



PROGRAM MATERIALS

Program #36225

October 13, 2026

AI Governance Basics

Copyright ©2026 by

- **Elizabeth Elices, Esq. - Caldwell University**

All Rights Reserved.

Licensed to Celesq®, Inc.

Celesq® AttorneysEd Center

www.celesq.com

5301 North Federal Highway, Suite 150, Boca Raton, FL 33487

Phone 561-241-1919

AI Governance Basics

Elizabeth Elices, Esq
October, 2026

Speaker Bio

Elizabeth Elices is an attorney licensed in New Jersey and New York. She currently serves as General Counsel for Caldwell University where she works on a wide range of issues including data privacy, employment matters, and contract law.

She is former judicial clerk and author of “Citizens United and the future of FCC Content Regulation” in the Hastings Communications and Entertainment Law Journal (Comm/Ent). She can be found online at

<https://www.linkedin.com/in/elizabeth-elices-esq-72baa114/> and
Elices.Esq@gmail.com

Anti-Bias Legislation and frameworks

How are different jurisdictions approaching the risks of AI in the context of existing anti-discrimination laws?

- NYC Automated Decision Tool law
- NJ Office of the Attorney General : Guidance on Algorithmic Discrimination
- NAIC AIS Governance Guidelines
- Connecticut CART Act
- Colorado AI Act

NYC Automated Decision Tool law (2021)

Governs software that:

“Uses machine learning, statistical modeling, data analytics, or artificial intelligence

AND

Helps employers and employment agencies make employment decisions.

AND

Substantially assists or replaces discretionary decision-making.”

NJ Office of the Attorney General : Guidance on Algorithmic Discrimination (2025)

Framework for use in assessing the use of AI in the context of potential discrimination

NAIC AIS Governance Guidelines

- “FACTS” approach
 - Fair and ethical (does it discriminate?)
 - Accountable (governance)
 - Compliant
 - Transparent (disclose data sources)
 - Secure and Safe

Connecticut CART Act (2026)

- Mandatory disclosures related to the use of automated employment-related decision technology

Colorado AI Act (2026)

- Risk-based approach
- “High risk artificial intelligence systems”
- Consequential decisions

Privacy

- Online Privacy Overview: key considerations
 - Best Practices in data collection online
 - COPPA: purpose and requirements
 - Emerging issues in social media regulation
 - Minors and AI
- GDPR considerations
 - The right to access personal data
 - The right to rectification (correction) of inaccurate data
 - The right to erasure (to be forgotten) under certain circumstances
 - The right to restrict processing of data
 - The right to data portability (to move data between service providers)

Privacy

- Biometric Privacy Regulations : faces, fingerprints, and the future
 - State examples: Illinois, California, Colorado, and Montana

Deepfakes

- What is a “deepfake”?
 - Synthetic images, sounds, and videos
- Federal TAKE IT DOWN Act
 - “Tools to Address Known Exploitation by Immobilizing Technological Deepfakes on Websites and Networks Act” (2025)
- State laws against online impersonation

Legislation and Emerging Issues

- Transparency as a focus
- California and Texas, compared
- Healthcare and other high-risk uses
- White House 2025 and 2026 AI Executive Orders

Ethics Considerations for Lawyers

- ABA and Model Rules 1.1, 1.4, 1.5, and 1.6
- Confidentiality
- Duty of Competency
- Appropriate Fees
- Accountability

Conclusion & CLE Credits

Presidential Documents

Title 3—

Executive Order 14409 of June 2, 2026

The President

Promoting Advanced Artificial Intelligence Innovation and Security

By the authority vested in me as President by the Constitution and the laws of the United States of America, it is hereby ordered:

Section 1. Purpose. The United States continues to lead the world in Artificial Intelligence (AI) because of the enormous talent and innovation of our AI industry, and because we refuse to stifle this innovation with overly burdensome regulation. My Administration has unleashed tremendous technological growth and economic investment in AI by slashing the bureaucratic constraints that the prior administration placed on America's AI developers and researchers, and by instead encouraging AI innovation and accelerating responsible AI adoption across government and industry.

Advanced AI capabilities make our Nation stronger, but also introduce new national security considerations that require coordinated action across executive departments and agencies (agencies), and components. As these capabilities evolve, my Administration will continue to work closely with industry to ensure that the best and most secure technology is deployed rapidly to confront any and all threats to our country. We will continue to lead an America First cybersecurity effort that enhances both our national security and our global AI dominance.

It is the policy of the United States to promote AI innovation and security by working collaboratively with the private sector to modernize government and private sector information systems and harden them against external threats; to protect American ingenuity and intellectual property from exploitation and theft by adversaries; and to cultivate America's advanced AI-enabled capabilities.

Sec. 2. Upgrading American Systems for Advanced AI. (a) Within 30 days of the date of this order, the Committee on National Security Systems shall prioritize the cyber defense of National Security Systems, as defined in 44 U.S.C. 3552(b)(6)(A), by taking appropriate and expeditious action consistent with the purpose of this order.

(b) Within 30 days of the date of this order, the Secretary of War shall prioritize the cyber defense of Department of War information systems by taking appropriate and expeditious action consistent with the purpose of this order.

(c) Within 30 days of the date of this order, the Secretary of Homeland Security, through the Director of the Cybersecurity and Infrastructure Security Agency (CISA), in consultation with the Director of the Office of Management and Budget (OMB), the Assistant to the President for National Security Affairs, and the National Cyber Director, shall release Binding Operational Directives and other guidance as appropriate to:

- (i) expedite and prioritize the cyber defense of civilian Federal Government information systems in order to protect our Nation's vital functions;
- (ii) establish or expand Federal programs and cybersecurity services that enhance AI-enabled defensive tools; and
- (iii) facilitate access to cybersecurity tools and services including, where appropriate, covered frontier models for agencies, State and local authorities, and operators of critical infrastructure such as rural hospitals, community banks, and local utilities.

(d) Within 30 days of the date of this order, the Secretary of the Treasury, in consultation with the National Cyber Director, the Secretary of War, through the Director of the National Security Agency (NSA), and the Secretary of Homeland Security, through the Director of CISA, shall form an AI cybersecurity clearinghouse, in voluntary collaboration with the AI industry and operators of critical infrastructure, that coordinates and deconflicts scanning for software vulnerabilities, discovers and validates such vulnerabilities, and coordinates and prioritizes remediation and distribution of vulnerability patches.

(e) Within 30 days of the date of this order, the Director of OMB, in coordination with the National Cyber Director and the Director of CISA, shall determine whether any Federal grant programs have available and relevant funding that can be directed toward applicants developing advanced AI vulnerability detection.

(f) Within 60 days of the date of this order, the Director of the Office of Personnel Management shall expand the United States Tech Force Information Cybersecurity Specialist hiring and placement pathways.

Sec. 3. *Secure Frontier Model Deployment.* Within 60 days of the date of this order, the Secretary of the Treasury, the Secretary of War, through the Director of NSA, and the Secretary of Homeland Security, through the Director of CISA, in consultation with the White House Chief of Staff, through the National Cyber Director, the Assistant to the President for Science and Technology (APST), and the Secretary of Commerce, through the Director of the National Institute of Standards and Technology, and in coordination with other agencies, as appropriate, shall:

(a) develop and maintain a classified benchmarking process to assess the advanced cyber capabilities of AI models and determine the threshold at which an AI model should be designated a “covered frontier model” for the purposes of this order, sharing such assessments with AI developers and researchers as appropriate. Such a determination shall be made by the Director of NSA, in consultation with the National Cyber Director, the APST, the Director of CISA, and other representatives of the Department of War, as appropriate.

(b) design a voluntary framework with AI developers through which developers would be able to:

(i) engage the Federal Government to determine whether model(s) under development meet the designation of “covered frontier model”;

(ii) provide the Federal Government with access to covered frontier models, subject to appropriate confidentiality, cybersecurity, insider-risk, and intellectual-property protection, use, and nondisclosure requirements, for a period of up to 30 days before they plan to release such models to other trusted partners; and


(iii) collaborate with the Federal Government to select trusted partners that will have early access to covered frontier models to promote secure innovation and strengthen the cybersecurity of critical infrastructure.

(c) Nothing in this section shall be construed to authorize the creation of a mandatory governmental licensing, preclearance, or permitting requirement for the development, publication, release, or distribution of new AI models, including frontier models.

Sec. 4. *Protection Against Criminal Actors.* The Attorney General shall prioritize the enforcement of 18 U.S.C. 1028, 18 U.S.C. 1030, 18 U.S.C. 1343, and all other applicable Federal criminal laws against anyone who utilizes AI to illegally access or damage a computer without authorization, or who utilizes AI while engaged in such illegal access to further any other crime. This includes breaching any public or private information technology system, or employing AI agents to unlawfully access data or information that is subsequently used for a criminal or unlawful purpose.

Sec. 5. *General Provisions.* (a) Nothing in this order shall be construed to impair or otherwise affect:

- (i) the authority granted by law to an executive department or agency, or the head thereof; or
- (ii) the functions of the Director of the Office of Management and Budget relating to budgetary, administrative, or legislative proposals.
- (b) This order shall be implemented consistent with applicable law and subject to the availability of appropriations.
- (c) This order is not intended to, and does not, create any right or benefit, substantive or procedural, enforceable at law or in equity by any party against the United States, its departments, agencies, or entities, its officers, employees, or agents, or any other person.
- (d) The costs for publication of this order shall be borne by the Department of War.

A handwritten signature in black ink, appearing to be a stylized name with a long, sweeping horizontal stroke at the bottom.

THE WHITE HOUSE,
June 2, 2026.



Guidance on Algorithmic Discrimination and the New Jersey Law Against Discrimination

January 2025

The New Jersey Office of the Attorney General and the Division on Civil Rights (DCR) issue this guidance to clarify how the New Jersey Law Against Discrimination (LAD) applies to algorithmic discrimination resulting from the use of new and emerging data-driven technologies, such as artificial intelligence (AI), by employers, housing providers, places of public accommodation, and other entities covered by the LAD.¹

In recent years, businesses and governments across the country have started using automated tools to make decisions that affect key aspects of our lives—who is hired or receives a promotion, who is selected for an apartment rental or obtains a mortgage to purchase a house, or who receives certain medical treatments or obtains insurance coverage for that treatment. These automated decision-making tools are reshaping modern society. These tools carry potential benefits for regulated entities and the public. But they also carry risks. If these tools are not designed and deployed responsibly, they can result in algorithmic discrimination. Algorithmic discrimination is discrimination that results from the use of automated decision-making tools.

While the technology powering automated decision-making tools may be new, the LAD applies to algorithmic discrimination in the same way it has long applied to other discriminatory conduct. As further discussed below, in New Jersey the LAD prohibits algorithmic discrimination in employment, housing, places of public accommodation, credit, and contracting on the basis of actual or perceived race, religion, color, national origin, sexual orientation, pregnancy, breastfeeding, sex, gender identity, gender expression, disability, and other protected characteristics. A covered entity—that is, an entity subject to the LAD’s requirements—that

¹ The purpose of this guidance document is to clarify and explain DCR’s understanding of existing legal requirements in order to facilitate compliance with the LAD. This guidance document does not impose any new or additional requirements that are not included in the LAD, does not establish any rights or obligations for any person, and will not be enforced by DCR as a substitute for enforcement of the LAD. This document does not provide legal advice and should not be treated as providing legal advice. Employees, employers, tenants, housing providers, places of public accommodation, and others with questions about the LAD are encouraged to speak with a qualified attorney to address their specific questions. Moreover, the examples discussed in this guidance are provided only for informational purposes to illustrate legal theories and concepts. Because claims of discrimination under the LAD require a fact-intensive inquiry based on the specific facts at issue, this guidance document does not purport to pre-judge the lawfulness of any particular factual scenario or any particular application of the automated decision-making tools discussed herein. The Law Against Discrimination, N.J.S.A. § 10:5-1 et seq., is available on the Division on Civil Rights’ (DCR) website at <https://www.njoag.gov/wp-content/uploads/2021/02/NJ-Law-Against-Discrimination.pdf>.

engages in algorithmic discrimination may be held liable for violating the LAD, even if the covered entity uses a tool it did not develop.

In light of the extensive use of automated decision-making tools in modern society, it is critical that all New Jerseyans understand what these tools are, how they are being used, and the risks and benefits associated with their use. To that end, this guidance provides an overview of automated decision-making tools, explores the risk of algorithmic discrimination posed by the use of these tools and their potential benefits, and outlines the LAD’s protections against algorithmic discrimination.

I. AUTOMATED DECISION-MAKING TOOLS AND ALGORITHMIC DISCRIMINATION

A. Automated Decision-Making Tools

The term “automated decision-making tool” refers to any technological tool, including but not limited to, a software tool, system, or process that is used to automate all or part of the human decision-making process. Automated decision-making tools may incorporate technologies at various levels of advancement and sophistication, such as generative AI, machine-learning models, traditional statistical tools, and decision trees.²

Automated decision-making tools are now used across a wide range of contexts. They are regularly used to help determine who sees a job advertisement, whether a human reviewer reads an applicant’s resume, whether an applicant is hired, whether an employee receives positive reviews or is promoted, and whether an employee is demoted or fired.³ Automated decision-making tools can determine whether someone is approved to rent an apartment or buy a home, what the terms of their mortgage are, the location of the homes they

A [recent survey](#) by Rutgers University revealed that New Jersey employers are increasingly using AI-enabled tools in hiring. Of employers surveyed, 63% use one or more tools to recruit job applicants and/or make hiring decisions. Nonetheless, 47% of surveyed employers agreed that using AI-enabled tools in hiring might lead to unfair outcomes.

Liu, et al., *Shaping the Future of Recruitment: A Survey on AI-enabled Hiring Tools*, 4 & 12 (2024).

² Not all automated decision-making tools are used in a way that causes algorithmic discrimination. The use of automated decision-making tools implicates the LAD only when that use results in unlawful discrimination. While different types of tools may pose different risks of algorithmic discrimination, a covered entity is liable for any unlawful discrimination that results from its use of any tool.

³ See generally U.S. Equal Emp. Opportunity Comm’n, “Select Issues: Assessing Adverse Impact in Software, Algorithms, and Artificial Intelligence Used in Employment Selection Procedures Under Title VII of the Civil Rights Act of 1964” (“EEOC Title VII Guidance”) (May 18, 2023), https://www.eeoc.gov/laws/guidance/select-issues-assessing-adverse-impact-software-algorithms-and-artificial#_edn7.

are shown, and whether they can tour specific homes.⁴ Schools use automated decision-making tools to monitor students' computer use, filter and block content available to students, aid in disciplinary decisions, predict students' likelihood of encountering the criminal justice system, and predict students' likelihood of graduation.⁵ In health care, automated decision-making tools help providers make diagnoses, decide the type of care a patient receives, allocate resources between patients, determine patients' health risks, and decide which medications to prescribe.⁶ Law enforcement agencies may use automated decision-making tools to identify suspects, read license plates, detect gunshots, and select areas to patrol,⁷ and courts may use automated decision-making tools to predict whether someone is likely to commit another crime, set bail, or decide if someone is eligible for pretrial release.⁸

Many automated decision-making tools accomplish their aims by using algorithms, or sets of instructions that can be followed, typically by a computer, to achieve a desired outcome. These algorithms analyze data, uncover correlations within that data, and make predictions, recommendations, or generate new data based on those correlations.⁹ In so doing, however, automated decision-making tools can create classes of individuals who will be either advantaged or disadvantaged in ways that may exclude or burden them based on their protected

⁴ See generally U.S. Dept. of Housing and Urban Dev., “Guidance on Application of the Fair Housing Act to the Screening of Applicants for Rental Housing” (“HUD ADM Housing Rental Screening Guidance”), at 1–4 (Apr. 29, 2024),

https://www.hud.gov/sites/dfiles/FHEO/documents/FHEO_Guidance_on_Screening_of_Applicants_for_Rental_Housing.pdf; U.S. Dept. of Housing and Urban Dev., “Guidance on Application of the Fair Housing Act to the Advertising of Housing, Credit, and Other Real Estate-Related Transactions through Digital Platforms” (“HUD ADM Advertising Guidance”), at 4–10 (Apr. 29, 2024), https://www.hud.gov/sites/dfiles/FHEO/documents/FHEO_Guidance_on_Advertising_through_Digital_Platforms.pdf.

⁵ See generally Ctr. for Democracy and Tech., *Off Task: EdTech Threats to Student Privacy and Equity in the Age of AI* (Sept. 2023), <https://cdt.org/wp-content/uploads/2023/09/091923-CDT-Off-Task-web.pdf>; U.S. Dept. of Ed. Office for Civil Rights, “Avoiding the Discriminatory Use of Artificial Intelligence” (Nov. 2024), <https://www.ed.gov/media/document/avoiding-discriminatory-use-of-ai>.

⁶ See e.g., Crystal Grant, “Algorithms Are Making Decisions About Health Care, Which May Only Worsen Medical Racism,” Am. Civil Liberties Union, (Oct. 2, 2022), <https://www.aclu.org/news/privacy-technology/algorithms-in-health-care-may-worsen-medical-racism>; Darshali A. Vyas, et al., *Hidden in Plain Sight—Reconsidering the Use of Race Correlation in Clinical Algorithms*, 9 N. ENGL. J. MED. 874–882 (Aug. 27, 2020).

⁷ See e.g., Cong. Research Serv., “Law Enforcement Use of Artificial Intelligence and Directives in the 2023 Executive Order” (Dec. 15, 2023), <https://crsreports.congress.gov/product/pdf/IN/IN12289>.

⁸ See e.g., Alexandra Chouldechova & Kristian Lum, *The Present and Future of AI in Pre-Trial Risk Assessment Instruments* (June 2020), https://www.ncsc.org/data/assets/pdf_file/0019/52516/AI-in-Pre-Trial-Risk-Assessment-Brief-June-2020-R2.pdf.

⁹ This generally refers to machine-learning models, which can in some cases constitute automated decision-making tools. See David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 671 (2017). Some machine-learning models are developed with the primary goal of assessing existing data. In comparison, “[g]enerative AI can be thought of as a machine-learning model that is trained to create new data, rather than making a prediction about a specific dataset.” Adam Zewe, “Explained: Generative AI,” MIT News, (Nov. 9, 2023), <https://news.mit.edu/2023/explained-generative-ai-1109>.

characteristics.¹⁰ Because automated decision-making tools use algorithms, discrimination stemming from the use of these tools is commonly referred to as “algorithmic discrimination.”

Some recent examples illustrate how automated decision-making tools can find and leverage correlations in the datasets they analyze in ways that may contribute to or amplify discriminatory outcomes:¹¹

- A 2019 study found that the use of an automated decision-making tool by hospitals to make recommendations about how to distribute care among patients resulted in discrimination based on race by underestimating the health risks of Black patients and, accordingly, under-recommending care for them.¹² According to the study, the tool was designed to predict a patient’s medical care needs by assigning each patient a health risk score; patients who scored above specified percentiles were automatically enrolled in or referred for consideration of hospital programs that provided access to more intensive resources to manage complex medical problems. However, the algorithm powering the automated decision-making tool was trained on medical claim insurance data (which did not explicitly include information on race) specifically to calculate a patient’s expected total medical expenditures for the next year, which were used by the algorithm as the proxy for health risk. That design choice was flawed because the model failed to account for historical disparities in access to healthcare for Black patients, which resulted in Black patients with high levels of medical need spending less than their white counterparts. As a result, the automated decision-making tool underestimated the health risks of Black patients and, accordingly, under-recommended care for them, thus amplifying racial bias in access to health care. Notably, after the study was published, the researchers worked with the developer of the tool to address this issue and rectify some of the biased outcomes stemming from the tool.¹³
- A 2024 study found that the use of an automated decision-making tool by employers to screen resumes and make hiring decisions resulted in discrimination based on race and gender.¹⁴ According to the study, the tool discovered correlations between protected classes and specific job types. In making its predictions, the tool assumed applicants’ race and gender based on their names and reproduced those correlations, inadvertently ranking equally-qualified applicants unequally for particular jobs based on their names. For example, the study found that the tool ranked Hispanic women more favorably for a

¹⁰ See U.S. Consumer Fin. Protection Bureau, et al., “Joint Statement on Enforcement of Civil Rights, Fair Competition, Consumer Protection, and Equal Opportunity Laws in Automated Systems” (hereinafter “2024 Joint Statement”) (Apr. 4, 2024) <https://www.justice.gov/crt/media/1346821/dl?inline>.

¹¹ As indicated above, the examples discussed in this guidance document are provided only for informational purposes and are not examples of per se violations of the LAD. See *supra* note 1.

¹² Ziad Obermeyer, et al., *Dissecting racial bias in an algorithm used to manage the health of populations*, 366 SCIENCE 447 (2019).

¹³ Starre Vartan, “Racial Bias Found in Major Health Care Risk Algorithm” (Oct. 24, 2019), <https://www.scientificamerican.com/article/racial-bias-found-in-a-major-health-care-risk-algorithm/> (quoting Ziad Obermeyer’s statement that the point of the paper was “to point out where the problem lies *and how to fix it*”) (emphasis added).

¹⁴ Leon Yin, et al., “OpenAI’s GPT is a Recruiter’s Dream Tool. Tests Show There’s Racial Bias,” Bloomberg (Mar. 7, 2024).

position as a human resources specialist than white men and ranked Asian women as the top candidates for a financial analyst role more than twice as often as Black men based on their names.

Not all automated decision-making tools are used in a way that results in algorithmic discrimination. Indeed, some automated decision-making tools offer the promise of reducing bias and discrimination.¹⁵ For example, in the housing context, researchers at Stanford University have developed an automated decision-making tool that can identify property records with racially restrictive language faster and more accurately than human searchers, enabling the removal of these racially restrictive covenants.¹⁶ In the credit context, in a joint letter to the U.S. Consumer Financial Protection Bureau, a group of nonprofit consumer advocates and for-profit fintech companies recently explained that automated decision-making tools can be used to assess non-traditional credit factors to improve approval odds for consumers with low traditional credit scores, increasing access to credit for marginalized communities and reducing discriminatory lending outcomes.¹⁷ And researchers at the University of Bath found that an automated decision-making tool used to optimize salesforce commissions on automobile loan decisions could be modified to prevent gender-based discrimination while still increasing overall profits.¹⁸

Ultimately, depending on how they are developed and deployed, automated decision-making tools can either contribute to or reduce the likelihood of discriminatory outcomes. If covered entities fail to account for or remedy bias in automated decision-making tools, the tools may contribute to discriminatory outcomes. But if covered entities incorporate fairness considerations in their use of automated decision-making tools, the tools may reduce discriminatory outcomes.¹⁹ It is, therefore, critical that covered entities understand the potential risks and benefits of their use of automated decision-making tools.

¹⁵ Mike Field, *Algorithms for a Fairer World*, Johns Hopkins Whiting Sch. of Eng'g Mag. (Spring 2024), <https://engineering.jhu.edu/magazine/2024/06/algorithms-for-a-fairer-world/>; see also New Jersey Artificial Intelligence Task Force, *2024 Report to the Governor*, (Nov. 2024), <https://innovation.nj.gov/news/NJ-AI-Task-Force-Report.pdf>.

¹⁶ Kaustuv Basu, “AI Seeks Out Racist Language in Property Deeds for Termination,” Bloomberg Law, (Oct. 17, 2024), <https://news.bloomberglaw.com/artificial-intelligence/ai-seeks-out-racist-language-in-property-deeds-for-termination>.

¹⁷ Nat'l Cmty Reinvestment Coal., et al., “Opportunities for the CFPB and FHFA to Reduce Discrimination in Financial Services, As Encouraged in the White House Executive Order on Artificial Intelligence,” (Sept. 25, 2024), <https://ncrc.org/wp-content/uploads/2024/09/NCRC-Innovation-Council-Joint-Letter-on-Fair-Lending-and-the-Executive-Order-on-AI-v2.pdf>.

¹⁸ Christopher Amaral, et al., *Optimizing Pricing Delegation to External Sales Forces via Commissions: An Empirical Investigation*, 33 (9) POMS 1839–1854 (Sept. 2024).

¹⁹ See generally White House, “Appendix: Examples of Automated Systems, White House Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People,” at 52–53 (Oct. 2022), <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.

B. How Automated Decision-Making Tools Can Contribute to Discriminatory Outcomes

Because the work performed by an automated decision-making tool can be invisible and not well understood, it can be difficult to detect if or how the tool contributes to discriminatory outcomes. Nonetheless, there are some common reasons why using automated decision-making tools may result in discriminatory outcomes. Most often, these discriminatory outcomes stem from the way these tools are designed, trained, or deployed.²⁰

Design: The choices a developer makes in designing an automated decision-making tool can skew the tool, either purposefully or inadvertently. These choices include how the developer translates a real-world problem into something—often numerical—that can be analyzed by an automated decision-making tool.²¹ Decisions regarding the output the tool provides, the model or algorithms the tool uses, and what inputs the tool assesses can introduce bias into the automated decision-making tool.²²

A [recent survey](#) by the Center for Democracy and Technology revealed that “[a]pproximately one-third of teachers agree that content associated with or about LGBTQ+ students and students of color is more likely to be filtered or blocked” by their school than other content.

Ctr. for Democracy and Tech., *Off Task: EdTech Threats to Student Privacy and Equity in the Age of AI*, 21 (2023).

- Automated decision-making tools used by the U.S. Internal Revenue Service (IRS) to select individuals to audit for underreporting their tax liability provide an example of how design choices can inadvertently contribute to a tool’s discriminatory impact.²³ The IRS first used a tool designed to predict the likelihood that someone misreported their tax liability. That tool disproportionately recommended lower-income and Black people for an audit. When the tool was redesigned to instead predict the *amount* of taxes underreported, the tool recommended more higher-income and more non-Black individuals for auditing. Changing the design of the automated decision-making tool—in other words, what the tool should predict and how it makes that prediction—changed the tool’s assessment and

²⁰ See Lehr & Ohm, *supra* note 9 at 701–05; see also 2024 Joint Statement; Harini Suresh & John Guttag, *A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle*, 21 EAAMO 1, 4–6 (Oct. 5–9, 2021) (describing seven sources of harm that can arise as an automated decision-making tool is designed, trained, and deployed, including historical bias, representation bias, measurement bias, aggregation bias, learning bias, evaluation bias, and deployment bias).

²¹ See Lehr & Ohm, *supra* note 9 at 668 (“An analyst must translate an abstract, often ill-defined goal into a highly specified outcome variable to be predicted by the algorithm.”); see also Emily Black, et al., *Less Discriminatory Algorithms*, 113 GEO. L. J. 53, 104 (2024); Nat’l Inst. of Standards and Tech., “Towards a Standard for Identifying and Managing Bias in Artificial Intelligence,” at 12 (Mar. 2022) <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf> (“Representing [] complex human phenomena with mathematical models comes at the cost of disentangling the context necessary for understanding individual and societal impact and contributes to a fallacy of objectivity.”).

²² See Lehr & Ohm, *supra* note 9 at 673–77; see also Suresh & Guttag, *supra* note 20 at 2–3 (described as “learning bias”).

²³ See Black, et al., *supra* note 21 (citing Emily Black et al., *Algorithmic Fairness and Vertical Equity: Income Fairness with IRS Tax Audit Models*, 2022 ACM Conference on Fairness, Accountability, and Transparency, at 1479–1503 (2022) and Hadi Elzayn, et al., *Measuring and Mitigating Racial Disparities in Tax Audits*, Stanford Inst. for Policy Research Working Paper (2023)).

recommendation for each individual and reduced the discriminatory impact of the automated decision-making tool.

- A U.S. Equal Employment Opportunity Commission (EEOC) enforcement action against a tutoring company, iTutorGroup, provides an example of how an automated decision-making tool can be purposely designed to discriminate.²⁴ In its complaint, the EEOC alleged that iTutorGroup used an automated decision-making tool to aid in hiring tutors. According to the EEOC, the tool was programmed to automatically reject women applicants aged 55 or older and men applicants aged 60 or older. The EEOC argued that the automated decision-making tool therefore had been purposely designed to discriminate against applicants based on their age.²⁵ In settling the case, iTutorGroup agreed not to ask for age-related information, such as birth date, from future applicants.

Training: An automated decision-making tool must be “trained” before it is ready for use in the real world. Training occurs by exposing the tool to training data from which the tool learns correlations or rules.²⁶ A developer can either create a new dataset to train an automated decision-making tool or can refine and reformat a pre-existing dataset. The training data may reflect the developer’s own biases, or it may reflect institutional and systemic inequities. The tool can become biased if the training data is skewed or unrepresentative, lacks diversity, reflects historical bias, is disconnected from the context the tool will be deployed in, is artificially generated by another automated decision-making tool, or contains errors.²⁷

- A U.S. Federal Trade Commission (FTC) enforcement action against Rite Aid provides an example of how the data used in developing an automated decision-making tool can potentially contribute to biased outcomes.²⁸ In its complaint, the FTC challenged Rite Aid’s practice of using facial recognition technology to identify customers who were likely to engage in shoplifting or other criminal behavior. According to the FTC, Rite Aid’s automated decision-making tool generated a disproportionate number of false positives, or falsely identified someone as likely to commit a crime, for Black, Latinx/e, Asian, and women consumers. The FTC alleged that Rite Aid failed to consider that facial recognition technologies are often less accurate when used on images of Black or Asian people

²⁴ See U.S. Equal Emp. Opportunity Comm’n Press Release, “iTutorGroup to Pay \$365,000 to Settle EEOC Discriminatory Hiring Suit” (Sept. 11, 2023), <https://www.eeoc.gov/newsroom/itutorgroup-pay-365000-settle-eeoc-discriminatory-hiring-suit>.

²⁵ Because the tool was programmed to automatically reject women applicants aged 55-59 but not men within that age range, the EEOC’s allegations suggest that iTutorGroup also may have discriminated based on gender.

²⁶ Different types of automated decision-making tools require different amounts and types of training data. For example, unlike traditional machine-learning models which can be trained on smaller datasets that are structured—or labelled and categorized, for example in an excel spreadsheet—generative AI is trained on massive unstructured datasets—for example, large volumes of text. Mani Pande, “Understanding the Difference Between GenAI & Traditional AI/ML Models,” Medium, (June 30, 2024), <https://medium.com/manipande/understanding-the-difference-between-gen-ai-traditional-ai-ml-models-cfd941296f23>; see also Cole Stryker & Mark Scapicchio, “What is generative AI?” IBM, (March 22, 2024), <https://www.ibm.com/topics/generative-ai>.

²⁷ See Black, et al., *supra* note 21 at 105–06.

²⁸ See *FTC v. Rite Aid Corp., et al.*, No. 23-05023, Complaint (E.D. PA. Dec. 19, 2023); Fed. Trade Comm’n Press Release, “Rite Aid Banned from Using AI Facial Recognition After FTC Says Retailer Deployed Technology Without Reasonable Safeguards” (Dec. 19, 2023), <https://www.ftc.gov/news-events/news/press-releases/2023/12/rite-aid-banned-using-ai-facial-recognition-after-ftc-says-retailer-deployed-technology-without>.

compared to white people and on images of women compared to men.²⁹ The FTC also alleged that Rite Aid used low-quality images in creating a database for its automated decision-making tool to use, which contributed to the tool's high false-positive rate. The FTC ultimately barred Rite Aid from continuing to use its facial recognition automated decision-making tool.

Deployment: Once an automated decision-making tool is deployed, algorithmic discrimination can occur for a number of reasons. The tool can be used in a purposely discriminatory manner—for example, if the tool is used to assess members of one protected class but not another. If a tool is used to make decisions that it was not designed to assess, its deployment may amplify any bias in the tool and systemic inequities that exist outside of the tool.³⁰ In some instances, deployment may reveal biases in the tool that were not apparent during testing.³¹ And deployment can lead to a feedback loop, where a biased tool contributes to discriminatory decisions that are then introduced back to the tool for continuous training. Each iteration of this loop exacerbates the tool's bias.

- Research on a now-discontinued predictive policing algorithm, PredPol, provides an example of how historical biases in training data can be amplified in a feedback loop upon deployment of an automated decision-making tool.³² PredPol was designed to help law enforcement agencies map “hotspots” of criminal activity to inform how they deployed police. However, the tool ultimately flagged for targeted policing locations that were already over-represented in police data, inadvertently reinforcing racial biases in historical policing practices. The researchers ran PredPol on a simulation of drug crimes that assumed the police would observe more crimes in these target neighborhoods compared to non-target neighborhoods because they had been sent there by the algorithm and would seek opportunities to make arrests during their patrols in the target neighborhoods. The researchers further assumed that these additional crimes, no matter how minor the offense, would be fed back into the algorithm for training. In the simulation, because of the increase in observed crimes in the target neighborhoods, PredPol predicted that even more crimes would be committed in those neighborhoods, leading to further targeted policing by law enforcement agencies. This created a feedback loop, where the targeted neighborhoods became increasingly likely to be flagged for additional policing. Following reports that

²⁹ See Nat'l Inst. of Standards and Tech, *Face Recognition Vendor Test Part 3: Demographic Effects*, at 71 (Dec. 2019), <https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf> (explaining that facial recognition tools developed in China were more accurate when used on Chinese faces, suggesting that the accuracy of facial recognition technologies on different races is dependent, at least in part, on the diversity reflected in the training data); Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, 81 *Proceedings of Machine Learning Research*, 1–15 (2018) (explaining that facial recognition tools trained on biased data have resulted in algorithmic discrimination).

³⁰ See Suresh & Gutttag, *supra* note 20 at 6 (describing “deployment bias” as “a mismatch between the problem a model is intended to solve and the way in which it is actually used”).

³¹ See Nat'l Telecomm. Info. Admin., *Artificial Intelligence Accountability Policy Report*, at 18–19 (Mar. 2024) <https://www.ntia.gov/sites/default/files/publications/ntia-ai-report-final.pdf> (explaining that “[n]ot all risks can be identified pre-deployment, and downstream developers/deployers may fine tune [automated decision-making tools] either to ameliorate or exacerbate dangers present in artifacts from upstream developers”).

³² Kristian Lum & William Isaac, *To Predict and Serve?*, 13 *SIGNIFICANCE* 1, 14–19 (2016).

PredPol was contributing to discriminatory policing practices, law enforcement agencies across the nation stopped using the tool.³³

Ultimately, bias can be introduced into automated decision-making tools if systemic racism, sexism, or other inequalities are not accounted for when designing, training, and deploying the tools. And this, in turn, can reinforce and exacerbate existing disparities, risking significant harm to marginalized populations. All of this underscores the importance of civil rights protections to prevent and address algorithmic discrimination.

II. THE LAD PROHIBITS ALGORITHMIC DISCRIMINATION

The LAD applies to discrimination stemming from the use of automated decision-making tools in the same way it has long applied to other forms of discriminatory conduct.

The LAD prohibits all forms of discrimination, irrespective of whether discriminatory conduct is facilitated by automated decision-making tools or driven by purely human practices. Indeed, consistent with its broad remedial purpose of eliminating discrimination in New Jersey,³⁴ the LAD draws no distinctions based on the mechanism of discrimination. Thus, a covered entity—that is, an employer, housing provider, place of public accommodation, credit provider, contractor, or any other party subject to the LAD’s requirements—is not immunized from liability for violating the LAD merely because its discriminatory policy or practice involves using or relying on an automated decision-making tool. A covered entity can violate the LAD even if it has no intent to discriminate, and even if a third-party was responsible for developing the automated decision-making tool.³⁵ In short, claims of algorithmic discrimination are assessed consistent with other claims of discrimination under the LAD.

The LAD prohibits algorithmic discrimination on the basis of actual or perceived race, religion, color, national origin, sexual orientation, pregnancy,³⁶ breastfeeding, sex, gender identity, gender expression, disability, and other protected characteristics.³⁷ When employers, housing providers, places of public accommodation, or other covered entities use automated decision-making tools, they may violate the LAD if those tools result in disparate treatment based on a protected characteristic or if those tools have a disparate impact based on a protected characteristic. The

³³ For example, while law enforcement in Plainfield, New Jersey obtained PredPol beginning in 2018, Plainfield police stated that they never used PredPol’s predictions to direct patrols and discontinued its use of the tool once concerns about bias were identified. See Aaron Sankin & Surya Mattu, *Predictive Policing Software Terrible at Predicting Crimes*, The Markup (Oct. 2, 2023) <https://themarkup.org/prediction-bias/2023/10/02/predictive-policing-software-terrible-at-predicting-crimes>.

³⁴ *Nini v. Mercer Cty. Cmty. Coll.*, 202 N.J. 98, 115 (2010).

³⁵ See *Lehman v. Toys ‘R’ Us, Inc.*, 132 N.J. 587, 604-05 (N.J. 1993) (“The LAD is not a fault- or intent-based statute. . . . The purpose of the LAD is to eradicate discrimination, whether intentional or unintentional. Although unintentional discrimination is perhaps less morally blameworthy than intentional discrimination, it is not necessarily less harmful in its effects, and it is at the effects of discrimination that the LAD is aimed.”).

³⁶ Consistent with the LAD, the term “pregnancy” in this document should be read broadly to include pregnancy-related conditions such as childbirth, and “breastfeeding” should be read broadly to include manually expressing milk; pumping milk; breastfeeding, chestfeeding, bodyfeeding, or otherwise feeding milk directly from an individual’s body to a child, as well as related conditions. N.J.S.A. § 10:5-12(s).

³⁷ N.J.S.A. § 10:5-12.

LAD also prohibits algorithmic discrimination when it precludes or impedes the provision of reasonable accommodations, or of modifications to policies, procedures, or physical structures to ensure accessibility for people based on their disability, religion, pregnancy, or breastfeeding status.³⁸

Disparate Treatment Discrimination

Disparate treatment discrimination involves conduct by a covered entity that treats a person differently because of their membership in an LAD-protected class.³⁹ Disparate treatment discrimination occurs if a policy or practice is intentionally discriminatory, or if a policy or practice is discriminatory on its face, even if a covered entity does not intend to discriminate.

With respect to automated decision-making tools, a covered entity engages in disparate treatment discrimination when it designs or uses automated decision-making tools to treat members of a protected class differently. For example, housing providers engage in disparate treatment discrimination based on source of lawful income, an LAD-protected characteristic, if they use automated decision-making tools designed to exclude prospective tenants who use housing vouchers.

Disparate treatment discrimination may also occur if covered entities selectively use automated decision-making tools to assess members of a protected class—for example, if a housing provider uses a tenant screening algorithm only to evaluate Black prospective tenants but not prospective tenants of other races. Further, the use of an automated decision-making tool may result in disparate treatment discrimination even if the tool does not directly consider a protected characteristic but makes recommendations based on a close proxy for a protected characteristic. For example, a housing provider that uses a housing screening tool designed to select applicants who provide individual taxpayer identification numbers instead of Social Security numbers—for instance, because the housing provider wishes to rent only to immigrants based on their belief that immigrants will be less likely to challenge deficiencies in a property—discriminates based on national origin. Even if the tool does not explicitly consider national origin, an individual taxpayer identification number, which is only available to certain immigrants, is so closely associated with national origin that using it to screen applicants would in effect screen applicants based on national origin.

Disparate Impact Discrimination

Even if automated decision-making tools do not expressly assess people differently based on a protected characteristic, the use of these tools may violate the LAD if the use has a disproportionately negative effect on members of an LAD-protected class. This form of discrimination is known as disparate impact discrimination.

³⁸ For more information on reasonable accommodations for pregnancy, breastfeeding, and related conditions, see DCR’s Guidance on Workplace Accommodations for Pregnant, Postpartum, Breastfeeding, and Lactating Employees (Dec. 2024), [DCR-Guidance-on-Pregnancy-Related-Workplace-Accommodations.pdf](#).

³⁹ N.J.S.A. § 10:5-12(a), (f), & (g); *Gerety v. Atl. City Hilton Casino Resort*, 184 N.J. 391, 398–99 (2005) (citing *Pepper v. Princeton Univ. Bd. of Trustees*, 77 N.J. 55, 81–82 (1978)).

Disparate impact discrimination occurs when policies or practices disproportionately affect members of an LAD-protected class in an unlawful manner.⁴⁰ Specifically, policies and practices that result in a disparate impact are prohibited under the LAD unless they are necessary to achieve a substantial, legitimate, nondiscriminatory interest and there is no less discriminatory alternative that would achieve the same interest.⁴¹ This is true even if policies and practices are not discriminatory on their face—in other words, even if they are facially neutral—and are not motivated by discriminatory intent.⁴²

Algorithmic discrimination constitutes disparate impact discrimination when an automated decision-making tool makes recommendations or contributes to decisions that disproportionately harm members of an LAD-protected class unless use of the tool serves a substantial, legitimate, nondiscriminatory interest. Even then, the use of the tool is prohibited if there is a less discriminatory alternative. In evaluating whether there are less discriminatory alternatives, whether the covered entity tested its automated decision-making tool for bias or evaluated alternatives may be considered as relevant evidence.⁴³

Disparate impact discrimination may arise, for example, if a company uses an automated decision-making tool to assess contract bids that disproportionately screens out bids from women-owned businesses. Likewise, disparate impact discrimination may occur if a store uses facial recognition technology to ban former shoplifters that disproportionately generates false positives for patrons who wear religious headwear compared to other patrons.

Reasonable Accommodations

Algorithmic discrimination may also violate the LAD if the use of automated decision-making tools precludes or impedes the provision of reasonable accommodations. This includes, but is not limited to, failing to make reasonable accommodations for a person's disability,⁴⁴ religion,⁴⁵ pregnancy,⁴⁶ or breastfeeding. The LAD requires a covered entity to make reasonable accommodations based on a person's membership in a protected class if the entity knew or should have known about the need for an accommodation and the accommodation will not cause the entity an undue hardship.⁴⁷

⁴⁰ 56 N.J.R. 969(a); *Gerety*, 184 N.J. at 398–399 (employment context); *In re Adoption of 2004 Low Income Hous. Tax Credit*, 369 N.J. Super. 2 (App. Div. 2004) (housing context).

⁴¹ 56 N.J.R. 969(a).

⁴² Nevertheless, a covered entity may also engage in disparate impact discrimination if it creates policies and practices that are intended to have a disproportionately discriminatory effect on members of a protected class. *See id.*

⁴³ *See generally id.* (indicating that “[a]n employer’s use of an automated employment decision tool that has not been adequately tested and shown not to adversely affect people in a protected class before its use may have a disparate impact on members of that protected class”).

⁴⁴ N.J.A.C. § 13:13-2.5 (employment context); N.J.A.C. § 13:13-3.4(f)(2) (housing context); N.J.A.C. § 13:13-4.11 (public accommodations context).

⁴⁵ N.J.S.A. § 10:5-12(q) (employment context).

⁴⁶ N.J.S.A. § 10:5-12(s) (employment context).

⁴⁷ *See* N.J.A.C. § 13:13-2.5; N.J.A.C. § 13:13-4.11; N.J.S.A. § 10:5-12(q); N.J.S.A. § 10:5-12(s).

A covered entity's use of an automated decision-making tool may implicate reasonable accommodations in several ways.⁴⁸ Reasonable accommodations may be necessary if an automated decision-making tool is inaccessible to individuals with disabilities or other individuals based on a protected characteristic. For example, reasonable accommodations may be necessary if an employer uses a tool to measure applicants' typing speed that cannot measure typing from a non-traditional keyboard. Such a tool would not fairly or accurately assess the typing speed of an applicant who uses a non-traditional keyboard because of a disability.

Reasonable accommodations may also factor into recommendations by an automated decision-making tool that result in disparate impact discrimination. If an automated decision-making tool is not trained on data that includes individuals who use an accommodation, the tool may not recognize that an accommodation is possible or may penalize individuals who have or need a reasonable accommodation. If a covered entity acts on the recommendation made by such a tool, the covered entity may engage in disparate impact discrimination. When used in hiring, for example, such a tool may disproportionately exclude applicants who could perform the job with a reasonable accommodation. By accepting the tool's recommendation to exclude these applicants, a covered entity could violate the LAD. Relatedly, some automated decision-making tools may fail to account for an existing reasonable accommodation. For example, if an employer uses a tool to monitor and track the productivity of its employees and the tool is programmed to flag atypical or unsanctioned breaks but is not programmed to consider reasonable accommodations, the tool may disproportionately flag for discipline employees who are allowed additional break time to accommodate a disability or to accommodate milk expression. The employer may violate the LAD if it accepts the recommendation from the tool to discipline these employees.

Who Is Liable for Discrimination Under the LAD

A covered entity is liable for any of its policies or practices that result in discrimination in violation of the LAD.⁴⁹ In many cases, employers, housing providers, places of public accommodation, creditors, and contractors use automated decision-making tools they did not help to develop. Often, these covered entities rely largely or entirely on a third-party developer to create or implement the automated decision-making tool. Nevertheless, a covered entity is not shielded from liability for algorithmic discrimination that results from the entity's use of an automated decision-making tool simply because the tool was developed by a third party or because the entity does not understand the inner workings of the tool.⁵⁰

⁴⁸ See U.S. Equal Emp. Opportunity Comm'n, "The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees" ("EEOC ADA Guidance"), at Questions 2 & 12 (May 12, 2022) <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence>; see also U.S. Dept. of Justice Civil Rights Division, "Algorithms, Artificial Intelligence, and Disability Discrimination in Hiring" (May 12, 2022), <https://www.ada.gov/resources/ai-guidance/>.

⁴⁹ See 56 N.J. R. 969(a) (proposed section 13:16-3.2(c)(3)).

⁵⁰ See EEOC ADA Guidance, *supra* note 48 at Question 3 (providing that employers may be responsible under the ADA for their use of automated decision-making tools designed or administered by another entity, such as a software vendor); see also EEOC Title VII Guidance, *supra* note 3 at Question 3 (providing same analysis pursuant to Title VII); Consumer Fin. Protection Bureau, "Circular 2022-03 Adverse Action Notification Requirements in Connection With Credit Decisions Based on Complex Algorithms," (May 26, 2022) <https://www.consumerfinance.gov/compliance/circulars/circular-2022-03-adverse-action-notification-requirements-in-connection-with-credit-decisions-based-on-complex-algorithms/> (providing that "[a] creditor cannot justify

It is critical that employers, housing providers, places of public accommodation, and other covered entities—as well as the developers and vendors of automated decision-making tools used by these entities—carefully consider and evaluate the design and testing of automated decision-making tools before they are deployed, and carefully analyze and evaluate those tools on an ongoing basis after they are deployed. Doing so is necessary to decrease the risk of discriminatory outcomes and thereby decrease the risk of possible liability under the LAD.⁵¹

The use of automated decision-making tools will continue to expand in key aspects of New Jerseyans' lives. As it does, the risk of algorithmic discrimination may also increase. The LAD prohibits algorithmic discrimination and protects New Jerseyans from discrimination no matter the type of technology covered entities use.

To find out more about the protections the LAD provides, go to www.NJCivilRights.gov. DCR also has several [resources](#) that provide an overview of the LAD. For more about algorithmic discrimination under the LAD, visit DCR's [website](#).

If you believe you have faced discrimination in violation of the LAD, you can file a complaint with DCR by visiting www.bias.njcivilrights.gov or calling 1-833-NJDCR4U (833-653-2748).

Sundeep Iyer
Director, New Jersey Division on Civil Rights
January 2025

noncompliance with [the Equal Credit Opportunity Act] and Regulation B's requirements based on the mere fact that the technology it employs to evaluate applications is too complicated or opaque to understand").

⁵¹ Because algorithmic bias can be introduced at one or more stages in the development of an automated decision-making tool, there is no one-size-fits-all solution to algorithmic discrimination. Rather, different solutions may be necessary to combat algorithmic discrimination based on where bias was introduced in the development and use of the tool. See Black et al., *supra* note 21 at 103–09; see also Suresh & Gutttag, *supra* note 20; Nat'l Fair Housing Alliance & FairPlay, *Improving Mortgage Underwriting and Pricing Outcomes for Protected Classes Through Distribution Matching* (Apr. 2024) <https://nationalfairhousing.org/wp-content/uploads/2024/04/Unlocking-Fairness-Final-April-2024.pdf>. A non-exhaustive list of steps covered entities can take to identify and mitigate algorithmic discrimination from their use of automated decision-making tools includes implementing quality control measures for any data used in designing, training, and deploying the tool; conducting impact assessments; having pre-and post-deployment bias audits performed by independent parties; providing notice of their use of an automated decision-making tool; involving people impacted by their use of a tool in the development of the tool; and red-teaming their automated decision-making tools—or purposely attacking the tools to search for flaws. See e.g., Off. of Mgmt. & Budget, Exec. Off. of the President, m-24-10, “Memorandum for the Heads of Executive Departments and Agencies, Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence,” (Mar. 28, 2024), <https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf>; see also U.S. Dept. of Labor, “Artificial Intelligence and Worker Well-Being: Principles and Best Practices for Developers and Employer,” (Oct. 2024), <https://www.dol.gov/sites/dolgov/files/general/ai/AI-Principles-Best-Practices.pdf>; HUD ADM Housing Rental Screening Guidance; *supra* note 4; HUD ADM Advertising Guidance, *supra* note 4.