without maintaining all records, or without providing all relevant information to an enforcement agency or auditor, in accordance with Standard 8; or
 10. That is a high-risk selection procedure.

- b. For any person to retaliate or otherwise discriminate against a worker for:
 - 1. Requesting a reasonable accommodation;
 - 2. Opting out of an automated selection procedure;
 - 3. Otherwise exercising their rights under these Standards;
 - 4. Opposing any practice or conduct that the worker reasonably believes to be prohibited by these Standards; or
 - 5. Making a charge, testifying, assisting, or participating in any manner in an investigation, proceeding, or hearing under these Standards.
- c. For an employer or employment agency to extract biometric data made available through a selection procedure or disseminate such data to unauthorized third parties for commercial purposes.
- d. For an employer, employment agency, or auditor to include knowingly false, materially misleading, or materially incomplete information in an audit summary, notice, enforcement agency filing, or other documentation required under these Standards.
- e. For an employer to use an automated selection procedure without first obtaining preclearance from the enforcement agency, as described in Standard 11(c).

Standard 10: Enforcement, remedies, and liability

Laws and regulations should be updated, amended, enacted, or promulgated to the extent necessary to ensure that the following enforcement mechanisms and remedies are available:

- a. The enforcement agency should have authority to bring an administrative or civil action against an employer or employment agency for any unlawful employment practice described in Standard 9, paragraph (a) or (b). If the enforcement agency proves that the alleged unlawful employment practice occurred, the agency or court should:
 1. Assess a civil penalty for each calendar day that an employer or employment agency used the selection procedure that was the subject of the unlawful employment practice;
 2. Enjoin the employer or employment agency from continuing to use the selection procedure, or any related or derivative selection procedure, that was the subject of the unlawful employment practice;
 3. Issue any other such orders as the agency or court deems necessary to eliminate the effects of the unlawful employment practice and prevent future violations, including but not limited to the payment of actual damages to affected candidates.
- b. The enforcement agency should have authority to bring an administrative or civil action against an auditor for any unlawful employment practice described in Standard 9, paragraph (d). If the enforcement agency proves that the alleged unlawful employment practice occurred, the agency or court should assess a civil penalty.
- c. The enforcement agency should have authority to bring a civil action in any court of competent jurisdiction to enjoin the use, sale, or marketing of any selection procedure that the enforcement agency has reasonable cause to believe has resulted, is resulting, or will result in an unlawful employment practice.
- d. Any person who has been subjected to any unlawful employment practice described in Standard 9, paragraph (a), subparagraphs (1), (2), (3), (4), (7), or (10), or described in Standard 9, paragraphs (b) or (c), should be able to file a complaint with the enforcement agency or a civil action in any court of competent jurisdiction.
 1. In any civil action under this paragraph, a prevailing plaintiff, or, in a successful class action, each member of the class, should be able to obtain, for each unlawful employment practice, the greater of:
 - A. \$500 per prevailing plaintiff or, in a successful class action, per each member of the class; or
 - B. Actual damages.

2. In any civil action under this paragraph, a court should:
 - A. Award a prevailing plaintiff costs and reasonable attorney fees;
 - B. Enjoin the employer from continuing to use the selection procedure, or any related or derivative selection procedure, that was the subject of the unlawful employment practice; and
 - C. Issue any other such orders as the court or agency deems necessary to eliminate the effects of the unlawful employment practice and prevent future violations.
- e. In any civil action claiming that an employment decision made using a selection procedure has violated an applicable law against employment discrimination, each employer and employment agency that used, sold, distributed, or developed the selection procedure used in the employment decision should be jointly and severally liable to a prevailing plaintiff for all damages awarded to that prevailing plaintiff for the unlawful discriminatory practice and any liquidated damages awarded under paragraph (b) of this Standard, except that in instances where an employment agency knowingly sells, provides, or distributes a selection procedure to a small employer, the small employer should not be liable for any unlawful discriminatory practices or other unlawful acts of the employment agency.

Standard 11: Enforcement agency authority and responsibilities

- a. The enforcement agency should promulgate rules:
 - 1. Detailing the degree and content of human review required when an employment decision is made by an automated selection procedure;
 - 2. Providing effective complaint processes for workers claiming to be aggrieved by an employer's use of a selection procedure;
 - 3. Establishing certification procedures for auditors and their methodologies for conducting pre-deployment and ongoing audits, as well as the standards that auditors' methodologies must meet to obtain certification;
 - 4. Establishing standards for what types and degrees of disparity constitute an adverse impact;
 - 5. Specifying any situations other than those described in the text of these Standards under which a candidate may choose to opt-out of a selection procedure;
 - 6. Specifying what additional information employers must include in the short-form disclosure described in Standard 4, and providing plain-language definitions of "audit," "adverse impact," and any other technical terms used in short-form disclosures;
 - 7. Specifying what findings, data, conclusions, and other information must be included in the summaries described in Standard 7;
 - 8. Specifying what information employers must retain as part of their recordkeeping obligations under Standard 8, as well as the cybersecurity and personal privacy standards governing any information retained under Standard 8;
 - 9. Designating additional selection procedures as high-risk selection procedures beyond those specifically enumerated in Standard 2(q), but the enforcement agency's rulemaking authority should not extend to removing the high-risk designation from the types of selection procedures specifically enumerated in Standard 2(q);
 - 10. Establishing reporting, investigation, and enforcement procedures that incentivize cooperation by employers, employment agencies, and other parties subject to these Standards and related laws against employment discrimination; and
 - 11. Providing all other guidance, procedures, and interpretations necessary or beneficial to implementing and enforcing the provisions of these Standards.
- b. The enforcement agency should publish example forms for the short-form disclosure, pre-deployment audit summary, ongoing audit summary, and any other documents, disclosures, or records that the enforcement agency deems necessary or beneficial, and should update such example forms when warranted.

- c. The enforcement agency should establish a preclearance process for automated selection procedures, under which employers must demonstrate that an automated selection procedure is the least discriminatory valid method of measuring essential job functions before it may be used for a particular position.
- d. The enforcement agency should update its rules as appropriate to keep pace with changes in employment discrimination laws, technological advancements, the nature and content of selection procedures, and the social science of test validation.

Standard 12: Construction

- a. The remedies described in these Standards are intended to be non-exclusive, and the provisions of these Standards should not be implemented in a manner that limits any preexisting right of any candidate to bring a civil action in any court of competent jurisdiction, limits the authority of appropriate agencies to enforce employment discrimination laws, or prohibits localities from enacting laws that provide greater or additional protection.
- b. The existence of Enforcement Agency preclearance to use an automated selection procedure under Standard 11(c) should not alter, eliminate, or otherwise affect any person's liability or remedies under other applicable law, including but not limited to laws relating to employment discrimination, nor should evidence of preclearance be admissible in any action involving a claim of an unlawful employment practice, including but not limited to actions for employment discrimination.

The Standards provide a concrete alternative to recent proposals that would set very weak notice, audit, and fairness standards for automated tools. They also map out a more rigorous and rights-focused approach as compared to the outdated rules that currently govern how employers assess whether their selection procedures are discriminatory and whether they actually measure the worker characteristics they claim to measure.

For more information on this work, contact the
Center for Democracy & Technology:

Matt Scherer

Senior Policy Counsel for Worker Privacy
mscherer@cdt.org

Ridhi Shetty

Policy Counsel, Privacy & Data Project
rshetty@cdt.org

cdt.org





Center for Industry
Self-Regulation



National
Programs

EPSTEIN
BECKER
GREEN

PRINCIPLES FOR

Trustworthy AI in Recruiting & Hiring

© The Center for Industry Self-Regulation, 2023. All Rights Reserved. BBB National Programs Charitable Foundation d/b/a Center for Industry Self-Regulation.

I. Goals for the Use of Artificial Intelligence (AI) in Recruiting and Hiring

Algorithmic systems using AI technologies, when responsibly designed, deployed, and monitored, can promote fairer outcomes and inclusivity by reducing the potential for bias that may exist in human-led decision-making during the recruiting and hiring process. The development and implementation of trustworthy AI¹ therefore should focus on the following goals:

- 1) Ensuring systems are valid and reliable
- 2) Promoting fair outcomes, with harmful bias managed
- 3) Increasing inclusivity
- 4) Facilitating compliance, transparency, and accountability
- 5) Striving for systems that are safe, secure, resilient, explainable, interpretable, and privacy-enhanced

II. Purpose of the Principles

Employers must be committed to robust, achievable, and fair standards for the use of AI in recruitment and hiring. Practical baseline standards for trusted business practices are necessary to provide appropriate guidance for Employers and Vendors seeking to properly leverage AI technology in recruitment and hiring processes, provide protection for Applicants and employees, send marketplace signals to showcase best-in-class examples of responsible implementation and utilization, and demonstrate to regulators that new technologies are being developed and deployed responsibly.

Employers will evaluate the potential benefits and societal risks associated with these methods, as they may have the ability to identify conscious and unconscious human biases and barriers to inclusion. The ultimate objective is to promote positive changes in the hiring process. Important decisions are subject to human direction and review as necessary, the extent of which depends on the application and the impact it has on the decision-making process (e.g., final selection, filtering selection, testing, etc.).

III. Scope

These self-regulatory principles serve as a global baseline standard for the use of artificial intelligence applications by employers for the purpose of recruitment and hiring. The Principles are meant to inform best practices when advanced adaptive algorithms are used as part of the employment selection process. Excluded from scope are other Employer activities that may make use of AI applications, including Passive Hiring, as defined below, and non-Selection Activities. Also excluded from scope are those Selection Activities that make use of simple algorithms that do not meet the definition of AI Application below (e.g., functions in a spreadsheet). The Principles do not supersede or replace national or local laws, regulations, or best practices.² Organizations (both Employers and Vendors) that commit to these principles agree to apply them as a consistent minimum set of practices across their global operations.

IV. Definitions

- A. Applicant: Individual who has expressed and maintained interest in a specific position in the manner specified by the Employer's application requirements.
- B. Artificial Intelligence (AI) Application: an algorithmic system that includes machine learning or deep learning techniques—or future technologies with similar complexity—and is used as part of a Selection decision.
- C. Artificial Intelligence (AI) Processing: the use of an AI Application.

¹ See National Institute of Standards and Technology Risk Management Framework 1.0 (Jan. 2023), <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>, at 12.

² The Principles are compatible with a holistic governance approach to AI risk management, such as that defined in the NIST AI RMF 1.0.

The Center for Industry Self-Regulation (CISR) is a 501(c)3 non-profit organization (FEIN: 84-5021924) and is a DBA of BBB National Programs Charitable Foundation. 1676 International Drive, Suite 550 | McLean, VA 22102

- D. Employer: organization seeking to hire individuals.
- E. Passive Hiring: the process of waiting for candidates to respond to job postings or career pages; the Employer does not affirmatively seek to identify relevant and qualified non-Applicants to fulfill a current or future staffing need.
- F. Personal Data: information about a natural person or data linked to such information, where processed to uniquely identify a natural person.
- G. Recruiting Activities (or “Sourcing”): the process of identifying, researching, and networking with potential job candidates with the intent to generate Applicants.
- H. Selection Activities: the process of evaluating Applicants for existing staffing needs and selecting which Applicants move forward in the hiring process or, ultimately, are selected for hire. (E.g., reviewing submitted resumes or screening candidates through mechanisms such as interviews and/or assessments.)
- I. Sensitive Personal Data: Personal Data that reveals information with enhanced legal protections (e.g., racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, genetic data, or biometric data.).
- J. Vendor: company that creates/markets AI Applications for end-user use by Employers.

V. Transparency

A. Objective

For Applicants: If AI Applications are being used, Applicants should be able to make an informed choice about whether to proceed at the appropriate stage in the recruiting or hiring process.

For Employers: Employers using AI should partner with Vendors to understand how the AI Application operates and to ensure the AI Application is designed and maintained consistent with the other self-regulatory principles.

B. Employers: General Notice if the Use of AI Processing

1. **Timing of the notice.** Employers should provide Applicants with an opt-out notice that AI Processing will be applied to personal data collected about them. Such notice may be provided to the Applicant at the time of application, or in any case prior to the Employer’s utilizing AI Processing to Select/Hire the Applicant.
2. **Content of the notice.** The initial notice of AI Processing should be written in clear, plain language and should link to the Employer’s privacy policy or other policy describing the AI Processing.³ Following is a model template that can be tailored by the Employer to fit specific use scenarios:

*To help us select the most qualified candidates from the many applications we receive, while reducing the potential risk of human bias, we use AI processing to evaluate your application. The AI processing used meets or exceeds the standards for fairness and inclusivity established in the **INDEPENDENT CERTIFICATION PROTOCOLS FOR AI-ENABLED HIRING AND RECRUITING TECHNOLOGIES**. We may use information you submit to consider you for other positions as well. For more information, see [company privacy*

³ The opt out described in this sample notice is meant to be separate and distinct from any ADA accommodation obligations. See Section VI(B).

policy]. You may opt-out of this process, but by doing so, you will not be considered for the position, unless otherwise required by relevant law.⁴

C. Employers: Detailed Notice of the Use of AI Processing

In its privacy policy, a dedicated notice, or communication from HR, the Employer should specify, in clear, plain language:

1. Whether its AI Processing meets the Principles for Trustworthy AI in Recruiting and Hiring and the **INDEPENDENT CERTIFICATION PROTOCOLS FOR AI-ENABLED HIRING AND RECRUITING TECHNOLOGIES**; or similar protocols as warranted.
2. The Personal Data elements, tasks, and inferences subject to AI Processing with sufficient detail to enable the Applicant to determine whether an accommodation is necessary and consistent with current legal accommodation requirements, and/or providing an opt-out option.
3. The sources of data elements subject to AI Processing, if not collected from the Applicant.
4. In accordance with disability regulations, the option for applicants to request an accommodation to proceed (separate from the general opt-out option).
5. Any choices the Employer provides to limit or review AI Processing, such as human review of outputs.
6. The functions of the AI Processing, such as the stages of the selection process at which it may be applied (e.g., resume review, assessment stage of the hiring process, etc.) the modalities through which it operates (text, images, audio), and the dimensions being assessed.
7. The nature of any specific automated decisions made through AI Processing (e.g., behavioral or skill assessments, video interviewing).
8. Whether the purpose of the AI Processing is limited to recruitment and selection.
9. Whether data will be used for future hiring decisions on other job openings.

D. Employers: Periodic Review of AI-Enhanced Outputs

1. Employers commit to appropriate examination of outputs where AI-enhanced methods are used for Selection Activities, including AI Processing to promote fairer and more inclusive outcomes and to reduce the potential for biased results, consistent with the U.S. Uniform Guidelines on Employee Selection Procedures (41 CFR 60-3) or other applicable legal requirements. The examination will be appropriately tailored to the Employer's individual circumstances and specific uses of AI Applications.
2. When making use of AI Applications developed by Vendors, Employers agree to appropriately vet Vendors using either:
 - a. Employer's own industry-standard⁵ internal validation processes or
 - b. A requirement for all Vendors to demonstrate:
 - (1) a third party-certified process validation or

⁴ If the Employer wishes to provide an alternative selection process for those who opt-out from AI processing, they may do so consistent with these Principles, but it is not required. At a minimum, Employers commit to provide the choice to opt-out from consideration for those who prefer their personal data not be subject to AI processing.

⁵ In conducting a review, employers are encouraged to utilize a third party, or, if using an internal process, implement safeguards to ensure independence of the review.

- (2) an internal validation by the Vendor designed to demonstrate basic fairness and to facilitate employer compliance with the U.S. Uniform Guidelines on Employee Selection Procedures (41 CFR 60-3) or other applicable legal requirements.

E. Vendors: Transparency about Commitment to Standards

In a dedicated notice to Employer or Applicant (as appropriate) the Vendor will specify:

1. Whether its AI Processing meets the Principles for Trustworthy AI in Recruiting and Hiring and the **INDEPENDENT CERTIFICATION PROTOCOLS FOR AI-ENABLED HIRING AND RECRUITING TECHNOLOGIES**;
2. A list of the steps it takes to assure the inputs and functions of AI Processing have been evaluated to ensure a level of data set bias that is minimal and job related;
3. A list of the techniques it deploys to promote fair, achievable, and bias-neutral outputs; and
4. The specific purpose(s) for which the AI Application is designed to be used by Employers.⁶ Vendors are not expected to validate an Employer's use of the AI Application, unless contractually obligated to do so.

VI. Fairness, Non-Discrimination, Technical Robustness, and Safety

Throughout the AI lifecycle (pre-design, design, development, and deployment) AI Processing should be evaluated with reference to three components: inputs, functions, and outputs. Optimally, organizations should rely on the scientific method and a holistic approach to testing, evaluation, validation, and verification (TEVV).

- A. Inputs of AI Processing should be evaluated with the following general goals as applicable to the technology used:
 1. Rely on scientifically rigorous approaches to AI Application development/techniques.
 2. Ensure model data set integrity to reduce the impact of systemic biases⁷ and increase probability of non-discriminatory results.
 3. To avoid the potential for discrimination by proxy, recognize that socially derived constructs (e.g., education, geography, group memberships, or resume gaps) may disadvantage groups.
 4. Treat cultural differences fairly (e.g., eye contact, visual cueing, and accents).
 5. Reduce computational and statistical biases,⁸ including selection bias (over- or under-representing populations) and reporting bias (under-reporting all available information).
 6. Anticipate accommodations needed for various physical or cognitive disabilities and ensure proper mechanisms are in place to allow for such accommodations and that such accommodations do not disadvantage the Applicant.
- B. Functions of AI Processing should be evaluated periodically for:
 1. Consistency with evolving best practices.

⁶ Clear documentation helps to mitigate deployment bias and emergent bias, National Institute of Standards and Technology Special Publication 1270, *Towards a Standard for Identifying and Managing Bias in Artificial Intelligence* (March 2022), <https://nvlpubs.nist.gov/NISTpubs/SpecialPublications/NIST.SP.1270.pdf> at 26.

⁷ For an explanation and comparison of systemic bias, computational and statistical bias, and human-cognitive bias, and the role of mitigation efforts throughout the AI lifecycle, see *id.* at 18.

⁸ *Id.*

The Center for Industry Self-Regulation (CISR) is a 501(c)3 non-profit organization (FEIN: 84-5021924) and is a DBA of BBB National Programs Charitable Foundation. 1676 International Drive, Suite 550 | McLean, VA 22102



2. How the AI Application will be deployed and used to determine if reasonable accommodation for a disability under the Americans with Disabilities Act (ADA) or other applicable law is appropriate on a case-by-case basis, as needed.
 3. Human-cognitive biases⁹ that can emerge while calibrating and interpreting models, such as confirmation bias (the tendency to search for, interpret, favor, and recall information in a way that confirms one's pre-existing beliefs or hypotheses).
 4. Providing sufficient safeguards to ensure data integrity and preventing accidental disclosure of Personal Data.
- C. Outputs of AI Processing should be evaluated for, and adjustments made, if necessary, for:
1. Discriminatory impact.
 2. Actual uses of the system conforming to expectations.
 3. Processes for monitoring, evaluating, and providing feedback including by engaging in ongoing monitoring, testing, and adjustment of AI Applications and processes, as described in Section VI above.
 4. Human decision-making controls at relevant stages.
 5. Safeguards in place to protect from disclosure consistent with jurisdictional privacy standards.

VII. Governance and Accountability

A. In General

Organizations that make use of AI Processing should implement robust mechanisms to manage risk¹⁰ and ensure data set, algorithm, and output integrity as well as rates of bias that meet or exceed baseline requirements.

Who is responsible?

1. Vendor: responsible for baseline reliable AI Application designed and tested to meet standards with clear disclosure of purposes for which the AI Application was designed.
2. Employer: responsible for proper AI Processing, execution, and validation of results, with the frequency of testing dependent on volume and other situational factors.

B. Vendor Responsibilities

1. Rigorous and industry-standard pre-testing and validation of the AI Application.
2. Appropriate confidence that AI Application meets or exceeds industry standards, including these Principles, based on, as appropriate and agreed by Employer and Vendor, internal testing, independent certification to the **INDEPENDENT CERTIFICATION PROTOCOLS FOR AI-ENABLED HIRING AND RECRUITING TECHNOLOGIES** and/or third-party audit.
3. Reasonable basis for achievable outputs, limited to the described purpose for which the AI Application designed.
4. Scheduled third-party review (essential for open-loop learning models; recommended for closed-loop models).

⁹ Id.

¹⁰ For an explanation of risk in relation to AI systems, see NIST AI RMF 1.0 at 38 (**Appendix B: “How AI Risks Differ from Traditional Software Risks”**).

The Center for Industry Self-Regulation (CISR) is a 501(c)3 non-profit organization (FEIN: 84-5021924) and is a DBA of BBB National Programs Charitable Foundation. 1676 International Drive, Suite 550 | McLean, VA 22102

5. Educate Employers about results of any prior validations or third-party reviews to ensure that the Employer is informed of best practices to operate the AI tool effectively and in accordance with these self-regulation principles.

C. Employer Responsibilities

1. Confirm that Vendor meets or exceeds industry standards, including these Principles, through contractual controls (specific conditions detailed prior to engagement as necessary for the type of AI Application being used).
2. Ensure that AI Applications are appropriately certified to the **INDEPENDENT CERTIFICATION PROTOCOLS FOR AI-ENABLED HIRING AND RECRUITING TECHNOLOGIES**, meet equivalent criteria, or other appropriate protocols under relevant law.
3. Employers are responsible for biased outcomes, but they can mitigate potential damages through a robust vetting and use of AI-based systems, including by engaging in ongoing monitoring and testing of AI Applications and organizational processes, as described in Section VI above.

D. Independent Certification

To ensure consistency and foster trust in the workforce, participating Employers may require Vendors to demonstrate independent certification of AI Applications against the parameters specified in the **INDEPENDENT CERTIFICATION PROTOCOLS FOR AI-ENABLED HIRING AND RECRUITING TECHNOLOGIES**.

The accompanying protocols specify the criteria for certification and the requirements for qualifying as a third-party certifying entity.

The independent certification process is designed to achieve the following goals:

1. AI Applications are only deployed after industry standard processes for testing and validation have been completed.
2. Uniformity in disclosures and communications between Vendors and Employers.
3. Employers that are properly informed about the purposes for which AI Applications have been developed and tested.
4. Enhanced levels of testing and accountability when appropriate (including when Sensitive Personal Data is processed in an AI Application).
5. Indicators to assist applicants in verifying whether Employers abide by industry standard practices for the use of AI Applications.



About the Center for Industry Self-Regulation

The Center for Industry Self-Regulation (CISR), BBB National Programs' 501(c)(3) nonprofit foundation, was created to harness the power of independent, industry self-regulation to address the marketplace trust challenges businesses face today. CISR is dedicated to education and research that supports responsible business leaders developing fair, future-proof best practices, and to the education of the general public on the conditions necessary for industry self-regulation.

About BBB National Programs

BBB National Programs, a non-profit organization, is the home of U.S. independent industry self-regulation, currently operating more than a dozen globally recognized programs that have been helping enhance consumer trust in business for more than 50 years. These programs provide third-party accountability and dispute resolution services that address existing and emerging industry issues, create a fairer playing field for businesses, and a better experience for consumers. BBB National Programs continues to evolve its work and grow its impact by providing business guidance and fostering best practices in arenas such as advertising, child-and-teen-directed marketing, data privacy, dispute resolution, automobile warranty, technology, and emerging areas. To learn more, visit bbbprograms.org.



Center for Industry
Self-Regulation



National
Programs

EPSTEIN
BECKER
GREEN

INDEPENDENT CERTIFICATION PROTOCOLS FOR

AI-Enabled Hiring & Recruiting Technologies

© The Center for Industry Self-Regulation, 2023. All Rights Reserved. BBB National Programs Charitable Foundation d/b/a Center for Industry Self-Regulation.

I. BACKGROUND

The Independent Certification Protocols for AI-Enabled Hiring and Recruiting Technologies (“Certification Protocols”) supplement the Principles for Trustworthy AI in Recruiting and Hiring (“Principles”) established by BBB National Programs and the Center for Industry Self-Regulation.

Capitalized terms in this document have the same meaning as in the Principles.

Participating Employers and Vendors may self-certify their compliance with the Principles through public notices and disclosures. Additionally, in order to enhance workforce trust, reduce risk,¹ and demonstrate accountability and transparency, Employers and Vendors may also undergo an independent certification through a third-party certifying organization (“Certifier”). Certified organizations may display an approved trust mark or seal provided by the Certifier. For such independent certification to be recognized, the Certifier and its procedures must meet or exceed the standards specified in these Certification Protocols.

To avoid confusion, the Certification Protocols refer to the organization seeking certification as the “Applicant Business.” Such organization may serve the role of Employer, Vendor, or both under the Principles. For purposes of certification, the organization’s role vis-à-vis AI-powered hiring tools and individual applicants for employment will be considered.

II. GOALS OF INDEPENDENT CERTIFICATION

- A. Promotes accountability by Employers and Vendors.
- B. Provides flexible response to technology changes.
- C. Protects intellectual property via confidential independent third-party review.
- D. Fosters process transparency and coalesces best practices.
- E. A trusted certification mechanism should:
 1. Require independent reviews of policies, procedures, and technical documentation;
 2. Ensure transparency via attestations from certified businesses in the notice provided at time of data collection, including explicit commitment to the Certification Protocols and Principles for Trustworthy AI;
 3. Clarify responsibilities between Employers and Vendors based on enforceable commitments to established best practices;
 4. Support rigorous algorithm testing using representative data sets;²
 5. Standardize understanding of factors that should trigger enhanced auditing;
 6. Normalize expectations for routine certification and risk-based approach to assessments and audits; and
 7. Engage in ongoing monitoring and/or redress for complaints.

¹ Independent certification reduces third-party risk and increases the feasibility of risk measurement when AI-enabled tools are used across organizations. This mitigates many challenges identified in NIST’s AI Risk Management Framework 1.0: “Risk can emerge both from third-party data, software or hardware itself and how it is used. Risk metrics or methodologies used by the organization developing the AI system may not align with the risk metrics or methodologies used by the organization deploying or operating the system. Also, the organization developing the AI system may not be transparent about the risk metrics or methodologies it used. Risk measurement and management can be complicated by how customers use or integrate third-party data or systems into AI products or services, particularly without sufficient internal governance structures and technical safeguards.” National Institute of Standards and Technology Risk Management Framework 1.0 (Jan. 2023), <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf> at 5.

² For examples of recommendations for AI bias audits, see the New York City Automated Employment Decision Tool Law (Local Law 144) and the White House Blueprint for an AI Bill of Rights, <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.

III. CERTIFIER REQUIREMENTS

A. Transparency

On its public website, the Certifier must include information regarding the Principles and the services it performs under the Certification Protocols.

This information must include:

1. information on or a link to the Principles and Certification Protocols;
2. detailed certification standards against which Applicant Businesses will be assessed, consistent with the Certification Protocols;
3. procedures for certification review and annual renewal, documenting how the Certifier meets the requirements outlined herein, including the possibility of denying and revoking certifications; and
4. a list of all currently certified organizations along with, for each, a description of the scope the certification and the most recent review date.

B. Impartiality

The Certifier must be free of actual or potential conflicts of interest. This means the ability of the Certifier to perform all tasks related to an Employer or Vendor certification and ongoing participation in the Principles free from influences that would compromise the Certifier's professional judgment, objectivity, and integrity.

To meet this requirement, the Certifier should have internal structural and procedural safeguards, such as:

1. Written policies for internal review of potential conflicts of interest with Applicant Businesses and certified organizations.
2. Written policies for disclosure of potential conflicts of interest and, where appropriate, withdrawal of the Certifier from particular engagements. Such withdrawal should be mandatory in cases where the Certifier is related to the Applicant Business or certified organization to the extent that it would give rise to a risk that the Certifier's professional judgment, integrity, or objectivity could be influenced by the relationship. At no time may a Certifier have a direct or indirect affiliation with any Applicant Business or certified organization that would prejudice the ability of the Accountability agent to render a fair decision with respect to their certification and ongoing participation.³
3. Written policies governing the separation of personnel handling certification functions from personnel handling sales and consulting functions.

C. Expertise & Substantive Review

The Certifier evaluates the written policies, procedures, and technical documentation of Applicant Businesses against a set of published program requirements that encompass the Principles. The scope of the review is focused on those documents necessary to validate processes for testing AI Applications for consistency with the Principles.

To ensure a fulsome review, the Certifier must:

1. Possess appropriate expertise to review and verify policies, procedures, and technical documentation against program requirements.

³ The Certifier may be engaged to perform consulting or technical services for an Employer or Vendor other than services relating to their certification and on-going participation in the Principles.

2. Engage in a comprehensive dialogue with the Applicant Business that includes, at a minimum, the production and review of all relevant policies, procedures, and technical documentation about AI Applications within scope of the certification, with an opportunity to remedy any gaps identified.
3. Any material statements and findings from the review should be signed by a senior officer of the organization with authority to bind it.

D. Confidentiality and Data Security

The Certifier must establish comprehensive written procedures designed to ensure the confidentiality of all information received from Applicant Businesses and certified organizations, including communications and documentation. The Certifier must also establish appropriate policies to ensure the security of all documents provided during the certification process, including a data retention policy.

E. Accountability

The Certifier must establish comprehensive written procedures designed to ensure the integrity of the certification process and to monitor the certified organization throughout the certification period to ensure compliance with the Principles. Such procedures should include, at minimum:

1. Mechanisms to enforce Certifier program requirements, whether by contract or other means, including the ability for the Certifier to revoke the certification at any time due to unremedied findings of non-compliance.
2. Ongoing monitoring of transparency and disclosure requirements, such as verifying that public notices include required information.
3. A mechanism for receiving and resolving individual complaints related to the certified organization's conduct under the Principles. Such mechanism must be monitored by personnel separate from those who conduct certifications.
4. Where there are reasonable grounds for the Certifier to believe that a certified organization has engaged in a practice that may constitute a breach of the Principles or Certification Protocols, an immediate review process will be triggered whereby verification of compliance will be carried out. Where non-compliance with any of the program requirements is found, the Certifier will notify the organization outlining the corrections it must make and a reasonable timeframe within which the corrections must be completed.

IV. CERTIFICATION PROTOCOL

A. Outline of Certification Steps

Stage 1 – Production of the application file, including organization self-assessment and supporting documentation such as written policies, procedures, and technical documentation of AI Application.

Stage 2 – Initial review of file against certification requirements.

Stage 3 – Feedback on non-conformity or anomalies, if any, and request to seek additional independent audit of AI Application, if necessary, based on risk.

Stage 4 – Secondary review of outputs/findings after opportunity to remedy identified gaps.

Stage 5 – Issue/deny certification and communicate results to organization in a comprehensive report, along with any other documentation needed to verify certified status. If approved, issue certification seal, verify public transparency and notice requirements are met, and list certified organization in public directory.

Stage 7 – Conduct annual re-certification reviews.

B. Application File

The application file submitted by the Applicant Business and reviewed by the Certifier must include the following elements:

I. Design process documentation:⁴

Applicant Business shall document the design, development, and testing protocols for all AI Applications within the scope of the certification. This documentation shall include internal policies, procedures, and records of assessment sufficient to demonstrate the requirements of these Certification Protocols.⁵ If the Applicant Business does not have oversight or control over a listed criterion, it may rely on the separate certification of the Vendor that controls that aspect of the design, development, and testing of the AI Application. Alternatively, it may provide documentation of industry-standard internal validation processes, consistent with Section V.D.2. of the Principles.

Provided documentation will specify, in detail, the following elements:

a. Assessment of inputs:

- (1) types, formats (images in png or jpg format, audio mpeg or mp3, specific structures in xml, json, csv, or specific binary format, etc.) and the compatibility of these formats with other solutions or environments (in the case where interoperability is an issue),
- (2) source/acquisition method (sensors embedded into the hardware solution, data imported from a database, etc.) used by the AI functionality (if the source depends on the use cases, the range should be specified),
- (3) volume and quality of training data sufficient for the performance requirements and risks associated with the use of the AI functionality,⁶
- (4) type(s) of algorithm(s) as well as the learning paradigm used by the AI functionality must be documented with respect to performance, maintenance and explainability constraints, frequency and output flow (the specifications of the learning paradigm (especially in the case where re-learning is possible after deployment),
- (5) the prototype of AI functionality whose normal operating domain of use and performance will be evaluated,
- (6) end-user requirements and expectations for AI functionality,
- (7) applicable regulatory requirements⁷ and how the AI Application meets these requirements,
- (8) complete description of how internal processes and procedures for reviewing inputs are consistent with requirements from similar activities and/or uses of AI Applications (standards, best practices etc.), and

⁴ Documentation can be specific, in the event the AI application is developed for a particular end-user, or generic, in the event it is expected to be widely distributed. A model can be built or adjusted based on data processed either manually by humans or using automated tools like machine learning algorithms, or both. Model Building often uses Historical Data/Memory to aggregate data automatically into the Model but can also use Expert Knowledge. Objectives (e.g. the output variables) and Performance Measures (e.g. accuracy, resources for training, representativeness of the dataset) guide the building process. OECD Expert Group on Artificial Intelligence, Nov. 2019

⁵ For a list of factors that written policies and procedures should address, see National Institute of Standards and Technology Special Publication 1270, *Towards a Standard for Identifying and Managing Bias in Artificial Intelligence* (March 2022), <https://nvlpubs.nist.gov/NISTpubs/SpecialPublications/NIST.SP.1270.pdf> at 43 “Policies and Procedures.”

⁶ See NIST Special Publication 1270 at 17 “Dataset Guidance.”

⁷ (Title VII of the Civil Rights Act of 1964 (Title VII), the Americans with Disabilities Act of 1990 (ADA), and the Age Discrimination in Employment Act of 1967 (ADEA) and other applicable regulations prohibit the use of discriminatory employment tests and selection procedures.

- (9) the presence of sensitive data, with discriminatory potential, should be checked and appropriate safeguards should be documented.
- b. Assessment of functions:
 - (1) Data quality. The collection and use of training, validation, and test data must comply with applicable regulations. For example, in the United States, compliance with the Uniform Guidelines on Employee Selection Procedures, 29 C.F.R. Part 1607.
 - (2) Segmentation. The data used for testing must be distinct from the data used for learning (training and validation/development). The segmentation methods (random, by date, by configuration, or based on objective criteria (date, age, etc.) or randomly etc.) of the training, validation, and test datasets must be documented.
 - (3) Training and test data must be unique and not poisoned. The methods used to ensure that the databases are not poisoned and do not contain duplicates must be documented.
 - (4) In the case of learning on structured data, an effective method of handling (detecting and dealing with) missing values among critical attributes in the training and test data shall be implemented and documented.
 - (5) The learning database must include a training set and a validation set. Both sets must individually ensure a sufficient coverage of the system's use cases with respect to the performance requirements and risks associated with the use of the AI functionality, limit selection biases, and be balanced with respect to the use cases and objects under study.
 - (6) The learning database shall incorporate rare events or, alternatively, the AI functionality shall be capable of detecting the occurrence of such events. The Vendor shall document the list of rare events incorporated into the database along with their frequency of occurrence and the method used to determine them.
 - (7) The test database shall ensure sufficient coverage of the system use cases with respect to the performance requirements and risks associated with the use of the AI functionality.
 - (8) The test database shall include rare events. The Vendor shall document the list of rare events included in the database as well as their frequency of occurrence and the method used to determine them.
 - (9) All data transformation steps must be documented. If not, a study must be performed to evaluate the bias introduced by the pre-processing of the test database and the input data of the AI functionality of the deployed system and/or to reduce the impact of this bias. For example, the data normalization process is an example of pre-processing to be listed.
 - (10) If the learning process can generate biases, debiasing algorithms in the pre-processing, processing, or post-processing phase must be used, and this use must be documented and the records retained in accordance with applicable record retention schedules.
 - (11) An evaluation protocol based on performance measures must be implemented during and at the end of the learning process and documented (including human intervention for performance checks).⁸
 - (12) Verification of compliance with the regulatory requirements established during the design process shall ensure that the AI functionality is not validated if any of the requirements are not met.

⁸ NIST recommends a holistic test, evaluation, validation, and verification (TEVV) approach. See generally NIST AI RMF 1.0 and NIST Special Publication 1270.

The Center for Industry Self-Regulation (CISR) is a 501(c)3 non-profit organization (FEIN: 84-5021924) and is a DBA of BBB National Programs Charitable Foundation. 1676 International Drive, Suite 550 | McLean, VA 22102

- (13) The evolutions and updates of the AI functionality, as well as the performance evaluations carried out following incremental or continuous learning, must be traced, archived, and documented, for a period of time adapted to the context of the application and to the life span of the AI functionality.

c. Assessment of outputs:

- (1) Specifications of the functionality for which the AI Application is designed and tested, documenting:
 - (a) the specifications of the system (e.g., for a resume filtering process, the description could be: “the AI Application aims to identify a range of candidates for a position based on a matched review of the required criteria with the data supplied on the resume”),
 - (b) the use of the AI functionality (the automated task and its purpose), for example, “the required position criteria were met or not met”,
 - (c) levels of autonomy (human actions and controls on automated tasks),
 - (d) levels of performance where appropriate, in terms of: accuracy, reliability, learning execution time / execution time of the automated task on target hardware, resilience,⁹ including to attacks and outliers,
 - (e) and other examples of macro tasks (classification, ranking, regression, etc.);
- (2) Requirements for associated documentation (for example, user documentation must be defined. It can also be defined that the test plans for the verification of the requirements concerning the AI functionality be included at specified intervals);
- (3) Justification of the acceptability criteria for each of the requirements defined;
- (4) Evaluation protocols, tools and metrics, the results of the evaluation process, the AI functionality whose field of use and performance will have been validated, and final risk analysis. The results of the evaluations (and any deviation from the requirements defined during the design process) must be documented; and
- (5) Detailed description of risk mitigation efforts used in prior testing to avoid or suppress biased outputs and detail follow-on process testing schedules to ensure mitigation efforts maintain rigor.

2. End-user (Employer) documentation

If the Applicant Business plans to serve as a Vendor under the Principles, offering covered AI Applications for use by other organizations, it shall provide to Certifier detailed documentation for communication with the Employer in standardized format, including:

- a. A description of the desired end-use of the AI Application along with instructions on its proper/stated use for end-users.
- b. Appropriate information about testing procedures deployed (for example regarding the origin and composition of the training databases, the results of the evaluations, certainty levels of the functionality, its open- or closed-source character, and the specifications of the input data).¹⁰

⁹ Resilience: The ability of an AI functionality to maintain compliance with expected requirements in the presence of inputs outside its use domain (e.g., in the event of failure, intentional or unintentional incident, and/or extreme stress).

¹⁰ See generally NIST Special Publication 1270. Robust documentation helps to mitigate deployment bias and emergent bias, see *id.* at 26.

- c. Descriptions of how the Principles have been met for the AI Application, including the following elements, as applicable:
 - (1) Confidentiality and respect for privacy, and access to data,
 - (2) Levels of transparency, in terms of:
 - (a) Explainability and interpretability (including what is required and how long data should be kept),
 - (b) Traceability, auditability of learning, and/or results,
 - (c) Non-discrimination, and bias mitigation,
 - (d) Results of evaluations,
 - (e) Regulatory constraints,
 - (f) Specifications related to the potential consequences of a failure of the AI functionality.
 - (3) Security of the AI Application.
- d. Templates for documenting and communicating with the Employer about the AI Application, to include all of the following:
 - (a) Clear indicators for validation.
 - (b) Description of the post-processing (suggested or built-in) for the use of the AI Application.
 - (c) Resources required and the operations to carry out calibration and maintenance (recommended or mandatory) after deployment.
 - (d) The availability of data needed for learning to achieve the targeted performance and confidence levels validated or mentioned as a necessary condition for using the AI Application.
 - (e) Contraindications (for example, if tests showed that the algorithm did not work for people with hyphenated surnames, then hyphenated surnames are a contraindication).

3. Risk analysis

The Applicant Business shall provide Certifier with a copy of the risk analyses performed at the pre-design, design, development, and deployment phases of covered AI Applications, as relevant and adapted to the functions of the AI Application. The risk analysis must identify, evaluate, and document the risks associated with the use of the AI Application and their potential impacts. This analysis must foresee the use of erroneous data that may be due to formatting errors, bugs in the data management system, or cyberattacks, and cover the components and sub-components as well as the interfaces between components of the AI functionality. The different failure modes of the AI functionality and their consequences must be established to allow the customers to be aware of the residual risks they are exposed to. The impacts can be quantified in terms of cost, safety, security, discrimination, etc.

C. Ongoing Commitments:

To receive and maintain certification, the Applicant Business and certified organization shall commit to the following practices:

- 1. To make available the specifications of functionality to anyone involved in the design, development, evaluation, or maintenance of the AI Application. For example, in the case of a Vendor, the specifications of the AI Application can be made available to Employers through the product sheet, available on the Vendor's website, in the customer documentation, etc.
- 2. To document the design assumptions made about the AI functionality (including statistical assumptions that may vary over time) and the approach taken to select model types and evaluate the AI functionality.

In the event of a change in the assumptions or requirements applicable to the AI functionality, the certified organization shall ensure that an impact analysis is performed to ensure that the changes do not negatively impact compliance with the requirements of the Principles or the Certification Protocols.

3. To conduct outside audits of AI Applications, if recommended by Certifier based on the identified risks,¹¹ with results reported during review.
4. To notify Certifier and provide fulsome documentation, if changes negatively impact compliance or otherwise would affect the disclosures and commitments made during certification.
5. To respond to individual inquiries or assist the Employer in responding to inquiries, as relevant, and cooperate with any dispute resolution procedures established by Certifier.
6. To maintain active certification by pursuing annual review and re-certification from Certifier.

¹¹ The Certification Protocols embrace a risk-based approach to the review of AI Applications. Independent audits always serve to increase accountability and transparency. However, they are not a necessary precondition for certification except in situations where the Certifier identifies a high level of risk, consistent with best practices and applicable regulations.

The Center for Industry Self-Regulation (CISR) is a 501(c)3 non-profit organization (FEIN: 84-5021924) and is a DBA of BBB National Programs Charitable Foundation. 1676 International Drive, Suite 550 | McLean, VA 22102



About the Center for Industry Self-Regulation

The Center for Industry Self-Regulation (CISR), BBB National Programs' 501(c)(3) nonprofit foundation, was created to harness the power of independent, industry self-regulation to address the marketplace trust challenges businesses face today. CISR is dedicated to education and research that supports responsible business leaders developing fair, future-proof best practices, and to the education of the general public on the conditions necessary for industry self-regulation.

About BBB National Programs

BBB National Programs, a non-profit organization, is the home of U.S. independent industry self-regulation, currently operating more than a dozen globally recognized programs that have been helping enhance consumer trust in business for more than 50 years. These programs provide third-party accountability and dispute resolution services that address existing and emerging industry issues, create a fairer playing field for businesses, and a better experience for consumers. BBB National Programs continues to evolve its work and grow its impact by providing business guidance and fostering best practices in arenas such as advertising, child-and-teen-directed marketing, data privacy, dispute resolution, automobile warranty, technology, and emerging areas. To learn more, visit bbbprograms.org.

Lee D. Winston
lwinston@winstoncooks.com
Roderick T. Cooks
rcooks@winstoncooks.com
Winston Cooks, LLC
420 20th Street North
Suite#2200
Birmingham, AL 35203
Telephone: (205) 502-0970
Facsimile: (205) 278-5876

Attorneys for the Plaintiff and the Proposed Class

**UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA
OAKLAND DIVISION**

DEREK L. MOBLEY, for and on behalf)	
of himself and other persons)	<u>CLASS ACTION</u>
similarly situated,)	
)	COMPLAINT
Plaintiffs,)	
)	JURY TRIAL DEMANDED
vs.)	
)	
WORKDAY, INC.)	
,)	
Defendant.)	
)	

NATURE OF COMPLAINT

Plaintiff, Derek L. Mobley (“Mobley” or “Representative Plaintiff”) brings this suit for injunctive, monetary, and declarative relief against Defendant Workday, Inc. (“Workday”) for engaging in a pattern or practice of illegal discrimination on the basis of race, age, and disability in violation of Title VII of the Civil Rights Act of 1964, the Civil Rights Act of 1866 (“42 U.S.C. § 1981”), the Age Discrimination in Employment Act of 1967, and the ADA Amendments Act of 2008 (“ADAAA”). Defendant Workday, Inc.’s (“Workday” or “Defendant”) continuous and systemic pattern or practice of discriminatory job screening—which disproportionately

1 disqualifies African-Americans, individuals over the age of forty (40) and individuals with
2 disabilities from securing gainful employment-resulted in the denial of employment to Mr.
3 Mobley.

4 Defendant Workday, Inc.’s artificial intelligence (“AI”) systems and screening tools rely
5 on algorithms and inputs created by humans who often have built-in motivations, conscious and
6 unconscious, to discriminate. Defendant Workday, Inc. unlawfully offers an algorithm-based
7 applicant screening system that determines whether an employer should accept or reject an
8 application for employment based on the individual’s race, age, and or disability.

9 10 11 **JURISDICTION AND VENUE**

12 1. The jurisdiction of this Court is invoked pursuant to 28 U.S.C. §§ 1331, 1343(3),
13 and (4), 2201 and 2202, 42 U.S.C. 2000d-2 and 2000e5(f), and 29 U.S.C. § 621, et seq.

14 2. This is a suit authorized and instituted pursuant to the Act of Congress known as
15 Title VII of the Civil Rights Act of 1964, 42 U.S.C. § 2000 et seq., as amended, “The Civil
16 Rights Act of 1866,” 42 U.S.C. § 1981, the Age Discrimination in Employment Act of 1967, 29
17 U.S.C. § 621, et seq., and the ADA Amendments Act of 2008 (“ADAAA”).

18 3. Venue is proper in the Northern District of California under 28 U.S.C. § 1391(B)
19 & (c) because Workday is located here and the acts complained of occurred in the Northern
20 District of California.

21 22 **PARTIES**

23 4. Plaintiff, Derek Mobley is an African -American male, over the age of forty (40)
24 and who suffers from depression and anxiety. Mr. Mobley is an applicant.

1 5. Defendant Workday is an employment agency pursuant to Section 703(b) of the
2 Act, 42 U.S.C. § 2000e-2(b). Defendant Workday's headquarters and principal place of business
3 is located at 6110 Stoneridge Mall Road, Pleasanton, California.

4 **CONDITIONS PRECEDENT TO SUIT UNDER**
5 **TITLE VII, THE ADEA AND THE ADAAA**

6 6. On June 3, 2021, Mr. Mobley filed a charge of discrimination with the Oakland
7 Field Office of the United States Equal Employment Opportunity Commission. On July 19,
8 2021, Mr. Mobley filed an amended charge of discrimination. On November 22, 2022, the
9 EEOC issued Mr. Mobley a Dismissal and Notice of Right to Sue, giving him ninety-days from
10 its receipt to file a case. Thus, Mr. Mobley has satisfied all prerequisites to bring this action
11 pursuant to Title VII, the ADEA, and the ADAAA.

12 7. Mr. Mobley's claims arising under 42 U.S.C. § 1981 do not require administrative
13 exhaustion and are subject to a four-year statute of limitations.

14 **CLASS ACTION ALLEGATIONS**

15 8. The Representative Plaintiff brings this action in his own behalf and on behalf of
16 all others similarly situated, pursuant to Rule 23 of the Federal Rules of Civil Procedure, and
17 seek to represent the following subclasses:
18

19 •All African-American applicants or former applicants who from June 3, 2019 to the
20 present who have not been referred and/or permanently hired for employment as a
21 result of the challenged discriminatory screening process.
22

23 •All applicants or former applicants over the age of forty (40) who from June 3, 2019 to
24 the present who have not been referred and/or permanently hired for employment as a result of
25 the challenged discriminatory screening process.
26

1 •All disabled applicants or former applicants who from June 3, 2019 to the present who
2 have not been referred and/or permanently hired for employment as a result of the challenged
3 discriminatory screening process.

4 Mr. Mobley in the case at bar challenges systemic discrimination by, and seeks class-
5 wide relief against, Workday for its administration and dissemination of discriminatory screening
6 products as part of its employment policies and procedures which constitute a pattern and
7 practice of discrimination on the basis of race, age, and disability with respect to selections.
8 These policies and procedures have been continuously utilized by the Defendant since at least
9 2018, and their implementation and use has personally harmed the named the Plaintiff, and the
10 putative class members he seeks to represent. Moreover, the selection tools marketed by
11 Workday to its customers allows these customers to manipulate and configure them in a
12 discriminatory manner to recruit, hire, and onboard employees. Workday's products process and
13 interpret an applicant's qualifications and recommend whether the applicant should be accepted
14 or rejected.

15 These tools when utilized by private companies as administered by Workday relies upon
16 subjective practices which have caused disparate impact and disparate treatment to applicants
17 who are African-American, over the age of forty (40) or and/or disabled. Applicants who are not
18 members of these protected groups and who are similarly situated to the Representative Plaintiff,
19 have not been subjected to such discriminatory treatment.

20 **A. COMMON QUESTIONS OF LAW AND FACT**

21 9. The prosecution of the claims of the Representative Plaintiff requires adjudication
22 of numerous questions of law and fact common to his individual claims and those of the putative
23 classes he seeks to represent. The common questions of law would include, inter alia: (a)
24 whether the Defendant's screening products cause African-American, individuals over the age of
25

forty (40), and/or individuals with a disability to be disproportionately and more likely denied employment; (b) whether the Defendant's doing so cannot be justified as a necessary business practice for evaluating potential employees; and (c) whether the Defendant's screening products have a disparate impact on applicants who are African-American, over the age of forty (40), and/or disabled in violation of the "Civil Rights Act of 1964," 42 U.S.C. § 2000 et seq., the "Civil Rights Act of 1866," 42 U.S.C. § 1981 and 1981a, the Age Discrimination in Employment Act of 1967, 29 U.S.C. § 621, et seq., and the ADA Amendments Act of 2008 ("ADAAA"). The common questions of fact would include, inter alia: (1) whether Workday's administration of its screening products discriminated against the aforementioned applicants because of their race, age, and/or disability with regards to hiring; (2) whether compensatory and punitive damages, injunctive relief, and other equitable remedies for the class are warranted; and (3) whether Workday discriminated against the aforementioned protected groups in other terms and conditions of employment. The details of the Representative Plaintiff's claims are encompassed within the claims prosecuted on behalf of the class and set forth in this Complaint.

B. TYPICALITY

10. The claims of the Representative Plaintiff are typical of those of the members of the class. The Representative Plaintiff and all class members have been and are similarly adversely affected by the systemic racially discriminatory practices complained of herein. Specifically, the representative claims, like those of the class members, arise out of Defendant's pervasive discriminatory conduct with regard to aforementioned discrimination in hiring and other terms and conditions of employment. The relief necessary to remedy the claims of the Representative Plaintiff is the same relief that is necessary to remedy the claims of the putative class members in this case. The Representative Plaintiff seeks the following relief for individual claims and class claims asserted herein: (1) declaratory judgment that Defendant has engaged in

1 systemic discrimination against African-Americans, individuals over the age of forty (40),
2 and/or the disabled; (2) a permanent injunction against such continuing discrimination; (3)
3 injunctive relief which reforms Workday's screening products, policies, practices and procedures
4 so that the Representative Plaintiff and the class members will be able to compete fairly in the
5 future for jobs and enjoy terms and conditions of employment traditionally afforded similarly
6 situated employees outside of the protected categories; (4) backpay, front pay, compensatory
7 damages, and other equitable remedies necessary to make the Plaintiff, and the class, whole from
8 Workday's past discrimination; and, (5) attorneys' fees, costs, and expenses.

10 **C. NUMEROSITY AND IMPRACTICABILITY OF JOINDER**

11 11. The class that the Representative Plaintiff seeks to represent is too numerous to
12 make joinder practicable. The proposed class consists of numerous former, current, and future
13 applicants who have been denied employment due to the discriminatory administration of
14 Workday's screening products. Workday's pattern or practice of discrimination also makes
15 joinder impracticable by making it impractical and inefficient to identify many members of the
16 class prior to the determination of the merits of Workday's class wide liability. Thus, the
17 number of Class members is currently indeterminate, but is certainly numerous.

19 **D. ADEQUACY OF REPRESENTATION**

20 12. The Representative Plaintiff will fairly and adequately protect the interests of the
21 class inasmuch as they are broadly representative, as reflected in the preceding paragraphs.
22 There are no conflicts of interest present with the members of the proposed class as each would
23 benefit from the imposition of a remedy for the Defendant's discriminatory employment
24 practices. The Representative Plaintiff has retained counsel experienced in litigating major class
25 actions in the field of employment discrimination, and who are prepared and able to meet the
26 time and fiscal demands of class action litigation of this size and complexity. The combined
27
28

1 interest, experience, and resources of the Representative Plaintiff and his counsel to litigate
2 competently the individual and class claims of employment discrimination at issue satisfy the
3 adequacy of representation requirement under Fed.R.Civ.P. 23(a)(4).

4 **E. EFFICIENCY OF CLASS PROSECUTION OF COMMON CLAIMS**

5
6 13. Certification of a class of similarly-situated aforementioned applicants is the most
7 efficient and economical means of resolving the questions of law and fact that are common to the
8 individual claims of the Representative Plaintiff and the proposed class. The individual claim of
9 the Representative Plaintiff requires resolution of the common question of whether Defendant
10 has engaged in a systemic pattern of discrimination against African-Americans, those over forty
11 (40) and the disabled. The Representative Plaintiff seeks remedies to undo the adverse effects of
12 such discrimination in his own life and career. The Representative Plaintiff has standing to seek
13 such relief because of the adverse effect that such discrimination has had on him individually and
14 on the putative classes he seeks to represent, in general. In order to gain such relief for himself,
15 as well as for the putative class members, the Representative Plaintiff will first establish the
16 existence of systemic discrimination as the premise of the relief they seek. Without class
17 certification, the same evidence and issues would be subject to re-litigation in a multitude of
18 individual lawsuits with an attendant risk of inconsistent adjudications and conflicting
19 obligations. Certification of the subclasses affected by the common questions of law and fact is
20 the most efficient and judicious means of presenting the evidence and arguments necessary to
21 resolve such questions for the Representative Plaintiff, the class and the Defendant. The
22 Representative Plaintiff's individual and class claims are premised upon the traditional
23 bifurcated method of proof and trial for systemic disparate treatment claims of the type at issue
24 in this complaint. Such a bifurcated method of proof and trial is the most efficient method of
25 resolving such common issues.
26
27
28

F. CERTIFICATION IS SOUGHT PURSUANT TO FED. R. CIV. P. 23(b)(2)

14. Workday has acted on grounds generally applicable to the Representative Plaintiff and the proposed class by adopting and following systemic practices and procedures that discriminate on the basis of race, age, and/or disability. Workday's screening products are regularly used to discriminate on the basis of race, age, and/or disability. Workday has refused to act on grounds generally applicable to the putative class by: (1) refusing to adopt or follow screening productions and selection procedures which do not systemically discriminate on the basis of race, age, and/or disability. Workday's discriminatory screening products have made appropriate final injunctive relief and declaratory relief with respect to the class as a whole. The injunctive relief and declaratory relief are the predominate reliefs sought because they are both the culmination of the proof of the Defendant's individual and class-wide liability at the end of Stage I of a bifurcated trial and the essential predicate for the Representative Plaintiff and the class members entitlement to monetary and non-monetary remedies at Stage II of such a trial. Declaratory and injunctive relief flow directly and automatically from proof of the common questions of law and fact regarding the existence of systemic discrimination against individuals on the basis of race, age, and/or disability. Such relief is the factual and legal predicate for the Representative Plaintiff's and the class members entitlement to injunctive and equitable remedies caused by such systemic discrimination.

**G. ALTERNATIVELY CERTIFICATION IS SOUGHT
PURSUANT TO FED. R.CIV. P. 23(b)(3)**

15. The common issues of fact and law affecting the claims of the Representative Plaintiff and proposed class members, including, but not limited to, the common issues identified above, predominate, over any issues affecting only individual claims. A class action is superior to other available means for the fair and efficient adjudication of the claims of the Representative

1 Plaintiff and members of the proposed class. The cost of proving the Defendant's pattern or
2 practice of discrimination makes it impracticable for the named Plaintiffs and members of the
3 proposed class to control the prosecution of their claims individually. The Northern District of
4 California is the most logical forum in which to litigate the claims of the Representative Plaintiff
5 and the proposed class in this case because the Defendant's home office is here and it engages in
6 or ratifies illegal conduct adversely affecting the Plaintiff here.

7
8 **H. ALTERNATIVELY, CERTIFICATION IS SOUGHT PURSUANT TO FED. R.
9 CIV. P. 23(c)(4) FOR INJUNCTIVE AND DECLARATORY RELIEF.**

10 16. Alternatively, claims for injunctive and declaratory relief for the Injunctive Relief
11 Class are properly certified under Federal Rule of Civil Procedure 23(c)(4) because such claims
12 present only common issues, the resolution of which would advance the interests of the parties in
13 an efficient manner.

14 **I. ALTERNATIVELY, CERTIFICATION IS SOUGHT PURSUANT TO FED. R.
15 CIV. P. 23(c)(4) FOR CLASS WIDE LIABILITY.**

16 17. Alternatively, class wide liability claims are properly certified under Federal Rule
17 of Civil Procedure 23(c)(4) for the Classes because such claims present only common issues, the
18 resolution of which would advance the interests of the parties in an efficient manner.

19 **J. PUNITIVE DAMAGES MAY ALTERNATIVELY BE CERTIFIED
20 PURSUANT TO FED.R.CIV.P. 23(b)(2).**

21 18. Punitive damages liability may alternatively be certified under Federal Rule of
22 Civil Procedure 23(b)(2) because such relief focuses on the conduct of Workday and not the
23 individual characteristics of the Plaintiff and are an allowable form of incidental monetary relief.
24
25
26
27
28

CLAIMS OF REPRESENTATIVE PLAINTIFFS**Derek Mobley**

19. Derek L. Mobley is an African-American male. Mr. Mobley possesses a Bachelor's degree in Finance from Morehouse College, and an Associate's Degree in Network Systems Administration, from ITT Technical Institute. Mr. Mobley is over the age of forty (40) and suffers from anxiety and depression.

20. Since 2018, Mr. Mobley has applied for at least 80-100 positions that upon information and belief use Workday, Inc. as a screening tool for talent acquisition and/or hiring. He has been denied employment each and every time.

21. Upon information and belief, Workday, Inc. provides screening tools that allow its customers to use discriminatory and subjective judgments in reviewing and evaluating employees for hire and allows the preselection of applicants outside of the protected categories. Indeed, because of the utilization of these screening products, there is no need in this case to distinguish one employer/client from another. Similarly, as the Representative Plaintiff here is challenging the use of these common discriminatory screening products per se, and not merely the individualized manifestations of their use, the fact that the common components may vary to some small degree or be applied by different customers is of no consequence.

22. If an individual does not make it past these Workday screening products, he/she will not advance in the hiring process.

23. Workday provides this service for hundreds if not thousands of companies, including numerous Fortune 500 firms. Individuals impacted the same way by these processes number in the thousands if not tens of thousands.

24. The selection tools and/or tests utilized by Workday, Inc. in making selection decisions-to include screening and hiring applicants discriminate on the basis of race in violation

1 of §703(k) of Title VII, 42 U.S.C. §2000e-2(k). Upon information and belief, these processes
2 disparately impact African-American applicants because the tests have the effect of
3 disproportionately excluding African-Americans. Furthermore, the tests or selection procedures
4 are not job-related, nor are they consistent with any business necessity.

5
6 25. Title VII prohibits discrimination by employment agencies. Section 703(b) of the
7 Act, 42 U.S.C. § 2000e-2(b), reads:

8 “it shall be an unlawful employment practice for an employment agency to fail or refuse to refer
9 for employment, or otherwise to discriminate against, any individual because of his race, color,
10 religion, sex, or national origin, or to classify or refer for employment any individual on the basis
11 of his race, color, religion, sex, or national origin. Section 701(c) of the Act, 42 U.S.C. §
12 2000e(c), defines the term “employment agency” as: any person regularly undertaking with or
13 without compensation to procure employees for an employer or to procure for employees
14 opportunities to work for an employer and includes an agent of such a person. Workday, Inc. Is
15 an employment agency under Title VII.

16
17 26. Upon information and belief, Mr. Mobley and other African-Americans have been
18 discriminated against because of their race (African-American), in violation of the Civil Rights
19 Act of 1964, as amended.

20
21 27. Furthermore, the screening tools marketed by Workday for the administration of
22 its products discriminate on the basis of disability in violation of the ADA Amendments Act of
23 2008 (ADAAA). Upon information and belief, these processes disparately impact disabled
24 applicants because the screening products have the effect of disproportionately excluding
25 individuals with disabilities. Furthermore, the screening products are not job-related, nor are
26 they consistent with any business necessity.

29. Upon information and belief, these screening tools disparately impact applicants over the age of 40 because the tests have the effect of disproportionately excluding them. Furthermore, the screening tools are not job-related, nor are they consistent with any business necessity.

A. COUNT ONE-Intentional Discrimination on the Basis of Race in Violation of Title VII of the Civil Rights Act of 1964 and 42 U.S.C. § 1981.

31. Workday has discriminated against the Representative Plaintiff and the class he
Seeks to represent with regards to selection procedures and other terms and conditions of
employment because of their race, in violation of Title VII of the Civil Rights Act of 1964 and
42 U.S.C. § 1981.

33. By reason of Workday's discriminatory employment practices, the Representative Plaintiff and the proposed class members have experienced extreme harm, including loss of compensation, wages, back and front pay, and other employment benefits, and, as such, are entitled to all legal and equitable remedies available under Title VII of the Civil Rights Act of 1964 and 42 U.S.C. § 1981.

B. COUNT TWO-Disparate Impact Discrimination on the Basis of Race and Disability in Violation of Title VII of the Civil Rights Act of 1964, and the ADA Amendments Act of 2008.

34. Representative Plaintiff restates and incorporates by reference all applicable paragraphs above as part of this Count of the Complaint.

35. The screening tools marketed Workday for making selection decisions-to include hiring discriminate on the basis of race and/or disability.

36. Workday's screening tools allow employers to consciously or unconsciously discriminate against individuals on the basis of race and/or disability in making hiring decisions. There is no business necessity justifying the disparate impact on individuals in the protective category.

37. Because of this, the decision-makers are free to exercise their discretion in an unguided, subjective manner that provides a ready mechanism for discrimination.

38. Workday's screening tools discriminated against the Representative Plaintiff and the proposed class both within and outside the liability period in this case.

39. As a direct result of Workday's discriminatory screening tools as described above, the Representative Plaintiff and the class he seeks to represent have suffered damages including, but not limited to, lost past and future income, compensation, and benefits.

**C. COUNT THREE- Intentional Discrimination
Age Discrimination in Employment Act of 1967, 29 U.S.C. §§ 623(a)(1)**

40. Representative Plaintiff restates and incorporates by reference all applicable paragraphs above as part of this Count of the Complaint.

41. This claim is brought by the Representative Plaintiff on behalf of himself and the collective he seeks to represent. Workday's screening tools discriminate against applicants aged

1 40 and older. Workday marketed applicant screening tools it knew intentionally discriminated
2 against the Representative Plaintiff and the collective in violation of the ADEA.

3 42. These screening tools have the effect of denying the Representative Plaintiff and
4 the collective of employment opportunities because of their age. The discriminatory acts that
5 constitute Workday's pattern and/or practice of discrimination have occurred both within and
6 outside the liability period in this case.

7
8 43. As a direct result of Workday's marketing its discriminatory screening tools as
9 described above, the Representative Plaintiff and the collective have suffered damages including,
10 but not limited to, lost past and future income, compensation, and benefits.

11 44. The foregoing conduct constitutes illegal, intentional discrimination and
12 unjustified disparate treatment prohibited by 29 U.S.C. § 623(a)(1).

13
14 **D. COUNT FOUR- Disparate Impact Discrimination**
Age Discrimination in Employment Act of 1967, 29 U.S.C. §§ 623(a)(2)

15 45. Representative Plaintiff restates and incorporates by reference all applicable
16 paragraphs above as part of this Count of the Complaint.

17
18 46. This Claim is brought by Representative Plaintiff on behalf of himself and the
19 collective he seeks to represent. Workday maintains discriminatory policies, patterns, and/or
20 practices that have an adverse impact on employees ages 40 and older in violation of the ADEA
21 and are not, and cannot be, justified by reasonable factors other than age. Specifically,
22 Workday's marketing of its screening tools have had a disparate impact on individuals over the
23 age of 40.

24 47. Workday has administered these discriminatory screening tools both within and
25 outside the liability period in this case.
26
27
28

1 48. As a direct result of Workday's discriminatory policies and/or practices as
2 described above, the Representative Plaintiff and the collective he seeks to represent have
3 suffered damages including, but not limited to, lost past and future income, compensation, and
4 benefits.

5
6 **PRAYER FOR RELIEF**

7 WHEREFORE, the Representative Plaintiff and the Proposed Class pray for relief as
8 follow:

- 9 1. Certification of the case as a class action on behalf the proposed subclasses;
- 10 2. Designation of Plaintiffs as representatives of the subclasses;
- 11 3. Designation of Plaintiffs' Counsel of record as Class Counsel;
- 12 4. A declaratory judgment that the practices complained of herein are unlawful and
13 violate Title VII, 42 U.S.C. § 1981, the ADEA, and the ADAAA;
- 14 5. A preliminary and permanent injunction against the Company and its officers,
15 agent, successors employees, representatives, and any and all persons acting in concert with them
16 from engaging in each of the unlawful policies, practices, customs, and usages set forth herein;
- 17 6. An order that the Company institute and carry out policies, practices, and
18 programs that provide equal employment opportunities for all minorities, and that it eradicate the
19 effects of its past and present unlawful employment practices;
- 20 7. For back pay, front pay and other monetary relief according to proof (including
21 interest and benefits);
- 22 8. For all damages sustained as a result of the Company's conduct according to
23 proof;
- 24 9. For compensatory damages, nominal damages, and liquidated damages according
25 to proof;
- 26
27
28

14. For such other and further relief as the Court deems proper.

Lee D. Winston
lwinston@winstoncooks.com
Roderick T. Cooks
rcooks@winstoncooks.com
Winston Cooks, LLC
420 20th Street North
Suite 2200
Birmingham, AL 35203
Telephone: (205) 502-0970
Facsimile: (205) 278-5876

JULIE A. TOTTEN (STATE BAR NO. 166470)
 jatotten@orrick.com
 ERIN M. CONNELL (STATE BAR NO. 223355)
 econnell@orrick.com
 KAYLA D. GRUNDY (STATE BAR NO. 300513)
 kgrundy@orrick.com
 ALEXANDRIA R. ELLIOTT (STATE BAR NO. 320293)
 aelliottt@orrick.com
 ORRICK, HERRINGTON & SUTCLIFFE LLP
 The Orrick Building
 405 Howard Street
 San Francisco, CA 94105-2669
 Telephone: (415) 773-5700
 Facsimile: (415) 773-5759

JUSTIN M. WASHINGTON (SBN 334389)
 justin.washington@orrick.com
 ORRICK, HERRINGTON & SUTCLIFFE LLP
 355 S. Grand Avenue, Suite 2700
 Los Angeles, CA, 90071
 Telephone: (213) 629-2020
 Facsimile: (213) 612-2499

Attorneys for Defendant
 WORKDAY, INC.

UNITED STATES DISTRICT COURT
 NORTHERN DISTRICT OF CALIFORNIA

DEREK L. MOBLEY, for and on behalf of
 himself and other persons similarly situated;

Plaintiff,

v.

WORKDAY, INC.

Defendant.

Case No. 4:23-cv-00770-YGR

**DEFENDANT WORKDAY, INC.'S
 MOTION TO DISMISS AND
 MEMORANDUM OF POINTS AND
 AUTHORITIES IN SUPPORT**

Date: September 12, 2023
 Time: 2:00 p.m.
 Courtroom: 1, 4th Floor
 Judge: Hon. Yvonne Gonzalez
 Rogers

Complaint Filed: February 21, 2023

TABLE OF CONTENTS

	Page
NOTICE OF MOTION AND MOTION TO DISMISS	1
MEMORANDUM OF POINTS AND AUTHORITIES	2
I. INTRODUCTION AND SUMMARY OF POINTS	2
II. ALLEGATIONS OF THE OPERATIVE COMPLAINT	4
A. Workday Provides Customers With Software That Assists In Screening Applicants For Employment.	4
B. Plaintiff Files A Charge Of Discrimination Before The EEOC Alleging Disparate Impact.	4
C. Mobley Brings This Class Action Alleging Claims Of Both Disparate Treatment And Disparate Impact.	5
III. APPLICABLE LEGAL STANDARD.....	6
IV. PLAINTIFF FAILS TO ALLEGE THAT WORKDAY IS AN ENTITY THAT CAN BE SUBJECT TO LIABILITY UNDER TITLE VII, THE ADEA, OR THE ADA (ALL COUNTS).....	6
A. An Employment Agency Must Be Engaged To A “Significant Degree” In Procuring Employment.	7
B. Plaintiff Fails To Allege That Workday Is An Employment Agency.....	9
V. PLAINTIFF’S INTENTIONAL DISCRIMINATION CLAIMS FAIL FOR MULTIPLE REASONS (COUNTS ONE AND THREE).	10
A. Plaintiff Failed To Exhaust Intentional Discrimination Claims Under Title VII And The ADEA.	11
B. Plaintiff Fails To Raise A Plausible Inference Of Intentional Discrimination On The Basis Of Race Or Age.	12
C. Plaintiff’s Allegations Of Discrimination Based On “Terms and Conditions Of Employment” And Invocation Of § 623(a)(1) Should Be Dismissed Or Struck.	15
VI. PLAINTIFF’S DISPARATE IMPACT CLAIMS FAIL FOR MULTIPLE REASONS (COUNTS TWO AND FOUR).	16
A. Neither Title VII Nor The ADEA Recognize A Disparate Impact Claim Against Employment Agencies.....	17
1. Disparate Impact Claims Lie Only Where The Applicable Provision Focuses On Effects And Not Just Motivation.....	17
2. Title VII And The ADEA’s Employment Agency Provisions Do Not Authorize Disparate Impact Claims.	18

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

TABLE OF CONTENTS
(continued)

	Page
B. Plaintiff Has Insufficiently Pled A Disparate Impact Claim Under All Of The Statutes On Which He Relies.....	21
1. Plaintiff Has Not Pled A Disparity.....	21
2. Plaintiff Has Not Pled A Specific Employment Practice.....	22
3. Plaintiff Has Not Sufficiently Pled Causation.	24
VII. CONCLUSION	25

TABLE OF AUTHORITIES

Page(s)

Cases

<i>Adams v. Vivo, Inc.</i> , 2012 WL 5525315 (N.D. Cal. Nov. 14, 2012).....	20
<i>Alhayoti v. Blinken</i> , 2022 WL 2869850 (N.D. Cal. July 21, 2022).....	22
<i>Ali v. ServiceNow, Inc.</i> , 2019 WL 11542365 (N.D. Cal. Sept. 23, 2019), <i>aff'd</i> , 797 F. App'x 322 (9th Cir. 2020)	14, 15
<i>Ansari v. Mnuchin</i> , 2017 WL 11696665 (N.D. Cal. Oct. 11, 2017).....	13, 14
<i>Ashcroft v. Iqbal</i> , 556 U.S. 662 (2009)	6
<i>Barnhart v. Sigmon Coal Co.</i> , 534 U.S. 438 (2002)	19
<i>Bell Atl. Corp. v. Twombly</i> , 550 U.S. 544 (2007)	6
<i>Blackman-Baham v. Kelly</i> , 2017 WL 679514 (N.D. Cal. Feb. 21, 2017).....	13, 15
<i>Borodaenko v. Twitter, Inc.</i> , 2023 WL 3294581 (N.D. Cal. May 5, 2023)	22
<i>Brown v. Puget Sound Elec. Apprenticeship & Training Trust</i> , 732 F.2d 726 (9th Cir. 1984).....	11
<i>Brush v. San Francisco Newspaper Printing Co.</i> , 315 F. Supp. 577 (N.D. Cal. 1970), <i>aff'd</i> , 469 F.2d 89 (9th Cir. 1972).....	8, 10
<i>Canatella v. Reverse Mortg. Sols. Inc.</i> , 2014 WL 7205146 (N.D. Cal. Dec. 17, 2014)	24
<i>Career Sys. Dev. Corp. v. Am. Home Ins. Co.</i> , 2010 WL 11636254 (N.D. Cal. Aug. 11, 2010).....	15
<i>Eldredge v. Carpenters 46 N. Cal. Ctys. Joint Apprenticeship & Training Comm.</i> , 833 F.2d 1334 (9th Cir. 1987).....	16, 20

- iii -

DEFENDANT WORKDAY, INC.'S MOTION TO DISMISS AND MEMORANDUM OF POINTS AND
AUTHORITIES IN SUPPORT

1	<i>Epstein v. Wash. Energy Co.</i> ,	
2	83 F.3d 1136 (9th Cir. 1996).....	15
3	<i>Flanagan v. Excel Staffing Sols., LLC</i> ,	
4	2018 WL 558499 (N.D. Ill. Jan. 25, 2018)	20
5	<i>Fonseca v. Sysco Food Servs. of Ariz., Inc.</i> ,	
6	374 F.3d 840 (9th Cir. 2004).....	12
7	<i>Freeman v. Oakland Unified Sch. Dist.</i> ,	
8	291 F.3d 632 (9th Cir. 2002).....	11
9	<i>Freyd v. University of Or.</i> ,	
10	990 F.3d 1211 (9th Cir. 2021).....	21
11	<i>Griggs v. Duke Power Co.</i> ,	
12	401 U.S. 424 (1971).....	18
13	<i>Grigorescu v. Bd. of Trustees of San Mateo Cnty. Cmty. Coll. Dist.</i> ,	
14	2019 WL 1790472 (N.D. Cal. Apr. 24, 2019)	13
15	<i>Hill v. Mississippi State Emp.t Serv.</i> ,	
16	918 F.2d 1233 (5th Cir. 1990).....	13
17	<i>Katz v. Regents of the Univ. of Cal.</i> ,	
18	229 F.3d 831 (9th Cir. 2000).....	21
19	<i>Knieval v. ESPN</i> ,	
20	393 F.3d 1068 (9th Cir. 2005).....	4
21	<i>Liu v. Uber Techs. Inc.</i> ,	
22	551 F. Supp. 3d 988 (N.D. Cal. 2021)	12, 21, 22, 24
23	<i>Lopez v. Pac. Mar. Ass’n</i> ,	
24	2009 WL 10680881 (C.D. Cal. Apr. 3, 2009), <i>aff’d</i> , 657 F.3d 762 (9th Cir.	
25	2011)	21
26	<i>Lucas v. Gold Standard Baking, Inc.</i> ,	
27	2014 WL 518000 (N.D. Ill. Feb. 10, 2014)	20
28	<i>McDonnell Douglas Corp. v. Green</i> ,	
	411 U.S. 792 (1973).....	13
	<i>Metro-Goldwyn-Mayer Studios Inc. v. Gorkster, Ltd.</i> ,	
	545 U.S. 913 (2005).....	13
	<i>Moussouris v. Microsoft Corp.</i> ,	
	2016 WL 4472930 (W.D. Wash. Mar. 7, 2016)	25

1	<i>Opara v. Yellen</i> ,	
2	57 F.4th 709 (9th Cir. 2023).....	13
3	<i>Radentz v. Am. Ass’n of Physician Specialists, Inc.</i> ,	
4	2014 WL 12601014 (C.D. Cal. Nov. 10, 2014).....	8, 10
5	<i>Ricci v. DeStefano</i> ,	
6	557 U.S. 557 (2009).....	18
7	<i>Riesgo v. Heidelberg Harris, Inc.</i> ,	
8	36 F. Supp. 2d 53 (D.N.H. 1997).....	16
9	<i>Rodriguez v. United States</i> ,	
10	852 F.3d 67 (1st Cir. 2017).....	11
11	<i>Rollins v. Traylor Bros., Inc.</i> ,	
12	2016 WL 258523 (W.D. Wash. Jan. 21, 2016).....	11
13	<i>Ross v. Lockheed Martin Corp.</i> ,	
14	267 F. Supp. 3d 174 (D.D.C. 2017).....	23
15	<i>Smith v. City of Jackson</i> ,	
16	544 U.S. 228 (2005).....	17, 18, 22
17	<i>Sprewell v. Golden State Warriors</i> ,	
18	266 F.3d 979 (9th Cir. 2001).....	15
19	<i>Stewart v. Am. Ass’n of Physician Specialists, Inc.</i> ,	
20	2015 WL 7722349 (C.D. Cal. Nov. 30, 2015).....	8
21	<i>Tex. Dep’t of Housing & Cmty. Aff. v. Inclusive Cmty. Project, Inc.</i> ,	
22	576 U.S. 519 (2015).....	17, 18
23	<i>Wal-Mart Stores, Inc. v. Dukes</i> ,	
24	564 U.S. 338 (2011).....	14, 23, 24
25	<i>Wards Cove Packing Co. v. Atonio</i> , 490 U.S. 642 (1989), <i>superseded by statute on</i>	
26	<i>other grounds</i> , 42 U.S.C. § 2000e-2(k).....	22
27	<i>Watson v. Fort Worth Bank & Trust</i> ,	
28	487 U.S. 977 (1988).....	22, 23
	<i>Wei v. Dep’t of Physics at Univ. of Cal. at Berkeley</i> ,	
	2018 WL 457211 (N.D. Cal. Jan. 18, 2018).....	15
	<i>Whitsitt v. Hedy Holmes Staffing Servs.</i> ,	
	2015 WL 5560119 (E.D. Cal. Sept. 21, 2015), <i>aff’d</i> , 671 F. App’x 1004 (9th	
	Cir. 2016).....	8

- v -

DEFENDANTS’ NOTICE OF MOTION AND MOTION TO DISMISS, STRIKE, OR ALTERNATIVELY, FOR MORE DEFINITE STATEMENT; SUPPORTING MEMORANDUM OF POINTS AND AUTHORITIES

1	<i>Whitsitt v. Hedy Holmes Staffing Servs.</i> ,	
2	671 F. App'x 1004 (9th Cir. 2016)	7

Statutes

4	29 U.S.C. § 623	<i>passim</i>
5	29 U.S.C. § 626(d)(1).....	11
6	29 U.S.C. § 630(c)	7
7	42 U.S.C. § 1981	2, 5, 12
8	42 U.S.C. § 2000e	<i>passim</i>
9	42 U.S.C. § 3604(a)	18
10	42 U.S.C. § 12111	6, 7

Other Authorities

12	Fed. R. Civ. P. 12(b)(6).....	1, 6
13	Fed. R. Civ. P. 12(f)	16
14	Fed. R. Civ. P. 23	14

NOTICE OF MOTION AND MOTION TO DISMISS

PLEASE TAKE NOTICE THAT at 2:00 p.m. on September 12, 2023, or as soon thereafter as counsel may be heard, in Courtroom 1 of this Court, 1301 Clay Street, Oakland, California 94612, before the Honorable Yvonne Gonzalez Rogers, Defendant Workday, Inc. (“Workday”) will and does move for an Order dismissing all causes of action alleged against it in Plaintiff Derek Mobley’s Complaint.

Workday moves to dismiss Plaintiff’s Complaint on the grounds that Plaintiff fails to state a claim upon which relief can be granted pursuant to Federal Rule of Civil Procedure 12(b)(6) and for failure to exhaust administrative remedies. This Motion is based on this Notice of Motion, the accompanying Memorandum of Points and Authorities, the Request for Judicial Notice, the complete pleadings and records on file herein, any matters of which the Court may take judicial notice, and such other evidence and arguments presented at the hearing on this Motion.

///

///

///

///

///

///

///

///

///

///

///

///

///

///

///

///

DEFENDANT WORKDAY, INC.’S MOTION TO DISMISS AND MEMORANDUM OF POINTS AND
AUTHORITIES IN SUPPORT

MEMORANDUM OF POINTS AND AUTHORITIES**I. INTRODUCTION AND SUMMARY OF POINTS**

This case attempts to turn a software provider into an employment agency. Defendant Workday, Inc. provides enterprise software to thousands of customers who then use that software to help run various aspects of their businesses— including payroll, financial planning, and human resources. Workday’s products are made to be adaptable, meaning customers configure them according to their business needs and change them as those needs evolve. Workday also builds artificial intelligence and machine learning into its platform, making its products nimble enough to provide a wide range of powerful business solutions that help improve the human experience by achieving better productivity, information, and problem-solving.

Plaintiff Derek Mobley brings this lawsuit under four federal employment discrimination laws—Title VII of the Civil Rights Act (Title VII), the Age Discrimination in Employment Act (ADEA), the Americans with Disabilities Act (ADA), and 42 U.S.C. Section 1981. But he does not allege that he is or ever was a Workday employee. He also does not allege that he ever applied to a position at Workday. Instead, he alleges that he applied to “80-100” unidentified positions at an unidentified number of unidentified employers who he claims used an unidentified “algorithm-based screening system” provided to them by Workday. He further alleges that these unidentified Workday customers “manipulate and configure” Workday’s unidentified tools “in a discriminatory manner,” Compl. ¶ 8, and “use discriminatory and subjective judgments” in screening applicants. *Id.* ¶ 21. In Plaintiff’s view, this makes Workday an “employment agency” covered by the federal employment discrimination laws, as he accuses Workday of employment discrimination on the basis of race, age, and disability under both disparate treatment and/or disparate impact theories. But Plaintiff’s claims sputter at the starting line.

Despite Plaintiff’s novel legal theory, the defects in the Complaint are basic. First, Plaintiff does not come close to pleading facts that plausibly demonstrate Workday is an “employment agency.” The relevant statutory definitions and controlling case law make clear an employment agency is an entity that “regularly undertak[es]” to “procure” employees for employers or employment opportunities for prospective employees. Here, the Complaint alleges only that

- 2 -

DEFENDANT WORKDAY, INC.’S MOTION TO DISMISS AND MEMORANDUM OF POINTS AND
AUTHORITIES IN SUPPORT

1 Workday provides and administers a software product that customers use themselves to screen
2 applicants. This deficiency alone requires dismissal of the entire Complaint. *Infra* § IV.

3 Plaintiff's claims each fail for multiple additional reasons as well. His intentional
4 discrimination claims must be dismissed as unexhausted because Plaintiff never raised a disparate
5 treatment theory in his Charge of Discrimination before the Equal Employment Opportunity
6 Commission (EEOC). More importantly, the Complaint contains no allegations that Workday acted
7 with discriminatory intent. Plaintiff says nothing at all about Workday's motivations and makes no
8 attempt to plead any element of the customary prima facie case of intentional discrimination,
9 including any information about whether he was qualified for the roles to which he applied, who
10 was ultimately hired instead of him, and their demographics. *Infra* § V.

11 Plaintiff's disparate impact claims fare no better. First, even if Plaintiff had sufficiently pled
12 that Workday is an employment agency, neither Title VII nor the ADEA authorize disparate impact
13 claims against employment agencies. Second, even if they did, the Complaint contains no facts that
14 would render plausible Plaintiff's bare assertion that Workday's screening tools result in a
15 statistical hiring disparity with respect to *any* of his disparate impact claims. Third, Plaintiff fails
16 to plead facts plausibly demonstrating that the alleged use of Workday's tools by unidentified
17 employers constitutes a single employment "policy" or "practice" that could ground a disparate
18 impact claim, let alone has Plaintiff sufficiently pled facts to plausibly show Workday's tools
19 caused any alleged disparate impact. Indeed, the Complaint is utterly devoid of factual allegations
20 specifying what Workday "screening tools" are at issue, how they work, or how they could possibly
21 create any statistical disparity with respect to a protected class. *Infra* § VI.

22 Simply put, Plaintiff's Complaint accuses a software provider of providing an *unidentified*
23 product to *unidentified* customers who allegedly used it in connection with *unidentified* jobs in
24 *unspecified* ways. Those facts do not amount to a plausible claim for employment discrimination.
25 The Complaint should be dismissed in its entirety.

26 ///

27 ///

28 ///

1 **II. ALLEGATIONS OF THE OPERATIVE COMPLAINT**

2 **A. Workday Provides Customers With Software That Assists In Screening**
 3 **Applicants For Employment.**

4 Based in Pleasanton, California, Workday is a leading provider of enterprise cloud
 5 applications for finance and human resources, bringing its software solutions to thousands of
 6 organizations. Request for Judicial Notice In Support Of Motion To Dismiss, Ex. A (Charge of
 7 Discrimination) at 2, ¶¶ 2, 4.¹

8 Plaintiff alleges that one of Workday’s many applications is a “screening tool” that
 9 customers “configure” to “review[] and evaluat[e] employees for hire.” Compl. ¶¶ 20-21. He
 10 further alleges Workday offers a “screening tool” or “tools,” “artificial intelligence (‘AI’) systems,”
 11 “algorithm-based applicant screening system,” “screening products, policies, practices, and
 12 procedures,” and “selection tools and/or tests.” *Id.* ¶¶ 10, 20-21, 27-29, 35-36, 38-39, 41-43, 46-
 13 47, and Nature of Complaint. According to Plaintiff, Workday “provides this service for hundreds
 14 if not thousands of companies, including numerous Fortune 500 firms.” *Id.* ¶ 23. Despite these
 15 allegations, the Complaint does not allege which of Workday’s alleged tool(s) are at issue, how
 16 they work, what they do, how customers use them, or how Workday “administers” them.

17 **B. Plaintiff Files A Charge Of Discrimination Before The EEOC Alleging**
 18 **Disparate Impact.**

19 Plaintiff Derek Mobley is an African-American male “over the age of forty.” *Id.* ¶ 19. He
 20 earned a bachelor’s degree in Finance and an associate’s degree in Network Systems
 21 Administration. *Id.* He also suffers from anxiety and depression. *Id.*

22 In July 2021, Plaintiff brought a Charge of Discrimination before the Equal Employment
 23 Opportunity Commission against Workday. Compl. ¶ 6; Ex. A (Charge of Discrimination) at 2.
 24 According to his EEOC Charge, “[s]ince 2018, [he] has applied for at least 80-100 . . . positions

25 _____
 26 ¹ The Court may consider materials attached to or incorporated by reference in the pleadings in
 27 deciding a motion to dismiss. *See Knievel v. ESPN*, 393 F.3d 1068, 1076 (9th Cir. 2005). Here,
 28 Plaintiff expressly references his charge of discrimination in his Complaint. *See* Compl. ¶ 6 (“On
 June 3, 2021, Mr. Mobley filed a charge of discrimination with the Oakland Field Office of the
 United States Equal Employment Opportunity Commission. On July 19, 2021, Mr. Mobley filed
 an amended charge of discrimination.”).

1 that use Workday, Inc. as a screening tool for talent acquisition and/or hiring” and “[h]e has been
2 denied employment each and every time.” *Id.* at 2. Plaintiff further alleges Workday is an
3 “employment agency” within the meaning of Title VII and that its “tools” discriminate on the basis
4 of race, age, and disability. *Id.* at 2-3. In each instance, Plaintiff alleges only that Workday’s tools
5 “disparately impact applicants” in a protected class. *Id.* On November 22, 2022, the EEOC issued
6 Plaintiff a Dismissal and Notice of Right to Sue. Compl. ¶ 6.

7 **C. Mobley Brings This Class Action Alleging Claims Of Both Disparate**
8 **Treatment And Disparate Impact.**

9 Plaintiff filed his Complaint in this lawsuit on February 22, 2023. He again alleges that
10 Workday is an “employment agency,” and that since 2018, he has “applied for at least 80-100
11 positions” that “use Workday, Inc. as a screening tool for talent acquisition and/or hiring” without
12 success. Compl. ¶ 20. Plaintiff does not identify any prospective employer or describe any position
13 for which he applied. He also fails to describe the required qualifications for any of the alleged
14 positions, and further fails to plead any details about who was hired instead of him, including the
15 successful applicants’ demographics.

16 The Complaint alleges no factual details regarding how customers purportedly “use”
17 Workday’s screening tools, or how this allegation is consistent with the allegation that Workday
18 “administers” the alleged tools for customers to “manipulate and configure” them in a discretionary,
19 discriminatory manner. Compl. ¶ 8. Instead, he baldly asserts that Workday’s “algorithm-based
20 applicant screening system . . . determines whether an employer should accept or reject an
21 application . . . based on the individual’s race, age, and or disability.” *Id.* at 2. He further asserts the
22 “tools . . . allow [Workday] customers to use discriminatory and subjective judgments in reviewing
23 and evaluating employees for hire and allows the preselection of applicants outside of the protected
24 categories.” *Id.* ¶ 21. He claims that “[i]f an individual does not make it past . . . Workday screening
25 products, he/she will not advance in the hiring process.” *Id.* ¶ 22.

26 Plaintiff asserts four causes of action against Workday: (1) intentional discrimination on the
27 basis of race under Title VII and 42 U.S.C. Section 1981; (2) disparate impact discrimination on
28 the basis of race and disability under Title VII and the ADA; (3) intentional discrimination under

the ADEA; and (4) disparate impact discrimination under the ADEA. *See generally* Compl. at 12-14. He also seeks to represent three subclasses:

- All African-American applicants or former applicants who, from June 3, 2019 to the present, have not been referred and/or permanently hired for employment as a result of the challenged discriminatory screening process.
- All applicants or former applicants over the age of forty (40) who, from June 3, 2019 to the present, have not been referred and/or permanently hired for employment as a result of the challenged discriminatory screening process.
- All disabled applicants or former applicants who, from June 3, 2019 to present, have not been referred and/or permanently hired for employment as a result of the challenged discriminatory screening process.

Compl. ¶ 8.

Plaintiff seeks certification of his alleged class action, as well as declaratory and injunctive relief, damages (including punitive damages), and costs and fees.

III. APPLICABLE LEGAL STANDARD

To survive a motion to dismiss, a complaint must “contain sufficient factual matter, accepted as true, to ‘state a claim to relief that is plausible on its face.’” *Ashcroft v. Iqbal*, 556 U.S. 662, 678 (2009). The factual allegations must allow “the court to draw the reasonable inference that the defendant is liable for the misconduct alleged.” *Id.* Allegations that are “merely consistent with” a defendant’s liability, *id.*, or that do not state a cognizable legal theory, cannot survive a motion to dismiss. Fed. R. Civ. P. 12(b)(6). “[A] plaintiff’s obligation to provide the ‘grounds’ of his ‘entitle[ment] to relief’ requires more than labels and conclusions, and a formulaic recitation of the elements of a cause of action will not do.” *Bell Atl. Corp. v. Twombly*, 550 U.S. 544, 555 (2007).

IV. PLAINTIFF FAILS TO ALLEGE THAT WORKDAY IS AN ENTITY THAT CAN BE SUBJECT TO LIABILITY UNDER TITLE VII, THE ADEA, OR THE ADA (ALL COUNTS).

A claim lies under Title VII, the ADEA, or the ADA only against an “employer,” “employment agency,” or “labor organization.” 42 U.S.C. § 2000e-2 (Title VII prohibitions); 29 U.S.C. § 623 (ADEA prohibitions); 42 U.S.C. § 12111 (ADA definition of “covered entity”).

- 6 -

DEFENDANT WORKDAY, INC.’S MOTION TO DISMISS AND MEMORANDUM OF POINTS AND
AUTHORITIES IN SUPPORT

1 Plaintiff does not allege, nor could he, that Workday was the prospective “employer” for any
2 position to which he applied or a “labor organization.” Instead, he labels Workday an “employment
3 agency.” Compl. ¶ 5. But labels alone do not suffice, and the Complaint is utterly devoid of any
4 factual assertions that could plausibly hold Workday accountable as an employment agency within
5 the meaning of the applicable statutes. This defect requires dismissal of Plaintiff’s entire case. *E.g.*,
6 *Whitsitt v. Hedy Holmes Staffing Servs.*, 671 F. App’x 1004, 1005 (9th Cir. 2016) (affirming
7 dismissal of claims where plaintiff failed to plausibly allege defendant was subject to ADEA).

8 A. **An Employment Agency Must Be Engaged To A “Significant Degree” In**
9 **Procuring Employment.**

10 “Employment agency” is defined in both Title VII and the ADA as:

11 Any person regularly undertaking with or without compensation to
12 procure employees for an employer or to procure for employees
13 opportunities to work for an employer and includes an agent of such
14 a person.

14 42 U.S.C. § 2000e(c); 42 U.S.C. § 12111(7) (“The term[] . . . ‘employment agency’ . . . shall have
15 the same meaning given such terms in section 2000e of this title.”). The ADEA definition similarly
16 embraces those who “regularly undertak[e]” to “procure employees for an employer,” but is
17 narrower in that it omits the phrase “or to procure for employees opportunities to work for an
18 employer,” 29 U.S.C. § 630(c); 42 U.S.C. § 2000e(c).

19 These statutory definitions are consistent with the plain meaning of the word “agency,”
20 which is defined as “an establishment engaged in doing business for another.” *Agency*, Merriam-
21 Webster Online Dictionary, <https://www.merriam-webster.com/dictionary/agency> (last visited
22 June 20, 2023). And while the statutes do not define the word “procure,” the most relevant definition
23 identifies it as meaning “to obtain (something) by particular care and effort.” *Procure*, Merriam-
24 Webster Online Dictionary, <https://www.merriam-webster.com/dictionary/procure> (last visited
25 June 20, 2023). Accordingly, the plain language of the relevant employment discrimination statutes
26 confirm they are intended to cover only the actions of entities that regularly and actively engage
27 with prospective employees to find them opportunities, or entities that actively engage with
28 employers and then go out and find them employees. Entities that do not engage in this type of

- 7 -

1 conduct—and are not employers or would-be employers themselves—cannot be liable for
2 employment discrimination.

3 Controlling case law further confirms that to be an employment agency, an entity must
4 actively engage with employees to obtain employment opportunities (e.g., a staffing agency or temp
5 agency), or actively engage with employers to find prospective employees (e.g., a headhunter or
6 executive search firm). For example, in *Brush v. San Francisco Newspaper Printing Co.*, a seminal
7 case interpreting the employment agency definition, the court explained that “the statutory
8 requirement that an employment agency be one that ‘regularly’ undertakes to procure employees
9 or employment opportunities indicates that the Congress had in mind to include only those engaged
10 to a significant degree in that kind of activity as their profession or business.” 315 F. Supp. 577,
11 580 (N.D. Cal. 1970), *aff’d*, 469 F.2d 89 (9th Cir. 1972). The court held that the fact that newspaper
12 organizations are “in the business of printing and publishing advertising copy presented by
13 employers, professional employment agencies and job seekers,” does not make them “employment
14 agencies” under the law. *Id.* The Ninth Circuit affirmed, agreeing that the “dictionary” could not
15 support a broader reading of the term “employment agency.” *Brush*, 469 F.2d at 90.

16 Courts since *Brush* have refined the “significant degree” requirement, explaining that a
17 defendant must “actively assist” in procuring employment opportunities or employees to be an
18 employment agency. *Stewart v. Am. Ass’n of Physician Specialists, Inc.*, 2015 WL 7722349, at *5
19 (C.D. Cal. Nov. 30, 2015); *see Radentz v. Am. Ass’n of Physician Specialists, Inc.*, 2014 WL
20 12601014, at *5 (C.D. Cal. Nov. 10, 2014) (finding that although defendant was “more significantly
21 involved in procuring employment opportunities than a newspaper that publishes classified ads,
22 such involvement does not rise to the level of a ‘significant degree’ of engagement in employment-
23 related activities” as defendant merely provided a “venue in which potential employers and
24 candidates can exchange information regarding employment opportunities” and was not “actively
25 involved in the application and interview process for any specific employers”); *Whitsitt v. Hedy
26 Holmes Staffing Servs.*, 2015 WL 5560119, at *3 (E.D. Cal. Sept. 21, 2015), *aff’d*, 671 F. App’x
27 1004 (9th Cir. 2016) (allegation of “screen[ing]” plaintiff’s employment application not sufficient
28 to allege defendant was an “employment agency”).

- 8 -

DEFENDANT WORKDAY, INC.’S MOTION TO DISMISS AND MEMORANDUM OF POINTS AND
AUTHORITIES IN SUPPORT

1 In short, Congress included “employment agenc[ies]” (and “labor organizations”) within
2 the federal antidiscrimination laws to embrace a defined class of entities that actively work with
3 the workforce and prospective employers to find a match. Adopting the expansive reading required
4 by the Complaint would dramatically widen the scope of the employment discrimination statutes
5 to embrace an enormous range of entities, from software providers to companies that perform
6 background checks to networking websites. Nothing in the statutory text or governing case law
7 suggests that Congress intended the definition of “employment agency” to broadly embrace anyone
8 that offers any product or service that is somehow related to the hiring process.

9 **B. Plaintiff Fails To Allege That Workday Is An Employment Agency.**

10 Rather than allege facts that plausibly could establish Workday comes within the statutory
11 definitions, Plaintiff instead makes only the bare assertion that “Workday is an employment
12 agency.” Compl. ¶ 5. Beyond that, the Complaint describes a software provider. Plaintiff alleges
13 merely that Workday “offers an algorithm-based applicant screening system that determines
14 whether an employer should accept or reject an application for employment.” Compl. at 2. He
15 further asserts Workday “market[s]” and “administer[s]” these software-based “screening tools” to
16 prospective employers. *See e.g.*, Compl. ¶¶ 8, 27-28, 35, 41, 43, 46, 47. Plaintiff further alleges
17 these tools “allow [Workday’s] customers to use discriminatory and subjective judgments in
18 reviewing and evaluating employees for hire and allow the preselection of applicants outside of the
19 protected categories.” *Id.* ¶ 21.

20 Even if Plaintiff could prove these allegations (which Workday denies), they do not describe
21 an entity remotely like those commonly understood to be employment agencies. Plaintiff does not
22 allege that Workday is involved in recruiting new employees, inquiring about job opportunities on
23 behalf of job seekers, assisting such job seekers in pursuing opportunities with new employers, or
24 negotiating and consummating employment agreements. Indeed, the Complaint does not allege that
25 Workday interacts with prospective employees at all—whether on behalf of employees to find them
26 positions, or on behalf of employers to find them employees. Nor does it allege that Workday itself
27 possesses or exercises control or decision-making authority over hiring decisions.

28 Workday is aware of no case that has ever found that an entity that neither interacts with

- 9 -

DEFENDANT WORKDAY, INC.’S MOTION TO DISMISS AND MEMORANDUM OF POINTS AND
AUTHORITIES IN SUPPORT

prospective employees nor makes hiring decisions is an “employment agency.” Indeed, cases like *Brush* and *Radentz* reject employment agency status as to entities with far more involvement in the procurement of prospective employees than Workday’s alleged involvement. In *Radentz*, the court found that a professional services agency that provided optional certification to its members, solicited and published job postings, and offered job-relevant training was not an employment agency under Title VII. *Radentz*, 2014 WL 12601014, at *1-3. In reaching that conclusion, the court compared the facts in *Radentz* to the facts of other cases where the defendant entity was found to be an employment agency, including a truck training program that required a certain number of job applications and invited recruiters onsite for interviews; a law school that was the primary source through which employers hired recent graduates and law school administrators spent significant time on student job placement; and a commission that provided lists of eligible employees and determined the criteria for who is eligible. *Id.* at *3-4. *Radentz* ultimately concluded that while the defendant was more involved in the procurement of employees than the newspaper at issue in *Brush*, such involvement did not rise to the level of “significant” because it merely “provide[d] a venue in which potential employers and candidates can exchange information regarding employment opportunities.” *Id.* at *5. The court reached this conclusion even though the defendant, on a recurring basis, was engaged in certifying prospective employees as “accomplished practitioner[s]” and presenting them with job postings. *Id.*

By contrast here, the Complaint comes nowhere close to alleging that level of active involvement with prospective employees. Instead, it alleges that Workday provides software to customers that customers can use to help them select candidates from an applicant pool that Workday had no part in generating. Plaintiff fails to allege that Workday “procures” anything—let alone that Workday procures to the “significant degree” required by *Brush* and its progeny. Because Plaintiff does not allege that Workday is an employer, employment agency, or labor organization within the meaning of any of the applicable statutes, all of his claims should be dismissed.

V. PLAINTIFF’S INTENTIONAL DISCRIMINATION CLAIMS FAIL FOR MULTIPLE REASONS (COUNTS ONE AND THREE).

Even if Plaintiff had pled facts showing Workday is an “employment agency,” Plaintiff’s

- 10 -

DEFENDANT WORKDAY, INC.’S MOTION TO DISMISS AND MEMORANDUM OF POINTS AND
AUTHORITIES IN SUPPORT

1 claims for intentional discrimination on the basis of race and age fail for multiple additional reasons.
2 As to claims brought under Title VII and the ADEA, Plaintiff failed to exhaust them before the
3 EEOC, as required to file a civil suit. *Infra* § A. On the merits, all of Plaintiff’s intentional
4 discrimination claims fail because Plaintiff fails to plead facts raising an inference of
5 discrimination, either through direct evidence of discriminatory intent or by trying to plead the
6 elements of a prima facie case. *Infra* § B. And in all events, Plaintiff’s (a) allegations of
7 discrimination in the “terms and conditions of employment” and (b) invocation of § 623(a)(1) of
8 the ADEA must be dismissed or struck because they are applicable only to employers, not
9 employment agencies. *Infra* § C.

10 **A. Plaintiff Failed To Exhaust Intentional Discrimination Claims Under Title**
11 **VII And The ADEA.**

12 Before a plaintiff may file a Title VII or ADEA suit in federal court, he must first file a
13 charge with the EEOC identifying the unlawful actions alleged. *See* 42 U.S.C. § 2000e-5(e)(1),
14 (f)(1); 29 U.S.C. § 626(d)(1). Although the express claim a plaintiff makes in a lawsuit need not
15 necessarily match one-to-one with the charge, there must be some link between the EEOC charge
16 and the complaint. *Brown v. Puget Sound Elec. Apprenticeship & Training Trust*, 732 F.2d 726,
17 729 (9th Cir. 1984). The question is whether the plaintiff’s claim would have fallen “within the
18 scope of . . . an EEOC investigation which can reasonably be expected to grow out of the charge of
19 discrimination.” *Freeman v. Oakland Unified Sch. Dist.*, 291 F.3d 632, 636 (9th Cir. 2002).

20 It is settled law in the Ninth Circuit that an administrative charge that, as here, is limited to
21 a disparate impact theory does not exhaust a disparate treatment claim. *Brown*, 732 F.2d at 730
22 (“Any investigation of whether [plaintiff’s employment] application was rejected as a result of
23 disparate *impact* would not have encompassed her subsequent claim that . . . she was subjected to
24 *intentional* sex discrimination.”); *e.g.*, *Rollins v. Traylor Bros., Inc.*, 2016 WL 258523, at *14
25 (W.D. Wash. Jan. 21, 2016) (“If an EEOC charge alleges only that an employer’s neutral policy
26 caused a disparate impact, it would be unreasonable for the EEOC to separately investigate whether
27 the claimant also suffered intentional discrimination.”). Other circuits are in accord. For example,
28 in *Rodriguez v. United States*, the First Circuit explained that “a disparate impact allegation” cannot

1 “somehow encompass an intentional discrimination claim on the theory that the agency would have
2 investigated intent in connection with the disparate impact claim” because “disparate treatment and
3 disparate impact . . . involve[e] different facts and evidence.” 852 F.3d 67, 80 (1st Cir. 2017)
4 (concluding charge alleging disparate impact did not exhaust claim for disparate treatment).

5 Plaintiff’s charge of discrimination before the EEOC alleged only disparate impact
6 discrimination—not disparate treatment. The charge (1) identifies three types of tests that
7 “impacted” applicants the same way; (2) cites an alleged violation of 42 U.S.C. § 2000e-2(k) (which
8 is a subsection of Title VII titled “Burden of Proof in Disparate Impact Cases”); and (3) alleges the
9 challenged tests disparately impact African-American, disabled, and over-40 applicants and are
10 neither “job-related [nor] . . . consistent with any business necessity.” 42 U.S.C. § 2000e-2(k).
11 There are *no* allegations asserting intentional discrimination or, indeed, addressing Workday’s
12 intent at all. Accordingly, Plaintiffs Title VII and ADEA claims must be dismissed.

13 **B. Plaintiff Fails To Raise A Plausible Inference Of Intentional Discrimination**
14 **On The Basis Of Race Or Age.**

15 Even if Plaintiff had exhausted his disparate treatment claims under Title VII and the
16 ADEA, his claims for intentional discrimination must nevertheless be dismissed because he fails to
17 plead facts that give rise to a plausible inference that Workday (or anyone else) acted with
18 discriminatory intent. All three of the statutes upon which Plaintiff relies require him to “plead[]
19 facts giving rise to an inference that the [defendant] intended to discriminate against the protected
20 group.” *Liu v. Uber Techs. Inc.*, 551 F. Supp. 3d 988, 992 (N.D. Cal. 2021); *see* 42 U.S.C. § 2000e-
21 2(b) (Title VII prohibits “fail[ure] or refus[al] to refer for employment . . . because of [] race, color,
22 religion, sex, or national origin”); *Fonseca v. Sysco Food Servs. of Ariz., Inc.*, 374 F.3d 840, 850
23 (9th Cir. 2004) (“Analysis of an employment discrimination claim under § 1981 follows the same
24 legal principles as those applicable in a Title VII disparate treatment case.”); 29 U.S.C. § 623(b)
25 (ADEA prohibits “fail[ure] or refus[al] to refer for employment . . . because of such individual’s
26 age”). The Complaint comes nowhere close to alleging that Workday acted with such motivations,
27 which is not surprising given that Workday, per Plaintiff’s *own* allegations, does nothing but
28 provide configurable software tools.

- 12 -

DEFENDANT WORKDAY, INC.’S MOTION TO DISMISS AND MEMORANDUM OF POINTS AND
AUTHORITIES IN SUPPORT

1 Indeed, although allegations establishing discriminatory motive can be direct or
2 circumstantial, *Opara v. Yellen*, 57 F.4th 709, 721-22 (9th Cir. 2023), Plaintiff includes neither
3 here. Instead, he makes only the threadbare assertion that Workday’s “conduct has been
4 intentional.” Compl. ¶ 32; *Blackman-Baham v. Kelly*, 2017 WL 679514, at *17 (N.D. Cal. Feb. 21,
5 2017) (dismissing complaint because “Plaintiff’s allegations concerning Defendant’s
6 discriminatory motives are entirely conclusory”). But the only Workday conduct Plaintiff alleges
7 is the provision of a software product. And Plaintiff’s theory is not that Workday’s product is
8 inherently discriminatory, but that its algorithm-based tools allow *customers* to “manipulate and
9 configure [the tools] in a discriminatory manner,” *id.* ¶ 8; “use discriminatory and subjective
10 judgments,” *id.* ¶ 21; and, ultimately, allows customers “to discriminate.” *Id.* ¶ 14. Many products,
11 software or otherwise, *can* be misused by another. Merely offering such a multi-use product,
12 however, does not establish culpable intent. *Cf. Metro-Goldwyn-Mayer Studios Inc. v. Gorkster*,
13 *Ltd.*, 545 U.S. 913, 933, 934 (2005) (noting patent and copyright law “limits [on] imputing culpable
14 intent as a matter of law from the characteristics or uses of a distributed product” or “understanding
15 that some of one’s products may be misused”).

16 Plaintiff also fails to raise an inference of discriminatory intent through the customary *prima*
17 *facie* elements of a failure to hire claim under any of the statutes on which he relies. Under Title
18 VII, § 1981, and the ADEA, Plaintiff must allege that (i) he is a member of a protected class; (ii)
19 he applied and was qualified for a job for which he sought to be placed; (iii) despite his
20 qualifications, he was not referred for employment; and (iv) the employment agency referred or
21 hired someone else to fill the position who was equally or less qualified and was outside the
22 protected class. *E.g., Ansari v. Mnuchin*, 2017 WL 11696665, *5 (N.D. Cal. Oct. 11, 2017) (Title
23 VII and ADEA); *Grigorescu v. Bd. of Trustees of San Mateo Cnty. Cmty. Coll. Dist.*, 2019 WL
24 1790472, at *10-11 (N.D. Cal. Apr. 24, 2019) (Section 1981); *Hill v. Mississippi State Emp.t Serv.*,
25 918 F.2d 1233, 1239 (5th Cir. 1990); *see generally McDonnell Douglas Corp. v. Green*, 411 U.S.
26 792, 802 (1973) (articulating elements of a *prima facie* case under Title VII). Although a plaintiff
27 need not “*prove* her *prima facie* discrimination case at the pleading stage, . . . [t]he elements of a
28 *prima facie* case . . . are relevant to determining whether a plaintiff has alleged sufficient facts to

1 state a plausible claim.” *Ansari*, 2017 WL 11696665 at * 5.

2 Once again, Plaintiff does not come close. Instead, Plaintiff alleges no more than the
3 following regarding the jobs he did not receive:

4 Mr. Mobley has applied for at least 80-100 positions that upon
5 information and belief use Workday, Inc. as a screening tool for
6 talent acquisition and/or hiring. He has been denied employment
each and every time.

7 Compl. ¶ 20. Plaintiff literally says nothing about the names of employers to whom he submitted
8 these “80-100” applications, the industry or industries allegedly involved, the job titles for the
9 positions to which he applied, whether the jobs are similar or different, or whether he met the
10 minimum and/or preferred qualifications for the positions at issue. Moreover, although Plaintiff
11 alleges his educational background, he never alleges whether that background was relevant with
12 respect to the jobs to which he applied. *See Ali v. ServiceNow, Inc.*, 2019 WL 11542365, at *4
13 (N.D. Cal. Sept. 23, 2019), *aff’d*, 797 F. App’x 322 (9th Cir. 2020) (dismissing complaint where
14 plaintiff “describe[d] a variety of his accolades and certifications, [but did] not explain how his
15 qualifications, or . . . past experience, [was] relevant to any of the positions to which he applied”).

16 Each of these problems become exponential when considering the classwide nature of
17 Plaintiff’s claims. Plaintiff attempts to excuse this failing by suggesting that because he is
18 “challenging the use of . . . common discriminatory screening products per se,” it is “of no
19 consequence” that he lumps together 80-100 unidentified positions with some number of
20 unidentified companies and untold class members. Compl. ¶ 21. While Plaintiff appears to make
21 this allegation in an attempt to support the “commonality” requirement for class adjudication under
22 Rule 23,² it confirms Plaintiff utterly fails to plead anything like a discriminatory motive on
23 Workday’s part with respect to any of the positions for which he claims he was not hired.

24
25 ² This allegation also fails to support any argument for class certification. To the contrary, Plaintiff
26 specifically alleges elsewhere in the Complaint that the purported problem with Workday’s “tools”
27 are that “customers . . . *manipulate and configure* them in a discriminatory manner,” Compl. ¶ 8
28 (emphasis added), which “allow[s] [customers] to use discriminatory and *subjective* judgments in
reviewing and evaluating employees for hire.” Compl. ¶ 21 (emphasis added). Accordingly,
Plaintiff alleges “just the opposite of a uniform employment practice that would provide the
commonality needed for a class action.” *Wal-Mart Stores, Inc. v. Dukes*, 564 U.S. 338, 355 (2011).

Plaintiff also fails to allege any facts that could permit the Court to conclude that whoever was selected for the jobs at issue were equally or less qualified than Plaintiff, and not members of the protected classes to which he asserts he belongs. Instead, he merely asserts that “[a]pplicants who are not members of these protected groups and who are similarly situated to the Representative Plaintiff, have not been subjected to such discriminatory treatment.” Compl. ¶ 8. Not only are those allegations legally insufficient, but they are too conclusory to be entitled to any weight at the pleading stage. *See Sprewell v. Golden State Warriors*, 266 F.3d 979, 988 (9th Cir. 2001) (“Nor is the court required to accept as true allegations that are merely conclusory. . . .”); *Epstein v. Wash. Energy Co.*, 83 F.3d 1136, 1140 (9th Cir. 1996) (“[C]onclusory allegations of law and unwarranted inferences are insufficient to defeat a motion to dismiss.”).

At bottom, Plaintiff fails to include *any* allegations that would suggest intentional discrimination by Workday. Accordingly, his disparate treatment claims must be dismissed. *See, e.g., Wei v. Dep’t of Physics at Univ. of Cal. at Berkeley*, 2018 WL 457211, at *3 (N.D. Cal. Jan. 18, 2018) (dismissing complaint because “Plaintiff’s allegation that Defendant declined to hire him on account of his age is unsupported and conclusory.”); *Kelly*, 2017 WL 679514, at *17.³ Indeed, courts have dismissed complaints with significantly more detail than what Plaintiff offers here. *See, e.g., Ali*, 2019 WL 11542365, at *4.

C. Plaintiff’s Allegations Of Discrimination Based On “Terms and Conditions Of Employment” And Invocation Of § 623(a)(1) Should Be Dismissed Or Struck.

Even if Plaintiff’s disparate treatment claims could survive in some form, this Court should address two pleading defects, both of which relate to the fact that Plaintiff alleges only that Workday is an employment agency, not an employer by dismissing or striking the allegations.

First, Plaintiff’s allegations that Workday discriminated with respect to the “terms and conditions of employment” cannot proceed. Compl. ¶¶ 9-10, 31. As a matter of law, such claims exist only against *employers*, not employment agencies. *Compare* 42 U.S.C. § 2000e–2(b) (“It shall

³ Plaintiff’s prayer for punitive damages must also be dismissed because Plaintiff fails to plead facts sufficient to support his allegation that Workday acted with malice or reckless indifference, particularly where, as here, Plaintiff has not stated a claim for intentional discrimination. *See Career Sys. Dev. Corp. v. Am. Home Ins. Co.*, 2010 WL 11636254, *2 (N.D. Cal. Aug. 11, 2010).

1 be an unlawful employment practice for an employment agency to fail or refuse to refer for
 2 employment, or otherwise to discriminate against, any individual because of [protected classes]”
 3 with 42 U.S.C. § 2000e–2(a) (“It shall be an unlawful employment practice for an employer to fail
 4 or refuse to hire or to discharge any individual, or otherwise to discriminate against any individual
 5 with respect to his compensation, terms, conditions, or privileges of employment because of
 6 [protected classes]”); *see also Eldredge v. Carpenters 46 N. Cal. Ctys. Joint Apprenticeship &*
 7 *Training Comm.*, 833 F.2d 1334, 1337–38 (9th Cir. 1987) (holding that defendant was not “in the
 8 same position as an employment agency” and could be held liable under Title VII for discrimination
 9 in providing employment); *Riesgo v. Heidelberg Harris, Inc.*, 36 F. Supp. 2d 53, 57 (D.N.H. 1997)
 10 (“42 U.S.C.A. § 2000e–2(b) . . . which catalogs the list of unlawful employment agency practices,
 11 does not bar discrimination with respect to the ‘terms, conditions, or privileges of employment.’”).

12 Second, Plaintiff may not invoke § 623(a)(1) of the ADEA as the basis for an intentional
 13 discrimination claim because that subsection applies only to employers. *See* Compl. at 13; 29
 14 U.S.C. § 623(a)(1) (“It shall be unlawful for an employer . . .”). To the extent that Plaintiff intended
 15 to invoke § 623(b), which applies to “Employment agency practices,” the grounds for dismissal
 16 Workday has identified above, *supra* § A, B, are fully applicable to an intentional discrimination
 17 claim under § 623(b) and therefore equally merit dismissal of such a claim.

18 Because Plaintiff attempts to plead claims against Workday solely as an “employment
 19 agency,” any invocation of statutory provisions applicable only to employers should be dismissed
 20 or struck. *See* Fed. R. Civ. P. 12(f).

21 **VI. PLAINTIFF’S DISPARATE IMPACT CLAIMS FAIL FOR MULTIPLE**
 22 **REASONS (COUNTS TWO AND FOUR).**

23 Plaintiff’s disparate impact claims, brought under Title VII, the ADA, and the ADEA, also
 24 fail for multiple reasons. With respect to Title VII and the ADEA, there is no such thing as a
 25 disparate impact claim against an employment agency as a matter of law, requiring dismissal of
 26 those claims with prejudice. *Infra* § A. In any event, Plaintiff’s allegations fall woefully short of
 27 stating a disparate impact claim. *Infra* § B.

28 ///

1 **A. Neither Title VII Nor The ADEA Recognize A Disparate Impact Claim**
 2 **Against Employment Agencies.**

3 As already explained, *supra* § V.C, Plaintiff’s attempt to hold Workday liable under the
 4 federal employment statutes is premised on Workday’s alleged status as an “employment agency.”
 5 Plaintiff does not (and cannot) allege that Workday was his prospective “employer.” As to Title
 6 VII, he specifically invokes § 2000e-2(b), Compl. ¶ 25, which applies only to “Employment agency
 7 practices.” As to the ADEA, Plaintiff again, perhaps inadvertently, invokes only the subsection of
 8 the statute applicable to employers, 29 U.S.C. § 623(a)(2). Compl. at 14. For reasons explained in
 9 connection with Plaintiff’s intentional discrimination claims, *supra* V.C, any claim premised on
 10 § 623(a) obviously must be dismissed because Plaintiff does not allege that Workday is an
 11 employer at all. Workday thus assumes for present purposes that Plaintiff intended to invoke
 12 § 623(b), the subsection applicable to employment agencies.

13 Even if Workday were an employment agency subject to § 2000e-2(b) of Title VII and
 14 § 623(b) of the ADEA, however, neither provision authorizes a disparate impact claim. That is
 15 because—in contrast to neighboring provisions that do authorize disparate impact claims—the
 16 employment agency provisions of Title VII and the ADEA contain no reference to the “*effects* of
 17 [an] action on the employee rather than the motivation for the action of the employer.” *Tex. Dep’t*
 18 *of Housing & Cmty. Aff. v. Inclusive Cmty. Project, Inc.*, 576 U.S. 519, 533 (2015). Although the
 19 availability of a disparate impact claim against an employment agency under Title VII and the
 20 ADEA appears to be a question of first impression, the answer flows directly from the text of the
 21 statute, read in light of the Supreme Court’s decisions in *Smith v. City of Jackson*, 544 U.S. 228
 22 (2005), and *Texas Department of Housing & Community Affairs v. Inclusive Communities Project,*
 23 *Inc.*, 576 U.S. 519 (2015).

24 **1. Disparate Impact Claims Lie Only Where The Applicable Provision**
 25 **Focuses On Effects And Not Just Motivation.**

26 Begin with *Smith*, which addressed whether the ADEA authorizes disparate impact claims
 27 under § 623(a)(2). That provision makes it unlawful for an employer “to limit, segregate, or classify
 28 his employees in any way which would deprive or tend to deprive any individual of employment

opportunities *or otherwise adversely affect* his status as an employee, because of such individual’s age.” 29 U.S.C. § 623(a)(2) (emphasis added). Those italicized words—“or otherwise adversely affect”—were all-important. 544 U.S. at 235 (plurality opinion). It was those same words in § 2000e-2(b) of Title VII, a plurality of the Court explained, that had supported the recognition of a disparate impact claim against employers under that statute. *Id.* at 235-36 (citing *Griggs v. Duke Power Co.*, 401 U.S. 424, 432-34 (1971)). The use of language that “focuses on the *effects* of the action” in § 623(a)(2) of the ADEA “strongly suggest[ed]” a disparate impact claim against employers under that provision as well. 544 U.S. at 236.

In *Texas Dept. of Housing*, the Court endorsed and adopted the same reasoning in analyzing the Fair Housing Act. That statute made it unlawful “[t]o refuse to sell or rent after the making of a bona fide offer, or to refuse to negotiate for the sale or rental of, *or otherwise make unavailable* or deny, a dwelling to any person because of race, color, religion, sex, familial status, or national origin.” 42 U.S.C. § 3604(a) (emphasis added). Expressly invoking *Smith*, the Court held “that antidiscrimination laws must be construed to encompass disparate-impact claims when their text refers to the consequences of actions and not just to the mindset of actors.” 576 U.S. at 533. The Court therefore found “the phrase ‘otherwise make unavailable’ . . . of central importance,” because it “refers to the consequences of an action rather than the actor’s intent.” *Id.* at 533-34.

By contrast, the Supreme Court has explained that neighboring provisions in the ADEA and Title VII that lack any reference to the effects of a practice do not embrace disparate impact liability. *Smith* observed that § 623(a)(1), which prohibits employers from “fail[ing] or refus[ing] to hire or to discharge any individual . . . with respect to his compensation, terms, conditions, or privileges of employment,” “does not encompass disparate-impact liability.” 544 U.S. at 236 n.6. And in *Ricci v. DeStefano*, the Court held that § 623(a)(1)’s Title VII counterpart, § 2000e-2(a)(1), can hold “employers liable only for disparate treatment.” 557 U.S. 557, 577 (2009).

2. Title VII And The ADEA’s Employment Agency Provisions Do Not Authorize Disparate Impact Claims.

The rule derived from *Smith* and *Texas Dep’t of Housing* is simply applied here. Section 623(b) of the ADEA and § 2000e-2(b) of Title VII lack any reference to the “effects” of an

1 employment agency’s actions, as compared with the agency’s motivations. The provisions are
2 materially identical. Both apply to employment agency practices and make it an unlawful
3 employment practice to “fail or refuse to refer for employment, or otherwise to discriminate against,
4 any individual because of [protected characteristic], or to classify or refer for employment any
5 individual on the basis of his [protected characteristic].” These prohibitions are solely focused on
6 the employment agency’s motivations for taking or not taking the relevant action. Absent from
7 either provision concerning employment agencies is the key phrase “otherwise adversely affect.”

8 The omission hardly could have been accidental. As explained above, in the employer-
9 focused provisions of each statute, Congress included reference to the effects of an action in one
10 subparagraph but not the others. *Compare* 42 U.S.C. § 2000e-2(a)(1) (no reference to effects) *and*
11 29 U.S.C. § 623(a)(1) (same) *with* 42 U.S.C. § 2000e-2(a)(2) (reference to effects) *and* 29 U.S.C.
12 § 623(a)(2) (same). Not only that, Congress carefully imported the “otherwise adversely affect”
13 language into the prohibitions that pertain to “labor organizations” as well, 42 U.S.C. §§ 2000e-2
14 (c)(2); 29 U.S.C. § 623(c)(2), without doing the same for employment agencies. “[W]hen Congress
15 includes particular language in one section of a statute but omits it in another section of the same
16 Act, it is generally presumed that Congress acts intentionally and purposely in the disparate
17 inclusion or exclusion.” *Barnhart v. Sigmon Coal Co.*, 534 U.S. 438, 452 (2002) (internal quotation
18 marks omitted).

19 Declining to extend disparate impact liability to employment agencies also makes good
20 sense. Employers have full control over their own employment practices, full perspective on the
21 effects those practices may have across their workforces, and a full understanding of when those
22 practices are justified by business necessity or factors other than some protected characteristic.
23 Similarly, labor organizations—which typically are made up of workers in a particular industry or
24 trade—have direct knowledge of the industries in which they operate, the needs of employers in
25 those industries, and the likely impact of particular practices used to match union members with
26 union jobs. The key factors governing disparate impact liability are thus within the knowledge and
27 control of both employers and labor organizations.

28 But that does not necessarily hold true for employment agencies, which are often just

1 intermediaries. An executive search firm that is provided non-discriminatory characteristics for an
2 ideal hire has no way of knowing if those neutral criteria, across a workforce, would have a
3 disparate impact. A temp agency tasked with finding dozens of employees with a particular
4 qualification has no direct means of confirming how essential that qualification is to the positions
5 being filled. Congress apparently appreciated that it would be unworkable and unfair to extend
6 disparate impact liability to entities that lack full knowledge of or control over the factors that could
7 yield liability. *See Eldredge*, 833 F.2d at 1337 (recognizing a union training organization “does not
8 stand in the same position as an employment agency” because it is “not simply an intermediary”).
9 And the problem would only be worse if, as Plaintiff suggests, the definition of employment
10 agencies extended to entities like Workday, whose knowledge of and control over the factors
11 bearing on disparate impact liability is likely to be even more tenuous.

12 Although courts have, in a few instances, entertained disparate impact claims against
13 employment agencies, none appear to have addressed the issue of whether the claim actually exists.
14 For example, in the sole Northern District of California case to previously address disparate impact
15 in the context of an employment agency, *Adams v. Vivo, Inc. and Genomic Health, Inc.*, Magistrate
16 Judge Ryu granted a motion to dismiss because the plaintiff failed to allege the existence of a
17 specific employment policy of the employment agency, rather than the employer. *Adams v. Vivo,*
18 *Inc.*, 2012 WL 5525315, *5 (N.D. Cal. Nov. 14, 2012). The defendants never contested whether a
19 disparate impact claim was proper as against the employment agency (as opposed to the employer,
20 who was the other named defendant in that case). *Id.* Other cases similarly had no occasion to
21 address the question. *Flanagan v. Excel Staffing Sols., LLC*, 2018 WL 558499, *4 (N.D. Ill. Jan.
22 25, 2018) (finding no “substantive difference between [plaintiff’s] disparate treatment theory and
23 his disparate impact theory,” both of which “say that [defendant] intentionally discriminated against
24 [the plaintiff] based on his race”); *Lucas v. Gold Standard Baking, Inc.*, 2014 WL 518000 (N.D.
25 Ill. Feb. 10, 2014) (permitting disparate impact claim to go forward on theory of joint employment).

26 At bottom, there is no basis in either statutory text or case law for the existence of a disparate
27 impact claim against an employment agency under Title VII and the ADEA. Plaintiff’s disparate
28 impact claims under those statutes must therefore be dismissed with prejudice.

- 20 -

DEFENDANT WORKDAY, INC.’S MOTION TO DISMISS AND MEMORANDUM OF POINTS AND
AUTHORITIES IN SUPPORT

1 **B. Plaintiff Has Insufficiently Pled A Disparate Impact Claim Under All Of The**
2 **Statutes On Which He Relies.**

3 To state a claim for disparate impact discrimination under the ADA (and also, to the extent
4 Plaintiff is permitted to bring a claim, under Title VII), “a plaintiff must allege (1) a significant
5 disparity with respect to employment for the protected group, (2) the existence of a specific
6 employment practice or set of practices, and (3) a causal relationship between the identified practice
7 and the disparity.” *Liu*, 551 F. Supp. 3d at 990 (citing *Freyd v. University of Or.*, 990 F.3d 1211,
8 1224 (9th Cir. 2021)); *Lopez v. Pac. Mar. Ass’n*, 2009 WL 10680881, at *5 (C.D. Cal. Apr. 3,
9 2009), *aff’d*, 657 F.3d 762 (9th Cir. 2011) (same elements as to ADA disparate impact claims). To
10 the extent Plaintiff could bring a claim under the ADEA, the elements are substantially similar: “a
11 plaintiff must demonstrate (1) the occurrence of certain outwardly neutral employment practices,
12 and (2) a significantly adverse or disproportionate impact on persons of a particular [age] produced
13 by the employer’s facially neutral acts or practices.” *Katz v. Regents of the Univ. of Cal.*, 229 F.3d
14 831, 835-36 (9th Cir. 2000) (internal quotation marks omitted). Here, Plaintiff’s Complaint fails in
15 every respect. It includes no allegations lending plausibility to his bare assertion of a statistical
16 disparity concerning protected classes. *Infra* § 1. Additionally, Plaintiff neither identifies a specific
17 employment practice or set of practices he is challenging, nor does he articulate *how* any such
18 practice(s) could cause a statistical disparity with the required amount of specificity. *Infra* § 2.

19 **1. Plaintiff Has Not Pled A Disparity.**

20 With respect to a disparity, Plaintiff alleges merely that “upon information and belief,
21 [Workday’s “selection tools”] disparately impact African-American, [disabled, and over-40]
22 applicants because the tests have the effect of disproportionately excluding African-American[],
23 [disabled, and over-40 applicants].” Compl. ¶¶ 24, 27, 29. Plaintiff does not specify whether this
24 alleged disparate impact exists with respect to any one Workday customer, across all of the
25 customers who posted the “80-100” positions to which Plaintiff applied, or something else. He
26 includes no allegations concerning why he believes this disparity exists, what investigation he
27 conducted into the matter, or the alleged magnitude of this purported disparity.

28 That is not enough. Although a plaintiff need not define an alleged disparity with great

1 precision at the pleading stage, he must do more than “merely describe[e] his own experience.” *Liu*,
2 551 F. Supp. 3d at 991. Instead, a plaintiff must at least make a “sophisticated effort . . . to develop
3 a plausible factual basis in support of his assertion” of a disparity. *Id.* Even where a plaintiff would
4 not be expected to have full information, he cannot simply “plow ahead to discovery” without
5 providing some basis for a plausible inference of a disparity. *Id.* Other cases are in accord. *See*
6 *Borodaenko v. Twitter, Inc.*, 2023 WL 3294581, at *4 (N.D. Cal. May 5, 2023) (“Plaintiff alleges
7 that ‘[m]any disabled employees . . . have felt that, because of their disability, they will not be able
8 to meet [Defendant’s] new heightened standard of performance and productivity,’ and ‘have felt
9 forced to resign.’ But Plaintiff has not alleged that such individuals have actually resigned or that
10 they have done so at higher rates than their non-disabled colleagues.”) (citation omitted); *Alhayoti*
11 *v. Blinken*, 2022 WL 2869850, at *9 (N.D. Cal. July 21, 2022) (“The plaintiff has not alleged any
12 facts showing that the Department’s distinction between local employees and direct hires had a
13 disproportionate impact on any protected group. The complaint does not cite any data, scientific
14 literature, or personal experiences to support the theory that the Department’s local-employee
15 hiring practice had a discriminatory impact on a specific protected group.”).

16 Because he “has offered no allegations about what is actually happening,” *Liu*, 551 F. Supp.
17 3d at 991, with respect to the hiring of Black, disabled, or over-40 applicants among customers who
18 allegedly use Workday’s software, Plaintiffs’ claim should be dismissed.

19 **2. Plaintiff Has Not Pled A Specific Employment Practice.**

20 Plaintiff’s disparate impact claims fail for the additional, independent reason that he fails to
21 identify a specific employment policy or practice that allegedly caused an unlawful disparity.

22 Time and again, the Supreme Court has recognized that it is “not enough to simply allege
23 that there is a disparate impact on workers, or point to a generalized policy that leads to such an
24 impact.” *Smith*, 544 U.S. at 241. Instead, Plaintiff must “isolate[e] and identify[] the *specific*
25 employment practices that are allegedly responsible for any observed statistical disparities.” *Id.*
26 (quoting *Watson v. Fort Worth Bank & Trust*, 487 U.S. 977, 994 (1988)); accord *Wards Cove*
27 *Packing Co. v. Atonio*, 490 U.S. 642, 657 (1989) (“[A] plaintiff must demonstrate that it is the
28 application of a specific or particular employment practice that has created the disparate impact.”),

- 22 -

1 *superseded by statute on other grounds*, 42 U.S.C. § 2000e-2(k); *Dukes*, 564 U.S. at 357
2 (emphasizing specificity requirement and explaining that it is “all the more necessary when a class
3 of plaintiffs is sought to be certified”).

4 Meeting this specificity requirement is also essential to pleading a causal connection
5 between an employment practice and an alleged disparity. As then-Judge Brown Jackson explained
6 after an in-depth analysis of disparate impact law, a plaintiff must ultimately “supply an account of
7 how [identified] procedures themselves resulted in . . . racially disparate outcomes.” *Ross v.*
8 *Lockheed Martin Corp.*, 267 F. Supp. 3d 174, 198 (D.D.C. 2017). Absent that, “[i]t is completely
9 unrealistic to assume that [employers’] unlawful discrimination is the sole cause of . . . statistical
10 imbalances in the composition of their work forces.” *Id.* at 181 (quoting *Watson*, 487 U.S. at 992
11 (plurality opinion)). Where a plaintiff fails even to identify the particulars of a challenged
12 employment practice, it cannot begin to articulate the mechanism by which that practice could
13 cause a disparity. *Id.* at 199.

14 Here, Plaintiff comes nowhere close to meeting these requirements of specificity. *First*,
15 Plaintiff fails to even identify the employment practice he purports to challenge. Instead, Plaintiff
16 makes various vague references to the following alleged Workday products:

- 17 • “artificial intelligence (“AI”) systems and screening tools,”
- 18 • “algorithm-based applicant screening system,”
- 19 • “discriminatory screening process,”
- 20 • “discriminatory screening products,”
- 21 • “selection tools marketed by Workday,”
- 22 • “Workday’s products,”
- 23 • “screening products, policies, practices, and procedures,”
- 24 • “selection tools and/or tests,”
- 25 • “selection procedures,” and
- 26 • “policies, patterns, and/or practices.”

27 It is impossible to discern any identifiable employment practice that could have caused an unlawful
28 disparity. Is it a single software product? Several different products? The algorithm embedded in

- 23 -

DEFENDANT WORKDAY, INC.’S MOTION TO DISMISS AND MEMORANDUM OF POINTS AND
AUTHORITIES IN SUPPORT

1 that product? The intended or ordinary use of that product? A service or practice Workday provides
2 with respect to that product? The process into which an alleged Workday screening product might
3 be incorporated? Does it differ based on the point in time, the client, or the job itself? The Complaint
4 is so devoid of detail regarding the alleged discriminatory employment practice that it fails even
5 the basic requirement of providing notice to Workday of conduct that is purportedly at issue in this
6 case. *See Canatella v. Reverse Mortg. Sols. Inc.*, 2014 WL 7205146, at *8 (N.D. Cal. Dec. 17,
7 2014) (plaintiff’s disparate impact claim failed where plaintiff did not plead facts demonstrating a
8 “specific, identified . . . practice or selection criterion” caused “a significant disparate impact”).

9 Indeed, the only information Plaintiff provides strongly suggests that he is not challenging
10 any specific employment practice, but a nebulously plural set of associated products or services
11 that are inevitably put to different uses by different customers. Plaintiff specifically alleges that
12 Workday’s screening tools “allow [Workday’s] customers to manipulate and configure them in a
13 discriminatory manner to recruit, hire, and onboard employees,” Compl. ¶ 8; “rel[y] upon
14 subjective practices [by private companies],” *id.*; and permit “decision-makers [] free[dom] to
15 exercise *their* discretion in an *unguided, subjective* manner that provides a ready mechanism for
16 discrimination” (Compl. ¶ 37) (emphasis added). In emphasizing this subjectiveness and variation,
17 Plaintiff has pled the opposite of a “specific employment practice or set of practices” responsible
18 for the alleged disparate impact and implemented by Workday. *Liu*, 551 F. Supp. 3d at 990; *Dukes*,
19 564 U.S. at 355. The only thing the Complaint tells us about what the challenged employment
20 practice *does* is that it almost certainly did it differently across the many positions Plaintiff applied
21 for, let alone across the positions that all members of the putative class applied to.

22 **3. Plaintiff Has Not Sufficiently Pled Causation.**

23 Even were the Court able to fashion a coherent challenged practice, Plaintiff provides no
24 detail concerning how the challenged employment practice could possibly cause any unlawful
25 disparate impact. Plaintiff says nothing of how Workday’s products work, what mechanism or
26 mechanisms they use to “screen” applicants, what inputs are relevant, and what outputs the products
27 provide. He also says nothing of how Workday’s customers make use of whatever it is he is
28 claiming Workday provides. Without this information, it is impossible for the Court to reach the

- 24 -

DEFENDANT WORKDAY, INC.’S MOTION TO DISMISS AND MEMORANDUM OF POINTS AND
AUTHORITIES IN SUPPORT

1 conclusion that the challenged practice *caused* the alleged disparate impact.

2 Plaintiff fails to explain *how* whatever practice he is challenging would work to create a
3 disparity concerning any protected class. He alleges no causal mechanism at all other than the
4 unspecific suggestion that customers could theoretically import subjectivity or bias into their use
5 of Workday’s tools. Courts have dismissed disparate impact claims based on far more detailed
6 accounts of challenged practices and their alleged effects. In *Moussouris v. Microsoft Corp.*, for
7 example, plaintiffs challenged a specific “‘stack-ranking’ employee review policy,” and articulated
8 how “the forced curve might lower an employee’s final score from the score the employee’s
9 manager felt the employee deserved before applying the curve.” 2016 WL 4472930, *9-10 (W.D.
10 Wash. Mar. 7, 2016). The plaintiffs alleged that these “uniform, unvalidated, and unreliable
11 procedures for evaluating employee performance . . . systematically undervalue female technical
12 employees relative to their similarly situated peers.” *Id.* But even these details were not enough,
13 because the plaintiffs did not supply the why: “[N]owhere does the operative complaint plausibly
14 allege why the system might ‘systematically undervalue[] female’ employees.” *Id.* at *9.

15 Because Plaintiff fails to allege what the challenged employment practice is or does, or how
16 it could cause a statistical disparity, his disparate impact claims should be dismissed.

17 **VII. CONCLUSION**

18 For the foregoing reasons, Defendant respectfully requests the Court grant Defendant’s
19 Motion to Dismiss.

20 Dated: July 17, 2023.

ORRICK, HERRINGTON & SUTCLIFFE LLP

21
22 By: /s/ Julie A. Totten
23 JULIE A. TOTTEN
24 ERIN M. CONNELL
25 KAYLA D. GRUNDY
26 ALEXANDRIA R. ELLIOTT
27 JUSTIN M. WASHINGTON
28 Attorneys for Defendant
WORKDAY, INC.



U.S. Equal Employment Opportunity Commission

Press Release

09-11-2023

iTutorGroup to Pay \$365,000 to Settle EEOC Discriminatory Hiring Suit

Settles Federal Charges Tutoring Provider Programmed its Online Software to Automatically Reject More Than 200 Older Applicants

NEW YORK – iTutorGroup, three integrated companies providing English-language tutoring services to students in China, will pay \$365,000 and furnish other relief to settle an employment discrimination lawsuit filed by the U.S. Equal Employment Opportunity Commission (EEOC), the federal agency announced today.

iTutorGroup, composed of iTutorGroup, Inc.; Shanghai Ping’An Intelligent Education Technology Co., Ltd.; and Tutor Group Limited hired tutors based in the United States to provide online tutoring from their homes or other remote locations. According to the EEOC’s lawsuit, iTutorGroup programmed their tutor application software to automatically reject female applicants aged 55 or older and male applicants aged 60 or older. iTutorGroup rejected more than 200 qualified applicants based in the United States because of their age.

iTutorGroup’s alleged conduct violates the Age Discrimination in Employment Act (ADEA), which prohibits employers from discriminating based on age. The EEOC filed suit (EEOC v. iTutorGroup, Inc., et al., Civil Action No. 1:22-cv-02565) after first attempting to reach a pre-litigation settlement through its conciliation process.

“Prohibitions on age and other types of discrimination do not stop at the border,” said Trial Attorney Daniel Seltzer. “Even companies doing business abroad will face serious consequences if they discriminate against U.S.-based employees.”

“Where, as alleged here, companies closely control the way fully remote workers perform their jobs, those workers are employees protected by federal anti-discrimination laws,” said Jeffrey Burstein, regional attorney for the EEOC’s New York District Office. “The EEOC will continue to enforce those protections for all covered employees.”

The decree settling the suit provides \$365,000 to be distributed to applicants who were automatically rejected due to age. Although iTutorGroup has ceased hiring tutors in the United States, the decree also provides for significant non-monetary relief designed to prevent discrimination should iTutorGroup ever resume its U.S. operations. That relief includes extensive and continuing training for those involved in hiring tutors, issuance of a robust new anti-discrimination policy, and strong injunctions against discriminatory hiring based on age or sex and requesting applicants’ birth dates. The EEOC will monitor iTutorGroup’s compliance with these obligations for at least the next five years or longer if iTutorGroup resumes hiring tutors in the United States, and if iTutorGroup does resume its U.S. operations, it must notify and interview those applicants allegedly rejected because of age.

“Everyone loses when employers engage in age discrimination,” said Timothy Riera, the acting director of the EEOC’s New York District. “Hundreds of applicants lost out on employment during a difficult time for job seekers, and iTutorGroup’s students lost the opportunity to learn English from highly qualified and experienced tutors.”

For more information on age discrimination, please visit [**https://www.eeoc.gov/age-discrimination**](https://www.eeoc.gov/age-discrimination) ([**https://www.eeoc.gov/age-discrimination**](https://www.eeoc.gov/age-discrimination)).

The EEOC’s New York District Office is responsible for addressing discrimination charges and conducting agency litigation in Connecticut, Maine, Massachusetts, New Hampshire, New York, northern New Jersey, Rhode Island and Vermont.

The EEOC advances opportunity in the workplace by enforcing federal laws prohibiting employment discrimination. More information is available at [**www.eeoc.gov**](http://www.eeoc.gov) ([**http://www.eeoc.gov**](http://www.eeoc.gov)). Stay connected with the latest EEOC news by subscribing to our [**email updates**](https://public.govdelivery.com/accounts/USEEOC/subscriber/new) ([**https://public.govdelivery.com/accounts/USEEOC/subscriber/new**](https://public.govdelivery.com/accounts/USEEOC/subscriber/new)).

B. Subject Matter Priorities for Fiscal Years 2024-2028

The Commission’s goal in identifying agency-wide subject matter priorities is to ensure that the agency’s resources are focused on preventing and remedying discrimination and advancing equal employment opportunity in circumstances where EEOC enforcement is most likely to achieve strategic impact. The EEOC will use all its tools, including enforcement (including investigations, settlements, and litigation), education and outreach, research, and policy development, to advance the agency’s priorities.

The Commission relied on the following criteria to identify subject matter priorities for this SEP:

1. Issues that will have **broad impact** because of the nature and scope of the employment practices addressed, the number of individuals impacted, or the employers or industries affected;
2. Issues **affecting workers who may be unaware of their legal rights** or reluctant or unable to exercise their rights;
3. Issues involving **developing areas of the law**, where the Commission’s expertise is particularly valuable;
4. Issues involving policies or practices that **impede or impair full enforcement of federal employment discrimination laws**; and
5. Issues that may be **best addressed by government action, including enforcement**, based on the nature of the claim, the types of relief available, practical or legal impediments to private enforcement, or the Commission’s access to information, data, and research.

C. Subject Matter Priorities

The following are the EEOC’s subject matter priorities for Fiscal Years 2024-2028:

1. Eliminating Barriers in Recruitment and Hiring

The EEOC will focus on recruitment and hiring practices and policies that discriminate on any basis unlawful under the statutes EEOC enforces, including sex, race, national origin, color, religion, age, and disability. These include:

- the use of technology, including artificial intelligence and machine learning, to target job advertisements, recruit applicants, or make or assist in hiring decisions where such systems intentionally exclude or adversely impact protected groups;
- job advertisements that exclude or discourage certain protected groups from applying;
- channeling, steering or segregating individuals into specific jobs or job duties based on protected characteristics;

- policies and practices that limit access to on-the-job training, pre-apprenticeship or apprenticeship programs, temp-to-hire positions, internships, or other job training or advancement opportunities based on protected characteristics;
- policies and practices that limit employees exclusively to temporary work on a basis prohibited by federal employment laws when permanent positions are available for which they are qualified;
- reliance on restrictive application processes or systems, including online systems that are difficult for individuals with disabilities or other protected groups to access; and
- the use of screening tools or requirements that disproportionately impact workers on a protected basis, including those facilitated by artificial intelligence or other automated systems, pre-employment tests, and background checks.

The continued underrepresentation of women and workers of color in certain industries and sectors (for example, construction and manufacturing, high tech, STEM, and finance, among others), are also areas of particular concern, especially in industries that benefit from substantial federal investment.

2. Protecting Vulnerable Workers and Persons from Underserved Communities from Employment Discrimination

The EEOC will focus on harassment, retaliation, job segregation, labor trafficking, discriminatory pay, disparate working conditions, and other policies and practices that impact particularly vulnerable workers and persons from underserved communities. With respect to employment discrimination, the Commission views the category of vulnerable workers as including:

- immigrant and migrant workers and workers on temporary visas;
- people with developmental or intellectual disabilities;
- workers with mental health related disabilities;
- individuals with arrest or conviction records;
- LGBTQI+ individuals;
- temporary workers;
- older workers;
- individuals employed in low wage jobs, including teenage workers employed in such jobs;
- survivors of gender-based violence;
- Native Americans/Alaska Natives; and
- persons with limited literacy or English proficiency.

These workers may be unaware of their rights under equal employment opportunity laws, may be reluctant or unable to exercise their legally protected rights, and/or have historically been underserved by federal employment discrimination protections. Factors such as immigration status, language barriers, education level, poverty and/or economic circumstances, geographic location, isolated work conditions, age, disability status, societal stigma, or lack of employment

experience can make these workers particularly vulnerable to discriminatory practices or policies.⁶

To implement this priority, district offices and the agency’s federal sector program will identify vulnerable workers and underserved communities in their districts or within the federal sector for focused attention, based on their assessment of how the EEOC can most effectively utilize its resources to address issues of concern for these groups. For example, employment discrimination against Native Americans/Alaska Natives, indigenous people from Latin America, agricultural workers, workers on temporary visas, or individuals with arrest or conviction records might be areas of focus as part of this priority.

3. Addressing Selected Emerging and Developing Issues

The EEOC will continue to prioritize issues that may be emerging or developing, including issues that involve new or developing legal concepts or topics that are difficult or complex. The agency is uniquely suited to address these issues given the EEOC’s research, data collection, receipt of charges in the private and public sectors, adjudication of complaints and oversight in the federal sector, and ongoing engagement with stakeholders.

Because of the nature of this priority category, the Commission may add or remove issues through interim amendments to the SEP. The following issues currently fall within this category:

- a) Qualification standards and inflexible policies or practices that discriminate against individuals with disabilities
- b) Protecting workers affected by pregnancy, childbirth, or related medical conditions under the Pregnancy Discrimination Act (PDA) and the Pregnant Workers Fairness Act (PWFA), as well as pregnancy-related disabilities under the Americans with Disabilities Act (ADA)
- c) Addressing discrimination influenced by or arising as backlash in response to local, national, or global events, including discriminatory bias arising as a result of recurring historical prejudices

For example, discrimination, bias, and hate directed against religious minorities (including antisemitism and Islamophobia), racial or ethnic groups, and LGBTQI+ individuals may fall

⁶ The agency’s 2013-2016 and 2017-2021 SEPs included sex discrimination against LGBTQI+ individuals as an emerging and developing issue. In *Bostock v. Clayton County*, 140 S. Ct. 1731 (2020), the Supreme Court affirmed—as the EEOC had held several years earlier—that Title VII of the Civil Rights Act of 1964 prohibits employment discrimination on the basis of sexual orientation or gender identity. Although Title VII’s coverage of sexual orientation and gender identity discrimination has been settled by the Supreme Court and thus is no longer an emerging legal issue, preventing and remedying discrimination against LGBTQI+ individuals remain a key priority for the EEOC, as such individuals are often vulnerable and have historically been underserved. The Commission has therefore included LGBTQI+ individuals in the vulnerable and underserved worker priority.

within this subcategory. The discriminatory practices or affected groups or individuals may change during the time period covered by this SEP.

- d) Employment discrimination associated with the long-term effects of the COVID-19 pandemic, including Long COVID

- e) **Technology-related employment discrimination**

The EEOC will focus on employment decisions, practices, or policies in which covered entities' use of technology contributes to discrimination based on a protected characteristic. These may include, for example, the use of software that incorporates algorithmic decision-making or machine learning, including artificial intelligence; use of automated recruitment, selection, or production and performance management tools; or other existing or emerging technological tools used in employment decisions.

4. Advancing Equal Pay for All Workers

The EEOC will continue to focus on combatting pay discrimination in all its forms—on the basis of sex under the Equal Pay Act and Title VII, on other protected bases covered by federal anti-discrimination laws, including race, national origin, disability, and age, and at the intersection of protected bases. Because many workers do not know how their pay compares to their coworkers' and, therefore, are less likely to discover and report pay discrimination, the Commission will continue to use directed investigations and Commissioner Charges, as appropriate, to facilitate enforcement.

The Commission will also focus on employer practices that may impede equal pay or contribute to pay disparities and may lead to violations of statutes the Commission enforces, such as pay secrecy policies, discouraging or prohibiting workers from asking about pay or sharing their pay with coworkers, and reliance on past salary history or applicants' salary expectations to set pay.

5. Preserving Access to the Legal System

The EEOC will focus on policies and practices that limit substantive rights, discourage or prohibit individuals from exercising their rights under employment discrimination statutes, or impede the EEOC's investigative or enforcement efforts. For example, this priority includes policies or practices that deter or prohibit filing charges with the EEOC or cooperating freely in EEOC investigations or litigation. Specifically, the EEOC will focus on:

- a) overly broad waivers, releases, non-disclosure agreements, or non-disparagement agreements;
- b) unlawful, unenforceable, or otherwise improper mandatory arbitration provisions;
- c) employers' failure to keep applicant and employee data and records required by statute or EEOC regulations; and
- d) retaliatory practices that could dissuade employees from exercising their rights under employment discrimination laws. This subcategory focuses on retaliatory practices that detrimentally impact or otherwise affect employees beyond those engaging in protected activity.

APPLYING OLD RULES TO NEW TOOLS: EMPLOYMENT
DISCRIMINATION LAW IN THE AGE OF ALGORITHMS

Matthew U. Scherer,* Allan G. King** & Marko N. Mrkonich***

I. INTRODUCTION.....	450
II. BACKGROUND.....	453
A. <i>Algorithmic Selection Tools</i>	453
1. <i>Machine Learning and Deep Learning</i>	453
2. <i>Algorithmic Employee Selection</i>	456
B. <i>Law of Discrimination</i>	457
1. <i>Disparate Treatment</i>	459
2. <i>Disparate Impact</i>	460
a. <i>The Development of the Disparate Impact Doctrine</i>	460
b. <i>Prima Facie Case</i>	462
c. <i>Business Necessity Defense</i>	464
d. <i>Least Discriminatory Alternative</i>	471
3. <i>Ricci v. DeStefano and the Interplay Between Disparate Treatment and Disparate Impact</i>	472
III. THE UNIQUE CHALLENGES FOR ALGORITHMIC SELECTION TOOLS.....	477
A. <i>Validation</i>	478
1. <i>The Difficult Path to Criterion-Related Validity Under the Guidelines</i>	478
a. <i>Job Analysis</i>	478
b. <i>Selecting Criteria</i>	479
2. <i>The Guidelines: Behind the Times</i>	482
B. <i>The Pitfalls of Correlations and Big Data</i>	484
1. <i>The Ubiquity and Meaninglessness of Statistical Significance in Large Data Sets</i>	484
2. <i>Construct-Irrelevant Variance and Construct Underrepresentation</i>	485
3. <i>Redundant Encodings</i>	491
C. <i>The Black Box Problem</i>	492
D. <i>A Clear Target</i>	494
E. <i>Disparate Treatment: A Brave New World</i>	496
IV. NEW RULES FOR THE NEW TOOLS: A PROPOSED LEGAL FRAMEWORK FOR ALGORITHMIC SELECTION TOOLS.....	499
A. <i>Overview</i>	499
B. <i>Disparate Impact</i>	502
1. <i>Prima Facie Disparate Impact</i>	502

2. <i>Business Necessity Defense</i>	505
3. <i>Alternative Selection Procedures</i>	509
C. <i>Disparate Treatment</i>	512
1. <i>Safe Harbor 1: Differential Validation</i>	513
2. <i>Safe Harbor 2: Statistical Independence</i>	517
3. <i>A Treacherous Harbor: Constrained Optimization</i>	518
V. CONCLUSION.....	520

I. INTRODUCTION

Companies, policymakers, and scholars alike are paying increasing attention to the use of machine learning (ML) in recruitment and hiring, most notably in the form of ML-based employee selection tools that use algorithms in place of traditional employment tests and the judgment of human recruiters.¹ To its advocates, ML-based selection processes can be more effective in choosing the strongest candidates, increasing diversity, and reducing the influence of human prejudices.² Many observers, however, express concern about other forms of bias that can infect algorithmic selection procedures, leading to fears regarding the potential for algorithms to create unintended discriminatory effects, reinforce existing patterns of discrimination, or mask more deliberate forms of discrimination.³

In the Authors' experiences, most employers very much want to improve diversity and inclusion, from company leadership down to the most junior hourly employees. Companies pursue these objectives not just to avoid legal liability for violating antidiscrimination statutes, but also because they have concluded that a more diverse and inclusive workforce is better from both a

* Analytics Associate, Littler Mendelson, P.C.

** Shareholder, Littler Mendelson, P.C.

*** Shareholder, Littler Mendelson, P.C.

1. See Gil Press, *120 AI Predictions for 2019*, FORBES (Dec. 9, 2018, 12:00 PM), <https://www.forbes.com/sites/gilpress/2018/12/09/120-ai-predictions-for-2019/#54aa4326688c> [<https://perma.cc/XTS4-EL9P>]; Neelie Verlinden, *Machine Learning in Recruitment & How to Do It Right*, HARVER (Oct. 23, 2018, 3:53 PM), <https://harver.com/blog/machine-learning-in-recruitment/> [<https://perma.cc/RPR3-T29M>] (discussing how machine learning uses algorithms).

2. See HAIYAN ZHANG ET AL., THE ROLE OF AI IN MITIGATING BIAS TO ENHANCE DIVERSITY AND INCLUSION 6 (2019).

3. See *id.* at 8.

business perspective⁴ and an ethical perspective.⁵ Indeed, many (if not most) of the employers who are turning to algorithmic and data-driven selection tools are doing so in part because they want to guard against human biases that can serve as barriers to employment for disadvantaged groups.⁶

Not coincidentally, the eradication of such barriers is, as the Supreme Court long ago recognized, the overarching objective of antidiscrimination laws.⁷ Because this is an area where the objectives of the law and of America's businesses are well-aligned,⁸ one would think that the law should serve as an inducement rather than a deterrent to companies who wish to deploy algorithmic selection tools that will allow them to improve both the quality and diversity of their employees. Unfortunately, that has not been the case.⁹

The rules governing employment tests and other employee selection procedures were developed in the 1970s and have remained largely unchanged in the decades since.¹⁰ Those rules, written as they were for paper-and-pencil tests and other in-person examinations, are ill-suited for selection procedures that rely on a candidate's historical data rather than real-time observations and firsthand assessments. Complicating matters further, the complexity of the algorithms that underlie ML-based selection tools makes it

4. See, e.g., Dennis Nally, *Five Reasons Why Diversity and Inclusion Matter to Every Business—and Every Employee*, PWC: CEO INSIGHTS (June 15, 2015), <https://pwc.blogs.com/ceoinsights/2015/06/five-reasons-why-diversity-and-inclusion-matter.html> [<https://perma.cc/N7W9-EPGJ>] (noting that 85% of surveyed CEOs whose companies have diversity and inclusion strategies say that it has “improved their bottom line”); DELOITTE, 2017 BOARD DIVERSITY SURVEY: SEEING IS BELIEVING 8 (2017) (noting that over 90% of respondents believe that greater diversity on a company's board of directors enables an organization to improve its ability to innovate, ability to manage disruptions, and overall business performance). See generally MARK KAPLAN & MASON DONOVAN, *THE INCLUSION DIVIDEND* (2013) (noting that diversity and inclusion should lead to a decrease of expenses or an increase of revenue).

5. See Nally, *supra* note 4 (“Diversity and inclusion are quite simply the right thing to do[.] It's about creating equal opportunities for everyone—and we can all see signs of progress. But the statistics make it equally clear that there's still a long way to go.”).

6. See Nicole Lewis, *Will AI Remove Hiring Bias?*, SHRM (Nov. 12, 2018), <https://www.shrm.org/resourcesandtools/hr-topics/talent-acquisition/pages/will-ai-remove-hiring-bias-hr-technology.aspx> [<https://perma.cc/FM3Y-4WL8>].

7. *Griggs v. Duke Power Co.*, 401 U.S. 424, 429–31 (1971) (“The objective of Congress in the enactment of Title VII is plain from the language of the statute. It was to achieve equality of employment opportunities and remove barriers that have operated in the past to favor an identifiable group of white employees over other employees . . . What is required by Congress is the removal of artificial, arbitrary, and unnecessary barriers to employment when the barriers operate invidiously to discriminate on the basis of racial or other impermissible classification.”).

8. See Lewis, *supra* note 6 (“A study from McKinsey and Company found that companies that have a diverse workforce financially outperform companies that don't.”).

9. See *id.* (“Discrimination in hiring, however, is proving difficult to reverse.”).

10. Michael A. McDaniel et al., *The Uniform Guidelines Are a Detriment to the Field of Personnel Selection*, 4 INDUS. & ORGANIZATIONAL PSYCHOL. 494, 507 (2011).

difficult for employers and employees alike to discern how and why an algorithm came up with its scores, rankings, or recommendations.¹¹ This Article seeks to both highlight the challenges employers, workers, courts, and agencies will face as companies develop and deploy algorithmic selection tools, and propose a framework through which courts and agencies can assess whether such tools comply with antidiscrimination laws.

Part II begins with a brief overview of the technological concepts that underlie algorithmic employee selection procedures. It continues with a discussion of the development of antidiscrimination laws, along with the broader philosophical and legal principles that animate the two major forms of employment discrimination—disparate treatment and disparate impact. Part III details why algorithmic selection procedures fit poorly into the legal framework that has developed around Title VII and similar antidiscrimination laws.

Part IV proposes a uniform analytical framework through which agencies and courts can analyze whether an employer using a particular algorithmic selection tool has engaged in disparate treatment or disparate impact. The proposed framework is built upon three unifying themes:

- In light of the low practical value of statistical significance tests in the age of Big Data, flexible tests of reasonableness should replace inflexible tests of statistical significance when assessing whether a correlation is legally meaningful;
- Courts and agencies should recognize certain forms of ML-based fairness techniques as acceptable ways for employers to mitigate disparate impacts without exposing themselves to disparate treatment liability; and
- Standards for assessing the validity of an algorithmic selection procedure should focus on whether the procedure is based on the essential and important job functions of a particular position, as identified through an adequate job analysis and incorporated into a properly constructed model.

This framework, we posit, would give full effect to the objectives of antidiscrimination laws without discouraging employers from using machine learning and Big Data not only to increase efficiency, but also to improve

11. See Gregory Barber, *Shark or Baseball? Inside the 'Black Box' of a Neural Network*, WIRED (Mar. 6, 2019, 12:00 PM), <https://www.wired.com/story/inside-black-box-of-neural-network/> [<https://perma.cc/43U5-TMVD>] (noting that how neural networks work is still a mystery).

diversity and reduce the effects of human biases on the recruitment and hiring process.

II. BACKGROUND

A. Algorithmic Selection Tools

Many of the principles discussed in this Article will be relevant to all forms of data-driven employee selection procedures. But the primary focus will be on algorithmic tools¹² that utilize machine learning with a particular emphasis on those that use deep learning. Tools that rely on these sophisticated algorithmic methods pose challenges that exceed those of earlier generations of data-driven employee selection tools.¹³

1. Machine Learning and Deep Learning

Machine learning is a branch of artificial intelligence consisting of algorithms that learn from data.¹⁴ In this context, “learn” means that the algorithm uses statistical methods and “data-driven insights” to allow an AI system to improve itself without human intervention.¹⁵ A learning algorithm uses training data to build a statistical model that can then be used to make predictions or other decisions about new data.¹⁶ Learning algorithms may entail varying levels of mathematical and computational complexity. The highest profile breakthroughs in artificial intelligence over the past several years have come from a subfield of machine learning known as deep

12. This Article will generally use the term “tool” to refer to such algorithmic programs both as a convenient shorthand and to suggest that for the foreseeable future, employers likely will be using algorithmic selection procedures primarily to supplement or improve their existing human-driven employee selection process. This Article will also use the terms “tool,” “test,” and “selection procedure” interchangeably to refer to algorithmic systems and programs that are used to make, or to help an employer make, personnel decisions. For the most part, this Article will focus on hiring, but in most cases, the same principles will apply to the use of algorithms to make decisions regarding placement, compensation, promotion, termination, transfer, or other actions that affect aspects of employment.

13. Cf. ZHANG ET AL., *supra* note 2 (noting that AI algorithms are reliant on human data and can lead to biases).

14. IAN GOODFELLOW ET AL., *DEEP LEARNING* 96 (2016); Verlinden, *supra* note 1.

15. Tobias Baer & Vishnu Kamalnath, *Controlling Machine-learning Algorithms and Their Biases*, MCKINSEY & CO. 1 (Nov. 2017), <https://www.mckinsey.com/business-functions/risk/our-insights/controlling-machine-learning-algorithms-and-their-biases#> [<https://perma.cc/745L-C45M>].

16. *Learning Algorithm*, TECHOPEDIA, <https://www.techopedia.com/definition/33426/learning-algorithm> [<https://perma.cc/53NE-ACH8/>].

learning.¹⁷ Deep learning involves the use of artificial neural networks that are inspired by how neurons in the human brain are thought to interact with each other.¹⁸ A neural network operates by taking certain data as an input and passing that data through one or more layers of artificial “neurons” that analyze and transform the data.¹⁹

Most machine learning approaches can be classified under one of two broad headings: supervised learning or unsupervised learning.²⁰ In supervised learning, the training data is labeled by humans.²¹ In unsupervised learning, by contrast, the algorithm proceeds by looking for patterns in unlabeled data.²² In general, supervised learning techniques are better suited for applications where the developers are interested in predicting a specific outcome.²³ For example, to build an algorithm that takes photographic images as inputs and that will output a prediction as to whether the image contains a cat, a sensible approach would be to use supervised learning where the training data consists of images that humans have reviewed and labeled as “cat” or “not a cat.” On the other hand, an unsupervised learning algorithm might be an appropriate choice for a more general object-recognition algorithm, where the algorithm would receive unlabeled images as input, examine the content of each image, and identify groups of images that it identifies as having shared characteristics.²⁴

In technical parlance, the data sets used to train learning algorithms are said to consist of “instances” (also known as examples, observations, subjects,

17. Tom Simonite, *The WIRED Guide to Artificial Intelligence*, WIRED (Feb. 1, 2018, 9:22 AM), <https://www.wired.com/story/guide-artificial-intelligence/> [<https://perma.cc/QR4E-CGXB>].

18. See Chris Nicholson, *A Beginner's Guide to Neural Networks and Deep Learning*, SKYMIND, <https://skymind.ai/wiki/neural-network> [<https://perma.cc/4YFU-XS7N>] [hereinafter *A Beginner's Guide*].

19. *Id.*

20. *Id.* A third type of machine learning is *reinforcement learning*. A reinforcement learning algorithm is conceptually similar to a supervised learning algorithm except that, rather than using human-labeled training data for optimization, it optimizes itself by updating its decision-making strategy in response to feedback received on prior decisions, usually via a reward function. See Chris Nicholson, *A Beginner's Guide to Deep Reinforcement Learning*, PATHMIND, <https://pathmind.com/wiki/deep-reinforcement-learning> [<https://perma.cc/4N7A-PMEP>]. This type of machine learning has become prominent in many spheres, but it seems unlikely to have much application in the context of recruitment and hiring algorithms in the near future and thus is not discussed further in this Article.

21. *Id.*

22. *Id.*

23. See Chris Nicholson, *A Beginner's Guide to Supervised Learning*, SKYMIND, <https://skymind.ai/wiki/supervised-learning> [<https://perma.cc/57JY-JAN3>].

24. See *A Beginner's Guide*, *supra* note 18.

or units) and “attributes” (also known as features or covariates).²⁵ Instances generally correspond to the rows on a spreadsheet²⁶ and, for purposes of the types of tools that are the subject of this Article, most often represent individual persons. Attributes are the measurable properties and characteristics of interest for each instance and are analogous to column headings in spreadsheets, such as “educational attainment” or “years of experience.”²⁷ The number of attributes included in a dataset is referred to as the “dimensionality” of the data set.²⁸

While a detailed description of deep learning architectures is beyond the scope of this Article, a few characteristics are notably relevant to the legal challenges that employers using algorithmic selection tools will likely face. Deep learning uses neural networks and various mathematical and statistical techniques to determine a set of parameters that an algorithm can use to make predictions based on a given set of input attributes.²⁹ To determine that optimal set of parameters, deep learning uses the neural network to combine, abstract (and recombine and re-abstract), and otherwise transform the input attributes as they pass through multiple layers of the neural network.³⁰ This process is repeated thousands or millions of times, with the algorithm making slight adjustments to the parameters during each iteration.³¹ The process continues until the model finds an optimal set of parameters—that is, until the model reaches a point where further slight adjustments to the parameters will no longer improve the model’s accuracy on the training data.³² The resulting parameters are what the algorithm ultimately uses to make predictions.³³

Importantly for legal purposes, the optimized parameters cannot be expressed easily and reliably in terms of the original attributes that were used as inputs, particularly if the algorithm regularly receives new training data. The complexity of the calculations embedded in the deep learning process

25. Jason Brownlee, *Data, Learning and Modeling*, MACHINE LEARNING MASTERY (Jan. 6, 2017), <https://machinelearningmastery.com/data-learning-and-modeling/> [https://perma.cc/3VTW-XK3P].

26. *Id.*

27. *See id.*

28. Stephanie Glen, *Dimensionality & High Dimensional Data: Definition, Examples, Curse of, STATISTICS HOW TO* (Oct. 10, 2016), <https://www.statisticshowto.datasciencecentral.com/dimensionality/> [https://perma.cc/F363-VY4S].

29. *See A Beginner’s Guide*, *supra* note 18.

30. *See id.*

31. *See id.*

32. *See id.* The optima generally referred to here are local optima, rather than global or absolute optima. This feature has important implications for how courts should assess algorithmic selection tools under antidiscrimination laws. *See infra* Section IV.E.

33. *See A Beginner’s Guide*, *supra* note 18.

means that the algorithm generates the parameters that will not be readily interpretable, and the exact path through which the algorithm arrived at those parameters might not be practically traceable or capable of reconstruction. Consequently, even if the developer of an algorithm knows and understands all of the input variables (hardly a given in the age of Big Data) and also knows the target variables (or criteria) on which the algorithm optimizes, the algorithmic tool may nevertheless be effectively opaque even to the developer, much less the broader public. That is why deep learning algorithms are often referred to as “black box” algorithms.³⁴ Once the developer has specified the target (or criterion) by which success is judged, and selected the attributes that are potential predictors, the means by which the algorithm determines the parameters that result in the most accurate predictions is opaque.³⁵

2. *Algorithmic Employee Selection*

This Article is focused on algorithmic tools designed to make predictions about job candidates’ suitability for particular jobs. Today, building such a prediction system is generally best accomplished through supervised learning. The training data for a particular job will generally consist of historical examples of employees who have held the same job or a similar job, and possibly candidates who have applied for such jobs but who were not ultimately hired. In such a data set, the instances in the training data are individual employees or candidates, while the attributes consist of data on various characteristics of those employees or candidates. The labels for this training data would be information indicating each employee or candidate’s actual or projected performance in the job.

As an example, say that an employer wants to predict job candidates’ future job performance based on their educational attainment and experience. For training data, the employer has a data set consisting of 100 current employees’ educational attainment and years of experience at the time of hire, with each employee labeled with their most recent performance rating. In this example, the 100 employees are the instances for the training data, whereas educational attainment, performance rating, and years of service are attributes. If this data were used to build a standard statistical model (not necessarily one that uses machine learning),³⁶ performance rating would be termed the target

34. See Barber, *supra* note 11.

35. See generally *id.* (discussing the difficulty of determining how neural networks generate complex outcomes).

36. In fact, it would make little sense to use deep learning, rather than simple regression analysis, to create a predictive model on such a simple data set.

variable while educational attainment and years of service would be termed the predictor variables.³⁷

But in deep learning, and as stated above, the algorithm ultimately makes its predictions by using the final set of parameters that the trained algorithm generates rather than directly using the original input attributes.³⁸ Those original inputs are the raw materials for the resulting model, but the algorithm transforms them into something unrecognizable when it actually constructs the model.³⁹ Consequently, the final parameters, rather than the original attributes that the employer included in the training data are, in some sense, the true predictors. For that reason, we more accurately refer to the original attributes as input variables rather than predictor variables for the remainder of this Article.

In recruitment and hiring, available attributes most often include job-relevant characteristics such as certifications and prior employers—i.e., information that can be drawn from a candidate’s resume or application. If it is being developed by a third party, the training data may include employees from several different companies. In either case, employers may have the ability to access or acquire data from other sources on many more attributes—which may or may not be job related—such as a candidate’s social media profiles, criminal history, and web browsing history. Consequently, the data sets on which the models are trained may have a very high dimensionality and include inputs with no obvious connection to job performance. Some may contain thousands of candidates with thousands of attributes (or more). This makes algorithmic selection procedures considerably more complex than aptitude tests and other traditional employee selection tools.

B. Law of Discrimination

The seminal event in the history of employment discrimination law was the passage of the Civil Rights Act of 1964. Title VII⁴⁰ of that statute made it unlawful for employers to, among other things, “fail or refuse to hire or to discharge any individual, or otherwise to discriminate against any individual with respect to his compensation, terms, conditions, or privileges of

37. See ANDREW BRUCE & PETER BRUCE, PRACTICAL STATISTICS FOR DATA SCIENTISTS ch. 1 (2017) (ebook) (describing the process of linear regression).

38. See John Villasenor, *Artificial Intelligence and Bias: Four Key Challenges*, BROOKINGS: TECHTANK (Jan. 3, 2019), <https://www.brookings.edu/blog/techtank/2019/01/03/artificial-intelligence-and-bias-four-key-challenges/> [https://perma.cc/3HMX-ADU3].

39. See *id.*

40. Civil Rights Act of 1964, Pub. L. No. 88-352, tit. VII, 78 Stat. 241, 253-66 (codified as amended at 42 U.S.C. §§ 2000e-2000e-17 (2012)).

employment, because of such individual's race, color, religion, sex, or national origin."⁴¹ Various other federal statutes have been passed over the years creating additional protected categories, including age (under the Age Discrimination in Employment Act, or ADEA) and disability (under the Americans with Disabilities Act, or ADA), and many states have their own antidiscrimination laws covering different or additional protected categories.⁴²

But more than half of a century after its enactment, Title VII remains the most prominent antidiscrimination law and has the most fully developed legal framework for assessing employee selection procedures. Title VII is generally described as having two basic prohibitions, termed "disparate treatment" and "disparate impact."⁴³ The Supreme Court introduced the disparate impact doctrine in 1971, framing it, in essence, as a logical corollary to the general bar on discrimination "because of" a protected characteristic.⁴⁴ But in the ensuing decades, courts drew increasingly stark contrasts between the disparate treatment and disparate impact theories of discrimination,⁴⁵ culminating in a 2009 Supreme Court decision, *Ricci v. DeStefano*, where the high court described the inclusion of both theories in Title VII as a "statutory conflict."⁴⁶ Navigating this intersection will be a key challenge for employers seeking to implement algorithmic selection procedures.

41. 42 U.S.C. § 2000e-2(a)(1) (2012).

42. Age Discrimination in Employment Act of 1967, 29 U.S.C. § 623(a) (2012) ("It shall be unlawful for an employer . . . to fail or refuse to hire or to discharge any individual or otherwise discriminate against any individual with respect to his compensation, terms, conditions, or privileges of employment, because of such individual's age . . ."); Americans with Disabilities Act of 1990, 42 U.S.C. § 12112(a) (2012) (stating that it is unlawful for an employer to "discriminate against a qualified individual on the basis of disability in regard to job application procedures, the hiring, advancement, or discharge of employees, employee compensation, job training, and other terms, conditions, and privileges of employment."). *See generally Discrimination – Employment Laws*, NCSL (July 27, 2015), <http://www.ncsl.org/research/labor-and-employment/discrimination-employment.aspx> [<https://perma.cc/VD6E-W79D>] (listing employment discrimination laws from the states).

43. JOSEPH A. SEINER, *EMPLOYMENT DISCRIMINATION: PROCEDURE, PRINCIPLES, AND PRACTICE* 80, 170 (2015) [hereinafter *EMPLOYMENT DISCRIMINATION*].

44. *See* George Rutherglen, *Disparate Impact, Discrimination, and the Essentially Contested Concept of Equality*, 74 *FORDHAM L. REV.* 2313, 2314 (2006) (citing *Griggs v. Duke Power Co.*, 401 U.S. 424, 429–30 (1971)); 42 U.S.C. § 2000e-2(a)(1).

45. *See* Rutherglen, *supra* note 44 ("A strictly chronological account of these developments would reveal a very checkered history, with decisions to adopt or reject liability for disparate impact soon followed by qualifications and limitations.").

46. *Ricci v. DeStefano*, 557 U.S. 557, 580 (2009).

1. *Disparate Treatment*

Title VII's prohibition against disparate treatment derives from the original text of § 703(a), which prohibits employers from taking any adverse action against an employee or applicant “because of” a protected characteristic.⁴⁷ Another provision in § 703 reinforces this primary prohibition by stating the following:

Nothing contained in this subchapter shall be interpreted to require any employer . . . to grant preferential treatment to any individual or to any group . . . on account of an imbalance which may exist with respect to the total number or percentage of persons of any race, color, religion, sex, or national origin.⁴⁸

These two provisions lie at the core of what became known as disparate-treatment discrimination, although that precise terminology did not become common until the Supreme Court recognized the disparate impact theory of discrimination.⁴⁹

The vast majority of disparate treatment case law focuses on intentional acts of discrimination. Courts generally follow the *McDonnell Douglas* burden-shifting framework to demonstrate circumstantial evidence of discriminatory intent—a near necessity in light of the fact that in most discrimination cases, there is no direct evidence of discriminatory intent.⁵⁰ Rare is the case where there is a “smoking gun” demonstrating that the employer used race or some other characteristic as the explicit justification for an adverse employment action. The *McDonnell Douglas* framework allows plaintiffs to create an inference of intent without such direct evidence of discriminatory animus.⁵¹

But on its face, § 703(a) does not actually require an intent to discriminate; it bars all discrimination made because of a protected characteristic.⁵² The absence of an explicit intent requirement created—and continues to generate—ambiguity regarding Title VII's scope.⁵³ The most

47. EMPLOYMENT DISCRIMINATION, *supra* note 43 (quoting 42 U.S.C. § 2000e-2(a)(1)).

48. 42 U.S.C. § 2000e-2(j).

49. See Joseph A. Seiner, *Disentangling Disparate Impact and Disparate Treatment: Adapting the Canadian Approach*, 25 YALE L. & POL'Y REV. 95, 104–05 (2006) [hereinafter *Disentangling Disparate Impact*].

50. See *id.* at 81, 84–85 (citing *McDonnell Douglas Corp. v. Green*, 411 U.S. 792 (1973)).

51. See *McDonnell*, 411 U.S. at 802 (describing the prima facie case that a plaintiff is required to establish in order to show an inference of intent where direct evidence of discrimination is lacking).

52. 42 U.S.C. § 2000e-2(a) (2012).

53. *Disentangling Disparate Impact*, *supra* note 49, at 96.

important consequence of the broad language of § 703(a) was the creation of the disparate impact doctrine.

2. *Disparate Impact*

a. *The Development of the Disparate Impact Doctrine*

The standards governing disparate-impact discrimination are considerably more complex and ambiguous than those governing disparate treatment. The Supreme Court first established the disparate impact doctrine in *Griggs v. Duke Power Co.*, a class action by a group of black employees challenging their employer's requirement that new employees, in all but the lowest paying departments, have a high school diploma or pass a general intelligence test.⁵⁴ Both requirements operated to disproportionately exclude black workers—an outcome that likely was intended, given that many of the new requirements were imposed immediately after the passage of the Civil Rights Act of 1964.⁵⁵ The court of appeals concluded that the education and intelligence test requirements did not violate Title VII because they were facially neutral—that is, that they made no express distinction between employees on the basis of race—and because there was “no showing of a racial purpose or invidious intent.”⁵⁶

The Supreme Court reversed with an opinion that reshaped the legal landscape for employment discrimination law.⁵⁷ The Supreme Court began by rejecting the court of appeals' holding that the absence of intent to discriminate insulates a facially neutral employment condition under Title VII:

The objective of Congress in the enactment of Title VII is plain from the language of the statute. It was to achieve equality of employment opportunities and remove barriers that have operated in the past to favor an identifiable group of white employees over other employees. Under the Act, practices, procedures, or tests neutral on their face, and even neutral in terms of intent, cannot be maintained if they

54. *Griggs v. Duke Power Co.*, 401 U.S. 424, 424–28 (1971).

55. *Id.* at 426–28.

56. *Id.* at 428–29. Apparently, the court of appeals did not perceive the close proximity between the enactment of Title VII and the imposition of the new requirements to be evidence of racially discriminatory intent.

57. *Id.* at 436.

operate to “freeze” the status quo of prior discriminatory employment practices.⁵⁸

In ruling that the intelligence test and high school diploma requirements were unlawful, the *Griggs* court emphasized the systemic disadvantages that African Americans face as a result of receiving “inferior education in segregated schools.”⁵⁹ The Court explained the rationale for its new doctrine by analogy to one of Aesop’s fables:

Congress has now provided that tests or criteria for employment or promotion may not provide equality of opportunity merely in the sense of the fabled offer of milk to the stork and the fox. On the contrary, Congress has now required that the posture and condition of the job-seeker be taken into account. It has—to resort again to the fable—provided that the vessel in which the milk is proffered be one all seekers can use. The Act proscribes not only overt discrimination but also practices that are fair in form, but discriminatory in operation.⁶⁰

The *Griggs* decision also announced what would become known as the business necessity defense: “The touchstone is business necessity. If an employment practice which operates to exclude Negroes cannot be shown to be related to job performance, the practice is prohibited.”⁶¹ The Court then concluded that “neither the high school completion requirement nor the general intelligence test is shown to bear a demonstrable relationship to successful performance of the jobs for which it was used.”⁶²

Four years after *Griggs*, the Supreme Court laid out what remains the basic framework for disparate impact litigation in *Albemarle Paper Co. v. Moody*.⁶³ In that case, the Court introduced a three-step rubric for disparate-impact cases that roughly corresponds to the *McDonnell Douglas* burden-shifting framework that it had adopted two years earlier for disparate treatment claims.⁶⁴ First, the complaining party must “[make] out a prima facie case of discrimination, i.e. . . . show[] that the tests in question select

58. *Id.* at 429–30, 436.

59. *Id.* at 430 (citing *Gaston County v. United States*, 395 U.S. 285, 287 (1969)).

60. *Id.* at 430–31.

61. *Id.*; see Susan S. Grover, *The Business Necessity Defense in Disparate Impact Discrimination Cases*, 30 GA. L. REV. 387, 387 (1996).

62. *Griggs*, 401 U.S. at 430–31.

63. *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425–36 (1975).

64. *Id.* at 425 (citing *McDonnell Douglas Corp. v. Green*, 411 U.S. 792, 801–02, 804–05 (1973)).

applicants for hire promotion in a racial pattern significantly different from that of the pool of applicants.”⁶⁵ If a *prima facie* case is established, the employer then can rebut by showing that the tests are “job related.”⁶⁶ Finally, if the defendant establishes job relatedness, the plaintiff may still prevail by demonstrating “that other tests or selection devices, without a similarly undesirable racial effect, would also serve the employer’s legitimate interest in ‘efficient and trustworthy workmanship.’”⁶⁷ These three stages of a disparate impact case are explored further below.

b. Prima Facie Case

Albemarle Paper states that a plaintiff makes out a *prima facie* case of disparate impact by showing “that the tests in question select applicants for hire or promotion in a racial pattern significantly different from that of the pool of applicants.”⁶⁸ The Court did not indicate, however, whether “significantly different” was intended to be a reference to significance in a formal statistical sense, or if it instead meant significant in some more colloquial sense.⁶⁹ This ambiguity has led to divergent interpretations of the nature and magnitude of the disparity necessary to establish a *prima facie* case of disparate-impact discrimination.

The Uniform Guidelines on Employee Selection Procedures (Guidelines) adopted the “four-fifths” or “80%” rule, under which:

A selection rate for any race, sex, or ethnic group which is less than four-fifths (4/5) (or eighty percent) of the rate for the group with the highest rate will generally be regarded by the Federal enforcement agencies as evidence of adverse impact, while a greater than four-fifths rate will generally not be regarded by Federal enforcement agencies as evidence of adverse impact.⁷⁰

65. *Albemarle Paper*, 422 U.S. at 425.

66. *Id.*

67. *Id.*; *Albemarle Paper* made no mention of the fact that the *McDonnell* framework was initially adopted in a case involving an allegation of intentional discrimination in what would today be termed a disparate treatment case. This further underscores the Court’s initial conception of the disparate impact doctrine as merely a corollary to—and not different-in-kind from—disparate treatment discrimination. See *Griggs*, 401 U.S. at 431 (describing scope of what Congress proscribed in Title VII).

68. *Albemarle Paper*, 422 U.S. at 425.

69. *Id.*

70. 29 C.F.R. § 1607.4(D) (2019).

At first glance, this rule appears to focus exclusively on differences in selection rates and examines only the magnitude of the differences rather than their statistical significance. But the Guidelines hedge this rule to the point of meaninglessness, noting that smaller differences “may nevertheless constitute adverse impact, where they are significant in both statistical and practical terms,” and that greater differences may not constitute adverse impact “where the differences are based on small numbers and are not statistically significant”⁷¹ The Guidelines offer no guidance on how enforcement agencies or the courts should determine whether an adverse impact exists where the four-fifths rule and a statistical significance test point in opposite directions.⁷²

Courts have generally shunned the four-fifths rule as a test for *prima facie* disparate impact, preferring instead to rely on statistical significance tests. In *Hazelwood School District v. United States*, the Supreme Court indicated in a footnote that a difference of “more than two or three standard deviations” between the expected and actual number of protected class employees selected would make “the hypothesis that [employees] were hired without regard to race . . . suspect.”⁷³ In the forty years since *Hazelwood*, courts have more often looked to the social science standard of statistical significance at the 5% level (1.96 standard deviations) than to *Hazelwood*’s less precise “two or three standard deviation” standard.⁷⁴ But no particular statistical method or threshold has been established as the *sine qua non* of disparate impact analysis.

Indeed, many courts have been openly hesitant to rely on statistical significance alone when attempting to assess adverse impact. Just as the Guidelines suggest that their four-fifths rule may be disregarded if observed disparities “are significant in both statistical and practical terms,”⁷⁵ courts have occasionally sought to inject a requirement of “practical” or “legal” significance—usually, akin to the four-fifths rule, by looking to the raw magnitude of the disparity—in addition to statistical significance.⁷⁶ But courts

71. *Id.*

72. *See* § 1607.4.

73. *Id.* at 308 n.14 (quoting *Castaneda v. Partida*, 430 U.S. 482, 496–97 n.17 (1977)).

74. *See, e.g.,* *Smith v. Xerox Corp.*, 196 F.3d 358, 366 (2d Cir. 1999) (“Courts generally consider [significance at the 5%] level sufficient to warrant an inference of discrimination.”); *Palmer v. Shultz*, 815 F.2d 84, 96 (D.C. Cir. 1987) (“[S]tatistical evidence must meet the 5% level . . . for it alone to establish a *prima facie* case under Title VII.”).

75. 29 C.F.R. § 1607.4(D).

76. *See, e.g.,* *Jones v. City of Boston*, 752 F.3d 38, 49 (1st Cir. 2014) (characterizing the district court’s use of the four-fifths rule as an examination of practical significance); *Apsley v. Boeing Co.*, 691 F.3d 1184, 1199–1201 (10th Cir. 2012). The *Apsley* court affirmed the district court’s rejection of plaintiff’s *prima facie* disparate impact claim due to a lack of practical significance. *See Apsley*, 691 F.3d at 1199–1200. The court also noted that: [T]he

have reached no consensus on what practical significance entails or even whether it need be examined at all.

The Supreme Court arguably closed the door on practical significance requirements in *Ricci v. DeStefano*, where the Court stated in passing that a prima facie case of disparate impact requires showing a statistically significant disparity “and nothing more.”⁷⁷ Although this statement is arguably dicta, it nevertheless suggests that statistical significance tests remains the primary means by which courts determine whether a prima facie case of disparate impact discrimination exists. This bodes ill for employers seeking to leverage Big Data because, as discussed in greater detail below, large data sets can render even the slightest differences in selection rates statistically significant, even if they have minimal real-world importance.⁷⁸

c. Business Necessity Defense

Under the amendments to § 703 enacted in the Civil Rights Act of 1991, employers faced with a prima facie case of disparate impact discrimination must “demonstrate that the challenged practice is job related for the position in question and consistent with business necessity” to escape liability.⁷⁹ The concepts of “job relatedness” and “business necessity” first appeared in *Griggs*,⁸⁰ but in the five decades since, courts, agencies, and Congress alike have struggled with the meaning, relative importance, and interplay between the two concepts.⁸¹

Employees’ own figures show that the Companies recommended and hired over 99% of the older employees they would have been expected to recommend and hire in the absence of any discrimination. While this disparity might still lead a social scientist to suspect that the divestiture process was not wholly free of age-based discrimination, it would not permit a jury to find that such discrimination was the Companies’ standard operating procedure.

Id. at 1200–01 (first citing *Apsley v. Boeing Co.*, 722 F. Supp. 2d 1218, 1239 (D. Kan. 2010); then citing *Castaneda v. Partida*, 430 U.S. 482, 496 n.17 (1977); and then citing *International Brotherhood of Teamsters v. United States*, 431 U.S. 324, 336 (1977)).

77. *Ricci v. DeStefano*, 557 U.S. 557, 587 (2009) (citing *Connecticut v. Teal*, 457 U.S. 440, 446 (1982)).

78. See discussion *infra* Section III.B.1.

79. 42 U.S.C. § 2000e-2(k)(1)(A)(i) (2012).

80. *Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971).

81. The *Principles for the Validation and Use of Personnel Selection Procedures* provides an oblique definition of job relatedness by stating how job relatedness must be shown:

The job relatedness of a selection procedure has been demonstrated when evidence supports the accuracy of inferences made from scores on, or evaluations derived from, those procedures regarding some important aspect of work behavior (e.g., quality or

In *Albemarle Paper*, the Court described the employer's burden as solely that of demonstrating job relatedness, without reference to business necessity.⁸² But the Supreme Court appeared to return to the concept of business necessity two years later in *Dothard v. Rawlinson*, which stated that "a discriminatory employment practice must be shown to be necessary to safe and efficient job performance to survive a Title VII challenge" under a disparate impact theory.⁸³ But courts in the late 1970s and 1980s generally continued to follow the *Albemarle Paper* approach, seemingly disregarding *Griggs*'s description of business necessity as the touchstone of the analysis and instead focusing on job relatedness.⁸⁴ The Supreme Court then attempted to put the final nail in the business necessity coffin in *Wards Cove Packing Co. v. Atonio*.⁸⁵ There, the Court held that "the dispositive issue is whether a challenged practice serves, in a significant way, the legitimate employment goals of the employer" and that "there is no requirement that the challenged practice be 'essential' or 'indispensable' to the employer's business for it to pass muster."⁸⁶

Just two years later, however, Congress overrode the Supreme Court's *Wards Cove* decision in the Civil Rights Act of 1991, which enshrined both job related and business necessity in the text of Title VII.⁸⁷ The latter term appeared not by itself, but instead as part of the phrase "job related for the position in question and *consistent with* business necessity."⁸⁸ The statutory text, like the case law that inspired it, does not clarify how job relatedness differs (if at all) from business necessity, nor does it indicate how demonstrating that something is "consistent with business necessity" differs

quantity of job performance; performance in training, advancement, tenure, turnover, or other organizationally pertinent behavior).

SOC'Y FOR INDUS. & ORGANIZATIONAL PSYCHOLOGY, PRINCIPLES FOR THE VALIDATION AND USE OF PERSONNEL SELECTION PROCEDURES 4 (5th ed. 2018) [hereinafter PRINCIPLES]. This definition may be adequate for social scientific purposes, but, as described further in the discussion of validation below, a test must measure a representative set of job behaviors and outcomes to satisfy the business necessity defense; it does not suffice that a test measures merely *some* important aspect of work behavior.

82. *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425 (1975).

83. *Dothard v. Rawlinson*, 433 U.S. 321, 331 n.14 (1977).

84. See LEX K. LARSON, 2 LARSON ON EMPLOYMENT DISCRIMINATION § 23.04[1] (2017), LEXIS.

85. *Wards Cove Packing Co. v. Atonio*, 490 U.S. 642 (1989), *superseded by statute*, Civil Rights Act of 1991, Pub. L. No. 102-166, 105 Stat. 1074, *as recognized in* *Raytheon Co. v. Hernandez*, 540 U.S. 44 (2003).

86. *Wards Cove Packing Co.*, 490 U.S. at 659.

87. Civil Rights Act of 1991, Pub. L. No. 102-166, § 105, 105 Stat. 1074 (codified as amended at 42 U.S.C. §§ 2000e–2000e-17 (2012)).

88. *Id.* (emphasis added).

(if at all) from demonstrating that it is an actual business necessity.⁸⁹ Further, Congress expressly limited the legislative history that may be used to elucidate these distinctions.⁹⁰

The original source for the pairing of *job related* with the phrase *consistent with business necessity* appears to be the Department of Labor regulations for federal contractors under the Rehabilitation Act, a predecessor to the ADA that applied to federal employees and contractors.⁹¹ The relevant Rehabilitation Act regulations, which predated the Civil Rights Act of 1991 by more than a decade, stated that “to the extent qualification requirements tend to screen out qualified handicapped individuals,” the requirement must be “job-related . . . and . . . consistent with business necessity and the safe performance of the job.”⁹²

The same wording that now appears in Title VII also appears, almost verbatim, in the Americans with Disabilities Act, which Congress enacted a year before the 1991 amendments to Title VII.⁹³ Specifically, the ADA prohibits employers from:

[U]sing qualification standards, employment tests or other selection criteria that screen out or tend to screen out an individual with a disability or a class of individuals with disabilities unless the standard, test or other selection criteria, as used by the covered entity, is shown to be *job-related for the position in question and is consistent with business necessity*.⁹⁴

89. *See id.* (emphasis added).

90. *Id.* § 105(b) (“No statements other than the interpretive memorandum appearing at Vol. 137 Congressional Record S 15276 (daily ed. Oct. 25, 1991) shall be considered legislative history of, or relied upon in any way as legislative history in construing or applying, any provision of this Act that relates to *Wards Cove*—Business necessity/cumulation/alternative business practice.”).

91. 41 C.F.R. § 60-741.44(c)(1) (2019); Exec. Order No. 11,758, 39 Fed. Reg. 2,075 (Jan. 15, 1974); Americans with Disabilities Act of 1990, Pub. L. No. 101-336, 104 Stat. 327; Rehabilitation Act of 1973, Pub. L. No. 93-112, 87 Stat. 355.

92. § 60-741.44(c)(1); *see also* 29 C.F.R. § 32.14(a) (2019) (stating that qualifications for jobs with programs receiving federal financial assistance must be “related to the performance of the job and . . . consistent with business necessity and safe performance” if they tend to exclude individuals with disabilities).

93. *See* 42 U.S.C. § 2000e-2(k)(1)(A)(i) (2012); Americans with Disabilities Act of 1990 § 12101, 42 U.S.C. § 12112(b)(6) (2012).

94. § 12112(b)(6) (emphasis added).

According to case law,⁹⁵ ADA regulations,⁹⁶ and Equal Employment Opportunity Commission (EEOC) guidance,⁹⁷ this provision is closely linked to the ADA's central inquiry into whether an individual can perform the "essential functions" of a position.

Title VII makes no explicit reference to the essential functions of a job, and the ADA's linking of essential functions to the business necessity defense remains mostly foreign to Title VII jurisprudence. Nevertheless, the general concept—that job relatedness and business necessity require linking the selection criteria to specific, articulable, and important job functions—is one of the few common themes pervading the scattershot judicial and administrative interpretations of Title VII's business necessity defense.⁹⁸ The Guidelines emphasize careful job analysis, with a particular focus on identifying the "critical or important job duties, work behaviors or work outcomes."⁹⁹ And courts have generally refused to countenance challenged selection procedures where the employer fails to demonstrate a connection

95. See, e.g., *Bates v. United Parcel Serv., Inc.*, 511 F.3d 974, 996 (9th Cir. 2007) ("To show 'job-relatedness,' an employer must demonstrate that the qualification standard fairly and accurately measures the individual's actual ability to perform the essential functions of the job."); *EEOC v. Exxon Corp.*, 203 F.3d 871, 875 (5th Cir. 2000) ("[T]he business necessity defense . . . involves whether the individual can perform the 'essential functions' of the job . . .").

96. See, e.g., 29 C.F.R. § 1630.14(b)(3) (2019) ("[I]f certain criteria are used to screen out an employee or employees with disabilities as a result of such an examination or inquiry, the exclusionary criteria must be job-related and consistent with business necessity, and performance of the essential job functions cannot be accomplished with reasonable accommodation as required in this part."); 29 C.F.R. § Pt. 1630, app. (2019) ("As part of the showing that an exclusionary criteria is job-related and consistent with business necessity, the employer must also demonstrate that there is no reasonable accommodation that will enable the individual with a disability to perform the essential functions of the job.").

97. See, e.g., EEOC, *The Americans With Disabilities Act: Applying Performance And Conduct Standards To Employees With Disabilities*, <https://www.eeoc.gov/facts/performance-conduct.html> [<https://perma.cc/NBS4-YECZ>] ("If an applicant or employee cannot meet a specific qualification standard because of a disability, the ADA requires that the employer demonstrate the importance of the standard by showing that it is 'job-related and consistent with business necessity.' This requirement ensures that the qualification standard is a legitimate measure of an individual's ability to perform an essential function of the specific position the individual holds or desires."); EEOC, A TECHNICAL ASSISTANCE MANUAL ON THE EMPLOYMENT PROVISIONS OF THE ADA § 9.4 (1992) ("If a worker has an on-the-job injury which appears to affect his/her ability to do essential job functions, a medical examination or inquiry is job-related and consistent with business necessity.").

98. See *Bates*, 511 F.3d at 996 (describing the relationship between essential functions of the job and job relatedness); *Exxon Corp.*, 203 F.3d at 875 (stating that essential job functions are part of the criteria for the "business necessity" defense); 29 C.F.R. § 1630.14(b)(3).

99. § 1607.14(B)(2).

between the selection procedure and specific, key aspects of job performance,¹⁰⁰ a process known as validation.

Establishing the *validity* of a selection procedure thus is the central task of employers faced with a *prima facie* case of disparate impact discrimination. The Supreme Court stated in *Griggs* 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.¹⁰¹ *Albemarle Paper* refined this rule by casting doubt on the usefulness of generic or subjective measures of performance to validate selection criteria.¹⁰² The Court refused to accept an employer’s attempt to validate its test by showing that the results correlated with supervisory ratings, holding that those ratings were “extremely vague and fatally open to divergent interpretations”:

There is no way of knowing *precisely what criteria of job performance* the supervisors were considering, whether each of the supervisors was considering the same criteria or whether, indeed, any of the supervisors actually applied a focused and stable body of criteria of any kind. There is, in short, simply no way to determine whether the criteria actually considered were sufficiently related to the Company's legitimate interest in *job-specific ability* to justify a testing system with a racially discriminatory impact.¹⁰³

Albemarle Paper narrowed the permissible focus of employment tests in other ways as well, effectively requiring employers to use tests that measure essential aspects of job performance.¹⁰⁴ The Court held that employers cannot use selection procedures that hold applicants to a higher standard than

100. See, e.g., *Ernst v. City of Chicago*, 837 F.3d 788, 805 (7th Cir. 2016) (“Chicago failed to establish that its physical-skills entrance test reflects ‘important elements of job performance.’ And this lack of connection between real job skills and tested job skills is, in the end, fatal to Chicago’s case.”) (quoting *EEOC v. Dial Corp.*, 469 F.3d 735, 743 (8th Cir. 2006)). Federal courts have at times endorsed, sometimes explicitly, an “essential job functions” interpretation of Title VII’s business necessity defense. Most notably, in *Dothard v. Rawlinson*, 433 U.S. 321, 331–2 (1977), the Supreme Court rejected a prison system’s argument that its minimum height and weight requirements for corrections counselors were job related. The Court acknowledged that height and weight requirements may “have a relationship to strength, a sufficient but unspecified amount of which is essential to effective job performance as a correctional counselor.” *Id.* at 331. But the Supreme Court found this insufficient, noting that the employer “produced no evidence correlating the height and weight requirements with the requisite amount of strength thought *essential to good job performance*.” *Id.* (emphasis added).

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

102. *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 432–33 (1975).

103. *Id.* (emphasis added).

104. *Id.* at 433–34.

successful people currently in the position by imposing requirements those current workers could not satisfy.¹⁰⁵ It also endorsed the contemporary EEOC guidelines' rule that employers could not test for criteria relevant only to higher level jobs unless the "job progression structures and seniority provisions are so established that new employees will probably, within a reasonable period of time and in a great majority of cases, progress to a higher level."¹⁰⁶ The Court reasoned that the flaw in such an approach is that:

The fact that the best of those employees working near the top of a line of progression score well on a test does not necessarily mean that that test, or some particular cutoff score on the test, is a permissible measure of the *minimal qualifications of new workers* entering lower level jobs.¹⁰⁷

While never using the term essential functions, the Court implied that the criteria used to assess job candidates must be based on key aspects of job performance.¹⁰⁸

Determining what those key aspects of job performance are—and demonstrating that the selection process effectively measures them—is the crux of test validation. The current version of the *Standards for Educational and Psychological Testing (Standards)* defines validity as "the degree to which *evidence* and *theory* support the interpretations of test scores for proposed uses of tests."¹⁰⁹ In the context of employee selection procedures,

105. *Id.* at 429, 435–36 ("The record shows that a number of white incumbents in high-ranking job groups could not pass the tests.").

106. *Id.* at 434 (quoting 29 C.F.R. § 1607.4 (c)(1) (2019)). This rule was later incorporated almost verbatim into the Uniform Guidelines. 29 C.F.R. § 1607.5(I) (2019).

107. *Albemarle Paper*, 422 U.S. at 434 (emphasis added). The Third Circuit has made a "minimal qualifications" requirement the crux of its construction of the business necessity defense, holding that "employers may not use criteria which have a discriminatory effect unless those criteria define the minimum qualifications necessary to perform the job." *NAACP v. N. Hudson Reg'l Fire & Rescue*, 665 F.3d 464, 477 (3d Cir. 2011). The Uniform Guidelines did not go quite so far, allowing employers to rely upon tests that measure important components of job performance. 29 C.F.R. § 1607.5(B) (2019) ("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 . . .").

108. *Albemarle Paper*, 422 U.S. at 433.

109. AM. EDUC. RESEARCH ASS'N ET AL., *STANDARDS FOR EDUCATIONAL AND PSYCHOLOGICAL TESTING* 11 (4th ed. 2014) (emphasis added). "Theory" is a key element of this formulation because data-driven algorithms are largely, if not entirely, empirical. *Id.* As a

the *evidence* is the information indicating how well the selection procedure actually measures the fitness of candidates for that particular job. The *theory* is the chain of logic that links the selection procedure to the job requirements. For example, there is a logical relationship between the requirement that a programmer be conversant with a particular computer language and the ability of that candidate to efficiently write code in that language. But no logic or theory suggests that the car one drives ought to predict a candidate's ability to succeed as a coder. Consequently, identifying the critical and important aspects of job performance—as well as metrics that have an evidentiary and theoretical connection to those aspects of job performance—lies at the heart of the validation process under the Guidelines.¹¹⁰

The Guidelines discuss three different types of validity, presenting each as an independent path through which an employer can establish the job relatedness of a selection procedure: criterion-related validity, which is based on correlations between performance on the test and performance on the job and is by far the validation method most frequently used for employee selection procedures; content validity, which requires designing a test that adequately simulates job performance; and construct validity, which is based on measuring more abstract characteristics that are important for successful job performance.¹¹¹ Of these, only criterion-related validation represents a plausible path to establishing the validity of an algorithmic selection procedure. Content validity is a poor match for most algorithmic selection tools, which do not attempt to directly test an applicant's job-related knowledge or ability to perform specific tasks central to the job. The Guidelines assume that evidence for construct validity will come from criterion studies,¹¹² because the Guidelines also recognize criterion-related studies alone as a basis for establishing the validity of a test, it rarely is efficient or even useful for an employer to pursue construct validation (at least as presented in the Guidelines)¹¹³ rather than criterion validation.

But even criterion-related validation is an arduous process under the Guidelines.¹¹⁴ Moreover, the Guidelines were promulgated in the 1970s¹¹⁵ and reflect half-century-old conceptions both of the nature and format of

result, this raises the question of whether the APA would endorse predictive methods that lack a theoretical foundation.

110. See 29 C.F.R. §§ 1607.14 B(2), C(4) (2019).

111. § 1607.5(A)–(B).

112. § 1607.14(B)(2)–(3).

113. See discussion *infra* Section III.A.2. As discussed in greater detail in Section III.A.2, scientific concept of construct validity has evolved considerably in the decades since the Guidelines were issued.

114. See discussion *infra* Section III.A.1.

115. McDaniel et al., *supra* note 10, at 507.

employment tests and of what makes a test valid. As discussed further below,¹¹⁶ this makes it difficult to predict how courts and agencies will assess the validity of algorithmic selection procedures.

d. Least Discriminatory Alternative

If an employer meets its burden in establishing the job relatedness of the selection procedure, the final stage of disparate impact analysis requires a plaintiff to demonstrate that a less discriminatory alternative was available that would meet the employer's business needs.¹¹⁷ This test traces its roots to *Albemarle Paper*, which stated that a plaintiff could prevail on a disparate impact claim by demonstrating "that other tests or selection devices, without a similar undesirable racial effect, would also serve the employer's legitimate interest in 'efficient and trustworthy workmanship.'" ¹¹⁸

A key question that remains largely unresolved is how effective the plaintiff's proposed alternative must be to defeat an employer's showing of business necessity. *Albemarle Paper*'s standard—that the procedure need only "serve the employer's legitimate interest in 'efficient and trustworthy workmanship'"—appeared to set the bar rather low, suggesting that the proposed alternative need not be *exactly* as effective as the challenged procedure, so long as it is adequate to meet the employer's needs.¹¹⁹ In *Wards Cove*, the Supreme Court attempted to reject this low bar, holding that an alternative practice "must be equally effective as [the employer's] chosen hiring procedures in achieving [the employer's] legitimate employment goals."¹²⁰ But as with *Wards Cove*'s alteration of the business necessity defense, Congress overrode the Court through the 1991 amendments to Title VII.¹²¹ In the Civil Rights Act of 1991, Congress explicitly restored the law governing alternative employment practices to "the law as it existed on June 4, 1984," the day before *Wards Cove* was decided.¹²²

Unfortunately, the exact nature of the "less discriminatory alternative" standard was far less than clear even before *Wards Cove*.¹²³ The only type of modification to a selection procedure that seems to have gained wide recognition as an adequate alternative is the practice of "banding" test scores,

116. See discussion *infra* Section III.A.1.

117. *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425 (1975) (quoting *McDonnell Douglas Corp. v. Green*, 411 U.S. 792, 801 (1973)).

118. See *id.* (quoting *McDonnell*, 411 U.S. 792, 801 (1973)).

119. *Id.*

120. *Wards Cove Packing Co. v. Atonio*, 490 U.S. 642, 661 (1989).

121. See *id.*; Civil Rights Act of 1991, Pub. L. No. 102-166, 105 Stat. 1074.

122. *Id.*; 42 U.S.C. § 2000e-2(k)(1)(C) (2012).

123. See *Disentangling Disparate Impact*, *supra* note 49, at 103–04.

where candidates are grouped together in bands based on differences between scores that are considered insignificant.¹²⁴ Because of the paucity of cases clarifying the standards by which alternative selection procedures should be judged, courts have generally been reluctant to decide cases on the basis of a plaintiff's showing of a less discriminatory alternative.¹²⁵

3. Ricci v. DeStefano and the Interplay Between Disparate Treatment and Disparate Impact

While *Griggs* cast the disparate impact theory as simply a logical corollary of the disparate treatment that Title VII clearly prohibited, these legal theories in fact spring from quite separate views on the thrust and purpose of antidiscrimination laws. Disparate treatment, as presently interpreted, reflects an anticlassification view of discrimination, which holds that the purpose of antidiscrimination laws is to prohibit classifying or differentiating between individuals on the basis of a protected characteristic.¹²⁶ Disparate impact, by contrast, reflects an antisubordination perspective on discrimination, under which the purpose of such laws is to “prohibit practices that enforce the social status of oppressed groups and allow practices that challenge oppression.”¹²⁷ The antisubordination roots of disparate impact theory can be seen in *Griggs*, where the Supreme Court emphasized the long-running and systemic disadvantages that blacks had endured, and rejected the notion that an employment test complies with Title VII so long as it is “fair in form.”¹²⁸

The conceptual tension between disparate treatment and disparate impact causes practical problems for employers who observe that their policies are having disparate impacts (or anticipate that they will have a disparate impact in the future) and perceive that the most logical way to stop such adverse impacts from arising is to take direct steps to correct for the disparate impact. But the very act of correcting disparate impacts may itself be a form of

124. See, e.g., *Officers for Justice v. Civil Serv. Comm'n*, 979 F.2d 721, 723–24, 728 (9th Cir. 1992) (“Today we hold that the banding process is valid as a matter of constitutional and federal law.”); *Chi. Firefighters Local 2 v. City of Chicago*, 249 F.3d 649, 656 (7th Cir. 2001) (“[Banding is] a universal and normally an unquestioned method of simplifying scoring by eliminating meaningless gradations.”).

125. LARSON, *supra* note 84, at § 24.02.

126. See Jack M. Balkin & Reva B. Siegel, *The American Civil Rights Tradition: Anticlassification or Antisubordination?*, 58 U. MIAMI L. REV. 9, 10 (2003).

127. Bradley A. Areheart, *The Anticlassification Turn in Employment Discrimination Law*, 63 ALA. L. REV. 955, 961 (2012).

128. *Griggs v. Duke Power Co.*, 401 U.S. 424, 430–31 (1971).

disparate treatment.¹²⁹ That dilemma made its way to the Supreme Court in the 2009 case *Ricci v. DeStefano*.¹³⁰

The plaintiffs in *Ricci* were white and Hispanic firefighters who had taken and passed an examination administered by the City of New Haven that determined the firefighters' eligibility for promotion to lieutenant or captain.¹³¹ The City worked with an outside consulting firm to develop the test over a period of several years.¹³² But the City's first real-life administration of the test showed that using the results of the exam would have an adverse impact on black and Hispanic firefighters; thirty-four of the seventy-seven firefighters who took the examination were black or Hispanic, but all ten of the candidates who scored high enough to be considered for promotion were white.¹³³ Based on these disproportionate outcomes, the City believed that using the results of the test would have an unlawful disparate impact and subject them to liability under Title VII.¹³⁴ Consequently, the City chose not to certify the examination results.¹³⁵

The firefighters who passed the test challenged the City's decision as expressly race based, and sought review by the Supreme Court.¹³⁶ The Court ruled for the firefighters and held that the City's decision, because it was driven by concern over the adverse impact on minority firefighters, was a decision made because of race in violation of Title VII's disparate treatment prohibition:

All the evidence demonstrates that the City chose not to certify the examination results because of the statistical disparity based on race—*i.e.*, how minority candidates had performed when compared to white candidates. As the District Court put it, the City rejected the test results because “too many whites and not enough minorities would be promoted were the lists to be certified.” Without some other justification, this express, race-based decision-making violates Title VII's command that employers cannot take adverse employment actions because of an individual's race.¹³⁷

129. See *Ricci v. DeStefano*, 557 U.S. 557, 593 (2009).

130. *Id.*

131. *Id.* at 562–63.

132. *Id.* at 564.

133. *Id.* at 566.

134. *Id.* at 563, 566.

135. *Id.* at 574.

136. *Id.* at 563, 574–75.

137. *Id.* at 579 (citation omitted).

Notably, the Supreme Court’s reasoning did not focus on racial animus or an intent to discriminate in the usual sense.¹³⁸ On the contrary, the Court acknowledged that the employer’s objective had been avoiding disparate-impact liability—in other words, to avoid committing unlawful discrimination.¹³⁹ But that objective did not insulate the employer from liability because it ignored “the City’s conduct in the name of reaching that objective.”¹⁴⁰ The Court reasoned:

Whatever the City’s ultimate aim—however well-intentioned or benevolent it might have seemed—the City made its employment decision because of race. The City rejected the test results solely because the higher scoring candidates were white. The question is not whether that conduct was discriminatory but whether the City had a lawful justification for its race-based action.¹⁴¹

The Court rejected the City’s argument that its violation of the disparate treatment prohibition should be excused because the City only did so to avoid the prospect of disparate impact liability.¹⁴² But in doing so, the Court explicitly left open the possibility that an employer, although not the employer in *Ricci* itself,¹⁴³ would be able to use the prospect of disparate impact liability as a defense to a disparate treatment claim.¹⁴⁴ The Court rejected the plaintiff firefighters’ blanket argument that “avoiding unintentional discrimination

138. *Id.* at 592.

139. *Id.* at 579.

140. *Id.*

141. *Id.* at 579–80.

142. *Id.* at 563.

143. As the dissent noted, despite its adoption of the “strong basis in evidence” standard discussed below, the *Ricci* majority did not remand the matter for further proceedings so that evidence could be presented on the strength of a potential business necessity defense and the availability (or not) of less discriminatory alternative selection procedures. *Id.* at 563, 631 (Ginsburg, J., dissenting). Instead, it ruled as a matter of law that plaintiffs were entitled to summary judgment. *Id.* at 592. The majority’s reasoning implies that it concluded either that the test was actually valid (which appeared to be a disputed factual question) or presumptively valid (which runs contrary to the Guidelines and the statute, which places the burden on the employer to establish validity). Given the absence in the record of criterion-related validity evidence, the basis of the Court’s validity finding is unclear. The Court seemed to imply that the rigorous job analysis that the City had performed, coupled with its efforts to craft a test based on that evidence, established the test’s validity. That line of reasoning most closely tracks a content validity argument, but a paper-and-pencil multiple-choice test would not be a direct test of job performance for a firefighter, as the dissenters in *Ricci* pointed out. See *id.* at 634 (Ginsburg, J., dissenting).

144. See *id.* at 593. The narrowness of the holding was to the apparent chagrin of Justice Scalia, who wrote a brief concurrence chiding the Court for declining to confront the constitutionality of the disparate impact doctrine directly. *Id.* at 594 (Scalia, J., concurring).

cannot justify intentional discrimination.”¹⁴⁵ Going a step further, it also declined to adopt a standard under which “an employer in fact must be in violation of the disparate-impact provision before it can use compliance as a defense in a disparate-treatment suit.”¹⁴⁶

Instead, borrowing from the Court’s constitutional Equal Protection Clause jurisprudence, the Court held that an employer could escape disparate treatment liability if it “can demonstrate a *strong basis in evidence* that, had it not taken the action, it would have been liable under the disparate impact statute.”¹⁴⁷ The City failed in this regard because it did not adequately consider evidence of the validity and job relatedness of the test—and job relatedness is a complete defense to a disparate impact claim.¹⁴⁸ After concluding—dubiously, given the case’s posture as an appeal from a summary judgment motion—that the City had failed to make such a “strong basis in evidence” showing, it ruled that the plaintiffs were entitled to summary judgment and, in effect, ordered that the City certify the examination results.¹⁴⁹

At first blush, *Ricci* seems a very ominous portent for employers considering whether and how to implement novel selection procedures—and it certainly is for employers who discover an adverse impact only after a selection procedure has been designed and administered. Such employers face a catch-22, where attempting to mitigate the disparate impact could subject them to disparate treatment liability, while inaction would leave them vulnerable to a disparate impact claim.¹⁵⁰ But the Court appeared to leave open an avenue through which employers could mitigate anticipated disparate impacts without necessarily violating Title VII.¹⁵¹

Specifically, the Court held that “Title VII does *not* prohibit an employer from considering, *before* administering a test or practice, how to design that test or practice in order to provide a fair opportunity for all individuals, regardless of their race.”¹⁵² Explaining the dissonance between that principle and the Court’s disposition of the firefighters’ examination results, the Court stated:

145. *Id.* at 580.

146. *Id.* at 580–81.

147. *Id.* at 563 (emphasis added).

148. *See id.* at 578.

149. *See id.* at 593 (holding that petitioners are entitled to summary judgment and remanding for further proceedings consistent with the opinion).

150. *See id.* at 629 (Alito, J., concurring).

151. *See id.* at 585.

152. *Id.* at 585 (emphasis added).

[W]e [do not] question an employer's affirmative efforts to ensure that all groups have a fair opportunity to apply for promotions and to participate in the process by which promotions will be made. But once that process has been established and employers have made clear their selection criteria, they may not then invalidate the test results, thus upsetting an employee's legitimate expectation not to be judged on the basis of race.¹⁵³

The City's actions, according to the Court, fell into the latter category.¹⁵⁴ The Court emphasized the "high, and justified, expectations of the candidates who had participated in the testing process on the terms the City had established for the promotional process," many of whom "had studied for months, at considerable personal and financial expense."¹⁵⁵ The unfairness of the City's decision lay not in its desire to avoid using a test that would have a disparate impact, but in the fact that the City only decided to discard the results after the examination design process was complete and the promotion candidates developed something akin to a reliance interest in having the examination used as a basis for promotion decisions.¹⁵⁶

The Court's reasoning seems consistent with the text of the most on-point provision in Title VII, § 703(I).¹⁵⁷ That provision makes it unlawful for employers to "adjust the scores of, use different cutoff scores for, or otherwise alter the results of, employment related tests" on the basis of protected class status.¹⁵⁸ Technically, designing a selection procedure to avoid disparate impacts would not be adjusting test scores or using different cutoffs because the scoring rubric for a selection procedure is not yet finalized during the test design stage.¹⁵⁹

That said, there is no case law squarely addressing the issue of how much license employers have to protect against disparate impacts by designing a selection procedure in a manner that explicitly takes protected class status into account. Is it permissible for employers to choose a suboptimal selection device, as measured by its accuracy, because it results in a more diverse workforce? It is safe to assume that there are limits—not least from § 703(a)'s general prohibition against making employment decisions because of protected class status—on the degree to which employers can be race or

153. *Id.* at 585.

154. *Id.* at 593.

155. *Id.*

156. *See id.*

157. *See* 42 U.S.C. § 2000e-2(I) (2012).

158. *Id.*

159. *See id.*

gender conscious when designing a selection procedure.¹⁶⁰ Using quotas or granting bonus points on the basis of protected class status, for instance, surely would not survive a disparate treatment challenge, even if an employer adds those features as part of initial test design.¹⁶¹ But it is not clear where courts will draw lines between permissible and impermissible methods of designing around disparate impacts.

This ambiguity is a source of concern for employees considering algorithmic selection procedures.¹⁶² Algorithms offer the potential for employers to design a selection procedure that reduce or eliminate disparate impacts using methods that are far more sophisticated and subtle than the blunt instruments available for traditional tests.¹⁶³ The degree to which those methods are deemed consistent with Title VII will likely determine how quickly employers adopt algorithmic selection procedures in the coming years.

III. THE UNIQUE CHALLENGES FOR ALGORITHMIC SELECTION TOOLS

Designing algorithmic selection tools that leverage the ability to generate unique data-driven insights while maintaining legal compliance will prove challenging under current law. A comprehensive treatment of all the practical and legal ambiguities surrounding algorithmic selection tools would be prohibitively lengthy, but sections A through D of Part III identify four overarching categories that encompass the most vexing legal compliance issues for algorithmic tools: challenges relating to the validation process; those stemming from algorithmic tools' reliance on correlation and use of Big Data; those relating to the opacity of models generated by deep neural networks; and those stemming from the bare fact that the deployment of algorithmic tools will provide plaintiffs' lawyers with a clear target for bringing discrimination claims.

A final issue that starkly illustrates the “square peg in a round hole” dynamic of algorithmic selection tools and current employment discrimination law is whether Title VII's disparate treatment doctrine can even be applied to machines that do not possess conscious intentions—or, indeed, consciousness at all. As explained in the final section of this part, despite the intent-focused tilt of case law on disparate treatment, the broad language of the statutory text and the equally broad early Supreme Court

160. See generally 42 U.S.C. § 2000e-2(a) (2012) (discussing unlawful employment practices for employers based on an “individual's race, color, religion, sex, or national origin”).

161. See NAT'L ARCHIVES, *EEO Terminology*, <https://www.archives.gov/eo/terminology.html> [<https://perma.cc/G4DQ-84JA>].

162. See ZHANG ET AL., *supra* note 2, at 8.

163. See *id.* at 10.

decisions interpreting it mean that employers are unlikely to escape disparate treatment liability if they deploy algorithms that make facially discriminatory classifications.

A. *Validation*

1. *The Difficult Path to Criterion-Related Validity Under the Guidelines*

The Guidelines establish rigorous standards for criterion-related validation studies.¹⁶⁴ These standards correctly ensure that a selection procedure has a demonstrable relationship to each job for which it is used, but the expense and complexity of establishing and maintaining the criterion-related validity of an algorithmic tool will blunt the efficiency gains that algorithmic tools promise.¹⁶⁵

a. *Job Analysis*

A criterion-related validation study begins with a careful job analysis conducted by industrial psychologists or other trained professionals to identify the critical and important elements of job performance.¹⁶⁶ Courts emphasize the thoroughness and attention to detail that a job analysis entails and often reject validation studies that are not supported by adequate job analyses.¹⁶⁷ One court described a job analysis: “A thorough survey of the relative importance of the various skills involved in the job in question and the degree of competency required in regard to each skill. It is conducted by interviewing workers, supervisors and administrators; consulting training manuals; and closely observing the actual performance of the job.”¹⁶⁸

164. See generally 29 C.F.R. § 1607.14 (2019) (“The following minimum standards, as applicable, should be met in conducting a validity study.”).

165. See 29 C.F.R. § 1607.14(B)(3) (2019).

166. See 29 C.F.R. § 1607.5(B)(2) (2019); § 1607.14(B)(2).

167. See, e.g., *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 432 (1975) (“The study in this case involved no analysis of the attributes of, or the particular skills needed in, the studied job groups. There is accordingly no basis for concluding that ‘no significant differences’ exist among the lines of progression, or among distinct job groupings within the studied lines of progression.”); *Rogers v. Int’l Paper Co.*, 510 F.2d 1340, 1351 (1975) (“The . . . absence of proper and careful job analyses . . . is fatal to the validation study.”).

168. *Guardians Ass’n of N.Y.C. Police Dep’t v. Civil Serv. Comm’n of N.Y.*, 633 F.2d 232, 242 (2d Cir. 1980), *aff’d sub nom*, *Guardians Ass’n of N.Y.C. Police Dept. v. Civil Serv. Comm’n of New York*, 463 U.S. 582 (1983) (quoting *Vulcan Soc’y of N.Y.C. Fire Dep’t v. Civil Serv. Comm’n of New York, Inc.*, 360 F. Supp. 1265, 1274 (S.D.N.Y. 1972)).

From these observations and information, the experts conducting the study then “break[] down an observed task into a set of component skills, abilities and knowledge,” and indicate what level of competence or proficiency is required for each component.¹⁶⁹ According to the American Psychological Association’s *Principles for the Validation and Use of Personnel Selection Procedures*, which apply the more generally applicable *Standards* to the employee selection setting, a proper job analysis “may include different dimensions or characteristics of work, including work complexity, environment, context, tasks, behaviors and activities performed, and worker requirements (e.g., KSAOs [Knowledge, Skills, Abilities, and Other Characteristics] or competencies).”¹⁷⁰

b. Selecting Criteria

From the critical and important job duties, work behaviors, and work outcomes identified during the job analysis, an employer must then select or develop measurable criteria that serve as metrics of how well an individual can perform the key functions of a job.¹⁷¹ Employees’ real-world performance with respect to those job related criteria then serve as the benchmarks for validation¹⁷²—and, in the case of algorithmic tools, as target variables for building a model.

Needless to say, criterion selection is crucial to a proper criterion-related validity study. “Criteria should be chosen on the basis of work relevance, freedom from contamination, and reliability rather than availability or convenience,”¹⁷³ and “should represent important or critical work behavior(s) or work outcomes,” as identified in the job analysis.¹⁷⁴ Where courts have refused to recognize proffered criterion validity studies, it has not usually been because the employer failed to show the proper correlation between the selection procedure and the criteria, but because the employer failed to select proper criteria in the first place.¹⁷⁵

169. See *Jones v. N.Y.C. Human Res. Admin.* 391 F. Supp 1064, 1080 (S.D.N.Y. 1975).

170. *PRINCIPLES*, *supra* note 81, at 1, 7.

171. See *id.* at 7.

172. See *id.* at 10.

173. *Id.* at 11.

174. 29 C.F.R. § 1607.14(B)(3) (2019).

175. *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 432–34 (1975) (stating that Albemarle Paper failed to show that its testing was job related when the criteria were subjective); see also, e.g., *Green v. U.S. Steel Corp.*, 570 F. Supp. 254, 273–77 (E.D. Pa. 1983) (stating that the hiring standard U.S. Steel adopted was a series of subjective criteria and found that because the employer failed to rebut plaintiffs’ prima facie case of disparate impact, the plaintiffs were entitled to judgment in their favor).

As a threshold matter, the criteria must be direct measures of job performance, and not separate on-the-job tests or assessments that have not themselves been validated.¹⁷⁶ Courts generally expect criteria to be specific and reasonably objective markers of job performance and frown on criteria that are vague, generic, or subjective. In *Albemarle Paper*, the Supreme Court found a purported criterion validity study inadequate in large part because the criteria consisted of subjective supervisory employee rankings that were made according to “a ‘standard’ that was extremely vague and fatally open to divergent interpretations.”¹⁷⁷ Consequently, supervisory ratings and assessments—which are often the only available measures of an employee’s on-the-job performance—may not be adequate to support criterion-related validity.¹⁷⁸

But it is difficult, and often impossible, to capture all essential and important job behaviors and job outcomes using readily available data. More general signals of employee performance such as statistics on hiring, retention, and tenure are generally available to employers. Formal performance reviews in some form may also be available, but if these include narrative sections or are not subject to a uniform rubric that ensures the reviews have consistent meaning, the reliability of the reviews (and the ability of an algorithm to make sense of them) as target variables might be limited.

Some jobs may have reasonably reliable performance metrics that seem to capture the essence of the job. But a closer examination often reveals that available metrics do not adequately measure job performance.¹⁷⁹ For example, a district attorney’s office may track the number of cases that its prosecutors try and the percentage of cases they win. These statistics, which can be tracked reliably at little or no cost, may make attractive target variables. But a prosecutor’s win-loss record may be a poor indicator of the quality of their lawyering. The best prosecutors might be the ones who take on the most difficult and time-intensive cases, and thus try fewer cases and have a lower

176. *Ernst v. City of Chicago*, 837 F.3d 788, 802 (7th Cir. 2016) (“Chicago created a skills test and a work-sample test, found a strong correlation between the skills test and the work-sample test, and thus concluded that the skills test is a good measure of job-related skills. As the plaintiffs argue, this is a statistical form of self-affirmation. There is no evidence that the work-sample test, which Chicago used to validate the skills test, is a proper validation of job skills.”).

177. *Albemarle Paper*, 422 U.S. at 432–33; see also, e.g., *Green*, 570 F. Supp. at 275–76 (discussing subjective hiring criteria). In *Green*, the court struck down a company’s “‘best qualified’ hiring standard,” which “consist[ed] of about twenty subjective criteria applied as an ‘amalgam,’” and where evidence showed that “each individual decision-maker essentially simply consulted his or her ‘gut level’ reaction to an individual applicant.” *Green*, 570 F. Supp. at 275–76.

178. See *Albemarle Paper*, 422 U.S. at 432–33.

179. Dick Grote, *The Myth of Performance Metrics*, HARV. BUS. REV. (Sept. 12, 2011), <https://hbr.org/2011/09/the-myth-of-performance-metric> [<https://perma.cc/RZ7R-HW7W>].

rate of positive outcomes than less skilled lawyers who shy away from such cases. But if the district attorney lacks the resources to closely observe the work of most prosecutors, the flawed trial statistics may be the only metrics available.¹⁸⁰

Similarly, employers are often tempted to search for readily observable characteristics that can serve as proxies for attributes essential to the job in question. But this too carries risk. In *Dothard v. Rawlinson*, an employer attempted to justify its minimum requirements for height and weight—metrics that were readily available—on the claimed basis that those requirements were meant to ensure that corrections counselors had the requisite physical strength, which was the job-relevant attribute of interest.¹⁸¹ The Court rejected this argument, reasoning that “[i]f the job-related quality that the appellants identify is bona fide, their purpose could be achieved by adopting and validating a test for applicants that measures strength directly.”¹⁸²

Having a representative set of participants is another key requirement for criterion-related validation.¹⁸³ The subjects must broadly reflect of the characteristics of the pool of actual applicants.¹⁸⁴ Thus, the sample must consist of entry-level employees if it is for an entry-level job; using employees from higher in the line of progression is not sufficient.¹⁸⁵ Representativeness across protected classifications is also required; the Guidelines state that the sample “should insofar as feasible include the races, sexes, and ethnic groups normally available in the relevant job market.”¹⁸⁶ Ultimately, an employer establishes criterion validity under the Guidelines by demonstrating that performance on the selection procedure correlates with a representative set of performance measures tied to the job criteria identified during the job analysis.¹⁸⁷

As the above discussion suggests, a proper criterion-related validity study is a major undertaking even for large and sophisticated employers. This may explain, in part, why most employers have shied away from using

180. The tendency to turn to easily available metrics as a substitute for deeper analysis is hardly unique to the hiring process. Before the recent explosion in sports analytics, Michael Lewis observed that “[f]or most of its history basketball has measured not so much what is important as what is easy to measure—points, rebounds, assists, steals, blocked shots—and these measurements have warped perceptions of the game.” Michael Lewis, *The No-Stats All-Star*, N.Y. TIMES MAG., Feb. 13, 2009, <https://www.nytimes.com/2009/02/15/magazine/15Battier-t.html> [<https://perma.cc/9LTB-ZRS7>].

181. *Dothard v. Rawlinson*, 433 U.S. 321, 331 (1977).

182. *Id.* at 332.

183. 29 C.F.R. § 1607.14(B)(1) (2019).

184. *Id.*

185. *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 434 (1975).

186. § 1607.14(B)(4).

187. § 1607.14(B)(2)–(3).

employment tests altogether;¹⁸⁸ relying on human judgment, however flawed, generates neither the cost nor the discoverable paper trail that validation entails.

Because many employers wishing to deploy an algorithmic selection procedure will not have ready access to a properly developed set of criteria that can serve as the basis for a criterion-related validity study, the process of developing and validating an algorithmic tool may take several years. That timetable that may prove problematic given the vintage of the Guidelines and the likelihood that courts and agencies will introduce new standards for validation in the coming years.

2. *The Guidelines: Behind the Times*

The long and difficult process of criterion-related validation under the Guidelines will be challenging enough for employers testing new hiring tools. But the Guidelines' forty-year-old standards are overdue for revamping or replacement to bring them in line with the modern social science of test validity, which has evolved considerably in the decades since the Guidelines first appeared. This adds an additional layer of legal uncertainty.

The EEOC and four other federal agencies and departments¹⁸⁹ jointly adopted the Guidelines in 1978. Recognizing that theories of test validity were still evolving, the Guidelines state that "[n]ew strategies for showing the validity of selection procedures will be evaluated as they become accepted by the psychological profession."¹⁹⁰ But the Guidelines' validation standards have, in fact, remained unchanged in the four decades since their promulgation. In the interim, the American Psychological Association (APA) has issued revised versions of the *Standards* three times (in 1985, 1999, and 2014). Starting with the 1985 edition, the *Standards* moved away from the Guidelines' trichotomous separation of test validity into content, criterion, and construct validity.¹⁹¹ Consequently, even before the advent of Big Data and the prospect of completely new types of employee selection procedures, many of the Guidelines' provisions and much of their terminology seemed dated.

Comparing the descriptions of construct validity in the Guidelines with those in modern scientific literature provides a stark example of how much the social science of test validation has evolved since the Guidelines were

188. See *infra* note 238 and accompanying text.

189. These federal agencies and departments include the Office of Personnel Management, Department of Justice, Treasury Department, and the Department of Labor's Office of Federal Contract Compliance Programs. See 29 C.F.R. § 1607.2(A) (2019).

190. § 1607.5(A).

191. See Samuel Messick, *Validity*, in EDUCATIONAL MEASUREMENT 13, 18–20 (Robert L. Linn ed., 3d ed. 1989).

issued. The Guidelines refer to construct validity as “a relatively new and developing procedure in the employment field,” for which there was, as of 1978, “a lack of substantial literature extending the concept to employment practices.”¹⁹² But the literature surrounding construct validity developed rapidly in the 1980s and 90s; today, far from an undeveloped and novel theory, construct validity is generally recognized as *the* overarching validity concept.¹⁹³ Where the Guidelines present construct and content validity as separate types of validity, modern social science treats test content and criterion relatedness as categories of evidence for demonstrating the broader concept of test validity.¹⁹⁴

Modern test literature treats test bias and fairness as potential threats to validity, and the vocabulary surrounding what constitutes test bias relies—sometimes explicitly—on the concept of the constructs that represent whatever the selection procedure is ultimately attempting to measure.¹⁹⁵ One specific threat to validity extensively studied by modern social scientists—construct-irrelevant variance—will take on particular importance in the age of Big Data and with the rise of algorithmic selection procedures, as discussed further below.¹⁹⁶ But the Guidelines and the existing case law on validation are bereft of meaningful discussion of these threats to validity, leaving employers to guess if, when, and how courts and agencies will take them into account.

Many courts continue to cite the Guidelines when discussing proper validation methods, and the EEOC’s Fact Sheet on Employment Tests and Selection Procedures still references the Guidelines as the primary source of regulatory guidance on validation of selection procedures.¹⁹⁷ Employers seeking to implement algorithmic selection procedures thus have little choice but to pursue validation that complies with the Guidelines. But the stringent requirements for criterion validation under the Guidelines can take many years to complete. The law may well change in the interim, which makes reliance on the Guidelines’ validity standards an inherently unstable proposition as long as they lag decades behind the prevailing social science.

192. § 1607.14(D)(1).

193. See, e.g., WILLIAM M. TROCHIM ET AL., RESEARCH METHODS 128–30 (2d ed. 2015); AM. EDUC. RESEARCH ASS’N ET AL., *supra* note 109, at 11 (“The term *construct* is used in the *Standards* to refer to the concept or characteristic that a test is designed to measure.”). But see Jerry A. Colliver et al., *From Test Validity to Construct Validity... and Back?*, 46 MED. EDUC. 366, 367–70 (2012) (criticizing the increasingly broad use of construct validity despite its rising popularity in social science).

194. AM. EDUC. RESEARCH ASS’N ET AL., *supra* note 109, at 14–19.

195. *Id.* at 5 (“Fairness and accessibility, the unobstructed opportunity for all examinees to demonstrate their standing on the construct(s) being measured, are relevant for valid score interpretations for all individuals and subgroups in the intended population of test takers.”).

196. See discussion *infra* Section III.B.2.

197. § 1607.14(B)(3).

B. The Pitfalls of Correlations and Big Data

The sheer size of data sets in the era of Big Data deepens the challenges that employers, agencies, and courts will face when attempting to analyze whether a particular algorithmic selection tool is legally compliant. Some of these challenges relate to algorithmic and data-driven selection tools' reliance on correlation rather than causation. In some ways, using correlative techniques across a huge number of attributes allows for a richer and more holistic analysis of candidates. But correlative techniques fit awkwardly (if at all) with existing legal frameworks, many of which—including antidiscrimination laws—rest on cause-and-effect relationships.¹⁹⁸ Reliance on correlation alone is also discouraged in modern test validity theory. This could complicate efforts to validate selection procedures that have an adverse impact.¹⁹⁹

A related challenge is that even fairly small gaps in selection rates will be statistically significant given a sufficiently large number of observations. The large number of attributes stored regarding candidates introduces additional dangers, most notably that the risks of construct-irrelevant variance and redundant encoding of protected class status, explained below, increase with the dimensionality of a data set.

1. The Ubiquity and Meaninglessness of Statistical Significance in Large Data Sets

If an employer uses an algorithmic tool to assess hundreds or thousands of candidates, rejected candidates who sue may find that the bar for making out a *prima facie* case of disparate impact discrimination under current law is remarkably low. Recall that the primary inquiry for *prima facie* disparate impact focuses on the differences in the rates at which members of protected class groups are selected, and that courts have most often focused on whether those differences are statistically significant.²⁰⁰ For selection procedures that are used on a few dozen candidates, the magnitude of the difference required for statistical significance is fairly large.

But, all else being equal, the magnitude of the difference necessary for statistical significance diminishes as the number of observations in a data set increases. If a data set has thousands of observations, even very small differences—say a 0.5% difference in selection rates between men and

198. See Allan G. King & Marko Mrkonich, “Big Data” and the Risk of Employment Discrimination, 68 OKLA. L. REV. 555, 563 (2016).

199. *Id.*

200. See discussion *supra* Section II.B.2.b.

women—may nevertheless be statistically significant. Under some interpretations of current law, such a statistically significant difference may, by itself, establish a prima facie case of disparate impact.²⁰¹

Consider the First Circuit’s 2014 decision in *Jones v. City of Boston*.²⁰² In that case, the First Circuit reversed a district court decision that had relied on the four-fifths rule in granting summary judgment to an employer, with the circuit court holding that the four-fifths rule cannot be used to “trump a showing of statistical significance,” particularly in cases with a large sample size.²⁰³ Indeed, the court ultimately rejected the notion of an additional “practical significance” requirement for prima facie disparate impact altogether, finding that “any theoretical benefits of inquiring as to practical significance outweighed by the difficulty of doing so in practice in any principled and predictable manner.”²⁰⁴

Employers seeking to leverage the power of Big Data at scale must either hope for a change in the prevailing winds of case law, or else find ways of eliminating statistically significant disparities between protected groups. But it may be devilishly difficult to reduce differences in selection rates to statistically insignificant levels without using techniques that make direct adjustments on the basis of protected characteristics—a technique that could constitute disparate treatment discrimination.²⁰⁵ Also, even if a selection procedure were designed and confirmed to have no disparate impacts during testing, disparate impacts may arise over time if the characteristics of the applicant pool diverge from the characteristics of the candidates in the training data. Current case law provides no clear guidance on whether making additional adjustments to the algorithm to reduce such later arising disparate impacts would constitute disparate treatment.

2. *Construct-Irrelevant Variance and Construct Underrepresentation*

The large number of attributes that are available in the age of Big Data will also present novel challenges as courts, agencies, and employers attempt to assess what a business necessity defense might look like in the context of algorithmic tools. A high-dimensionality data set presents an increased risk of construct-irrelevant variance, that is, nonrandom differences in test results that

201. *Ricci v. DeStefano*, 557 U.S. 557, 587 (2009) (citing *Connecticut v. Teal*, 447 U.S. 440, 446 (1982)) (“[A] prima facie case of disparate-impact liability [is] essentially, a threshold showing of a statistical disparity . . . and nothing more.”).

202. *Jones v. City of Boston*, 752 F.3d 38 (1st Cir. 2014).

203. *Id.* at 46, 49, 52–53.

204. *Id.* at 53.

205. See discussion *infra* Sections III.E, IV.C.

are the result of factors unrelated to the intended construct.²⁰⁶ This can happen for a variety of reasons, including when a criterion or predictor measures something more or different than the target construct (e.g., if the scores on a mathematical aptitude test are affected by a test-taker's proficiency in written English); or when scores reflect cultural differences rather than (or in addition to) differences in job related competencies. The inverse of construct-irrelevant variance is construct underrepresentation or construct deficiency.²⁰⁷ This occurs when criterion measures or predictors fail to reflect construct-relevant sources of variance because the criteria or predictors are unrepresentative or otherwise do not capture important aspects of the target construct.²⁰⁸ Both construct-irrelevant variance and construct deficiency can generate adverse impacts if members of certain subgroups perform differently on the improperly included or excluded aspects of job performance.²⁰⁹

The manner in which predictors and the test sample are selected in an algorithmic selection tool creates a risk of construct deficiency and introduces a potential source of construct-irrelevant variance in addition to those that affect traditional employment tests. According to modern test validation literature, the proper method for selecting predictors involves not just searching for statistical relationships between predictors and criteria, but also examining whether there are theoretical and logical reasons to suppose that the predictors are related to the criterion—in other words, that they are related in more than a mere correlational sense.²¹⁰

This was not a major issue for the sorts of employee selection procedures that existed at the time the Guidelines were promulgated because having a conceptual basis for predictor selection is a practical necessity for paper-and-

206. See AM. EDUC. RESEARCH ASS'N ET AL., *supra* note 109, at 12–13; Messick, *supra* note 191, at 34.

207. See AM. EDUC. RESEARCH ASS'N ET AL., *supra* note 109, at 12–13; Messick, *supra* note 191, at 34.

208. See AM. EDUC. RESEARCH ASS'N ET AL., *supra* note 109, at 12–13; Messick, *supra* note 191, at 34; PRINCIPLES, *supra* note 81, at 11–12.

209. See PRINCIPLES, *supra* note 81, at 11–12.

210. See, e.g., Messick, *supra* note 191, at 17 (“[E]mpirical relationships between the predictor scores and criterion measures should make theoretical sense in terms of what the predictor test is interpreted to measure and what the criterion is presumed to embody . . . [E]ven for purposes of applied decision making, reliance on criterion validity or content coverage is not enough. The meaning of the measure, and hence its construct validity, must always be pursued”); PRINCIPLES, *supra* note 81, at 12 (“The rationale for a choice of predictor(s) should be specified. A predictor is more likely to provide evidence of validity if there is good reason or theory to suppose that a relationship exists between it and the behavior it is designed to predict. A clear understanding of the work (e.g., via results of a work analysis), the research literature, or the logic of predictor development provides this rationale. This principle is not intended to rule out the application of serendipitous findings, but such findings, especially if based on small research samples, should be verified through replication with an independent sample.”).

pencil employment tests; it would be inefficient, to say the least, for the developers of such a test to provide a sample of hundreds or thousands of random questions to current employees and blindly search the results to see which questions correlate with performance on the criterion measures of interest. Instead, the designers of traditional employment tests select or develop questions because they have a prior reason to believe that there is a relationship between the proposed test questions and the criterion of interest. Choosing predictors based on their theoretical relationship with the target construct thereby allows test designers to be alert to potential sources of construct-irrelevant variance and to ensure that the test is measuring a sufficiently representative set of job related criteria. The hypothesized relationship between predictors and criteria is then tested by analyzing the results of the validation study to see if the test responses correlate with the criterion measures.²¹¹

But the algorithms that drive ML-based selection procedures do not consider theoretical or logical relationships between variables, or whether the training data includes attributes that constitute a representative set of predictors. The training algorithm instead examines numerous individual attributes and combinations of the attributes available in the training data and then develops a model based on correlations with the criterion measures—without regard to whether there was a prior reason to suppose that the attributes would have predictive value with respect to the criterion. This is both a blessing and a curse. It is a blessing because it has the potential to unearth job related predictors that would not have been obvious to humans. But it also creates a heightened risk that an algorithm will discover and capitalize on chance correlations.²¹² That risk that is heightened further when data sets contain a large number of observations (because small differences can constitute statistically significant correlations given a large enough sample size) or attributes (because more attributes also means more opportunities for chance correlations).

In the science world, the tendency of algorithmic tools—particularly those that utilize deep learning—to “discover” chance correlations is already

211. Cf. AM. EDUC. RESEARCH ASS'N ET AL., *supra* note 109, at 17 (“[T]he test is not a measure of a criterion, but rather is a measure hypothesized as a potential predictor of that targeted criterion. Whether a test predicts a given criterion in a given context is a testable hypothesis.”).

212. PRINCIPLES, *supra* note 81, at 13 (“In cases where scores from . . . algorithms are used as part of the selection process, the conceptual and methodological basis for that use should be sufficiently documented to establish a clear rationale for linking the resulting scores to the criterion constructs of interest. In addition, when some form of empirical keying is used, clear evidence of cross-validity should be provided prior to operational use to guard against empirically driven algorithms’ propensity to capitalize on chance.”).

causing a “reproducibility crisis,” in the words of Rice University statistician Dr. Genevera Allen.²¹³ Allen discovered in her research several instances where scientists using deep learning algorithms claimed to have identified previously unknown associations between variables, only to find that other researchers were unable to reproduce the results when applying the same techniques to different data sets.²¹⁴ They discovered associations between variables that existed only in the particular samples available to the researchers, but those associations had no generality because these correlations were absent from different sets of similar data.²¹⁵

Similar phenomena pose a substantial threat to validity for users of algorithmic employee selection tools. First, as with the genomic and health research that was the focus of Allen’s study,²¹⁶ there is a risk that algorithmic selection tools will discover correlations between variables in the training data that do not actually exist in the broader real-world applicant pool. While machine learning offers a number of well-accepted techniques for cross-validation, those methods may not be adequate to weed out all of the construct-irrelevant associations between variables in large data sets, particularly if a data set contains information on thousands (or tens or hundreds of thousands) of attributes.

There is another type of correlation that can also afflict employee selection procedures—associations between attributes that do hold in the population at large but that are nevertheless construct irrelevant. The number of such correlations may increase if the training examples tend to come from individuals from the same demographic group or groups, and who therefore share non-job-related attributes in the data. For example, if musical tastes differ by race, and the best incumbent job performers for a particular position are predominantly from a given race, then a high correlation between musical taste and job performance may exist—but only due to demographics, and not because musical taste is an accurate and generalizable predictor of job performance. The less representative the training data are of the population at large, the higher the risk that a deep learning model will identify and create a model that relies upon such demographics-dependent correlations.

213. Pallab Ghosh, *AAAS: Machine Learning ‘Causing Science Crisis’*, BBC (Feb. 16, 2019), <https://www.bbc.com/news/science-environment-47267081> [<https://perma.cc/VT3H-KQHW>].

214. *Id.*

215. *Id.* The *Principles* allude to this potential problem when discussing validation in the context of algorithmic selection procedures. See *PRINCIPLES*, *supra* note 81, at 13 (“[W]hen some form of empirical keying is used, clear evidence of cross-validity should be provided prior to operational use to guard against empirically driven algorithms’ propensity to capitalize on chance.”).

216. See Ghosh, *supra* note 213.

An example of this phenomenon can be seen in the results of the MIT Media Lab Gender Shades study.²¹⁷ That study examined the accuracy of gender classification systems—that is, machine learning software that takes a photograph of a person as its input and outputs a predicted classification of that person’s gender as male or female.²¹⁸ The MIT study used the gender classification systems on photographs of Northern European and African politicians.²¹⁹ The study showed that each of the three facial recognition platforms was more accurate in classifying the European legislators than their African counterparts.²²⁰ Not only that; the study also indicated that the accuracy of the tool was generally better for people with skin types typically associated with moderately dark skin than those with very dark skin.²²¹

The authors hypothesized that this may be because darker skinned individuals may have been “less represented in the training data.”²²² If so, the tool’s accuracy might have been diminished either because of the dissimilarity of darker subjects’ skin from those that dominated the training data set or because darker skin may be highly correlated with other gender-distinctive attributes that were also underrepresented in the training data.²²³ The tool may thus have learned attributes useful for distinguishing white males and white females, while devaluing gender-distinctive attributes present in individuals with darker skin, and underweighting those attributes that actually are useful predictors across the population as a whole.

This is an illustration of a broader challenge with correlation-based selection: the more dissimilar an individual is from the population that served as training examples, the less reliable the tool’s output will be for that individual. That could lead to undesirable—and perhaps unlawful—outcomes with algorithmic employee selection tools.²²⁴ In the employment setting, if the positive examples used in the training data are predominantly individuals with a certain set of protected class characteristics, the data may tell the tool that

217. JOY BUOLAMWINI & TIMNIT GEBRU, GENDER SHADES: INTERSECTIONAL ACCURACY DISPARITIES IN COMMERCIAL GENDER CLASSIFICATION (Sorelle A. Friedler & Christo Wilson eds., 2018).

218. *Id.* at 1.

219. *Id.* at 5.

220. *Id.* at 8, 10.

221. *Id.* at 7 (noting the difference in the distributions of lighter and darker skinned subjects, labeled according to the Fitzpatrick classification system).

222. *Id.* at 10.

223. *See id.*

224. Among psychologists, the term “subgroup validity” refers to the different validity coefficients that can arise between tested subgroups, and the differences in those coefficients are termed “differential validity.” *See* Richard J. Klimoski & Lori B. Zukin, *Psychological Assessment in Industrial/Organizational Settings*, in 10 HANDBOOK OF PSYCHOLOGY: ASSESSMENT PSYCHOLOGY 317, 324 (Irving B. Weiner ed. 2003). The issue of differential validity is discussed further in Section IV.C.1.

those individuals' attributes—whether construct-relevant or not—are associated with success for the position in question. The more highly qualified candidates' attributes differ from the training benchmarks, the more the algorithm's ability to identify those candidates would diminish.

As an example, say 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 study 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.

If courts and agencies reassess the legal standards of employee selection procedures to bring them in line with modern scientific standards, the resulting new standards will likely include a requirement that an employer demonstrate some level of construct relevance—as opposed to relevance in the correlational sense—for algorithmic selection procedures. In either case, employers may find conducting a legally compliant validation cumbersome at best and infeasible at worst, given the sheer number of attributes that would need to be reviewed. The task would be doubly challenging in the context of a deep learning tool, which may transform the input variables into representations that are not human interpretable.²²⁵ Because courts have never ruled on the requirements of validity studies in the context of algorithmic selection procedures that utilize thousands of features, it simply is not clear

225. See discussion *infra* Section III.C.

how courts will treat such tools if they produce a disparate impact and the employer is unable to explain how and why the variables considered and constructed by the tool were relevant to the job in question.

3. *Redundant Encodings*

On the surface, it may seem easy for the developer of an algorithmic selection tool to design around disparate treatment—simply ensure that gender, race, and other protected class status information is not made available to the selection tool during training. But in the age of Big Data, it may not be that simple. First, it may be difficult to reliably excise protected class status information if the training data pools information on candidates from a variety of sources, each of which may encode the sensitive characteristic differently. Even if employers overcome that hurdle, however, a tool trained on data sets of high dimensionality could effectively reconstruct a protected characteristic from other attributes with which it is correlated, a problem called redundant encoding.²²⁶ When that occurs, the redundant encoding effectively creates a reliable proxy for the protected characteristic, even if it does not use the characteristic itself.²²⁷

If the tool is able to reconstruct the protected characteristic, has the tool engaged in disparate treatment? Or does the fact that it did not explicitly consider the candidate's gender mean that the redundant encoding is facially neutral, such that disparate impact provides the proper analytical rubric? Unsurprisingly, this issue is not addressed in antidiscrimination case law, meaning that courts and agencies will have to decide which rubric to use when faced with redundant encodings.

Say that redundant encoding allows the algorithm to reconstruct a person's sex with 99.9% accuracy—say, by using the candidate's height, weight, college attended, and recent clothing purchases—and uses the resulting proxy for sex as part of the model. If the model then systematically disfavors women, women may plausibly argue that they were rejected because of their sex. Such a ruling would be consistent with the prevailing trend in case law, under which courts have increasingly held that, because Title VII prohibits discrimination because of sex, the prohibition against disparate treatment covers “not just discrimination based on sex itself, but also discrimination based on traits that are a function of sex.”²²⁸ Thus, courts have

226. See CYNTHIA DWORK ET AL., *FAIRNESS THROUGH AWARENESS* 22 (2011).

227. Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CAL. L. REV. 671, 695 (2016).

228. *Zarda v. Altitude Express, Inc.*, 883 F.3d 100, 111–12 (2d Cir. 2018); see also *Hively v. Ivy Tech Cmty. Coll.*, 853 F.3d 339, 339 (7th Cir. 2017). See generally Mary Stuart King,

held that using attributes related to sex, such as life expectancy,²²⁹ conformance to gender norms,²³⁰ and sexual preference²³¹ constitutes disparate treatment.

But it is not clear how far disparate treatment liability may extend when the discrimination is based on proxy characteristics. One court attempted to draw a distinction between characteristics that are a “proxy” for a protected characteristic and those that merely “correlate” with it.²³² But it is unclear where the line between proxy and correlate lies. It is difficult to imagine a court countenancing a model that uses a predictor variable that perfectly correlates with a protected characteristic. But what about a predictor variable with an R-squared value of 0.99 with respect to the protected characteristic? Or 0.8? Or 0.5? Until these questions are resolved, employers cannot afford to assume that they can insulate themselves from disparate treatment liability risk simply by removing demographic data and related information from the training data.²³³

C. *The Black Box Problem*

Perhaps the issue that legal commentators raise most frequently when discussing algorithmic selection tools is the black box problem—that is, that it may be difficult or impossible for a human to reconstruct or interpret the logical steps that the tool took when assessing the fitness of a candidate for a particular job. In this way, ML-powered selection tools share much in common with human decision makers, whose reasoning behind a particular selection decision may not be apparent to outside observers. But human

Note, *To Protect or Not to Protect: An Empirical Approach to Predicting Where the Fourth Circuit Would Stand on Coverage for Sexual Orientation Discrimination Under Title VII*, 705 S.C. L. REV. 1075, 1076–80 (discussing the expansion and development of Title VII in regards to sex as a protected class).

229. See *City of L.A. Dep’t of Water & Power v. Manhart*, 435 U.S. 702, 711 (1978) (finding Title VII violation where employer “require[d] 2,000 individuals to contribute more money into a fund than 10,000 other employees simply because each of them is a woman, rather than a man,” even though contributions were based on observed actuarial differences between the sexes in longevity).

230. *Zarda*, 883 F.3d at 112 (citing *Price Waterhouse v. Hopkins*, 490 U.S. 228, 250–51 (1989)).

231. *Id.* at 113–15.

232. *Bowers v. Nat’l Collegiate Athletic Ass’n*, 563 F. Supp. 2d 508, 517–18 (D.N.J. 2008).

233. Even if the effect of redundant encodings was subject only to disparate impact analysis, the presence of redundant encodings would still pose liability risks to employers. The presence of such encodings could lead to gaps in selection rates for protected groups. And if the attributes that generated the redundant encodings are construct-irrelevant, an employer would likely be unable to establish a business necessity defense.

decision makers can be put on the witness stand and forced to explain their reasoning. Their underlying motivations for a particular decision may also be illuminated by other evidence, such as emails, text messages, conversations with friends, and social media activity. Machines are, for now at least, not able to testify regarding their decisions, and because ML algorithms are effectively built on the closed universe of their training data, little other evidence will likely be available that could shed light on how an algorithmic selection tool arrived at a particular score or recommendation for a particular candidate.

With the rise of deep learning, this inscrutability is not simply a problem for plaintiffs and courts. One result of the complexity of deep neural networks is that the precise inner workings of an algorithm may be indecipherable even to the algorithm's designers.²³⁴ While Title VII does not prohibit opaque selection procedures per se, the potential opacity of algorithmic tools will present considerable practical challenges for both plaintiffs and employers in discrimination suits based on the use of such tools once adverse impact is established.

For example, consider what would happen if redundant encodings of protected characteristics allowed algorithmic tools to essentially reconstruct the protected characteristics themselves, with discriminatory effects on certain protected groups. Regardless of whether courts characterize any resulting discrimination as disparate treatment or disparate impact, the employer may have difficulty deciphering whether—much less how—redundant encoding arose. This would complicate both efforts to rectify the discrimination and preparation of an adequate legal defense.

Of course, plaintiffs would have difficulty determining how the discriminatory output had been generated as well. In a disparate impact case, plaintiffs are responsible for identifying the subset of attributes responsible for the redundant coding, unless they can prove the attributes are “not capable of separation for analysis.”²³⁵ That may seem to suggest that employers who use such systems may escape liability for discrimination. But if the tool is as opaque to the employer as it is to the employee, it is difficult to predict whether employers or employees will suffer the greater disadvantage from the tool's opacity.

If courts view the discrimination through a disparate impact lens, the employer would seem to be at a greater strategic disadvantage than the plaintiff. Because the final output of the tool is not a black box, a plaintiff would have little difficulty determining whether the ultimate effect of the tool

234. See generally Will Knight, *The Dark Secret at the Heart of AI*, MIT TECHNOLOGY REVIEW (Apr. 11, 2017), <https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/> [<https://perma.cc/F249-2N9D>] (discussing the difficulty in assessing results from AI technology).

235. 42 U.S.C. § 2000e-2(k)(1)(B)(i) (2012).

was to disproportionately disfavor a protected class, as necessary to establish a *prima facie* case. On the other hand, the employer's efforts to establish the validity of the procedure would be complicated by the impracticability of tracing the neural network's transformation of the original input attributes into the final parameters used by the model. This is particularly true if courts require validation of the individual components of an algorithmic selection procedure. In such a situation, the employer might find itself hamstrung by its inability to identify and validate the components of the model that are having an adverse impact. This problem becomes even more serious—and perhaps intractable—if an algorithmic tool is updated frequently or continuously as new data is received. In such situations, the employer may not have a practical way of reconstructing the algorithm's parameters at the relevant time(s). If, as the Supreme Court has held, selection procedures are inadequate when their validation studies rely on ratings that are “vague and fatally open to divergent interpretations,”²³⁶ it is unlikely that a court will be satisfied by a selection procedure whose standards are completely opaque and not open to any human-decipherable interpretation.

If courts hold that the use of a reconstructed protected characteristic constitutes disparate treatment rather than disparate impact, it is not clear that employers would fare much better. While a plaintiff might find it impossible to explain how an algorithmic tool discovered redundant encodings of a protected characteristic, it is not difficult to imagine courts taking a *res ipsa loquitur* attitude if it appears obvious that a tool is employing an effective proxy for a protected characteristic.²³⁷ If so, current law does not appear to provide employers with an easily identifiable defense; the *McDonnell Douglas* framework is inapplicable if courts determine that the use of a redundant encoding is tantamount to use of the protected characteristic itself, and therefore direct evidence of discrimination.

D. A Clear Target

One of the most striking consequences of the *Griggs* decision and the subsequent development of disparate impact litigation has been the deformatization of employee selection procedures. As Lex Larson has observed, starting with *Griggs*, fear that testing would generate liability for disparate impact has driven many employers toward increased reliance on subjective decision-making:

236. See *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 433 (1975).

237. See RESTATEMENT (SECOND) OF TORTS § 328D (AM. LAW INST. 1965) (describing tort doctrine of *res ipsa loquitur* as a basis for negligence liability).

This dramatic reversal in business' attitude toward testing was tinged with irony; for the most part, businesses had moved toward the use of tests as a way to lend objectivity to the selection process to select the best-qualified personnel. Starting with *Griggs*, the courts began telling employers that these devices, too, could result in discrimination. As a result, many employers went back to using subjective judgment in making employment decisions.²³⁸

Of course, reliance on human judgment can lead to adverse impacts as well—which is precisely why algorithmic tools represent an appealing alternative. But subjective human judgments leave a lesser paper trail than more formal hiring practices. It is also harder to cast such subjective decision-making by numerous different decision makers as a unified employment practice that could serve as the basis for a class action disparate impact suit.²³⁹ These characteristics make relying on the humans in human resources more appealing, particularly in comparison to the lengthy, costly, and uncertain process of designing and validating a formal selection procedure.

These drawbacks are equally, if not more, apparent in the specific context of algorithmic selection procedures. The very essence of an algorithmic selection procedure is to take the observable characteristics of a candidate and reduce them to rows of data. The output of the selection procedure is essentially a function of complex mathematical formulae. The process simply does not work unless both the candidates and the selection procedure that assesses them are formalized and ultimately reduced to computer code, and the procedure loses its value if it is not used consistently for all candidates under consideration for a given position. When disparate impacts arise, algorithmic selection procedures give potential plaintiffs an obvious target.

The inherent explicitness will also muddy the waters in disparate treatment cases. It is trivial for an employer to ensure that an algorithm does not use a protected characteristic as an input when assessing a candidate. But if the algorithm reconstitutes the protected characteristic through redundant encodings, and if courts hold that using such redundant encodings constitutes disparate treatment, it will be equally trivial for a plaintiff to demonstrate that the algorithmic selection procedure is the source of the disparate treatment. The inner workings of the algorithm may be opaque, which will hinder plaintiffs' ability to demonstrate precisely how a protected characteristic was

238. LARSON, *supra* note 84, at § 25.02.

239. See *Wal-Mart Stores, Inc., v. Dukes*, 564 U.S. 338, 356–57 (2011) (rejecting class certification because “[r]espondents have not identified a common mode of exercising discretion that pervades the entire company” and different managers would likely say they were using different criteria when making decisions).

reconstituted. But as discussed elsewhere in this Article,²⁴⁰ that may not provide employers with an escape route.

The explicitness of algorithmic selection procedures will also complicate employers' efforts to navigate the intersection of disparate impact with disparate treatment, as considered in the *Ricci* case. In the algorithmic age, it will be easier than ever for employers to eliminate disparate impacts in their selection procedures—but race norming, boosting algorithmic scores of candidates from disadvantaged groups, and other preferential practices risk disparate treatment liability. Less direct methods of eliminating disparate impacts remain untested in court, leaving employers with no clear options regarding how to cure disparate impacts when they arise.

And it is almost inevitable that at least some disparate impacts will arise. Even if an employer succeeds in designing an algorithmic selection procedure that has no disparate impacts during initial training, adverse impacts may creep in as the characteristics of candidates and successful employees in a given position change. Making changes after a tool has already been deployed is problematic under *Ricci*, which held that such modifications may be made only prospectively.²⁴¹ Employers will then be forced to make conscious decisions about how to manage those adverse impacts, and any adjustments made to the model in response will themselves have to be reduced to computer code and explained during the course of litigation. Faced with this morass of legal uncertainty, many employers may prefer to continue to rely on subjective human judgment—and with it, the potential effects of human prejudice—rather than risk getting bogged down in the marsh of an unsettled area of law.

E. Disparate Treatment: A Brave New World

Given the manner in which disparate-treatment case law has developed, concerns have been raised regarding whether companies might be effectively immune from disparate treatment liability if they use algorithmic selection devices that learn, without any express human programming, to classify workers in a discriminatory manner on the basis of protected characteristics.²⁴² The premise is the belief that because machines cannot have “intent” in the human sense, there can be no liability for their actions under Title VII unless the machine was intentionally programmed to discriminate.²⁴³ This concern seems misplaced.

240. See discussion *supra* Section III.C; *infra* Section III.E.

241. See *Ricci v. DeStefano*, 557 U.S. 557, 585 (2009).

242. See Charles A. Sullivan, *Employing AI*, 63 VILL. L. REV. 395, 404–10 (2018).

243. See *id.*

First, fixating on intent means ignoring the clear anticlassification rule set forth in the statutory text. Under the plain text of § 703(a), a Title VII violation occurs whenever an adverse employment or hiring action is because of a protected characteristic.²⁴⁴ That is language of causation, not intent. The disparate impact theory of discrimination itself first arose out of this language, with the Supreme Court explicitly holding that employers cannot escape liability under § 703(a) for practices with discriminatory effects simply by pleading lack of intent:

[G]ood intent or absence of discriminatory intent does not redeem employment procedures or testing mechanisms that operate as ‘built-in headwinds’ for minority groups and are unrelated to measuring job capability.

The Company’s lack of discriminatory intent is suggested by special efforts to help the undereducated employees through Company financing of two-thirds the cost of tuition for high school training. But Congress directed the thrust of the Act to the consequences of employment practices, not simply the motivation.²⁴⁵

The frequent connection of disparate treatment with intent in the case law has never been cast as mandated by the statutory text. More likely, it is a consequence of the fact that, up to now, hiring practices have been driven by human decision makers.

To that point—who is to say that courts would necessarily conclude that machines cannot possess intent? True, many definitions of intent reference a “state of mind” or a “conscious” desire to bring about a particular result, terms that seem to refer to distinctly human traits.²⁴⁶ But other definitions are far broader, focusing only on the party’s “objective” or “purpose.”²⁴⁷ Under the criminal laws of many states, an entire category of “general intent” exists where the defendant’s state of mind is not relevant so long as the defendant

244. 42 U.S.C. § 2000e-2(b) (2012).

245. *See* *Griggs v. Duke Power Co.*, 401 U.S. 424, 432 (1971).

246. *See, e.g., Intent*, BLACK’S LAW DICTIONARY (11th ed. 2019) (defining intent, in part, as “[t]he state of mind accompanying an act . . .”).

247. *See, e.g.,* WASH. REV. CODE ANN. § 9A.08.010 (West, Westlaw through 2019 Reg. Sess.) (“A person acts with intent or intentionally when he or she acts with the objective or purpose to accomplish a result which constitutes a crime.”); *State v. Salinas*, 423 P.3d 463, 465 (Idaho Sup. Ct. 2018) (“Intent ‘is the purpose to use a particular means to effect a certain result.’”) (citing *State v. Stevens*, 454 P.2d 945, 950 (Idaho Sup. Ct. 1969)).

acted volitionally as opposed to accidentally.²⁴⁸ Tort law treats an act as “intentional” when the actor believes that the consequences of his act are “substantially certain” to result from it.²⁴⁹ Such definitions of intent could easily apply to decisions made by machines. It will not do to simply assume that because machines generally are not considered to have consciousness in the metaphysical sense, they necessarily cannot possess intent in the legal sense or that intent cannot be imputed to those who deploy them.

Moreover, intent has proven to be quite a malleable concept in the context of Title VII, as in other areas of law. The *Ricci* majority stated that disparate treatment requires intent but, at the same time, acknowledged that the employer’s objective in that case was avoiding legal liability;²⁵⁰ to the extent that race factored into the decision, the employer’s intent was not to treat workers differently on the basis of race but rather to avoid discrimination on the basis of race. Nevertheless, citing Title VII’s use of the broad term *because of*, the Court treated that motivation as itself a form of disparate treatment.²⁵¹

Employers can also be held liable for sexual harassment even if the harassment was committed by nonemployees and even if the employer had no actual knowledge of the harassment.²⁵² A number of courts have also upheld the “cat’s paw” theory of discrimination,²⁵³ under which “an employer who acts without discriminatory intent can be liable for a subordinate’s discriminatory animus if the employer uncritically relies on the biased subordinate’s reports and recommendations in deciding to take adverse employment action.”²⁵⁴ If a court is willing to find intent based on a decision maker’s uncritical reliance on another person’s biased recommendation, it seems highly unlikely it would excuse an employer for uncritically relying on the recommendation of a machine it chose to use, regardless of the metaphysics of whether an algorithm can have intent.

Lastly, even if some element of human intent were an absolute requirement for disparate treatment liability, algorithmic selection tools will

248. *See, e.g.*, *United States v. Lamott*, 831 F.3d 1153, 1156 (9th Cir. 2016) (“In a crime requiring ‘specific intent,’ the government must prove that the defendant subjectively intended or desired the proscribed act or result. By contrast, a general intent crime requires only that the act was volitional (as opposed to accidental), and the defendant’s state of mind is not otherwise relevant.”).

249. RESTATEMENT (SECOND) OF TORTS § 8A (AM. LAW INST. 1965).

250. *See Ricci v. DeStefano*, 557 U.S. 557, 579–80 (2009).

251. *Id.*

252. 29 C.F.R. § 1604.11(e) (2019).

253. *Thomas v. Berry Plastics Corp.*, 803 F.3d 510, 514 (10th Cir. 2015); *see also Vasquez v. Empress Ambulance Serv., Inc.*, 835 F.3d 267, 271–73 (2d Cir. 2016); *Lust v. Sealy, Inc.*, 383 F.3d 580, 584 (7th Cir. 2004); *cf. Staub v. Proctor Hosp.*, 562 U.S. 411, 421–22 (2011) (upholding a court’s use of the cat’s paw theory of discrimination, albeit in a case brought under the Uniformed Services Employment and Reemployment Rights Act, not under Title VII).

254. *Thomas*, 803 F.3d at 514.

very much be the product of human motivations and intentions. Because of the need to validate selection procedures that may have a disparate impact, a topic discussed further below, algorithmic selection tools will rely on data that is labeled by humans—ideally, managers or HR employees for the company seeking to use the tool—tasked with assessing the fitness of candidates in the training data for a particular job. Those labelers’ motivations and intentions are incorporated, however indirectly, into the final selection procedure. Similarly, the training data itself will ideally include employee performance data, such as supervisor ratings. Because the input of human decision makers will be baked into the algorithm, it is difficult to imagine courts and enforcement agencies shrugging their collective shoulders and holding that employers who rely on the recommendations of algorithmic selection tools are immune from disparate treatment liability.

For these reasons, the disparate treatment doctrine will not fade into legal obscurity in the age of algorithms. But it is true that courts developed the prevailing judicial interpretations of the disparate treatment doctrine with human decision-making in mind and that the contours of disparate treatment liability in the context of algorithmic tools have yet to be established. This means that courts and agencies will have to consider the meaning of the statutory text afresh if or when they are faced with algorithmic selection tools that classify candidates on the basis of a protected characteristic—regardless of whether such classification was intended by the algorithmic tool’s designers or users.

IV. NEW RULES FOR THE NEW TOOLS: A PROPOSED LEGAL FRAMEWORK FOR ALGORITHMIC SELECTION TOOLS

A. Overview

Title VII requires (1) that employers avoid making employment decisions because of protected characteristics and (2) that employers establish the job relatedness of any selection tool that has an adverse impact on one or more protected groups. Although the legal regime governing employee selection tools was developed without algorithmic selection tools in mind, the broad principles set forth in the statutory text certainly can be applied to algorithmic selection procedures. What is required is not so much a new legal framework as a new conceptual approach to assessing employee selection procedures in the age of algorithms.

In particular, algorithmic selection procedures require taking the fundamental principles of Title VII, and the landmark Supreme Court cases interpreting them, and developing a set of standards that address the unique challenges posed by AI and Big Data discussed in Part III. The ultimate goal

should be to allow employers to find innovative ways of uncovering talent and building a diverse workforce—objectives fully consistent with Title VII—while remaining true to the purpose of Title VII itself. It should not be difficult to reconcile these objectives because selecting the highest quality candidate for a job while ensuring broad participation by disadvantaged groups is, as the Supreme Court held in *Griggs*, the very essence of Title VII.²⁵⁵ Algorithmic selection tools create an unprecedented opportunity to advance these goals by excising human prejudice and bias from personnel decisions.²⁵⁶

Our proposal weaves the disparate impact and disparate treatment inquiries into a single analytical framework:

Step 1: Determine whether the algorithmic procedure had an unlawful disparate impact.

- a. Prima facie disparate impact: Determine whether the gap between protected groups is large enough to give a reasonable employer concern that the algorithmically generated model is unreasonably disadvantaging members of a protected group.
 - i. *If no unreasonable gap exists, skip to Step 2.*
- b. Employer's defense: If a disparate impact exists, determine whether the tool has been properly validated.
 - i. *If the tool has not been validated, the employer is liable for disparate impact, and the court or agency should skip to Step 2 to determine whether the employer also engaged in disparate treatment.*
- c. Less discriminatory alternative: Determine whether the employer considered and rejected an alternative modeling method that would have affected a reasonable reduction in adverse impact but would have continued to meet the employer's legitimate objectives.

Step 2: Determine whether the ML tool used any methods that make prohibited classifications or that otherwise constitute disparate treatment.

255. *Griggs v. Duke Power Co.*, 401 U.S. 424, 436 (1971) (“Far from disparaging job qualifications as such, Congress has made such qualifications the controlling factor, so that race, religion, nationality, and sex become irrelevant.”).

256. See Jon Kleinberg et al., *Discrimination in the Age of Algorithms*, 10 J. LEGAL ANALYSIS 113, 163–64 (2018).

In broad strokes, with details to follow below, the first step is to determine whether use of the tool has a *prima facie* disparate impact on one or more protected groups. Because the currently favored statistical significance approaches to *prima facie* disparate impact would sweep too broadly in the age of Big Data, however, a modified approach to disparate impact analysis is required. Instead of focusing on the presence or absence of statistical significance, the inquiry should be one of reasonableness—a plaintiff can establish a *prima facie* case of disparate impact demonstrating by producing evidence demonstrating that the gap between protected groups is large enough to give a reasonable employer concern that the algorithmically generated model may be disproportionately disadvantaging members of a protected group.

The next step in the analytical process depends on whether a *prima facie* disparate impact exists. If it does not, the disparate impact inquiry ceases, and the only remaining issue is whether the algorithmic tool used techniques that constitute disparate treatment.

If, on the other hand, the plaintiff does present *prima facie* proof of disparate impact, the inquiry would instead progress to whether the algorithmic assessment is job related and consistent with business necessity. The key inquiry here would be whether the criteria that serve as target variables for the training algorithm represent “essential” and “important” job functions, as identified through standard job analysis.²⁵⁷ Essential functions can be used as screening criteria for an algorithmic selection tool; that is, employers can use algorithmic selection tools to screen out candidates the tool identifies as lacking the ability to perform essential job functions. Important job functions can also be used as target criteria to be optimized, but they cannot be used as hard screening devices. As in ADA cases, the employer’s designation of essential and important job functions would be entitled to some deference. An algorithmic tool would be considered job related if the criteria meet these requirements, if the outputs of the algorithmic assessment are significantly correlated with adequate measures of those criteria, and if the employer demonstrates that it took reasonable steps to guard against construct-irrelevant variance in the results. The employee can rebut this showing with proof that the employer used criteria that were not job related or failed to model these dimensions correctly.

The third step of the current disparate impact analysis—the plaintiff’s burden of demonstrating the existence of a less discriminatory alternative—would, in the case of algorithmic tools, require a plaintiff to show that the employer considered and rejected an alternative modeling method that would

257. See 29 C.F.R. § 1607.14(B)(2) (2019); § 1630.14(b)(3).

have effected a reasonable reduction in adverse impact but would have continued to meet the employer's legitimate objectives. Once again, we eschew the requirement of a statistically significant reduction because of the likelihood that any reduction in adverse impact, in a Big Data world, would meet that criterion. In the same vein, any reduction in the accuracy with which this alternative modeling method selected the best employees also would be deemed statistically significant in a world of Big Data.

After the disparate impact analysis concludes, attention should turn to whether the algorithm used any methods that constitute unlawful disparate treatment. The framework identifies two techniques through which employers, during the development and training process, should be permitted to take measures to prevent disparate impacts without exposing themselves to disparate treatment liability.

B. Disparate Impact

1. Prima Facie Disparate Impact

The standard approach to determining whether prima facie evidence of a disparate impact exists relies on formal statistical tests, with most courts relying on a bright-line rule that statistically significant differences in selection rates between favored and disfavored groups suffice to prove the first element of a disparate impact claim.²⁵⁸ In the era of Big Data, this criterion is no longer appropriate because, all else equal, the larger the sample, the smaller the differences that will be deemed statistically significant. At a certain point—which we are fast approaching for practical purposes—all differences, no matter how small, will be statistically significant. That means that a statistical significance requirement will be meaningless. How then should courts assess disparities when statistical significance no longer is a useful criterion for distinguishing discriminatory from nondiscriminatory assessment methods?

One possible policy response to the diminishing meaningfulness of statistical significance would be to abolish the disparate impact doctrine altogether. The doctrine has been criticized by some legal commentators and jurists on constitutional grounds, including Justice Scalia in his *Ricci* concurrence.²⁵⁹ And in practical terms, one could argue that in the age of Big Data, which allows for a richer analysis of candidates while reducing the practical significance of statistical tests, the doctrine has simply outlived its usefulness. But it usually is not possible to design an employment test, whether algorithmic or not, that is so comprehensive that it captures all

258. See discussion *supra* Section III.B.1.

259. See *Ricci v. DeStefano*, 557 U.S. 557, 594–96 (Scalia, J., concurring).

characteristics predictive of good job performance. Moreover, in the context of algorithmic selection tools, the effects of past discrimination may be baked into training data, meaning that unchecked reliance on existing data sets could repeat and reinforce existing patterns of discrimination. Similarly, the amount of statistical noise inherent in large data sets create too many opportunities for an algorithm to settle on parameters that relate more to demographic characteristics than to ability to perform the job. As a result, the concept of disparate impact discrimination still has a place in the age of algorithms.

But rather than rigidly relying upon statistical significance—which is not, in any event, mandated by any statute—courts and agencies should substitute a less formal reasonableness criterion when assessing whether a *prima facie* disparate impact.²⁶⁰ In other words, policies and practices would be deemed to have a disparate impact only when selection rates between groups differ unreasonably. Although this dispenses with the certainty that a purely statistical rule provides, the loss of that certainty is more than offset by the benefits of adopting a more flexible standard that can be adapted to the changing nature of algorithmic tools and the data sets that they use.

Applying a more flexible test should not be especially difficult; courts have, after all, hardly adhered to a uniform, bright-line rule with respect to statistical tests in the context of disparate impact suits. The Supreme Court’s “two or three standard deviations” formulation is not a bright-line rule, and the Court’s endorsement of this standard was arguably in dictum and is weaker than generally supposed.²⁶¹ And while courts have generally preferred to use tests of statistical significance, a substantial number of courts have looked to the Guideline’s four-fifths rule or otherwise examined the magnitude of the

260. See generally Allan G. King, “Two or Three Standard Deviations” from What?: *How Gross v. FBL Financial Services Changes the Statistical Benchmark in ADEA Collective Actions*, 37 EMP. REL. L. J. 17 (2011) (describing a system for finding a reason alone, as long as that reason taints the employers decision-making).

261. The Court observed in *Castaneda v. Partida* that the statistical disparity at issue, in excess of 12 standard deviations, was probative because scientists routinely consider differences of just two or three standard deviations sufficient to reject a null hypothesis. See *Castaneda v. Partida*, 430 U.S. 482, 496 n.17 (1977) (“Thus, in this case the standard deviation is approximately 12. As a general rule for such large samples, if the difference between the expected value and the observed number is greater than two or three standard deviations, then the hypothesis that the jury drawing was random would be suspect to a social scientist.”). Strictly speaking, this is not the case’s holding because its affirmance would be equally consistent with a rule that 10 standard deviations was required to prove discrimination. Indeed, Justice O’Connor subsequently noted in *Watson*: “Our formulations, which have never been framed in terms of any rigid mathematical formula, have consistently stressed that statistical disparities must be sufficiently substantial that they raise such an inference of causation.” *Watson v. Fort Worth Bank & Tr.*, 487 U.S. 977, 994–95 (1988).

disparity rather than applying a rigid statistical significance rule.²⁶² The theoretical certainty that mathematical tests provide thus has not been consistently attained in practice.

In any event, reasonableness tests are eminently workable, as their continuing popularity and ubiquity in law indicate. Criminal law, tort law, contract law, and, indeed, employment law are all replete with reasonableness tests that courts interpret and apply on a regular basis. In employment law, courts routinely assess whether a proposed accommodation for an employee with a disability is reasonable,²⁶³ whether an employment decision in an age discrimination case was motivated by reasonable factors other than age,²⁶⁴ and what amount of attorney fees are reasonable for a prevailing plaintiff,²⁶⁵ among many other examples. Reasonableness standards give courts the ability to avoid the unjust results that can accompany hard-and-fast rules.

In the context of assessing whether a *prima facie* disparate impact exists, the inquiry into whether a gap in selection rates is unreasonably large should not focus on whether the gap is sufficiently justified or explained by the criteria that underlie the selection procedure; that falls more properly within the realm of the business necessity defense. Rather, the test should be whether, in light of the magnitude of the difference in selection rates and the size of the affected candidate pool, is the gap large enough to permit a reasonable fact finder to conclude that the test systematically disadvantages members of a protected group. If the gap raises such a concern, then the employer would be required to demonstrate that the selection procedure is job related and consistent with business necessity. The *prima facie* case would therefore serve a gatekeeping function, protecting employers from having to validate gaps that, while significant in the statistical sense, are meaningless in practical economic and legal terms.

In assessing whether a gap is reasonable, the statistical significance of a gap would be one factor, but it would be assessed alongside indicators of the magnitude of the gap, such as an odds ratio or other measures of effect size. Courts and agencies could substitute other rules of thumb to serve as benchmarks for magnitude, as the Guidelines did with the four-fifths rule. This would allow courts and agencies to recoup some of the lost efficiencies that come with a bright-line rule.

262. See, e.g., *M.O.C.H.A. Soc’y, Inc., v. City of Buffalo*, 689 F.3d 263, 274 (2d Cir. 2012) (“Consistent with our precedent, the district court properly deferred to [the four-fifths rule] in finding M.O.C.H.A. to have carried its *prima facie* burden”); *Allen v. City of Chicago*, 351 F.3d 306, 310–12, n.5 (7th Cir. 2003) (holding that “promotions made on the basis of the assessment exercise did have a disparate impact on African-American and Hispanic officers” after applying the four-fifths rule).

263. See 42 U.S.C. § 12112(b)(5) (2012).

264. 29 U.S.C. § 623(f)(1) (2012).

265. 42 U.S.C. § 2000e-5(k).

2. *Business Necessity Defense*

If a plaintiff does establish a *prima facie* case of disparate impact, the burden shifts to the employer to show that it has validated the selection procedure and demonstrated its job relatedness. Because content-related evidence will not be sufficient to validate a selection procedure based on passive data,²⁶⁶ the most plausible route to validation for algorithmic tools will rely on criterion-related evidence of validity.

Under the Guidelines, the criterion validation process must begin with a careful job analysis to “determine measures of work behavior(s) or performance that are relevant to the job.”²⁶⁷ These measures can then be used as criteria in the validation study if they “represent important or critical work behavior(s) or work outcomes.”²⁶⁸ There is no reason to depart from these basic principles when validating an algorithmic selection procedure. But a slight change in wording would help ensure consistency across discrimination laws and obviate the need to select different criterion measures for different protected classifications. Specifically, and borrowing from the statutory language of the ADA, the criteria should reflect essential or important job functions.

The EEOC’s interpretive “Questions and Answers” on the Guidelines support this substitution. That guidance document states that if a particular work behavior is essential to the performance of a job, that behavior is “critical” within the meaning of the Guidelines, even if a worker does not spend much work time engaged in that behavior.²⁶⁹ The Q&As use the example of a machine operator for whom the ability to read is “essential” because the worker must be able to read simple instructions, even though the reading of those instructions “is not a major part of the job.”²⁷⁰ The essential nature of being able to read instructions is thus a critical task for purposes of the Guidelines.

The concept of essential job functions is central in the ADA, where it is closely identified with that statute’s version of the job relatedness test.²⁷¹

266. Content-related evidence would be a more persuasive indicator of validity could be a viable path for interactive algorithmic tools that incorporate chatbots or other means of directly eliciting information from candidates rather than using static data testing aspects of job performance, but such interactive assessment tools are beyond the scope of this Article.

267. 29 C.F.R. § 1607.14(B)(2) (2019).

268. § 1607.14(B)(3); *see also* § 1607.14(B)(2) (“[M]easures or criteria [of work behavior or performance] are relevant to the extent that they represent critical or important job duties, work behaviors or work outcomes as developed from the review of job information.”).

269. *See* Adoption of Questions and Answers to Clarify and Provide a Common Interpretation of the Uniform Guidelines on Employee Selection Procedures, 44 Fed. Reg. 11,996, 12,005 (Mar. 2, 1979).

270. *Id.*

271. 42 U.S.C. § 12112(b)(6) (2012).

Given that algorithmic employee selection procedures will have to comply with the ADA no less than Title VII, it would be logical to ensure that criterion standards for purposes of validation have a consistent meaning in both ADA and Title VII cases. Thus, a criterion should be acceptable for purposes of an algorithmic selection procedure if that criterion represents an essential job function, as that term is defined in the ADA.

The Guidelines also provide that noncritical but nevertheless important job duties can also serve as criteria for purposes of establishing the job relatedness of a selection procedure. The question is how an important job function differs from a critical or essential one. Here too, the ADA provides a useful framework. The ADA's job relatedness requirements apply only to criteria that "screen out or tend to screen out an individual with a disability or a class of individuals with disabilities."²⁷² The ADA does not prohibit an employer from taking important but nonessential job functions into account when designing a selection procedure, but it may not use the ability to perform such functions as a screening device or otherwise apply them in a manner that would effectively bar individuals with disabilities from the position in question.²⁷³ In other words, the ability to perform important, but nonessential, job functions can be a factor—just not an inherently decisive one.

Consistent with this principle, the rule should be that employers can use both essential and important job duties as part of an algorithmic employee selection procedure, but only attributes that strongly correlate with essential job functions can be used in algorithms that act, in form or effect, as screening devices. That is, if a validation study shows that the presence or absence of certain attributes is strongly predictive of a candidate's ability to perform one or more essential functions of a job, then the algorithmic tool can use those attributes to remove candidates from consideration for a position. Important job functions can be used as criteria and serve as target variables and used to score or rank candidates so long as they are not used to screen out candidates altogether. In addition, if the algorithm's target variable represents a composite of multiple criteria, its validity does not rest solely on the proper selection of criteria. Rather, the employer must also assign reasonable weights to the criteria in accordance with their relative importance to the performance of the job in question, as revealed by job analysis.²⁷⁴

As in ADA cases, an employer's assessment of which job criteria are essential and important should generally be entitled to deference. The same

272. *Id.*

273. See Adoption of Questions and Answers to Clarify and Provide a Common Interpretation of the Uniform Guidelines on Employee Selection Procedures, 44 Fed. Reg. at 12,005.

274. See PRINCIPLES, *supra* note 81, at 14 ("If the testing professional combines scores from several criteria into a composite, there should be a rationale to support the rules of combination, and the rules of combination should be described.").

rule should apply to the employer's identification of and assignment of weights to the various job functions that serve as the basis for criteria; a plaintiff should not be able to defeat a finding of job relatedness simply by quibbling about the precise weights the employer chose. As long as the employer demonstrates that the selected criteria and weights are reasonable in light of an adequate job analysis, the employer will have satisfied its burden on criterion selection, and the only remaining question would be whether the test results correlate with those criteria.

In accordance with the evolution of the social science of test validity, the rules governing validation of algorithmic selection procedures should also reflect the need to avoid contamination and reduce construct-irrelevant variance. Even if the chosen job criteria are limited to essential and important job functions, there still may be attributes that correlate with the performance of those functions in the training data simply because those attributes are more prevalent among the demographic groups that predominate in the training data. Here, eliminating differential validity, requiring statistical independence, or both could ensure that predictor-criterion relationships do not tend to unfairly exclude members of protected groups for construct-irrelevant reasons.²⁷⁵

Differential validity occurs when a test has substantially greater validity for some tested subgroups than for others.²⁷⁶ For example, a test that accurately predicts job performance for men but not for women has differential validity. The Gender Shades study showed differential validity for the gender classification systems—the tool predicted gender almost perfectly for light-skinned individuals but was noticeably less accurate for darker

275. Of course, it would be impractical to eliminate all sources of construct-irrelevant variance in the uncontrolled setting of recruitment and hiring. But eliminating differential validation and requiring statistical independence across protected groups would at least help ensure that such variance does not stem from protected group membership itself.

276. See generally JOHN W. YOUNG, DIFFERENTIAL VALIDITY, DIFFERENTIAL PREDICTION, AND COLLEGE ADMISSION TESTING: A COMPREHENSIVE REVIEW AND ANALYSIS (2001).

skinned individuals.²⁷⁷ Differential validity and its cousin, differential prediction,²⁷⁸ are well-recognized threats to validity in test design.²⁷⁹

Two variables are said to be *statistically independent* if knowing the value of one of the variables does not provide any information about the value of the other variable.²⁸⁰ In the context of employee selection procedures, race would be statistically independent of the outcome of the selection procedure if knowing an individual's race would not help someone ascertain that individual's performance on the selection procedure.

Viewed through a job-relatedness lens, the concepts of differential validation and statistical independence are intertwined; if an attribute is predictive of the criteria only for certain demographic groups, then the attribute will both have differential validity between demographic groups and not be statistically independent from membership in those groups. In theory, an adversarial learning process should allow the algorithm to tune the model's use of those attributes so that they are no longer dependent on the sensitive characteristic. That, in turn, should help ensure that the algorithmic tool is assessing candidates on the basis of characteristics that relate to job

277. See BUOLAMWINI & GEBRU, *supra* note 217, at 5–10.

278. Differential validity and differential prediction are both forms of test bias that create differences in the meaning of test results for different subgroups. Christopher M. Berry, *Differential Validity and Differential Prediction of Cognitive Ability Tests: Understanding Test Bias in the Employment Context*, 2 ANN. REV. ORGANIZATIONAL PSYCHOL. & ORGANIZATIONAL BEHAV. 435, 436 (“[T]wo forms of test bias especially relevant to personnel selection [are] differential validity (subgroup differences in test validities) and differential prediction (subgroup differences in test–criterion regression equations).”). The formal difference between differential validity and prediction is that the former refers to differences in correlation coefficients between subgroups, while the latter refers to differences in the regression line equations between subgroups. See JOHN W. YOUNG, DIFFERENTIAL VALIDITY, DIFFERENTIAL PREDICTION, AND COLLEGE ADMISSION TESTING: A COMPREHENSIVE REVIEW AND ANALYSIS 4 (2001); YOUNG, *supra* note 276, at 4 (“[D]ifferential validity refers to differences in the magnitude of the correlation coefficients for different groups of test-takers, and differential prediction refers to differences in the best-fitting regression lines or in the standard errors of estimate between groups of examinees.”). Less formally, differential validity can be thought of as differences in the *magnitude* of the relationship between predictor and criterion, whereas differential prediction refers to differences in the *nature* of that relationship.

279. See also PRINCIPLES, *supra* note 81, at 24 (“[P]redictive bias analysis should be undertaken when there are compelling reasons to question whether a predictor and a criterion are related in a comparable fashion for specific subgroups, given the availability of appropriate data.”). See generally YOUNG, *supra* note 276, at 4–5 (discussing how differential validity and differential prediction affect validation).

280. In the terminology of information theory, independence means that there is no mutual information between the variables. Peter E. Latham & Yasser Roudi, *Mutual Information*, SCHOLARPEDIA (2009), http://www.scholarpedia.org/article/Mutual_information [<https://perma.cc/RJ89-YJHD>].

performance and not to membership in a protected group. These techniques are discussed in greater detail below in section IV.C.

In sum, an employer using a algorithmic selection procedure that adversely impacts one or more protected groups would bear the burden of showing (1) that the chosen criteria are representative of essential and important job functions identified through an adequate job analysis; (2) that criteria reflecting nonessential job functions were not used to screen candidates; (3) that the employer assigned reasonable weights to the identified criteria in constructing the selection tool's ultimate target variable; (4) that the output of the selection procedure are correlated with performance on the chosen criteria; and (5) that the employer made reasonable efforts to ensure that predictors and criteria are not contaminated by construct-irrelevant factors that are correlated with protected-class status.

3. *Alternative Selection Procedures*

Under the longstanding framework codified in the 1991 amendments to Title VII, the third and final stage of the disparate impact analysis is the employee's effort to rebut the employer's showing of job relatedness by demonstrating the existence of a less discriminatory alternative selection procedure.²⁸¹ For at least two reasons—one affecting plaintiffs and the other affecting employers—this framework will prove a misfit for algorithmic selection procedures. The challenge for plaintiffs relates to the black box problem: if a deep learning algorithm is particularly opaque or complex, a plaintiff may not be able to gain the level of understanding necessary to mount an effective rebuttal.

From the employers' perspective, the major problem with the current framework is uncertainty surrounding the legal standards. Courts have generally avoided deciding Title VII disparate impact cases at the third stage of the analysis.²⁸² This has resulted in Title VII jurisprudence that lacks clear standards on how a proffered less discriminatory alternative should be judged, particularly on the key point of how available and effective a proposed alternative selection procedure must be to satisfy the plaintiff's burden. Must the employer provide its algorithm to the plaintiff, who then might attempt to reengineer it to reduce the adverse impact? That prospect will deter many employers from using algorithmic selection tools, perhaps even more so than for prior generations of employee selection procedures.

281. *See Ricci v. DeStefano*, 557 U.S. 557, 578 (2009) (citing 42 U.S.C. §§ 2000e–2(k)(1)(A)(ii), (C) (2012)).

282. *See id.* at 589–91 (rejecting three arguments of an “equally valid, less discriminatory testing alternative” brought by plaintiffs that defendant would necessarily have refused to adopt).

For example, many deep learning algorithms in use today rely on mathematical techniques that are guaranteed only to find a locally optimal model rather than the most accurate and effective model possible. That is, from a given set of initial conditions and parameters, the algorithm makes small adjustments until it reaches a point where further small adjustments will reduce rather than improve the accuracy of the model. Algorithm designers use this approach because performing a comprehensive search for a globally optimal model is computationally complex for even a modestly large data set, and wholly impractical for the high-dimensionality data sets that are available in the age of Big Data.

This has two important consequences. First, the process is not guaranteed to find the globally optimal set of parameters for a particular model. Second, two neural networks using the same data may generate different sets of locally optimal parameters, depending on the starting points specified for the parameters at the beginning of the training process.

Because absolute optimization cannot be guaranteed, there is always a risk that a plaintiff will be able to generate a model with equal or better accuracy that has less of a disparate impact. More generally, it is not feasible for employers to know in advance which machine learning algorithm will be most effective in identifying a globally optimal, nondiscriminatory model, or to test every conceivable type of algorithm to discover which one provides the most accurate predictions. If a plaintiff develops or identifies during litigation a better performing algorithm that employer had not previously considered, it seems reasonable for a court to order the tool to be modified, going forward; but it would be punitive for the court to provide retrospective relief based on an algorithm of which the employer had not previously been aware.

The legal system could use the fact that the outputs of algorithmic tools are the result of the mechanistic application of mathematical optimization techniques to greatly simplify the less discriminatory alternative legal standards. The essential decisions in designing an algorithmic employee selection tool are the selection and weighting of the criteria identified in the job analysis. If the criteria are properly identified and incorporated into a single target variable on which to optimize, the algorithmic procedure will find a model that is at least locally optimal. Consequently, and assuming the employer selected and weighted the target criteria properly,²⁸³ a plaintiff's later identification of a model with equal or better accuracy and less disparate impact should not, by itself, suffice to defeat an employer's assertion of business necessity. Instead, the employee should only be able to prevail in the face of an otherwise-valid selection procedure if the employer actually

283. However, if the criteria were not weighted or selected properly, then the test itself is invalid—i.e., the employer would have failed to establish the business necessity defense, rendering the third stage of the analysis moot.

considered and rejected an alternative algorithm that would have resulted in a reasonable reduction in adverse impact and, to paraphrase *Albemarle Paper*, would still have served the employer's legitimate objectives in selecting well-qualified candidates for a particular position.²⁸⁴ Otherwise, the discovery of a more optimal and less discriminatory model should only bind the employer prospectively.

To defeat an employer's job relatedness defense, then, a plaintiff should be required to demonstrate that the following:

- (1) The employer considered and rejected an alternative modeling methodology;
- (2) The alternative modeling methodology would have served the employer's legitimate interests in selecting suitable candidates for a particular position;
- (3) The alternative methodology would have resulted in a reasonable improvement in the selection rate for the plaintiff's protected class; and
- (4) The alternative modeling methodology would not have unreasonably lowered the employer's ability to select the best candidates for the particular position in comparison to the modeling methodology that the employer ultimately chose.

Note that this test does not require the employer to identify the globally optimal model. Given the impracticality of searching for a global optimum for large high-dimensionality data sets, employers should not be penalized for failing to do so. If there is evidence that an employer, with intent to discriminate, selected or manipulated the initial conditions or consciously chose not to search for a global optimum where it would have been reasonable to do so, then the plaintiff would have a claim for disparate treatment. But an employer should not be subject to disparate impact liability for a valid test constructed using well-established optimization methods simply because the plaintiff chances onto a more accurate and less discriminatory model later.

Components (3) and (4) of this analysis are intertwined and, as with the proposed test for prima facie disparate impact, eschew tests of significance in favor of tests of reasonableness. An improvement in adverse impact is more likely to be reasonable if adopting the alternative methodology would have caused little or no reduction in model performance, and a reduction in model

284. *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425 (1975) (citing *McDonnell Douglas Corp. v. Green*, 411 U.S. 792, 801 (1973)).

performance is more likely to be deemed unreasonable if the alternative modeling methodology would have only slightly improved selection rates for an adversely impacted group. The crux of these final two elements of the employee's rebuttal is demonstrating that the employer's rejection of the alternative modeling methodology was objectively unreasonable.²⁸⁵

C. Disparate Treatment

The Supreme Court's decision in *Ricci* greatly curtailed employers' ability to use race-conscious methods to mitigate statistical imbalances and, despite Justice Ginsburg's prediction that the *Ricci* decision would lack staying power,²⁸⁶ the current composition of the Court makes it unlikely that the decision will be overturned in the foreseeable future. Employers seeking to remove potential biases from algorithmic selection tools must do so while complying with *Ricci*'s strictures. Fortunately, employers could use machine learning and statistical techniques to prevent disparate impacts from arising while remaining within the boundaries set by *Ricci*. These methods could serve as safe harbors for employers seeking to correct disparate impacts in algorithmic selection procedures without running afoul of the prohibition against disparate treatment.

One strategy would be to engage in new forms of differential validation; that is, ensuring that a selection procedure has validity not only within the

285. Even if a group of plaintiffs makes this showing, however, a largely unresolved question in the context of disparate impact discrimination is how to calculate damages in cases involving competitive selection processes. Say, for example, that 50 white and 50 Hispanic candidates apply for 10 open positions with a company, and the company, based on an algorithmic selection procedure using facially neutral criteria, selects white candidates for all 10 positions. Even if the Hispanic candidates brought suit and established that the selection procedure violated Title VII's disparate impact prohibition, it would have been impossible for all 50 Hispanic candidates to be hired into just 10 open positions. The appropriate method of calculation of damages in such cases remains an open legal question. Judge Posner suggested in *Doll v. Brown* that courts use the probabilistic, tort-based recovery theory of "loss of a chance" as a basis for calculating damages in discrimination cases involving competitive settings. See *Doll v. Brown*, 75 F.3d 1200, 1205–06 (7th Cir. 1996). While *Doll* was a disparate treatment case, its "loss of a chance" logic has been generalized to disparate impact cases where the number of open positions is less than the number of affected plaintiffs. See, e.g., *Howe v. City of Akron*, 801 F.3d 718, 752 (6th Cir. 2015) (remanding disparate impact for further proceedings to determine whether loss of a chance is an appropriate method of calculating back pay); *Biondo v. Chicago*, 382 F.3d 680, 688 (7th Cir. 2004) (approving use of lost chance calculation in a disparate impact case, and stating "[i]f four people competing for one position lost an equal chance to get it, then each should receive 25% of the benefits available"). A full exploration of this issue is beyond the scope of this Article, but the applicability of "loss of a chance" theory in disparate impact cases appears to remain underexplored by courts.

286. See *Ricci v. DeStefano*, 557 U.S. 557, 609 (2009) (Ginsburg, J., dissenting).

dominant demographic groups, but also across all demographic groups. Deep learning gives employers the ability to differentially validate selection procedures without resorting to the blunt instrument of race norming, a procedure outlawed by the Civil Rights Act of 1991 that previously had been the primary method that employers used for differential validation. Another de-biasing technique involves using adversarial learning to ensure that the outputs of a selection procedure are statistically independent from protected class status.

The final segment of this section addresses what often is developers' first instinct when seeking to design an algorithm that avoids disparate impacts: imposing explicit constraints on the model to ensure that selection rates are roughly equal across protected classes. Imposing such constraints sits less easily, however, with the spirit and text of antidiscrimination laws than eliminating differential validity or enforcing statistical independence.

1. *Safe Harbor 1: Differential Validation*

In the Title VII case law, differential validation—that is, eliminating differential validity from a test²⁸⁷—has been conflated with the practice of race norming, which was the most common method that employers used to correct for differential validity between whites and nonwhites.²⁸⁸ Race norming ordinarily involves using different cutoff scores for members of different subgroups or adjusting test scores so that the highest performing members of one subgroup receive the same final scores as other subgroups. The 1991 amendments to Title VII explicitly prohibited those practices.²⁸⁹

But differential validation is not the same as—and need not involve—norming. At its root, differential validation simply means making sure that a procedure can distinguish between higher and lower performers not only in the majority groups but also within and across all demographic groups of interest. Some methods for correcting differential validity, such as the race norming of test scores, may constitute disparate treatment, but that does not mean that differential validation itself is inconsistent with Title VII. On the contrary, ensuring that a selection procedure is useful for all applicants, not just for applicants in certain groups, is precisely the sort of “removal of

287. See *supra* notes 276-279, and accompanying text.

288. LARSON, *supra* note 84, at § 27.12 (“[T]he concept of ‘differential validation,’ leading to the practice of adding points to or otherwise adjusting the scores of protected group members, has come under attack by scholars.”).

289. 42 U.S.C. § 2000e-2(I) (2012).

artificial, arbitrary, and unnecessary barriers to employment” that *Griggs* recognized as the very crux of Title VII.²⁹⁰

Any employment test that truly measures job performance should be able to clear that bar. The Guidelines require that subjects of a criterion-related validation study be representative of the relevant labor market precisely because unrepresentative samples can result in a test that does not accurately measure competence for the actual applicant pool.²⁹¹ And the Guidelines specifically warn of the need to check for bias and relevance when a criterion results in “significant differences in measures of job performance for different groups.”²⁹²

Differential validation need not involve norming if it is done as part of the test-design process, rather than as a post hoc adjustment to test scores. But incorporating differential validation into the test-design process would have been impractical for written employment tests and other traditional selection procedures. Eliminating differential validity requires complex analyses of how different groups performed on different proposed components of the test to determine which components should be selected and weighted so the test as a whole is comparably accurate within and across groups. That process would have been exceedingly time consuming and costly with traditional examination-based selection procedures. Thus, employers’ only practical option for eliminating differential validity was to use the blunt instrument of norming at the back end. When the 1991 amendments to Title VII prohibited that practice, it had the practical effect of eliminating employer efforts to engage in differential validation; only three cases even mentioned the terms *differential validation* or *differential validity* since 1991 and none have done so since 2005.²⁹³

But in the context of data-driven selection procedures and with the advent of deep learning, the complexity of differential validation is manageable. Using well-established machine learning techniques, an algorithmic tool could be designed to check different combinations of attributes, test them for

290. *Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971).

291. See 29 C.F.R. § 1607.14(B)(4) (2019) (“[T]he sample subjects should insofar as feasible be representative of the candidates normally available in the relevant labor market for the job or group of jobs in question, and should insofar as feasible include the races, sexes, and ethnic groups normally available in the relevant job market.”); see also *Albemarle Paper Co., v. Moody*, 422 U.S. 405, 435 (1975) (discussing validation study). The Court rejected employer’s validation study in part because “Albemarle’s validation study dealt only with job-experienced, white workers; but the tests themselves are given to new job applicants, who are younger, largely inexperienced, and in many instances nonwhite.” *Id.*

292. 29 C.F.R. § 1607.14(B)(2).

293. The Authors conducted a Westlaw search on January 17, 2019, of all federal cases (“All Federal”) and all state cases (“All States”) for the terms “differential validation” or “differential validity.”

validity within each subgroup, and make minute adjustments to the weights on the components until the model has comparable predictive validity across different protected classes. The resulting model then can serve as the basis for the selection procedure.

For at least two reasons—one temporal and one teleological—this type of differential validation would not run afoul of either the text of Title VII or the Supreme Court’s holding in *Ricci*. From a temporal standpoint, § 703(l) only prohibits “adjust[ing]” scores, “us[ing] different cutoff scores,” or “otherwise alter[ing]” the scores of a test.²⁹⁴ By its own terms, and consistent with *Ricci*’s design versus post-design distinction, this prohibition against norming in § 703(l) applies only to selection procedures whose content has already been determined. Because the whole purpose of algorithmic differential validation is to decide on a scoring system in the first instance, there simply are no scores to adjust.

But perhaps more importantly, the differential validation process has as its objective not achieving equal score performance across protected groups, but rather equal predictive performance—that is, an equal ability to distinguish between high- and low-performing future employees within and across protected groups. The resulting model would not necessarily—or even usually—achieve roughly equal selection rates or test performance across protected groups. It would instead ensure that the model does not give undue weight to characteristics that are only associated with good job performance among certain subgroups. This helps the model focus on characteristics that are tied to the underlying job-related constructs, rather than construct-irrelevant attributes that happen to be more prevalent within specific groups. When conducted using deep learning, differential validation thus is both different in time and in kind from the types of protected class-driven adjustments that Title VII’s race norming prohibition targets.

United States v. City of Erie provides some legal precedent for this distinction.²⁹⁵ At issue was the validity of a physical agility test administered to entry-level candidates to the police department.²⁹⁶ The City had required all candidates to complete seventeen push-ups as part of a broader physical fitness test.²⁹⁷ One of the expert witnesses for the U.S. Department of Justice, which brought the suit, testified that this test suffered from differential validity because “if a man and a woman obtained the same score on the push-ups test,

294. 42 U.S.C. § 2000e-2(l) (2012).

295. See generally *United States v. City of Erie*, 411 F.Supp.2d 524 (W.D. Pa. 2005) (noting that different standards when applied to different genders can function as the same standard in terms of predicted success for a particular job).

296. *Id.* at 524.

297. *Id.* at 532–33.

the woman's predicted job performance would be better than the man's."²⁹⁸ Another expert testified that in physical fitness tests recognized by the American College of Sports Medicine, "the format for women is typically modified, requiring them to push-up from the knees" rather than with their entire body outstretched.²⁹⁹

The City argued that allowing women to pass with a lower number of push-ups would constitute unlawful gender norming, but the Court firmly rejected this argument: "[I]n this circumstance, requiring that men and women complete different numbers of push-ups to pass the test is not 'gender-norming,' and it is not using 'different standards' for males and females. Rather it is using the same standard in terms of predicted success on the job task at issue."³⁰⁰ Machine learning can similarly be used to adjust model parameters so that the output of an algorithmic tool has equal predictive power among different protected groups.

Modern social science recognizes eliminating construct-irrelevant variance between groups—including variance that arises as a result of differential validation or prediction—as an essential part of the validation process.³⁰¹ Differential validation or some other comparable measure of bias control may be especially needed in the context of algorithmic selection tools because of well-recognized problems that machine learning systems encounter when they are used on groups that were underrepresented in the training data. Take the example of Beauty.AI, billed as the world's first AI-judged beauty contest.³⁰² Similar to the gender classification tool used in the Gender Shades study,³⁰³ the Beauty.AI judge was trained on a data set where darker skinned individuals were underrepresented.³⁰⁴ As a result, the system was less accurate in rating photographs of nonwhites and ended up picking winners that were mostly white and, to a lesser extent, Asian³⁰⁵—groups for which the algorithm was more confident in its rating because of their greater representation in the training data.³⁰⁶ Given that at least some protected class groups are certain to be underrepresented in any given data set, differential

298. *Id.* at 560.

299. *Id.* at 549.

300. *Id.* at 560.

301. See PRINCIPLES, *supra* note 81, at 4, 24.

302. Sam Levin, *A Beauty Contest Was Judged by AI & the Robots Didn't Like Dark Skin*, GUARDIAN (Sept. 8, 2016), <https://www.theguardian.com/technology/2016/sep/08/artificial-intelligence-beauty-contest-doesnt-like-black-people> [<https://perma.cc/9SAC-79EJ>].

303. See *supra* text accompanying notes 234–40.

304. Levin, *supra* note 302.

305. See *id.* (noting that out of forty-four winners, nearly all were white, a few were Asian, and only one had dark skin).

306. See *id.* (noting that when a minority group is underrepresented in a data set, "algorithms can reach inaccurate conclusions for those populations").

validation is likely to prove essential for employers seeking to build an unbiased and legally compliant algorithmic selection tool.

Unfortunately, differential validation alone will not eliminate disparate impacts in all cases. Due to inequalities in education and socioeconomic status as well as the effects of prior discrimination, members of different protected groups will sometimes differ in construct-relevant attributes as well. For example, in industries where women have been historically underrepresented, men may have had more opportunities to receive job-relevant training and to gain experience performing high-level tasks. Those attributes may be equally predictive of actual job performance for both men and women, but male applicants would be more likely to have those attributes. That is certainly a type of “artificial, arbitrary, and unnecessary barrier[] to employment”³⁰⁷ that Title VII was intended to eliminate, but differential validation would be of little assistance in eliminating that barrier.

2. *Safe Harbor 2: Statistical Independence*

The concept of statistical independence³⁰⁸ supplies a basis upon which employers seeking to implement algorithmic selection procedures may be able to greatly reduce disparate impacts—whether from construct-relevant or construct-irrelevant sources—without running afoul of the prohibition against disparate treatment. Of course, statistical independence could conceivably be achieved through means that do nothing to advance the employer’s objective of identifying the best candidates for a particular position; randomly assigning test scores to test-takers would result in independence, but it would be useless as the basis for a selection procedure.

The type of statistical independence that could serve as a benchmark for disparate treatment is conditional independence. Two variables x and y are conditionally independent given a third variable z if, once the value of z is known, knowing the value of y provides no additional information about x (and vice versa). For employee selection procedures, therefore, the relevant inquiry would be whether protected class status and the outcome of the selection procedure are independent given the values of the independent variables ultimately used in the procedure.

As with differential validity, designing a traditional examination-based selection procedure to have such conditional independence would be impractical. But modern machine learning techniques provide a potential path through which statistical independence can be achieved algorithmically. Harrison Edwards and Amos Storkey of the University of Edinburgh demonstrated how this can be accomplished through adversarial learning,

307. *Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971).

308. *See supra* note 280 and accompanying text.

which is a machine learning technique in which a digital “adversary”—essentially, a second training algorithm—is programmed to disrupt the performance of the predictor algorithm in some way.³⁰⁹

In Edwards and Storkey’s technique, the adversary is fed a representation of the data on a particular candidate and attempts to predict a sensitive attribute such as race or gender.³¹⁰ If the adversary’s prediction is correct, the predictor algorithm is penalized and the adversary is rewarded.³¹¹ Over many iterations, the predictor algorithm reduces the weights of attributes that carry substantial information about a person’s protected class status, while increasing the weight of attributes that correlate well with the target variables that do not reveal information about protected class status.³¹² Eventually, these adjustments should result in model outcomes that are independent of the sensitive attribute.³¹³

Conditional independence is, in many ways, the antithesis of disparate treatment. If, in the words of *Griggs*, the goal of § 703(a) is to make “race, religion, nationality, and sex become irrelevant,”³¹⁴ that is precisely what conditional independence ensures. True, if characteristics relevant to the job constructs are, in fact, unequally distributed between protected groups, statistical independence may reduce the selection procedure’s overall predictive accuracy. But for many companies looking to leverage the combined power of Big Data and deep learning, that loss in accuracy would likely be a price worth paying if the law recognized statistical independence as providing employers with a safe harbor from disparate treatment liability while still benefiting from the efficiencies and predictive power of algorithmic selection tools.

3. *A Treacherous Harbor: Constrained Optimization*

Correcting differential validity and achieving statistical independence are indirect methods of preventing or correcting disparate impacts because they do not alter selection rates in and of themselves. A more direct—but possibly unlawful—approach to correcting disparate impacts would be constrained

309. Harrison Edwards & Amos Storkey, *Censoring Representations with an Adversary* (Int’l Conf. on Learning Representations Conference Paper) (Mar. 4, 2016), <https://arxiv.org/pdf/1511.05897.pdf> [<https://perma.cc/4X8H-ED8M>].

310. *See id.* at 1, 5.

311. *See id.* at 1–2.

312. *See id.*

313. Edwards and Storkey use the term “fairness” to describe this statistical independence. *See id.* at 1 (“Here, fairness means that the decision is not-dependent on (i.e. marginally independent of) the sensitive variable.”).

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

optimization. In this context, constrained optimization means finding a model that maximizes predictive accuracy (optimization) but limits the search space by requiring the model to satisfy certain conditions (constraints).³¹⁵ In the context of an algorithm designed to avoid disparate impact liability, the constraints could be catered to the rules governing a *prima facie* case of disparate impact. Thus, under the Guidelines' four-fifths rule, the algorithm could find an optimal model subject to the constraint that the selection rate for each protected group can be no lower than 80% of the selection rate of any other group in the same protected category.

Enticingly, *Ricci* rejected the proposition that employers may not take disparate impact into account—even where this means being race conscious—when designing a selection procedure to ensure that the procedure provides a fair chance for all individuals.³¹⁶ *Ricci* thus suggests that there is a distinction between designing a selection procedure in a way that checks for and mitigates bias on one hand, and post-design test score adjustments and conscious decisions to incorporate protected class preferences into a model on the other.³¹⁷ Technically, using constrained optimization during the design phase would be consistent with this principle.

That said, the fact that constrained optimization explicitly examines and makes adjustments based on the selection rates of different groups distinguishes it from the approaches geared toward achieving differential validation and statistical independence. Differential validation aims to ensure that the metrics have comparable accuracy across protected groups, and statistical independence ensures that the selection tool is not encoding protected class information as part of its model. Any effect on selection rates is a beneficial side effect of these techniques, rather than their objective. A constrained optimization approach, by contrast, makes equalizing group selection rates a direct and explicit goal.

If Title VII could be viewed through a purely antisubordination lens, such a direct approach would be unproblematic. But it runs contrary to the law's anticlassification strictures. True, an employer may not know beforehand which precise groups will see their selection rates improve with constrained

315. See Introduction to Constrained Optimization in the Wolfram Language, Wolfram Language & System Documentation Center, <https://reference.wolfram.com/language/tutorial/ConstrainedOptimizationIntroduction.html> (defining constrained optimization as a problem for which a function must be minimized or maximized subject to constraints).

316. *Ricci v. DeStefano*, 557 U.S. 557, 585 (2009) (“Title VII does not prohibit an employer from considering, before administering a test or practice, how to design that test or practice in order to provide a fair opportunity for all individuals, regardless of their race.”). Because design-stage modifications were not at issue in *Ricci* and the Court’s articulation of the design versus postdesign distinction was not essential to its holding, this statement from *Ricci* is arguably dictum, but at least one court has relied on it as precedent. See *Maraschiello v. City of Buffalo Police Dep’t*, 709 F.3d 87, 95 (2d Cir. 2013).

317. See *Ricci*, 557 U.S. at 584–85.

optimization as compared to an unconstrained modeling approach. But that logic would not save an employer that normalizes test scores after administration of a traditional employment test; § 703(l)'s prohibition against norming prohibits all score adjustments made "on the basis of" a protected classification without regard to whether the employer knew in advance which groups would benefit from such norming. Even at the design stage, making adjustments explicitly based on protected characteristics sits uneasily with § 703(a)'s broad prohibition against employment decisions made because of such characteristics.

The Authors would encourage courts and agencies to adopt *Ricci*'s design versus post-design distinction, thus giving employers the freedom to fashion unbiased algorithmic selection tools without risking disparate treatment liability. Indeed, the only published case interpreting this passage from *Ricci* relied on this distinction to affirm a grant of summary judgment in favor of an employer that may have been "motivated in part by its desire to achieve more racially balanced results" when it adopted a new employment test.³¹⁸ This rule is fully consistent with the objectives of Title VII, as elucidated in *Griggs* and *Albemarle Paper*. But a rule permitting constrained optimization will likely face greater resistance than one permitting employers to use more sophisticated machine learning approaches that avoid such explicit reliance on protected class status.

V. CONCLUSION

For now, the above proposal is just that—a proposal, albeit one strongly rooted in the text of Title VII and the case law interpreting it. At this point, employers considering implementing AI-powered recruitment and hiring at scale simply do not know how a court would analyze an algorithmic selection procedure under Title VII. That is one reason the EEOC should act quickly to clarify the legal standards by which it will assess algorithmic selection procedures. Employers will undoubtedly be wary of developing (or at least implementing) such procedures in the meantime.

The framework offers two routes by which employers can avoid liability for the inevitable adverse impacts that algorithmic selection tools will generate: (1) correcting any disparate impacts by using one of the disparate treatment safe harbors; or (2) satisfying the business necessity defense by conducting a proper job analysis followed by criterion validation. That presents a conundrum for employers wishing to use algorithmic selection procedures today. The most efficient and practical way to achieve legal compliance under the above proposal—using one of the proposed disparate

318. *Maraschiello*, 709 F.3d at 95.

treatment safe harbors—is the path that carries the greatest legal uncertainty because using algorithms that use machine learning to eliminate differential validation or achieve statistical independence have never been tested in court. Conversely, the path to compliance that would provide the greatest legal certainty—following the validation standards described in the framework, which adhere closely to the Guidelines and existing case law—may be neither efficient nor practical, and may prove to be a wasted effort in any event if the Guidelines receive a long-overdue overhaul to bring them in line with the modern social science.

The difficulty of validation is partially just a function of the difficulty of validating large high-dimensionality data sets. But perhaps even more fundamentally, employers may find it extremely difficult to build a sufficiently representative set of measurable job behaviors and outcomes to serve as a proper set of validation criteria. Traditional employee selection procedures were actual examinations whose content an employer could cater to the actual skills and knowledge relevant to job performance. Algorithmic selection tools, by contrast, generally rely primarily on passive analysis of historical (and therefore static) data, which often cannot easily be crafted to fit the job functions of a particular position. Many employers have only very general or subjective measures of job performance available, such as tenure or impressionistic supervisor ratings, which courts have historically disfavored for purposes of validation when examining an employer's proposed business necessity defense.

These factors, coupled with the availability of data sets containing hundreds or thousands of attributes, will make it increasingly difficult for employers to validate employee selection tools in accordance with the Guidelines. That underscores the need for policymakers and courts to both adopt new standards for validation and establish clear safe harbors that allow employers to prevent disparate impacts from arising without exposing themselves to disparate treatment liability. Modifying the traditional framework by eliminating tests of statistical significance and replacing them with reasonableness standards is also necessary to avoid missing opportunities to materially improve the diversity and inclusiveness of an employer's workforce with minimal sacrifice in quality.

Companies today are leveraging algorithmic tools powered by machine learning and built on Big Data to enhance every aspect of their business activities. Despite the fact that algorithmic tools offer employers a vehicle for more effective and inclusive HR selection decisions, with less discrimination and more long-term accountability, the use of such tools to improve recruitment, hiring, and other human resources decisions has lagged behind their use in other business operations. The levee eventually must break, and legal standards will have to evolve quickly to stem the tide. For now, courts,

agencies, and employers alike must be attuned to the growing mismatch between the state of technology and existing legal standards so that the promise of these technologies is not squandered.