

Algorithmic Discrimination: Legal Principles

October 20, 2026

Congressional Research Service

<https://crsreports.congress.gov>

R49490



R49490

October 20, 2026

Abigail A. Graber
Legislative Attorney

Algorithmic Discrimination: Legal Principles

As automated tools for decisionmaking, including those using artificial intelligence (AI), have grown in popularity and sophistication, they raise questions about how their use intersects with federal antidiscrimination laws. Federal laws address various kinds of discrimination, including disparate treatment and disparate impact. *Disparate treatment* generally occurs when a person makes a decision or classifies people on the basis of a protected trait, such as race, sex, religion, or another characteristic. *Int'l Bhd. of Teamsters v. United States*, 431 U.S. 324, 335 n.15 (1977). *Disparate impact* occurs when a person makes a decision using neutral criteria, but the decision adversely and disproportionately affects members of a protected group, and either the decisionmaker cannot put forth a legally adequate justification for using the criteria, or an alternative exists that would serve the decisionmaker's purposes with less of a disparate impact. *Tex. Dep't of Hous. & Cmty. Affs. v. Inclusive Cmty. Project, Inc.*, 576 U.S. 519, 524–25 (2015).

How disparate treatment and disparate impact frameworks apply to decisionmaking via automated, algorithmic tools (including AI) is not always clear. Some situations involving algorithmic decisionmaking would clearly be disparate treatment under present law. For example, a person covered by antidiscrimination law who intentionally designs an algorithm to discriminate on the basis of a protected characteristic, or who selects an algorithm because of its discriminatory effect on a protected group, engages in disparate treatment. *See Pers. Adm'r of Mass. v. Feeney*, 442 U.S. 256, 279 (1979). However, the law is unsettled on how to evaluate cases where no human consciously chooses to discriminate, and yet an automated system classifies people on the basis of protected characteristics. There is also substantial uncertainty over when an algorithmic system could be said to engage in disparate treatment via proxy—that is, classification based on a facially neutral trait that is closely correlated with a protected characteristic. The law is not clear on whether proxy discrimination always requires human intent to use a proxy to discriminate, or whether some characteristics are so closely related to protected traits that they are inherent proxies. It is also not clear how one could prove that certain AI systems are engaging in proxy discrimination if there is no way to know how some systems arrive at their conclusions. How laws handle such situations, and how difficult it is for plaintiffs to prove their case, may also depend on the causation and intent standards specific to each law.

Many unanswered questions persist regarding when and how the use of algorithms can give rise to disparate impact claims, as well. Using algorithmic tools can result in disparate impacts in several ways. An algorithm could use decision rules or targeted outcomes that proxy for protected characteristics or that are more or less favorable to certain groups. Algorithms could rely on criteria that are less predictive for certain groups than others, or fail to consider criteria that would make them more predictive for particular groups. AI systems may be trained on data that are affected by past discrimination, or that are over- or under-representative of particular groups, potentially leading to biased or less accurate outcomes for members of those groups. *See Solon Barocas & Andrew D. Selbst, Big Data's Disparate Impact*, 104 CAL. L. REV. 671, 677–94 (2016).

There are already gray areas in legally significant elements of disparate impact doctrine. These uncertainties make it hard to predict how courts will respond when algorithmic decisionmaking is added to the mix. Courts may have to decide, for example, whether algorithmic systems should be analyzed as a whole or whether plaintiffs can, or must, challenge particular elements of how a system works. They may have to determine whether a decisionmaker can rely on the fact that an algorithm's predictions are accurate, or whether the decisionmaker has to posit or prove how the algorithm works or why it is accurate. Courts may also need to decide how much of a premium decisionmakers may place on cost savings or efficiency when justifying the use of an algorithm that causes a disparate impact. To evaluate plaintiffs' proposed less discriminatory alternatives, courts may need to grapple with whether plaintiffs must put forward alternatives that are as efficient and cost-effective as a challenged algorithm and what standards govern showing that an alternative—including an alternative algorithm—is available.

Congress has many options if it chooses to regulate algorithmic decisionmaking to prevent discrimination or bias. It may consider amending current antidiscrimination laws to better take into account the unique features of algorithmic decisionmaking, or it can pass laws specifically targeted at potential algorithmic bias. It may prohibit discrimination or require other, affirmative conduct by those that develop or deploy algorithmic systems, such as greater transparency, bias audits, or guaranteed human participation in certain decisions. Some potential pathways to mitigating the risks of algorithmic discrimination may raise constitutional concerns under the Takings Clause, First Amendment, and Equal Protection Clause. Congress may also consider how to balance the risks of discrimination with the risks of regulation, which could stifle the growth of automated tools, including those that may advance less discriminatory decisionmaking.

Contents

Glossary of Abbreviations	1
Introduction	2
Background	2
Algorithmic Decisionmaking and AI	2
Antidiscrimination Legal Frameworks.....	4
Fundamentals of Disparate Treatment Doctrine	5
Fundamentals of Disparate Impact Doctrine	9
Tensions Between Disparate Treatment and Disparate Impact	14
Applying Antidiscrimination Legal Frameworks to Algorithmic Decisionmaking.....	15
Disparate Treatment	16
How Algorithms May Result in Disparate Treatment.....	16
Open Questions on Disparate Treatment and Algorithms.....	17
Disparate Impact	23
How Algorithms May Cause Disparate Impacts.....	23
Open Questions on Disparate Impact and Algorithms.....	25
Considerations for Congress.....	29
Regulating Algorithmic Decisionmaking Under General Antidiscrimination Laws or	
Using Standalone Antidiscrimination Rules	31
Adopting Broad Versus Tailored Requirements	31
Deciding Who Should Be Regulated.....	32
Considering Requirements Beyond Nondiscrimination.....	33
Transparency	33
Audits.....	35
Design Requirements	36
Role of Human Decisionmakers	37
Examining Disparate Impact.....	38
Balancing the Risks of Discrimination with the Risks of Regulation	39

Contacts

Author Information.....	41
-------------------------	----

Glossary of Abbreviations

AI	Artificial Intelligence
ADA	Americans with Disabilities Act
ADEA	Age Discrimination in Employment Act
CFPB	Consumer Financial Protection Bureau
DOJ	Department of Justice
ECOA	Equal Credit Opportunity Act
EEOC	Equal Employment Opportunity Commission
FHA	Fair Housing Act
FTC	Federal Trade Commission
GAO	Government Accountability Office
GINA	Genetic Information Nondiscrimination Act
HUD	Department of Housing and Urban Development
NIST	National Institute of Standards and Technology
OLC	Office of Legal Counsel
OMB	Office of Management and Budget
Title IX	Title IX of the Education Amendments of 1972
Title VI	Title VI of the Civil Rights Act of 1964
Title VII	Title VII of the Civil Rights Act of 1964

Introduction

Automated decisionmaking and artificial intelligence (AI) tools have generated interest from lawmakers about the effect of new technology and “big data” on privacy, security, and opportunity, among other areas. One area that has seen limited federal legislative development is in guarding against algorithmic discrimination. At various points in time, the executive branch,¹ lawmakers,² and others³ have voiced concern about the use of “biased” algorithms and AI systems in important areas of public life, including employment, housing, financial transactions, and beyond.⁴ At the same time, these and other commentators have championed the potential of algorithmic decisionmaking to root out discrimination and make recommendations or decisions beyond the reach of human bias.⁵

This report explores legal frameworks for evaluating claims of algorithmic discrimination. It reviews the broad legal principles that animate much of federal antidiscrimination law and how those principles may, or may not, apply when people use algorithms to support decisionmaking.

Background

Algorithmic Decisionmaking and AI

There is a lack of consensus on much of the terminology used to describe different forms of algorithmic decisionmaking and AI. These terms are defined below as used in this report. More detail on types and applications of AI can be found in other CRS products.⁶

¹ E.g., Exec. Order No. 14110, 88 Fed. Reg. 75191, 75211 (Nov. 1, 2023); EXEC. OFF. OF THE PRESIDENT, BIG DATA: A REPORT ON ALGORITHMIC SYSTEMS, OPPORTUNITY, AND CIVIL RIGHTS (2016), https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf [<https://perma.cc/4JNR-LDNS>].

² E.g., STAFF OF THE BIPARTISAN H. TASK FORCE ON AI, 118TH CONG., REPORT ON ARTIFICIAL INTELLIGENCE: GUIDING PRINCIPLES, FORWARD-LOOKING RECOMMENDATIONS, AND POLICY PROPOSALS TO ENSURE AMERICA CONTINUES TO LEAD THE WORLD IN RESPONSIBLE AI INNOVATION 79–87 (2024) [hereinafter BIPARTISAN HOUSE TASK FORCE REPORT], https://www.govinfo.gov/content/pkg/GOVPUB-Y1_2-PURL-gpo236622/pdf/GOVPUB-Y1_2-PURL-gpo236622.pdf [<https://perma.cc/MB4R-KY8J>]; STAFF OF THE BIPARTISAN S. AI WORKING GRP., 118TH CONG., REPORT ON DRIVING U.S. INNOVATION IN ARTIFICIAL INTELLIGENCE: A ROADMAP FOR ARTIFICIAL INTELLIGENCE POLICY IN THE UNITED STATES SENATE 11 (2024), https://www.govinfo.gov/content/pkg/GOVPUB-Y1_3-PURL-gpo229697/pdf/GOVPUB-Y1_3-PURL-gpo229697.pdf [<https://perma.cc/HG8B-QKEN>].

³ E.g., Jo Gasior-Kavishe, *AI Is Infringing on Your Civil Rights. Here's How We Can Stop That*, AM. C.L. UNION (Dec. 3, 2025), <https://www.aclu.org/news/racial-justice/ai-is-infringing-on-your-civil-rights-heres-how-we-can-stop-that> [<https://perma.cc/FBW6-RPHJ>]; Chiraag Bains, *The Legal Doctrine That Will Be Key to Preventing AI Discrimination*, BROOKINGS (Sep. 13, 2024), <https://www.brookings.edu/articles/the-legal-doctrine-that-will-be-key-to-preventing-ai-discrimination> [<https://perma.cc/3WDC-3A4P>].

⁴ For a discussion of some potential ways to understand “bias” in the context of AI, see NAT’L INST. OF STANDARDS AND TECH. (NIST), ARTIFICIAL INTELLIGENCE RISK MANAGEMENT FRAMEWORK (AI RMF 1.0), at 18 (2023) [hereinafter NIST AI RMF 1.0], <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf> [<https://perma.cc/S5NF-STDD>], and REVA SCHWARTZ ET AL., NIST, NIST SPECIAL PUBLICATION 1270, TOWARDS A STANDARD FOR IDENTIFYING AND MANAGING BIAS IN ARTIFICIAL INTELLIGENCE 6–9 (2022), <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf> [<https://perma.cc/3L2R-WKBZ>].

⁵ See, e.g., EXEC. OFF. OF THE PRESIDENT, BIG DATA, *supra* note 1, at 5–6; Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 WM. & MARY L. REV. 857, 871 (2017); *infra* “Balancing the Risks of Discrimination with the Risks of Regulation.”

⁶ See CRS Report R48909, *Artificial Intelligence: CRS Products*, by Laurie Harris and Rachael D. Roan (2026).

In general, an *algorithm* is “a set of step-by-step instructions for solving a problem or completing a task.”⁷ In the context of computing, an algorithm tells “a computer how to process inputs such as data to produce an output, such as a decision, prediction, or sorted list.”⁸ While people can use algorithms without computers, in this report the term *algorithmic decisionmaking* means the use of automated, machine-based processes to make or assist with decisions, predictions, or recommendations. Algorithmic decisionmaking can be quite basic, using rules programmed by human experts. Software programmed to pass all college applicants with a grade point average above 3.0 to an admissions officer for further consideration and reject all others is an example of simple algorithmic decisionmaking. An algorithmic system could also use a more complex statistical formula and consider multiple variables (such as the difficulty of a college applicant’s high school course of study, the quality of their high school, their extracurriculars, standardized test scores, or other factors).

Algorithmic systems may rely on AI. While definitions of AI vary,⁹ this report uses *AI* to refer to “algorithms and techniques that aim to give computer systems the ability to learn new concepts or tasks and solve complex problems in a manner that mimics human intelligence.”¹⁰ When experts discuss *AI models*, they often mean computer programs engaged on some level in *machine learning*, a process in which a program uses algorithms and mathematical functions to process data and discover patterns and statistical relationships between variables. AI models can then use these discovered relationships to classify new inputs, that is, to make recommendations or predictions using previously unencountered data.¹¹ An AI system may include an underlying AI model and a user-facing application that is powered by the model. Such systems include AI virtual assistants, chatbots, and screening tools, among others.

There are many different types of AI models, varying in, among other things, how they are trained, the mathematical functions that power their operations, and the outputs they generate. For example, AI models may be trained using *supervised learning*, *unsupervised learning*, or *reinforcement learning*.¹² A subset of AI, known as *generative AI*, can be used to “generate new content,” such as “text, images, videos, computer code, or music.”¹³ Other AI models may

⁷ *What Is an Algorithm?*, STAN. U., HUMAN-CENTERED A.I., <https://hai.stanford.edu/ai-definitions/what-is-an-algorithm> [<https://perma.cc/2B55-4VMS>] (last visited Sep. 18, 2026).

⁸ *Id.*

⁹ See, e.g., CRS Report R46795, *Artificial Intelligence: Background, Selected Issues, and Policy Considerations*, by Laurie Harris (2021) (discussing debate over whether “robotic process automation” uses AI or whether “rule-based systems” are AI). Compare, e.g., 15 U.S.C. § 9401(3) (defining *AI* in the National Artificial Intelligence Initiative Act of 2020, H.R. 6216, 116th Cong. (2020)), with John S. McCain National Defense Authorization Act for Fiscal Year 2019, Pub. L. No. 115-232, § 238(g), 132 Stat. 1636, 1696 (2018) (defining *AI* for purposes of certain activities within the Department of Defense), and CAL. GOV’T CODE § 11546.45.5(a) (West 2026) (defining *AI* under California law), and CONN. GEN. STAT. § 4a-2e(a)(1) (2026) (defining *AI* under Connecticut law). For a visual taxonomy of machine-based systems, see CRS Infographic IG10077, *Artificial Intelligence (AI) Taxonomy*, by Laurie Harris and Nora Wells (2025).

¹⁰ CRS Infographic IG10077, *Artificial Intelligence (AI) Taxonomy*, by Laurie Harris and Nora Wells (2025).

¹¹ See, e.g., *id.*; *AI Models & Deployments*, MICROSOFT AI TOUR, <https://microsoft.github.io/Workshop-Interact-with-OpenAI-models/ai-models> [<https://perma.cc/48M5-SA7R>] (last visited Sep. 22, 2026); *What Is an AI Model?*, IBM, <https://www.ibm.com/think/topics/ai-model> [<https://perma.cc/EA9W-9E2L>] (last visited Sep. 22, 2026); see also Solon Barocas & Andrew D. Selbst, *Big Data’s Disparate Impact*, 104 CAL. L. REV. 671, 677 (2016) (describing “data mining”).

¹² See *What Is an AI Model?*, *supra* note 11, for explanations of these terms.

¹³ CRS In Focus IF12426, *Generative Artificial Intelligence: Overview, Issues, and Considerations for Congress*, by Laurie Harris (2025).

classify data.¹⁴ This report uses the term *AI* in a very general sense, without differentiating between the various kinds of AI models available today.

As a hypothetical example of the use of one type of AI system, consider an employer using an AI system to screen resumes. The employer might feed the system data about employees the employer considers “successful.” “Successfulness” is the employer’s *target variable*, or the outcome a person seeks to predict using AI.¹⁵ The AI system may then discover statistical relationships between data points about these employees—perhaps that “successful” employees are statistically likely to have a certain employment background, use a particular computer browser, or own a particular car.¹⁶ The AI system may assign weights to the elements it finds to be statistically related to success or lack thereof, and use the resulting algorithms to rank new applicants (i.e., the program creates internal *decision rules*).¹⁷ In this scenario, the employer may or may not tell the program what “success” means, and the employer may or may not know which input data the AI system considers salient and what weight it has assigned them. It may not be possible to determine the reason for some AI systems’ classifications due to the scale and complexity of certain AI models, a scenario often referred to as the “black box” problem.¹⁸ Additionally, there may not be a single, discoverable reason; the decision rules may change with each set of inputs.¹⁹

Antidiscrimination Legal Frameworks

Federal antidiscrimination laws have different, though sometimes overlapping, scopes. The laws vary in whom they cover, the decisions they govern, and the characteristics they protect. The Fifth Amendment’s Due Process Clause and the Fourteenth Amendment’s Equal Protection Clause, for example, apply only to government actors and subject laws that classify on certain bases, including race, national origin, and sex, to more rigorous scrutiny by courts.²⁰ Other federal antidiscrimination schemes are differently targeted. Title VII of the Civil Rights Act of 1964

¹⁴ *Id.*

¹⁵ See Anya E.R. Prince & Daniel Schwarcz, *Proxy Discrimination in the Age of Artificial Intelligence and Big Data*, 105 IOWA L. REV. 1257, 1273 (2020); Barocas & Selbst, *supra* note 11, at 677.

¹⁶ See Matthew U. Scherer, Allan G. King & Marko J. Mrkonich, *Applying Old Rules to New Tools: Employment Discrimination Law in the Age of Algorithms*, 71 S.C. L. REV. 449, 457 (2019) (describing data potentially available to employers, “which may or may not be job related—such as a candidate’s social media profiles, criminal history, and web browsing history”); cf. Matthew Adam Bruckner, *The Promise and Perils of Algorithmic Lenders’ Use of Big Data*, 93 CHI.-KENT L. REV. 3, 15 (2018) (describing data used by modern “algorithmic lenders,” including “a consumer’s email addresses, brand of car, Facebook friends, educational background and college major, even whether he or she sends text messages in all capital letters or in lower case” (quoting Gregory Roberts, *Regulator Wades into Big Data Credit Swamp*, BLOOMBERG LAW: BANKING (Apr. 20, 2017), <https://www.bna.com/regulator-wades-big-n57982086887> [<https://perma.cc/T8VP-JPZR>])).

¹⁷ Carey K. Morewedge, Haaya F. Naushan & Sendhil Mullainathan, *Decisions with Algorithms*, in THE HANDBOOK OF SOCIAL PSYCHOLOGY ch. 44 (D.T. Gilbert, S.T. Fiske & W.E. Mendes eds. 2025), <https://openpublishing.princeton.edu/read/decisions-with-algorithms/section/6092ce7a-c408-4eaa-ac59-d94baf7c5edf> [<https://perma.cc/5789-RZNK>].

¹⁸ See CRS Report R47644, *Artificial Intelligence: Overview, Recent Advances, and Considerations for the 118th Congress*, by Laurie Harris (2023) (describing the “‘black box’ problem”); Scherer, King & Mrkonich, *supra* note 16, at 455–56.

¹⁹ See Talia B. Gillis & Jann L. Spiess, *Big Data and Discrimination*, 86 U. CHI. L. REV. 459, 475–77 (2019). Similarly, changes in the data on which an AI model is trained can produce different decision rules in ways that are difficult for developers to predict. Emily Black et al., *Less Discriminatory Algorithms*, 113 GEO. L.J. 53, 67 n.42 (2024).

²⁰ U.S. CONST. amend. V, cl. 5 (governing the federal government); *id.* amend. XIV, § 1 (governing state and local governments); see *Clark v. Jeter*, 486 U.S. 456, 461 (1988) (describing the different levels of scrutiny that apply to laws based on the nature of the classifications they use). The Supreme Court has held that the Fifth Amendment’s Due Process Clause incorporates the same equal protection principles as the Fourteenth Amendment’s Equal Protection Clause. *Weinberger v. Wiesenfeld*, 420 U.S. 636, 638 n.2 (1975).

(Title VII), for example, generally prohibits “employers,” “employment agencies,” and “labor organizations” from making employment-related decisions on the basis of race, color, national origin, religion, and sex.²¹ The Equal Credit Opportunity Act (ECOA) forbids “creditors” from discriminating in credit transactions on the basis of race, color, religion, national origin, sex, marital status, age, or source of income.²² Other antidiscrimination laws prohibit discrimination on the basis of race,²³ age,²⁴ or disability²⁵ by recipients of federal funds.

In addition to who and what is covered by an antidiscrimination law, the types of claims available, and how to prove them, vary from law to law. Two of the primary forms of discrimination recognized in federal antidiscrimination law are disparate treatment and disparate impact.²⁶ The specifics of disparate treatment and disparate impact vary depending on the law; this report largely focuses on big-picture principles common to many approaches.

The next sections provide an overview of the fundamentals of disparate treatment and disparate impact claims, followed by a discussion of how these legal frameworks may apply to allegations of algorithmic discrimination.

Fundamentals of Disparate Treatment Doctrine

Intentional Classifications

Disparate treatment occurs when someone makes a decision based at least in part on another’s protected trait, such as race or sex.²⁷ Disparate treatment can involve formal classifications, such

²¹ Pub. L. No. 88-352, tit. VII, § 703, 78 Stat. 241, 255 (codified as amended at 42 U.S.C. § 2000e-2).

²² Pub. L. No. 93-495, tit. V, § 503, 88 Stat. 1521 (1974) (codified as amended at 15 U.S.C. § 1691(a)).

²³ Civil Rights Act of 1964, Pub. L. No. 88-352, tit. VI, 78 Stat. 241, 252 (codified as amended at 42 U.S.C. §§ 2000d–2000d-7).

²⁴ Age Discrimination Act of 1975, tit. III, Pub. L. No. 94-135, 89 Stat. 728 (codified as amended at 42 U.S.C. §§ 6101–6107).

²⁵ Rehabilitation Act of 1973, Pub. L. No. 93-112, § 504, 87 Stat. 355, 394 (codified as amended at 29 U.S.C. § 794).

²⁶ People can sometimes claim denial of reasonable accommodations or reasonable modifications on the basis of religion or disability, *see, e.g.*, 42 U.S.C. §§ 2000e(j), 12112(b)(5)(A), 12131(2), 12182(b)(2)(A)(ii), a type of claim not addressed here.

²⁷ *E.g.*, *Int’l Bhd. of Teamsters v. United States*, 431 U.S. 324, 335 n.15 (1977). While disparate treatment often involves treating a person worse than others on the basis of a protected trait, *e.g.*, *Cicalese v. Univ. of Tex. Med. Branch*, 924 F.3d 762, 766 (5th Cir. 2019), separating or classifying people on the basis of protected traits is also often disparate treatment even if those groups are ultimately treated the same. *See, e.g.*, *Anderson v. Martin*, 375 U.S. 399, 404 (1964) (invalidating state requirement that candidates’ race be indicated on election ballots); *Brown v. Bd. of Ed.*, 347 U.S. 483, 495 (1954) (“Separate educational facilities are inherently unequal.”). Not every decision that may involve considering a protected trait is disparate treatment, however. For example, regulated entities often may, and sometimes are required to, collect and report demographic data without running afoul of antidiscrimination law. *See, e.g.*, 42 U.S.C. § 2000e-8(c) (requiring entities covered by Title VII to engage in recordkeeping as the Equal Employment Opportunity Commission (EEOC) requires to track unlawful discrimination); *Caulfield v. Bd. of Ed.*, 583 F.2d 605, 611 (2d Cir. 1978) (holding that “the collection of racial and ethnic data is authorized by Title VI [of the Civil Rights Act of 1964]” in the context of an administrative investigation into potential discrimination). *But see* *MD/DC/DE Broads. Ass’n v. FCC*, 236 F.3d 13, 19, 22 (D.C. Cir. 2001) (holding that rules requiring broadcasters to undertake certain demographic data collection violated the Equal Protection Clause when an agency used the data to pressure broadcasters to recruit more minority job applicants); *Removal of Reporting Requirements*, 91 Fed. Reg. 46332, 46332 (proposed July 23, 2026) (to be codified at 29 C.F.R. pt. 1602) (proposing to rescind EEOC’s data collection and reporting requirements as “inconsistent with equal employment opportunity law and potentially unconstitutional”). Though this area of the law is unsettled, parties may also at times be able to consider the effect on particular groups when deciding on a course of action. *See, e.g.*, *Parents Involved in Cmty. Sch. v. Seattle Sch. Dist. No. 1*, 551 U.S. 701, 787–89 (2007) (Kennedy, J., concurring in part); *Raso v. Lago*, 135 F.3d 11, 17 (1st Cir. 1998); (continued...)

as when a school district assigns students to schools based on race.²⁸ It can also involve the intentionally discriminatory application of facially neutral policies, that is, policies that do not overtly classify on the basis of a protected characteristic. For example, the selective enforcement of a neutral school dress code against only members of a protected group is a form of disparate treatment.²⁹

A party also engages in disparate treatment when it adopts a facially neutral policy with discriminatory intent.³⁰ To plead a disparate treatment claim on such facts, a plaintiff must show that a defendant adopted the policy “[b]ecause of,” not merely “in spite of,” its adverse effects upon an identifiable group.³¹ For example, an apartment complex that enacts a rule against placing decorations on doors and doorposts may be engaging in disparate treatment if the company adopted the rule to harass or drive away Jewish tenants (who may want to place mezuzot on their doorways).³²

This third type of intentional discrimination can take the form of proxy discrimination. *Proxy discrimination* occurs when an entity classifies people based on facially neutral traits that are closely correlated with protected characteristics.³³ An example is *redlining*, where financial institutions discriminate against consumers based on zip codes, which, due to residential segregation, can be an effective proxy for race.³⁴ Understanding proxy discrimination is relevant to understanding algorithmic discrimination because, as discussed below, algorithms that are programmed to avoid considering protected characteristics might nevertheless rely on proxies.³⁵

The law on proxy discrimination as disparate treatment is not entirely clear.³⁶ The Supreme Court has at times held that a regulated entity must be “actually motivated” by the relationship between the proxy and the protected characteristic for proxy discrimination to support a disparate treatment claim.³⁷ The Court may presume such motivation when a policy targets traits or activities that are “an irrational object of disfavor . . . if they also happen to be engaged in exclusively or predominantly by a particular class of people.”³⁸ For example, “[a] tax on wearing yarmulkes is a tax on Jews.”³⁹ At other times, the Supreme Court has suggested that some

South-Suburban Hous. Ctr. v. Greater S. Suburban Bd. of Realtors, 935 F.2d 868, 883–84 (7th Cir. 1991); CRS Report R47471, *The Constitution and Race-Conscious Government Action: Narrow Tailoring Requirements*, by April J. Anderson (2023); *infra* “Tensions Between Disparate Treatment and Disparate Impact.”

²⁸ See *Parents Involved*, 551 U.S. at 709–11.

²⁹ E.g., *Frith v. Whole Foods Mkt., Inc.*, 38 F.4th 263, 271 (1st Cir. 2022); *Sewell v. Monroe City Sch. Bd.*, 974 F.3d 577, 583 (5th Cir. 2020); *Proctor ex rel. K.R. v. Duluth Pub. Schs. Acad.*, 591 F. Supp. 3d 418, 428 (D. Minn. 2022).

³⁰ E.g., *Inclusive Cmtys. Project, Inc. v. Lincoln Prop. Co.*, 920 F.3d 890, 910 (5th Cir. 2019); *Bloch v. Frischholz*, 587 F.3d 771, 785 (7th Cir. 2009) (en banc).

³¹ *Pers. Adm’r of Massachusetts v. Feeney*, 442 U.S. 256, 279 (1978).

³² *Bloch*, 587 F.3d at 785. “The mezuzah is a central aspect of the Jewish religious tradition—many Jews believe they are commanded by God to affix mezuzot on the exterior doorposts of their dwelling . . .” *Id.* at 772.

³³ See, e.g., *Holland v. Elevance Health, Inc.*, 171 F.4th 126, 131 (1st Cir. 2026); *Davis v. Guam*, 932 F.3d 822, 833, 837–38 (9th Cir. 2019).

³⁴ See *Prince & Schwarcz*, *supra* note 11, at 1268–69; see also *Saint-Jean v. Emigrant Mortg. Co.*, 129 F.4th 124, 139 (2d Cir. 2025) (defining “redlining”), *cert. denied*, 146 S. Ct. 1501 (2026) (mem.).

³⁵ See *infra* “How Could Plaintiffs Prove an AI Model Engaged in Disparate Treatment Via Proxy?”

³⁶ Proxy discrimination can also give rise to disparate impact claims. See *Prince & Schwarcz*, *supra* note 11, at 1260; Ifeoma Ajunwa, *The Paradox of Automation As Anti-Bias Intervention*, 41 CARDOZO L. REV. 1671, 1679–80, 1679 (2020).

³⁷ *Hazen Paper Co. v. Biggins*, 507 U.S. 604, 610, 612–13 (1993).

³⁸ *Bray v. Alexandria Women’s Health Clinic*, 506 U.S. 263, 270 (1993).

³⁹ *Id.*

classifications might proxy for protected characteristics without considering motivation at all.⁴⁰ However, one recent case adds weight to the view that a close correlation between two characteristics is not enough, standing alone, to show intentional discrimination via proxy. In *United States v. Skrametti*, the Court held that laws that prohibit certain medical care to treat gender dysphoria, but allow the same treatments for other conditions, do not classify on the basis of transgender status, despite the close link between having gender dysphoria and being transgender.⁴¹

On the other hand, some circuit courts have suggested that certain characteristics may be considered proxies for protected traits because of their inherent connection to those traits, without requiring evidence of motivation.⁴² Some of the proxy classifications cited as examples by these courts could fall within the category of “irrational object[s] of disfavor” that “also happen to be engaged in exclusively or predominantly by a particular class of people”⁴³—for example, discrimination based on “gray hair”⁴⁴—but not all of them necessarily would be considered “irrational.” The Third Circuit,⁴⁵ for example, has held that discrimination in the terms of employee health benefits based on Medicare eligibility is age discrimination because “Medicare eligibility ‘follow[s] ineluctably upon attaining age 65.’”⁴⁶ Not all circuit courts consider characteristics with a close fit to a protected trait to be inherent proxies.⁴⁷

⁴⁰ See *Hernandez v. New York*, 500 U.S. 352, 371–72 (1991) (plurality opinion) (“It may well be, for certain ethnic groups and in some communities, that proficiency in a particular language, like skin color, should be treated as a surrogate for race under an equal protection analysis.”). In *Guinn v. United States*, 238 U.S. 347 (1915), the Court invalidated a state constitutional provision that required would-be voters to pass a literacy test but that exempted descendants of people who could vote on or prior to January 1, 1866—in other words, people whose ancestors could vote before the Fifteenth Amendment went into effect. Acknowledging that the provision did not facially discriminate on the basis of “race, color, or previous condition of servitude prohibited by the Fifteenth Amendment,” the Court observed that “the standard itself inherently brings that result into existence.” *Id.* at 364. *Guinn* suggests that this inherent connection between the proxy (ancestry) and the protected characteristic (race) may be enough to show intentional discrimination. *Guinn*, however, is an old case, and the Court additionally reasoned that the irrational nature of the classification merited the same result: “Certainly it cannot be said that there was any peculiar necromancy in the time named which engendered attributes affecting the qualification to vote which would not exist at another and different period unless the Fifteenth Amendment was in view.” *Id.* at 365.

⁴¹ 605 U.S. 495, 517–19 (2025).

⁴² See, e.g., *Pac. Shores Props., LLC v. City of Newport Beach*, 730 F.3d 1142, 1160 n.23 (9th Cir. 2013) (“Proxy discrimination . . . arises when the defendant enacts a law or policy that treats individuals differently on the basis of seemingly neutral criteria that are so closely associated with the disfavored group that discrimination on the basis of such criteria is, constructively, facial discrimination against the disfavored group.”); *Erie Cnty. Retirees Ass’n v. County of Erie*, 220 F.3d 193, 211 (3d Cir. 2000).

⁴³ *Bray*, 506 U.S. at 270.

⁴⁴ *Pac. Shores Props.*, 730 F.3d at 1160 n.23 (“For example, discriminating against individuals with gray hair is a proxy for age discrimination because ‘the “fit” between age and gray hair is sufficiently close.’” (quoting *McWright v. Alexander*, 982 F.2d 222, 228 (7th Cir. 1992)).

⁴⁵ For purposes of brevity, references to a particular circuit in the body of this report (e.g., the Third Circuit) refer to the U.S. court of appeals for that circuit.

⁴⁶ *Erie Cnty. Retirees Ass’n*, 220 F.3d at 211 (alteration in original) (quoting *Erie Cnty. Retirees Ass’n v. County of Erie*, 91 F. Supp. 2d 860, 867 (W.D. Pa. 1999)). In *Erie County Retirees Association*, the Third Circuit held that an employer had treated some of its retirees differently because of their age, and may have violated the Age Discrimination in Employment Act (ADEA), when it shifted Medicare-eligible retirees to a health care plan that those retirees alleged was inferior to the plan available to non-Medicare-eligible retirees. *Id.* at 213, 217.

⁴⁷ *Farhan v. 2715 NMA LLC*, 161 F.4th 475, 483 & n.5 (7th Cir. 2025) (holding that an apartment complex did not discriminate against a tenant on the basis of national origin when it punished her for displaying a Palestinian flag); see also *Holland v. Elevance Health, Inc.*, 171 F.4th 126, 131 (1st Cir. 2026) (defining “proxy discrimination” with reference to the fit between the criteria and the disfavored group, but stating that “[t]his Circuit has never recognized a proxy discrimination theory, whether framed as a theory of disparate impact or facial discrimination”).

Mixed Motives and “Rational” Discrimination

In some cases, a defendant may have multiple motives for a decision, one or more of which may be discriminatory while others are not. Federal laws take different approaches to these so-called “mixed-motive” cases.⁴⁸ Under Title VII, for example, so long as the protected characteristic is a “motivating factor” for the decision, the fact that other factors may have been determinative does not matter; the decision is still illegal.⁴⁹ Courts applying the Equal Protection Clause take a different approach: Once a plaintiff shows that a decision was at least in part motivated by a discriminatory purpose, the burden shifts to the defendant to prove it would have made the same decision without the illegal consideration (and thereby avoid liability).⁵⁰

Not all laws recognize “motivating factor” discrimination or shift the burden to the defendant once some discriminatory intent is shown. Under the Age Discrimination in Employment Act (ADEA),⁵¹ for example, private-sector employees must prove that age (greater than forty) was a “but-for” cause of any adverse treatment. They cannot bring disparate treatment claims if their employers improperly considered their age but would have reached the same decisions regardless, and plaintiffs always bear the burden of showing age was determinative.⁵²

A defendant’s motive is generally discriminatory even if the defendant’s reliance on a protected trait is arguably “rational.”⁵³ In other words, even if a protected trait is strongly correlated with or has a causal relationship to a factor a person may legitimately consider, the decisionmaker often may not take the protected trait into account.⁵⁴ For example, the Supreme Court has held that Title

⁴⁸ Whether “mixed-motive” claims are available under any given antidiscrimination law, and the standards that govern such claims, are not always clear. *See, e.g.*, ROBERT G. SCHWEMM, *HOUSING DISCRIMINATION LAW AND LITIGATION* § 10:3 (2026) (discussing the Fair Housing Act (FHA)).

⁴⁹ 42 U.S.C. § 2000e-2(m); *see Univ. of Tex. Sw. Med. Ctr. v. Nassar*, 570 U.S. 338, 343, 347–49 (2013). The available remedies may be different, however, if the employer proves that the protected characteristic did not play a determinative role in the decision. 42 U.S.C. § 2000e-5(g)(2)(B). “Motivating-factor” claims are not available for Title VII plaintiffs in retaliation cases. *Nassar*, 570 U.S. at 352, 359–60.

⁵⁰ *Village of Arlington Heights v. Metro. Hous. Dev. Corp.*, 429 U.S. 252, 265–66, 271 n.21 (1977).

⁵¹ Pub. L. No. 90-202, 81 Stat. 602 (codified as amended at 29 U.S.C. §§ 621–634).

⁵² *Gross v. FBL Fin. Servs., Inc.*, 557 U.S. 167, 175–78, 180 (2009). Motivating-factor liability is available for claims against federal employers. In these cases, unlike Title VII cases, *see* 42 U.S.C. § 2000e-5(g)(2)(B), the burden remains with the plaintiff to prove that age was a but-for cause of the adverse action if the plaintiff seeks remedies such as reinstatement, backpay, and compensatory damages. *Babb v. Wilkie*, 589 U.S. 399, 402, 407–08, 413–14 (2020).

⁵³ *See, e.g.*, Samuel R. Bagenstos, “*Rational Discrimination, Accommodation, and the Politics of (Disability) Civil Rights*,” 89 VA. L. REV. 825, 837–38, 848–52 (2003) (discussing reasons for illegal discrimination that are grounded in an “employer’s pursuit of his ultimate end of maximizing profit”); Christine Jolls, *Antidiscrimination and Accommodation*, 115 HARV. L. REV. 642, 687 (2001) (discussing potential real costs imposed on employers by antidiscrimination law).

⁵⁴ *E.g.*, FED. TRADE COMM’N (FTC), *BIG DATA: A TOOL FOR INCLUSION OR EXCLUSION? UNDERSTANDING THE ISSUES* 18 (2016) [hereinafter *FTC, BIG DATA*], <https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf> [<https://perma.cc/V5Y3-FFN7>]; *see* sources cited *supra* note 53. This rule is not universal. Its application depends on the protected trait and the law. For example, under legal regimes that target racial discrimination, such discrimination is almost always illegal no matter how strong a correlation between race and a legitimate consideration. *See, e.g.*, *Students for Fair Admissions, Inc. v. President & Fellows of Harvard Coll.*, 600 U.S. 181, 208 (2023) (stating that “race-based state action” is unconstitutional in all but “the most extraordinary case”); *Ferrill v. Parker Grp., Inc.*, 168 F.3d 468, 473 (11th Cir. 1999) (observing that while Title VII allows employers to consider religion, sex, or national origin in “extremely narrow” cases where those traits are a “bona fide occupational qualification,” even that exception is not available for race). The law is generally more tolerant of disability discrimination. Under the Americans with Disabilities Act of 1990 (ADA), Pub. L. No. 101-336, 104 Stat 327, for example, covered employers generally may refuse to employ a person with a disability if their disability renders them unqualified for the job and no reasonable accommodation can solve the problem. *See* 42 U.S.C. § 12112(a), (b)(5) (prohibiting disability discrimination against “qualified individual[s]” and requiring “reasonable (continued...)”).

VII does not allow employers to require female employees to make higher contributions to their pension funds, even if one accepts that women generally live longer than men.⁵⁵ Employers also cannot refuse to hire women for physically demanding jobs, even if sex is correlated with certain physical aptitudes.⁵⁶ Strict limitations on “rational” discrimination are a significant consideration when examining algorithmic discrimination. AI models may discover hidden relationships, including correlations, between input data and desired outcomes. The AI model may then offer recommendations, make predictions, or take action based on those machine-identified patterns and trends. Under federal antidiscrimination law, the fact that a correlation is real does not make acting on that correlation legal.

Fundamentals of Disparate Impact Doctrine

Disparate impact discrimination occurs when a protected group experiences disproportionately adverse effects from a facially neutral practice (that is, a practice that does not take protected characteristics into account), and the regulated entity lacks an appropriate justification.⁵⁷ Not all federal antidiscrimination laws provide for disparate impact liability. Under the Genetic Information Nondiscrimination Act of 2008 (GINA) and Title VI of the Civil Rights Act of 1964 (Title VI), for example, plaintiffs may bring only disparate treatment claims.⁵⁸ Under laws that allow for disparate impact claims, such as Title VII,⁵⁹ the ADEA,⁶⁰ and the Fair Housing Act

accommodations” for individuals who are “otherwise qualified”). The law does, however, require individualized determinations of whether a person with a disability is qualified; a general correlation between disability and lack of qualification at a population level is likely insufficient in most cases. *See, e.g., Taylor v. Pathmark Stores, Inc.*, 177 F.3d 180, 192 (3d Cir. 1999) (“[U]nder the ADA, it is the employer’s burden to educate itself about the varying nature of impairments and to make individualized determinations about affected employees.”). The law also does not allow for all “rational” forms of disability discrimination: for example, the ADA does not allow covered entities to take into account the costs of accommodating a person’s disability unless such costs pose an “undue hardship,” a relatively high threshold. 42 U.S.C. § 12112(b)(5)(A).

⁵⁵ *L.A. Dep’t of Water & Power v. Manhart*, 435 U.S. 702, 707–08 (1978); *see also, e.g., J.E.B. v. Alabama ex rel. T.B.*, 511 U.S. 127, 140 n.11 (1994) (“We have made abundantly clear in past cases that gender classifications that rest on impermissible stereotypes violate the Equal Protection Clause, even when some statistical support can be conjured up for the generalization.”).

⁵⁶ *See Int’l Union, United Auto., Aerospace & Agric. Implement Workers of Am. v. Johnson Controls, Inc.*, 499 U.S. 187, 202 (1991) (stating that the Court would not “allow employers to use sex as a proxy for strength although it might be a fairly accurate one”).

⁵⁷ *E.g., Tex. Dep’t of Hous. & Cmty. Affs. v. Inclusive Cmty. Project, Inc.*, 576 U.S. 519, 524–25 (2015); *Dothard v. Rawlinson*, 433 U.S. 321, 329 (1977); *Int’l Bhd. of Teamsters v. United States*, 431 U.S. 324, 335 n.15 (1977); *see CRS In Focus IF13057, What Is Disparate-Impact Discrimination?*, by April J. Anderson (2026).

⁵⁸ Pub. L. No. 110-233, tit. II, § 208, 122 Stat. 881, 917 (codified at 42 U.S.C. § 2000ff-7(a)); *Alexander v. Sandoval*, 532 U.S. 275, 280, 285 (2001) (interpreting Title VI to provide a private right of action for disparate treatment claims only).

⁵⁹ 42 U.S.C. § 2000e-2(a)(2), (k). The Office of Legal Counsel (OLC) at DOJ opined in June 2026 that EEOC’s interpretation and application of Title VII’s disparate impact provisions is unconstitutional. *Constitutionality of Disparate-Impact Liability Under Title VII*, 50 Op. O.L.C. (June 9, 2026) [hereinafter OLC Disparate Impact Op.], <https://www.justice.gov/olc/media/1444871/dl?inline> [<https://perma.cc/R4D7-3QAC>]. This section discusses Title VII’s statutory requirements and court interpretations, not EEOC’s Title VII regulations, guidance, or application.

⁶⁰ *Smith v. City of Jackson*, 544 U.S. 228, 240 (2005) (plurality opinion); *id.* at 243 (Scalia, J., concurring).

(FHA),⁶¹ a party may be liable without any intent to discriminate;⁶² disparate impact is therefore sometimes referred to as “unintentional discrimination.”⁶³ The purpose of disparate impact law, the Supreme Court has said, is to eliminate “artificial, arbitrary, and unnecessary barriers to” opportunity “when the barriers operate invidiously to discriminate on the basis of racial or other impermissible classification.”⁶⁴ Disparate impact doctrine also “plays a role in uncovering discriminatory intent” by “permit[ing] plaintiffs to counteract unconscious prejudices and disguised animus that escape easy classification as disparate treatment.”⁶⁵ At the same time, the Court has said that disparate impact law must allow regulated entities to exercise appropriate judgment without undue second-guessing from judges.⁶⁶

The precise formulation of disparate impact doctrine varies somewhat depending on the applicable law and facts, but generally speaking, a basic disparate impact case proceeds through three inquiries:

1. whether the plaintiff can show that a regulated party engages in a specific practice that causes a significant disparate effect on the basis of a protected trait;
2. whether the regulated party can advance a legally adequate justification for the practice; and
3. if there is such a justification, whether the plaintiff can prove there is an alternative practice that would serve the regulated party’s legitimate interest with less of a disparate impact.⁶⁷

⁶¹ *Inclusive Cmty. Project*, 576 U.S. at 546–47; see 24 C.F.R. § 100.500 (2025). The Department of Housing and Urban Development (HUD) has proposed to rescind its FHA disparate impact regulations. HUD’s Implementation of the Fair Housing Act’s Disparate Impact Standard, 91 Fed. Reg. 1475 (proposed Jan. 14, 2026) (to be codified at 24 C.F.R. pt. 100); cf. HUD’s Implementation of the Fair Housing Act’s Disparate Impact Standard; Amendments to HUD’s Title VI Regulations, 91 Fed. Reg. 51416 (proposed Aug. 10, 2026) (to be codified at 24 C.F.R. pt. 1) (proposing to rescind HUD’s Title VI disparate impact regulations and reopening the comment period for the rescission of the FHA disparate impact regulations).

⁶² *E.g.*, *Wards Cove Packing Co. v. Atonio*, 490 U.S. 642, 646 (1989), *superseded by statute on other grounds*, Civil Rights Act of 1991, Pub. L. No. 102-166, tit. 1, § 105, 105 Stat. 1071, 1074–75; *Int’l Bhd. of Teamsters*, 431 U.S. at 335 n.15.

⁶³ *E.g.*, *Ricci v. DeStefano*, 557 U.S. 557, 580 (2009); *Bos. Parent Coal. for Acad. Excellence Corp. v. Sch. Comm. for Bos.*, 89 F.4th 46, 57 (1st Cir. 2023).

⁶⁴ *Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971); see also *Inclusive Cmty. Project*, 576 U.S. at 539 (stating that disparate impact doctrine targets policies and practices that “function unfairly to exclude minorities . . . without any sufficient justification”); *Connecticut v. Teal*, 457 U.S. 440, 449 (1982) (“Congress’ primary purpose was the prophylactic one of achieving equality of employment ‘opportunities’ and removing ‘barriers’ to such equality.” (quoting *Griggs*, 401 U.S. at 429–30)).

⁶⁵ *Inclusive Cmty. Project*, 576 U.S. at 540. The Court has also described disparate impact doctrine as a way to eliminate practices that “operate to ‘freeze’ the status quo of prior discriminatory . . . practices.” *Griggs*, 401 U.S. at 429–30. An entity may be liable for disparate impact discrimination, however, without any showing of prior intentional discrimination. *Watson v. Ft. Wor. Bank & Tr.*, 487 U.S. 977, 988 (1988).

⁶⁶ See, e.g., *Inclusive Cmty. Project*, 576 U.S. at 533.

⁶⁷ *E.g.*, 42 U.S.C. § 2000e-2(k) (Title VII); 24 C.F.R. § 100.500 (FHA regulation); *Sw. Fair Hous. Council, Inc. v. Maricopa Domestic Water Improvement Dist.*, 17 F.4th 950, 960–61 (9th Cir. 2021). This formulation is common to most, but not all, antidiscrimination laws that cover disparate impact. For example, under the ADEA, a court asks only whether the practice causing the disparate impact is “based on reasonable factors other than age.” 29 U.S.C. § 623(f)(1). The court does not inquire into whether alternatives exist that would lessen the disparate impact. *Smith v. City of Jackson*, 544 U.S. 228, 243 (2005). (Certain plaintiffs may not be able to bring ADEA disparate impact claims. See discussion *infra* note 84). There are other forms of disparate impact not discussed in this report. For example, some courts have held that under the FHA, plaintiffs may challenge a zoning regulation based on “harm to the community generally by the perpetuation of segregation.” *Huntington Branch, N.A.A.C.P. v. Town of Huntington*, 844 F.2d 926, 937 (2d Cir. 1988), *superseded by regulation on other grounds*, 24 C.F.R. § 100.500(c)(3), *aff’d in part*, 488 U.S. 15 (continued...)

Disparity Caused by a Specific Practice

Typically, plaintiffs in disparate impact cases rely on statistical disparities between outcomes for members of a protected group versus others.⁶⁸ Plaintiffs might attempt to show, for example, that a bank approves credit applications at lower rates for members of a particular race, or that a school suspends or expels students of one sex at higher rates. On its own, however, a statistical disparity in outcomes does not make out an initial disparate impact claim; the plaintiff must also identify a specific practice or policy of the defendant's that causes the disparity.⁶⁹ The Supreme Court has repeatedly cautioned that plaintiffs must tie disparities to causes within defendants' control so that defendants are held liable only for their own acts, rather than general social conditions; otherwise, the Court has warned, defendants would be forced to turn to forbidden preferences and quotas to avoid disparate impact liability.⁷⁰

Disparate impact plaintiffs commonly challenge the use of particular criteria or tests.⁷¹ Plaintiffs may also challenge defendants' subjective decisionmaking processes in addition to more concrete policies or criteria.⁷²

Legally Adequate Justification

Once a plaintiff has shown that a defendant's policy or practice causes a disproportionate impact, the burden generally shifts to the defendant to give its reasons for the challenged practice. The standard for a legally adequate reason varies by statute and is often unclear. In its first decision developing the concept of disparate impact liability, the Supreme Court articulated the standard in different ways. In that case, *Griggs v. Duke Power Co.*, the Court left unclear whether Title VII required employers to show a challenged practice was truly "necessary" or only "related" to the job in question, stating: "The touchstone is business *necessity*. If an employment practice which operates to exclude [a protected group] cannot be shown to be *related* to job performance, the practice is prohibited."⁷³ Later in the same opinion, the Court said that any employment

(1988) (per curiam); see *Mt. Holly Gardens Citizens in Action, Inc. v. Twp. of Mount Holly*, 658 F.3d 375, 385 (3d Cir. 2011).

⁶⁸ See, e.g., *Watson*, 487 U.S. at 994–95 (plurality opinion); *N.Y.C. Transit Auth. v. Beazer*, 440 U.S. 568, 584 (1979); *Mt. Holly*, 658 F.3d at 382. Statistical disparities may not be necessary under every disparate impact law. Regulations implementing Title II of the ADA, for example, prohibit public entities from using "eligibility criteria that screen out or tend to screen out *an individual with a disability* or any class of individuals with disabilities . . ." 28 C.F.R. § 35.130(b)(8) (2025) (emphasis added).

⁶⁹ E.g., *Inclusive Cmty. Project*, 576 U.S. at 542; *Smith*, 544 U.S. at 241; *Watson*, 487 U.S. at 994–95 (plurality opinion).

⁷⁰ *Inclusive Cmty. Project*, 576 U.S. at 542; *Wards Cove Packing Co. v. Atonio*, 490 U.S. 642, 646 (1989), *superseded by statute on other grounds*, Civil Rights Act of 1991, Pub. L. No. 102-166, tit. 1, § 105, 105 Stat. 1071, 1074–75; see also *Ricci v. DeStefano*, 557 U.S. 557, 581–82 (2009) (warning that employers may take "race-based action at the slightest hint of disparate impact").

⁷¹ See, e.g., *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425 (1975); *Louis v. Saferent Sols., LLC*, 685 F. Supp. 3d 19, 37 (D. Mass. 2023); *Conn. Fair Hous. Ctr. v. Corelogic Rental Prop. Sols., LLC*, 369 F. Supp. 3d 362, 378 (D. Conn. 2019).

⁷² See, e.g., *Watson*, 487 U.S. at 988–91; *Ramirez v. GreenPoint Mortg. Funding, Inc.*, 633 F. Supp. 2d 922, 927–28 (N.D. Cal. 2008); *Hoffman v. Option One Mortg. Corp.*, 589 F. Supp. 2d 1009, 1011 (N.D. Ill. 2008).

⁷³ 401 U.S. 424, 431 (1971) (emphasis added).

requirement “must have a manifest relationship to the employment in question.”⁷⁴ Subsequent Title VII cases also used different formulations.⁷⁵

Then, in 1989, the Court in *Wards Cove Packing Co. v. Atonio* held that “the dispositive issue is whether a challenged practice serves, in a significant way, the legitimate employment goals of the employer.”⁷⁶ Congress responded with the Civil Rights Act of 1991, which codified the concepts of *business necessity* and *job relatedness* in Title VII.⁷⁷ Congress declared its intent to capture those concepts as “enunciated by the Supreme Court in *Griggs v. Duke Power Co.*, 401 U.S. 424 (1971), and in the other Supreme Court decisions prior to *Wards Cove Packing Co. v. Atonio*, 490 U.S. 642 (1989),” a case Congress thought had “weakened” Title VII.⁷⁸ The statutory text of Title VII now requires a challenged practice to be “job related for the position in question and consistent with business necessity.”⁷⁹ The Supreme Court’s different formulations prior to *Ward’s Cove*, however, leave in doubt whether “job relatedness” and “business necessity” are the same or separate standards and whether a practice must be truly “necessary,” or meet some lower bar, to be legally sufficient.⁸⁰ The Court has injected further variability into how to evaluate the legitimate reason by highlighting the importance of deferring to defendants’ expertise.⁸¹

Moreover, while Title VII often influences how courts interpret disparate impact liability under other laws,⁸² not all laws apply the same standards. Some statutes and regulations implementing disparate impact liability refer, like Title VII, to practices that are “necessary” to the defendant’s

⁷⁴ *Id.* at 432.

⁷⁵ For example, the Court required a challenged practice in one case to be “necessary to safe and efficient job performance.” *Dothard v. Rawlinson*, 433 U.S. 321, 331 n.14 (1977). In another, the Court required only that an employer’s “goals” be “significantly served” by a practice. *N.Y.C. Transit Auth. v. Beazer*, 440 U.S. 568, 587 n.31 (1979). *See, e.g.*, Scherer, King & Mrkonich, *supra* note 16, at 464–65 (recounting the Court’s different formulations of a standard).

⁷⁶ 490 U.S. 642, 659 (1989), *superseded by statute*, Civil Rights Act of 1991, Pub. L. No. 102-166, tit. 1, § 105, 105 Stat. 1071, 1074–75.

⁷⁷ Civil Rights Act of 1991 § 3(2); *id.* § 105 (codified at 42 U.S.C. § 2000e-2(k)).

⁷⁸ *Id.* § 2(2).

⁷⁹ 42 U.S.C. § 2000e-2(k)(1)(A)(i).

⁸⁰ *See, e.g.*, *Chi. Tchrs. Union v. Bd. of Educ. of Chi.*, 419 F. Supp. 3d 1038, 1049 (N.D. Ill. 2020), *aff’d*, 14 F.4th 650 (7th Cir. 2021); *Hack v. President & Fellows of Yale Coll.*, 237 F.3d 81, 95 (2d Cir. 2000) (Moran, J., dissenting in part); GEORGE RUTHERGLEN & EARLE K. SHAW, *FED. JUD. CTR., MAJOR ISSUES IN THE FEDERAL LAW OF EMPLOYMENT DISCRIMINATION* 32 (Kris Markarian ed. 5th ed. 2012); Scherer, King & Mrkonich, *supra* note 16, at 465–66. For example, in a 2009 case, the Court asked only that the practice be “job-related.” *Ricci v. DeStefano*, 557 U.S. 557, 578 (2009). One prominent commentator has stated that part of the difficulty in evaluating the defendant’s justification stems from the fact that disparate impact doctrine largely evolved in response to a narrow set of practices, namely seniority systems and written tests. Clear standards have not evolved in other contexts, this scholar has asserted. Michael Selmi, *Was the Disparate Impact Theory a Mistake?*, 53 UCLA L. REV. 701, 705, 716 (2006). Professor Selmi also contends that most of the seminal disparate impact cases involved employers that previously engaged in intentional discrimination. *Id.* at 714. In June 2026, OLC interpreted Title VII to establish a “low threshold” for business necessity, requiring only that “a challenged practice . . . be a ‘reasonable’ way of accomplishing a ‘valid’ interest.” OLC Disparate Impact Op., *supra* note 59, at 15 (quoting *Tex. Dep’t of Hous. & Cmty. Affs. v. Inclusive Cmty. Project, Inc.*, 576 U.S. 519, 541 (2015)).

⁸¹ *See Inclusive Cmty. Project*, 576 U.S. at 540–41; *Watson v. Ft. Wor. Bank & Tr.*, 487 U.S. 977, 999 (1988) (plurality opinion); *Furnco Const. Corp. v. Waters*, 438 U.S. 567, 578 (1978).

⁸² *See, e.g.*, *Inclusive Cmty. Project*, 576 U.S. at 533–34; *Farhan v. 2715 NMA LLC*, 161 F.4th 475, 485 n.5 (7th Cir. 2025); *Sw. Fair Hous. Council, Inc. v. Maricopa Domestic Water Improvement Dist.*, 17 F.4th 950, 960 (9th Cir. 2021). For example, in one of the Supreme Court’s most recent disparate impact cases, the Court described the disparate impact liability standard under the FHA as “analogous to the business necessity standard under Title VII.” *Inclusive Cmty. Project*, 576 U.S. at 541.

operations or the advancement of its interests.⁸³ The ADEA, by contrast, requires only that an employer's practice be "reasonable."⁸⁴ Complicating the picture even more, while Congress rejected parts of *Wards Cove* in amendments to Title VII, some courts have applied all of *Wards Cove* to disparate impact claims under other laws, such as the FHA.⁸⁵ In these cases, defendants must prove that challenged practices "serve[], in a significant way, the[ir] legitimate . . . goals."⁸⁶

Less Discriminatory Alternative

If the defendant proves a legally adequate justification for a challenged practice, the third and final inquiry in many disparate impact cases is whether the plaintiff can show that there is an alternative practice that would cause less of a disparate impact. An academic review of Title VII disparate impact cases observed that "relatively few cases have explored the concept" of the less discriminatory alternative.⁸⁷

Generally, courts appear to require plaintiffs to propose an alternative to the challenged practice that is "available," "valid," and "less discriminatory."⁸⁸ Formulations of these factors vary from court to court and law to law. Broadly speaking, to be "available," a practice must be concrete, not "vague" or "theoretical,"⁸⁹ and the defendant must have "kn[o]w[n] or should have known" about it.⁹⁰ The alternative must be at least "substantially equally valid," meaning, in the employment context, that it "would lead to a workforce [that is] substantially equally qualified."⁹¹ Some courts, including a plurality of the Supreme Court, have said the alternative should be "equally as effective as the challenged practice in serving the [defendant's] legitimate business

⁸³ *E.g.*, 42 U.S.C. § 12182(b)(2)(A)(i) (ADA Title III); 24 C.F.R. § 100.500(b)(1)(i) (FHA regulation); 28 C.F.R. § 35.130(b)(8) (ADA Title II regulation).

⁸⁴ 29 U.S.C. § 623(f)(1) ("It shall not be unlawful for an employer, employment agency, or labor organization to take any action otherwise prohibited [by the ADEA] . . . where the differentiation is based on reasonable factors other than age . . ."); *see* *Smith v. City of Jackson*, 544 U.S. 228, 240, 243 (2005) (clarifying that an employer's burden in an ADEA disparate impact case is lower than its burden under Title VII). Some circuits have held that job applicants cannot bring ADEA disparate impact claims. *Kleber v. CareFusion Corp.*, 914 F.3d 480, 481 (7th Cir. 2019) (en banc); *Villarreal v. R.J. Reynolds Tobacco Co.*, 839 F.3d 958, 961 (11th Cir. 2016). *But see* *Rabin v. PricewaterhouseCoopers LLP*, 236 F. Supp. 3d 1126, 1128 (N.D. Cal. 2017) (holding that job applicants can bring ADEA disparate impact claims). At least one court has held that federal employees cannot bring ADEA disparate impact claims. *Allard v. Holder*, 840 F. Supp. 2d 269, 279–80 (D.D.C. 2012).

⁸⁵ *E.g.*, *Hardie v. Nat'l Collegiate Athletic Ass'n*, 876 F.3d 312, 319 (9th Cir. 2017) (applying *Wards Cove* to Title II of the Civil Rights Act of 1964); *Maricopa Domestic Water Improvement Dist.*, 17 F.4th at 967–68 (applying *Wards Cove* to FHA); *Pickett v. City of Cleveland*, No. 19-CV-2911, 2023 WL 11987525, at *12 (N.D. Ohio Sep. 29, 2023) (same).

⁸⁶ *Hardie*, 876 F.3d at 320 (third alteration in original) (quoting *Wards Cove Packing Co. v. Atonio*, 490 U.S. 642, 659 (1989)); *see* *Reyes v. Waples Mobile Home Park Ltd. P'ship*, 91 F.4th 270, 277 (4th Cir. 2024), *cert. denied*, 145 S. Ct. 172 (2024) (mem.) (citing same standard in FHA case without citing *Wards Cove*).

⁸⁷ *Selmi*, *supra* note 80, at 782 n.221; *see* *Scherer, King & Mrkonich*, *supra* note 16, at 471–72, 509.

⁸⁸ *Adams v. City of Chicago*, 469 F.3d 609, 613 (7th Cir. 2006); *see* *Johnson v. City of Memphis*, 770 F.3d 464, 472 (6th Cir. 2014).

⁸⁹ *Chi. Tchrs. Union v. Bd. of Educ. of Chi.*, 419 F. Supp. 3d 1038, 1052 (N.D. Ill. 2020) (quoting *Puffer v. Allstate Ins. Co.* 675 F.3d 709 (7th Cir. 2012)), *aff'd*, 14 F.4th 650 (7th Cir. 2021).

⁹⁰ *Johnson v. City of Memphis*, No. 00-CV-2608, 2006 WL 3827481, at *14 (W.D. Tenn. Dec. 28, 2006), *rev'd on other grounds*, 770 F.3d 464 (6th Cir. 2014). Title VII requires a plaintiff to show that "the respondent refuses to adopt such alternative employment practice." 42 U.S.C. § 2000e-2(k)(1)(A)(ii). Some courts have held that plaintiffs must "demonstrate a viable alternative and give the employer an opportunity to adopt it." *Allen v. City of Chicago*, 351 F.3d 306, 313 (7th Cir. 2003); *accord, e.g.*, *Lopez v. City of Lawrence*, 823 F.3d 102, 120 (1st Cir. 2016). Whether this holding requires a plaintiff to propose the alternative to a defendant is not clear.

⁹¹ *Erdman v. City of Madison*, 91 F.4th 465, 476 (7th Cir. 2024) (alteration in original) (quoting *Allen*, 351 F.3d at 314); *see* *Ricci v. DeStefano*, 557 U.S. 557, 589–90 (2009) (rejecting proposed alternative because evidence did not show it was "equally valid"); *Johnson*, 770 F.3d at 472.

goals,” taking into account “factors such as the cost or other burdens.”⁹² Little guidance is available on how burdensome on the defendant a less discriminatory alternative can be before it is not equally effective.⁹³ Finally, as mentioned, the alternative must cause less of a disparate impact.⁹⁴

Plaintiffs must support their showing with specific evidence as to each point; they may not meet their burden with speculation that a different approach would have less of an impact or would effectively accomplish the defendant’s legitimate objectives.⁹⁵

Tensions Between Disparate Treatment and Disparate Impact

There are limits on how far people can go to eliminate disparate impacts without running afoul of prohibitions on disparate treatment. The Supreme Court has cautioned that disparate impact liability may not be used to require “numerical quotas,” or to peg compliance to eliminating racial imbalance.⁹⁶ In *Ricci v. DeStefano*, the Court went further and held that New Haven, CT, violated Title VII when it invalidated the results of a firefighter promotion process that had ended in no eligible Black candidates.⁹⁷ While the city argued that it was lawfully attempting to avoid disparate impact liability, the Court held that it had intentionally discriminated against White and Hispanic officers when it acted “because of the statistical disparity based on race.”⁹⁸ Title VII allowed such race-based decisionmaking after the results of a selection process were known, the Court held, only in the face of a “strong basis in evidence of disparate-impact liability.”⁹⁹

How far *Ricci* goes to limit attempts to reduce disparate impacts is not fully known. In *Ricci*, an employer voided the results of a process that had concluded, resulting in identifiable people losing a benefit because of their race.¹⁰⁰ Parties may have more flexibility to modify their policies or practices to eliminate *prospective* disparate impacts.¹⁰¹ Post-*Ricci*, the Supreme Court opined

⁹² *Watson v. Ft. Wor. Bank & Tr.*, 487 U.S. 977, 998 (1988) (plurality opinion); see *Sw. Fair Hous. Council, Inc. v. Maricopa Domestic Water Improvement Dist.*, 17 F.4th 950, 970 (9th Cir. 2021) (requiring alternative to be “‘equally effective’ as the defendant’s chosen policy at serving the defendant’s interest(s), taking into account . . . ‘the cost or other burdens.’” (first alteration in original) (quoting *Hardie v. Nat’l Collegiate Athletic Ass’n*, 876 F.3d 312, 320 (9th Cir. 2017))).

⁹³ *Erdman*, 91 F.4th at 477. Compare *Hardie*, 876 F.3d at 322–23 (describing costs as “only a ‘factor’ in the equal effectiveness analysis” (quoting *Wards Cove Packing Co. v. Atonio*, 490 U.S. 642, 661 (1989))), with *Nat’l Fair Hous. All. v. Bank of Am., N.A.*, No. 18-CV-01919, 2025 WL 2030225, at *22 (D. Md. July 21, 2025) (faulting plaintiffs for failing to show how their alternative “can be accomplished at the *same* or *similar* price point as Defendants’ current policies” (emphases added)).

⁹⁴ See *Chi. Tchrs. Union*, 419 F. Supp. 3d at 1053 (rejecting suggestion that transferring teachers to vacant positions instead of laying them off was a less discriminatory alternative, because transfers may themselves have been adverse employment actions).

⁹⁵ See, e.g., *Maricopa Domestic Water Improvement Dist.*, 17 F.4th at 971 & n.18; *Johnson*, 770 F.3d at 477; *Allen*, 351 F.3d at 314.

⁹⁶ *Tex. Dep’t of Hous. & Cmty. Affs. v. Inclusive Cmty. Project, Inc.*, 576 U.S. 519, 542 (2015) (quoting *Wards Cove*, 490 U.S. at 653).

⁹⁷ 557 U.S. 557 (2009).

⁹⁸ *Id.* at 579.

⁹⁹ *Id.* at 583; see *id.* at 585.

¹⁰⁰ *Id.* at 583–84.

¹⁰¹ See *id.* at 585 (“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.”); Barocas & Selbst, *supra* note 11, at 725–26 (describing how *Ricci* may apply to efforts to design algorithms to minimize disparate impacts); Richard A. Primus, *The Future of Disparate Impact*, 108 MICH. L. REV. 1341, 1345 (2010) (laying out different readings of *Ricci* and positing that “within the category of formally race-neutral actions (continued...)”).

in dicta in an FHA case that government may “foster diversity and combat racial isolation with race-neutral tools.”¹⁰²

Ricci is also a statutory decision. The Court declined to issue a general rule on whether and when policies undertaken to avoid disparate impacts, or government mandates for private parties to avoid disparate impacts, constitute unconstitutional intentional discrimination.¹⁰³ Some Supreme Court Justices, both then and now, have expressed concern that public actors violate the Constitution when they reformulate policy to avoid prospective disparate racial impacts or increase racial diversity.¹⁰⁴ Concurring in *Ricci*, Justice Scalia wrote that the case’s

resolution . . . merely postpones the evil day on which the Court will have to confront the question: Whether, or to what extent, are the disparate-impact provisions of Title VII of the Civil Rights Act of 1964 consistent with the Constitution’s guarantee of equal protection? . . . Title VII’s disparate-impact provisions place a racial thumb on the scales, often requiring employers to evaluate the racial outcomes of their policies, and to make decisions based on (because of) those racial outcomes. That type of racial decisionmaking is, as the Court explains, discriminatory.¹⁰⁵

While that view has not garnered a majority on the Supreme Court, it was echoed by the Department of Justice (DOJ) in a 2026 opinion from the Office of Legal Counsel (OLC), which contended that Title VII disparate impact liability suffers from a “fundamental problem” in that it “tends to incent—and even coerce—employers to make race-based decisions.”¹⁰⁶

Applying Antidiscrimination Legal Frameworks to Algorithmic Decisionmaking

Algorithms have been used in at least some decisions covered by antidiscrimination law for many years. Lenders, for example, have long used credit scores to determine applicants’ creditworthiness.¹⁰⁷ Technological advances, particularly in data analysis, machine learning, and AI, have raised new questions under antidiscrimination law. Commentators have observed that machine-supported or AI-powered decisionmaking (including the data, algorithms, and models involved) can be too complex for humans to understand; that advances have broadened the data that can be considered in decisionmaking; and that automated decisionmaking is being used in an increasing number of domains and has the potential to cause much larger-scale impacts than

intended to improve the position of disadvantaged racial groups, equal protection doctrine may well distinguish between those that have visible victims and those whose costs are more diffuse”).

¹⁰² *Tex. Dep’t of Hous. & Cmty. Affs. v. Inclusive Cmty. Project, Inc.*, 576 U.S. 519, 545 (2015).

¹⁰³ *Ricci*, 557 U.S. at 584 (“Our statutory holding does not address the constitutionality of the measures taken here in purported compliance with Title VII. We also do not hold that meeting the strong-basis-in-evidence standard would satisfy the Equal Protection Clause in a future case.”).

¹⁰⁴ *See Coal. for TJ v. Fairfax Cnty. Sch. Bd.*, 146 S. Ct. 541, 544–45 (2024) (mem.) (Alito, J., dissenting from denial of certiorari) (arguing that a school district may have engaged in intentional discrimination when it changed its admissions policies to increase diversity); *Ricci*, 557 U.S. at 594 (Scalia, J., concurring).

¹⁰⁵ *Ricci*, 557 U.S. at 594 (Scalia, J., concurring).

¹⁰⁶ OLC Disparate Impact Op., *supra* note 59, at 8.

¹⁰⁷ Michael Griffith, Note and Comment, *AI Lending and the ECOA: Avoiding Accidental Discrimination*, 27 N.C. BANKING INST. 349, 356 (2023). Credit scores are calculated from a formula that considers certain financial data, including a person’s payment history, debt, current credit, and other factors. *Id.*

human decisionmaking.¹⁰⁸ Technological advances and their consequences raise novel questions as to how antidiscrimination law may operate when algorithms make or contribute to decisions.

Disparate Treatment

How Algorithms May Result in Disparate Treatment

Discriminatory intent could make its way into algorithmic decisionmaking in a number of ways. The most straightforward example may be when someone intentionally designs an algorithm to rely on a protected characteristic in decisionmaking. In one matter, for example, the Equal Employment Opportunity Commission (EEOC) alleged that an employer violated Title VII and the ADEA when it programmed software to reject applicants older than certain ages, with different thresholds for men and women.¹⁰⁹ In another, the Department of Housing and Urban Development (HUD) and DOJ alleged that Meta violated the FHA by allowing housing advertisers to use Facebook to target or screen out consumers based on protected characteristics, such as sex, religion, and family status.¹¹⁰ Housing advertisers could, for example, elect to screen out people Meta had identified as “Christian” or “moms of grade school kids.”¹¹¹ Black patients have sued hospitals and the nonprofit managing the national organ transplant waiting list for disparate treatment, claiming that until recently, the algorithm that ranks candidates for kidney transplants expressly factored in race and caused them to wait longer for transplants.¹¹²

Algorithms may facilitate intentional discrimination by proxy as well. In the case against Meta, for example, HUD and DOJ alleged that Facebook’s algorithm also allowed advertisers to target close proxies for, among other things, disability and religion, such as users’ interest in “service animals” or “Hijab Fashion.”¹¹³ Former Meta employees have separately alleged that Meta used AI to evaluate employee productivity to make decisions about layoffs, and the AI system used metrics that proxied, at least in part, for employees’ protected leave and disability

¹⁰⁸ See, e.g., U.S. GOV’T ACCOUNTABILITY OFF. (GAO), GAO-25-107201, PROPERTY TECHNOLOGY FOR HOMEBUYING: PRODUCTS PRESENT BENEFITS AND RISKS AMID EVOLVING FEDERAL OVERSIGHT 21–22, 25–26 & n.45 (2025) [hereinafter GAO, PROPERTY TECHNOLOGY], <https://www.gao.gov/assets/gao-25-107201.pdf> [<https://perma.cc/V3ZJ-9VMA>] (discussing the opacity of automated valuation and underwriting models); NIST AI RMF, *supra* note 4, at 18 (“AI systems can potentially increase the speed and scale of biases and perpetuate and amplify harms to individuals, groups, communities, organizations, and society.”); FTC, BIG DATA, *supra* note 54, at 1–2 (discussing how technological advances allow companies to more quickly and easily process larger volumes and a greater variety of data); Griffith, *supra* note 107, at 358–59; Ajunwa, *supra* note 36, at 1699; W. Nicholson Price II, *Risks and Remedies for Artificial Intelligence in Health Care*, BROOKINGS (Nov. 14, 2019), <https://www.brookings.edu/articles/risks-and-remedies-for-artificial-intelligence-in-health-care> [<https://perma.cc/C6F3-YY96>].

¹⁰⁹ EEOC, ARTIFICIAL INTELLIGENCE AND ALGORITHMIC FAIRNESS: REPORT OF THE U.S. EQUAL EMPLOYMENT OPPORTUNITY COMMISSION TO THE COMMITTEES ON APPROPRIATIONS 4 (2024), <https://aboutblaw.com/bfLr> [<https://perma.cc/QXR9-ZWTU>]. (This report is no longer available on EEOC’s website.)

¹¹⁰ Complaint ¶¶ 35–36, *United States v. Meta Platforms, Inc.*, No. 22-CV-05187 (S.D.N.Y. June 21, 2022), Dkt. No. 1, 2022 WL 2835379; Charge of Discrimination ¶¶ 14, 16, *HUD v. Facebook, Inc.*, FHEO No. 01-18-0323-8 (ALJ Mar. 28, 2019), https://archives.hud.gov/news/2019/HUD_v_Facebook.pdf. Private litigants have made similar allegations against Meta. See *Vargas v. Facebook, Inc.*, No. 21-16499, 2023 WL 6784359 (9th Cir. Oct. 13, 2023).

¹¹¹ Charge of Discrimination, *supra* note 110, ¶ 14.

¹¹² E.g., *Maddox v. United Network for Organ Sharing (UNOS)*, 773 F. Supp. 3d 1109, 1120–21 (W.D. Wash. 2025) (allowing case to proceed against both the hospital and the nonprofit that manages the transplant list); *Welch v. UNOS*, 767 F. Supp. 3d 746, 778–82 (M.D. Tenn. 2025) (allowing case to proceed against the hospital but not the nonprofit), *adhered to on reconsideration*, No. 24-CV-00422, 2025 WL 824137 (M.D. Tenn. Mar. 14, 2025); *Randall v. UNOS*, 720 F. Supp. 3d 864, 879–80, 883 (C.D. Cal. 2024) (allowing claims to proceed against both nonprofit (under Title VI and state law) and hospital (under state law)).

¹¹³ Complaint, *supra* note 110, ¶ 40; Charge of Discrimination, *supra* note 110, ¶ 14.

accommodations.¹¹⁴ In another reported incident, a company vetting a resume-screening tool determined that the algorithm favored applicants named “Jared,” a potential proxy for male identity.¹¹⁵

A party could also select a facially neutral algorithm intending to harm a protected group. In one lawsuit, for example, a fair housing nonprofit alleged that a developer of tenant-screening software designed its program to consider applicants’ criminal records with the intent to allow housing providers to discriminate based on race.¹¹⁶ Although the system did not directly consider race, the plaintiffs were able to plead a disparate treatment claim by alleging the defendant’s discriminatory intent.¹¹⁷

Open Questions on Disparate Treatment and Algorithms

The foregoing examples represent relatively straightforward applications of the disparate treatment framework to potential algorithmic discrimination. Yet using algorithms and AI can also introduce new variables into these frameworks that courts have yet to consider in any depth. Two such examples are addressed below.

Can Defendants Who Don’t Understand an Algorithm Be Liable?

The law is not clear on how the disparate treatment framework applies when a party uses an algorithmic tool into which they have little or no insight. Courts have not, until recently, had to consider situations in which a defendant does not know the rules governing a challenged decision. While federal statutes and the Constitution do not use these terms, courts often discuss disparate treatment by referencing the defendant’s “purpose,” “motive,” or “intent,” words that suggest a level of consciousness by a human decisionmaker.¹¹⁸ Indeed, the term “disparate treatment” is often used synonymously with “intentional discrimination.”¹¹⁹ To be clear, courts do not require a malicious motive: people are liable if they use prohibited classifications for ostensibly benevolent reasons.¹²⁰ Courts have also frequently stated that there is no need to inquire into motive when confronted with explicit classifications.¹²¹ In such cases, however, the defendant has

¹¹⁴ The plaintiffs contend that Meta used AI to track employee work output, including everything from keystroke data to how often employees used AI tokens, and that Meta did not account for time employees were on protected leave. Complaint ¶¶ 3–2, 47–53, *Doe v. Meta Platforms Inc.*, No. 26-CV-07122 (N.D. Cal. filed July 13, 2026), Dkt. No. 1.

¹¹⁵ Dave Gershgor, *Companies Are on the Hook If Their Hiring Algorithms Are Biased*, QUARTZ (July 20, 2022), <https://qz.com/1427621/companies-are-on-the-hook-if-their-hiring-algorithms-are-biased> [<https://perma.cc/D3C3-JSL5>].

¹¹⁶ *Conn. Fair Housing Ctr. v. CoreLogic Rental Prop. Sols.*, 369 F. Supp. 3d 362, 376–77 (D. Conn. 2019).

¹¹⁷ *Id.*; see also *Conn. Fair Hous. Ctr. v. CoreLogic Rental Prop. Sols., LLC*, 478 F. Supp. 3d 259, 304 (D. Conn. 2020) (allowing disparate treatment claims to proceed to trial), *aff’d in part, vacated in part on other grounds*, 167 F.4th 605 (2d Cir. 2026).

¹¹⁸ *E.g.*, *Ricci v. DeStefano*, 557 U.S. 557, 577 (2009) (“A disparate-treatment plaintiff must establish ‘that the defendant had a discriminatory intent or motive’ for taking a job-related action.” (quoting *Watson v. Ft. Wor. Bank & Tr.*, 487 U.S. 977, 986 (1988) (plurality opinion))); *Washington v. Davis*, 426 U.S. 229, 239 (1976) (requiring a “racially discriminatory purpose” for equal protection claims).

¹¹⁹ *E.g.*, *Ricci*, 557 U.S. at 580; *Saint-Jean v. Emigrant Mortg. Co.*, 129 F.4th 124, 153 (2d Cir. 2025), *cert. denied*, 146 S. Ct. 1501 (2026) (mem.); *Frith v. Whole Foods Mkt., Inc.*, 38 F.4th 263, 270 (1st Cir. 2022); *Pandazides v. Va. Bd. of Educ.*, 13 F.3d 823, 830 n.9 (4th Cir. 1994).

¹²⁰ *E.g.*, *Ricci*, 557 U.S. at 579–80; *City of Richmond v. J.A. Croson Co.*, 488 U.S. 469, 500 (1989).

¹²¹ *E.g.*, *Parents Involved in Cmty. Schs. v. Seattle Sch. Dist. No. 1*, 551 U.S. 701, 720 (2007) (“It is well established that when the government distributes burdens or benefits on the basis of individual racial classifications, that action is reviewed under strict scrutiny.”); *Hunt v. Cromartie*, 526 U.S. 541, 546 (1999) (“When racial classifications are (continued...)”).

presumptively created or known about the classification.¹²² As one scholar puts it, courts assume that “decisionmakers possess ‘transparency of mind’—that they are aware of the reasons why they are about to make, or have made, a particular . . . decision.”¹²³ That assumption may be untenable in many situations where defendants use algorithms that they do not understand.¹²⁴

Whether federal antidiscrimination law covers someone’s unknowing use of a discriminatory algorithm, even an algorithm that expressly classifies using protected characteristics, is uncertain. As referenced above, at times courts have stated that using a prohibited classification alone creates disparate treatment liability.¹²⁵ Courts have made these statements, however, in the context of defendants accused of intentionally making the relevant classifications.¹²⁶ It is difficult to predict how courts will resolve cases absent that context. Scholars disagree as to whether current law protects against discrimination when a regulated party does not know that an algorithm relied on a prohibited trait.¹²⁷ Several commentators posit that express classifications violate

explicit, no inquiry into legislative purpose is necessary.”); *UAW v. Johnson Controls, Inc.*, 499 U.S. 187, 199 (1991) (“Whether an employment practice involves disparate treatment through explicit facial discrimination does not depend on why the employer discriminates but rather on the explicit terms of the discrimination.”); *Cnty. Servs., Inc. v. Wind Gap Mun. Auth.*, 421 F.3d 170, 177 (3d Cir. 2005) (“Where a regulation or ‘policy facially discriminates on the basis of the protected trait, in certain circumstances it ‘may constitute per se or explicit . . . discrimination’ because ‘the protected trait by definition plays a role in the decision-making process, inasmuch as the policy explicitly classifies people on that basis.’” (alteration in original) (quoting *DiBiase v. SmithKline Beecham Corp.*, 48 F.3d 719, 726 (3d Cir. 1995))); *Santos v. UNOS*, No. 24-CV-11692, 2025 WL 2625321, at *5 (D. Mass. Sep. 10, 2025) (“Plaintiffs have sufficiently pled intentional discrimination by UNOS where they have alleged that UNOS adopted a policy that on its face encourages transplant hospitals to treat kidney transplant candidates differently based on their race. Race-based classifications are harmful in and of themselves.”).

¹²² In his dissent in *Guardians Association v. Civil Service Commission*, for example, Justice Marshall explained that the Court had not needed to determine “whether proof of discriminatory intent is necessary to establish a violation of Title VI” in a previous case, because the affirmative action program challenged in that prior case “*deliberately* used racial criteria.” 463 U.S. 582, 617 (1983) (Marshall, J., dissenting) (emphasis added). As the Third Circuit explained, plaintiffs bringing disparate treatment cases need not prove an “invidious motive,” only that an “actor *meant* to single out a plaintiff because of the *protected characteristic* itself. In a school-segregation case, for instance, ‘the “intent” which triggers a finding of unconstitutionality is not an intent to harm black students, but simply an intent to bring about or maintain segregated schools.’” *Hassan v. City of New York*, 804 F.3d 277, 298 (3d Cir. 2015), *as amended* (Feb. 2, 2016) (citations omitted) (quoting *United States v. Sch. Dist. of Omaha*, 521 F.2d 530, 535 (8th Cir.1975)).

¹²³ Linda Hamilton Krieger, *The Content of Our Categories: A Cognitive Bias Approach to Discrimination and Equal Employment Opportunity*, 47 STANFORD L. REV. 1161, 1185 (1995).

¹²⁴ In some cases, AI or algorithmic products may be proprietary, such that parties who did not develop the system cannot investigate how they work. See, e.g., Ziad Obermeyer et al., *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, 366 SCIENCE 447, 447 (2019), <https://www.science.org/doi/epdf/10.1126/science.aax2342> [<https://perma.cc/7EU3-V92X>].

¹²⁵ See sources cited *supra* note 121 and accompanying text.

¹²⁶ For example, in *Bostock v. Clayton County*, the Supreme Court held that an employer that discriminated on the basis of sexual orientation discriminated on the basis of sex, even if the employer never knew the employee or applicant’s sex. The Court held that employers may be liable for “intentionally setting out a rule” that makes their decisions turn on a protected characteristic. 590 U.S. 644, 668 (2020). In *International Brotherhood of Teamsters v. United States*, the Court stated that one could, at times, infer a “discriminatory motive”—which it described as “critical” to a disparate treatment case—from the “mere fact of differences in treatment.” 431 U.S. 324, 335 n.15 (1977). See also Scherer, King & Mrkonich, *supra* note 16, at 499 (“[C]ourts 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.”).

¹²⁷ Compare, e.g., Barocas & Selbst, *supra* note 11, at 697–700 (concluding that disparate treatment under Title VII requires “conscious intent” and captures discrimination through data mining only when an employer selects a model because it generates biased results), with, e.g., James Grimmelman & Daniel Westreich, *Incomprehensible Discrimination*, 7 CAL. L. REV. ONLINE 164, 164, 176 (2017) (expressing that the authors are “slightly more optimistic” than Professors Barocas and Selbst “that the law already has the doctrinal tools it needs to deal appropriately with cases of this sort” and contending that an employer who unknowingly uses a model that selects applicants based on a (continued...)

antidiscrimination law regardless of intent.¹²⁸ Others think it likely that, doctrinally, “an AI cannot intentionally discriminate independently of any human.”¹²⁹

At least one court seems to favor the view that human intent must be involved. In *Mobley v. Workday, Inc.*, one of the few cases alleging algorithmic discrimination to have generated judicial guidance to date, a job applicant alleged discrimination by the developer of resume-screening software.¹³⁰ The district court acknowledged that the challenged algorithm might use proxy information to classify applicants based on race and age¹³¹ but stated that the plaintiff could not plead a disparate treatment claim against the software developer unless the developer subjectively “intended” disparate outcomes.¹³² Discussing a hypothetical, the district court also suggested that an employer would not be liable for intentional discrimination if it used an algorithm that, unbeknownst to the employer, a third party designed to intentionally discriminate against graduates of historically Black colleges.¹³³

On the other hand, a different federal district court ruled that the government was liable for intentional discrimination under the Equal Protection Clause even when it outsourced decisionmaking to an AI system. In *American Council of Learned Societies v. National Endowment for the Humanities (NEH)*, plaintiffs challenged the terminations of their NEH grants.¹³⁴ As described by the court, discovery revealed that government staff had selected which grants to terminate in large part by instructing ChatGPT to identify grants involving “DEI,” without elaborating further on the meaning of that term, without understanding the AI’s methodology, and without reviewing the AI’s results for accuracy.¹³⁵ The court found that the results expressly classified on the basis of protected characteristics, including race, sex, and religion.¹³⁶ For example, the court found that the government had made the “chilling” decision to terminate a grant to produce a Holocaust documentary because the project “centered on ‘Jewish cultures’ and ‘female’ voices.”¹³⁷ The court ruled that regardless of who, or what, generated the classifications, “[w]hen the Government adopts AI-generated classifications as part of its official actions, those classifications are the Government’s own for constitutional purposes.”¹³⁸ In this

protected characteristic is liable because “Title VII does not permit an employer to do indirectly what it could not do directly”).

¹²⁸ See Scherer, King & Mrkonich, *supra* note 16, at 496–97; Primus, *supra* note 101, at 1350; Kathleen M. Sullivan, *After Affirmative Action*, 59 OHIO STATE LAW J. 1039, 1046–47 (1998).

¹²⁹ Prince & Schwarcz, *supra* note 11, at 1263; see Charles A. Sullivan, *Employing AI*, 63 VILL. L. REV. 395, 404–05 (2018); see also *Chapter One: Resetting Antidiscrimination Law in the Age of AI*, 138 HARV. L. REV. 1562, 1565–66 (2025) (describing why the authors think that “[w]hen it comes to AI, the intent requirement could render the [Equal Protection Clause] a dead letter”); Black et al., *supra* note 19, at 72 (“[T]he concept of intent does not fit easily with algorithms . . .”).

¹³⁰ 740 F. Supp. 3d 796 (N.D. Cal. 2024).

¹³¹ *Id.* at 812.

¹³² *Id.* Under the facts pled, the court held that the software company was potentially liable as multiple employers’ “agent.” *Id.* at 806.

¹³³ See *id.* at 806.

¹³⁴ No. 25-CV-3657, 2026 WL 1256545 (S.D.N.Y. May 7, 2026), *appeals filed*, No. 26-1829 (2d Cir. July 8, 2026), No. 26-1832 (2d Cir. July 8, 2026), No. 26-1996 (2d Cir. July 22, 2026).

¹³⁵ *Id.* at *11–14, 40, 46. ChatGPT is “an artificial-intelligence tool that generates text in response to user prompts.” *Id.* at *11. “DEI” stands for “diversity, equity, and inclusion,” *id.* at *8, although it does not appear from the court’s description that this explanation was provided to ChatGPT.

¹³⁶ *Id.* at *55 (“[T]he ChatGPT-generated ‘DEI Rationales’ repeatedly classified grants as DEI precisely because they involved specific races, ethnicities, religions, or genders.”).

¹³⁷ *Id.* at *56.

¹³⁸ *Id.* at *61.

case, ChatGPT generated explanations for its classifications that were visible to the government decisionmakers and clearly implicated protected characteristics (although it is not clear whether the decisionmakers reviewed those explanations).¹³⁹ Many of those explanations, the court found, showed that ChatGPT’s conclusions “rest[ed] solely on the project’s focus on a particular race,” such as an explanation classifying a grant as “DEI” because the project “invit[ed] reflection on the Asian American experience.”¹⁴⁰ Whether a court would reach a similar result in a case where plaintiffs could prove an AI system relied on discriminatory classifications, but indicia of discrimination were less available to the defendant, is unknown.

Beyond cases directly addressing the unknowing use of an allegedly discriminatory algorithm, courts at times have considered cases where someone who influences a decision exhibits prejudice, but the ultimate decisionmaker lacks conscious discriminatory intent. In the employment discrimination context, for example, courts have considered cases where the ultimate decisionmaker unknowingly acts as a “conduit” for the discriminatory animus of a lower-level supervisor or coworker.¹⁴¹ For example, an employer may fire an employee based on a supervisor’s poor performance review, not knowing that the supervisor was biased. Courts refer to these types of matters as “cat’s paw” cases.¹⁴²

The Supreme Court has ruled that employers may be liable in cat’s paw cases when they act as a conduit for the discriminatory intent of low-level supervisors.¹⁴³ The employer is liable only if the supervisor with discriminatory animus acted with the intent that the employer take some adverse action against the victim.¹⁴⁴ The Court has not decided whether employers are liable when a coworker’s discriminatory animus influences the decisionmaking process.¹⁴⁵

These facets of the doctrine may make it difficult to analogize an employer’s use of a discriminatory algorithm to a cat’s paw case. First, a court would have to decide whether an algorithm was more like a biased low-level supervisor or more like a coworker. Second, these cases still hinge on intentional human actions. An algorithm cannot intend an adverse action against an employee, which may make it impossible to use precedent from cat’s paw cases in algorithmic discrimination matters.

How Could Plaintiffs Prove an AI Model Engaged in Disparate Treatment Via Proxy?

AI developers concerned about discrimination could further complicate the question of what constitutes disparate treatment by excluding or hiding information about protected characteristics from AI training data, or by programming AI models to not consider protected characteristics in

¹³⁹ See *id.* at *55–56; see also *id.* at *47–48 (reviewing grants flagged by ChatGPT as relating to “DEI” with explanations that were not “coherent”).

¹⁴⁰ *Id.* at *55–56.

¹⁴¹ See, e.g., *Marshall v. Rawlings Co. LLC*, 854 F.3d 368, 378 (6th Cir. 2017); *Vasquez v. Empress Ambulance Serv., Inc.*, 835 F.3d 267, 272 (2d Cir. 2016).

¹⁴² For an explanation of the “only marginally” clear reason these matters are called cat’s paw cases, see *Staub v. Proctor Hospital*, 562 U.S. 411, 416 n.1 (2011).

¹⁴³ *Id.* at 422.

¹⁴⁴ *Id.* at 419–22.

¹⁴⁵ *Id.* at 422 n.4. For example, a company may discipline an employee based on complaints by a coworker, not knowing that the complaints are fabricated and motivated by the coworker’s discriminatory intent. See, e.g., *Vasquez*, 835 F.3d at 269–71 (employee fired after being falsely accused of sexual harassment by coworker whose sexual advances she had rejected).

their decision rules.¹⁴⁶ An AI model's decision rules may still reflect protected characteristics, contrary to its developers' intent, because protected characteristics may be correlated with other data points.¹⁴⁷ Some commentators posit that given a rich enough dataset, an AI model will invariably create decision rules that use proxies to incorporate the protected characteristics of people whose data are included, so long as the protected characteristics are correlated with the algorithm's target variable.¹⁴⁸

This concept is called *omitted variable bias*. Omitted variable bias occurs when a predictive variable is omitted from the data on which a statistical model is built, leading the model to misattribute the omitted variable's determinative effects to other variables included in the data. In that case, "the included variable will partially proxy for the omitted characteristic and the estimated impact of the included variable will be biased."¹⁴⁹ Two law professors use the example of an AI model that attempts to select police recruits and favors recruits who listen to certain pop songs. It could be that a preference for a particular style of music is proxying for some forbidden demographic characteristic, like race, sex, or age. On the other hand, perhaps a preference for this musical style is correlated with "motivation to exercise," and such motivation could be the "relevant hidden variable" that the model cannot directly measure.¹⁵⁰ The hidden variable could even be a factor or factors not known by social scientists or understood by humans.

To the extent users or a court would be able to identify any explanation for an AI system's output, it may not be possible to determine what, if any, omitted variables are at play.¹⁵¹ It would matter to the law whether the hypothetical AI model described above had identified a relationship between music preferences and suitability as a police officer separate from a relationship between music preference and a forbidden demographic trait. As previously discussed, under many

¹⁴⁶ See, e.g., Alexandra Chouldechova & Daniel J. Hemel, *Race-Conscious Admissions Algorithms and the Law*, 123 PROC. NAT'L ACAD. SCI. AM. 1, 4 (2026), <https://www.pnas.org/doi/epdf/10.1073/pnas.2509764123>. A number of commentators have opined that excluding information about protected characteristics from training data could make algorithms more discriminatory. See, e.g., BRIAN CHRISTIAN, THE ALIGNMENT PROBLEM: MACHINE LEARNING AND HUMAN VALUES 64–65 (2020); Kim, *supra* note 5, at 917–18; Cynthia Dwork & Deirdre K. Mulligan, *It's Not Privacy, and It's Not Fair*, 66 STAN. L. REV. ONLINE 35, 37 (2013).

¹⁴⁷ E.g., Barocas & Selbst, *supra* note 11, at 695.

¹⁴⁸ See, e.g., Gillis & Spiess, *supra* note 19, at 469; Joshua A. Kroll et al., *Accountable Algorithms*, 165 U. PA. L. REV. 633, 695 (2017); Barocas & Selbst, *supra* note 11, at 695 ("[T]he use of protected class as an input is usually irrelevant to the outcome in terms of discriminatory effect, at least given a large enough number of input features. The target variable will, in reality, be correlated to the membership in a protected class somewhere between 0 percent and 100 percent. . . . Membership in the protected class will prove relevant to the exact extent it is already redundantly encoded. Given a rich enough set of features, the chance that such membership is redundantly encoded approaches certainty.").

¹⁴⁹ Devin G. Pope & Justin R. Sydnor, *Implementing Anti-Discrimination Policies in Statistical Profiling Models*, 3 AM. ECON. J.: ECON. POL'Y 206, 207 (2011).

¹⁵⁰ Grimmelmann & Westreich, *supra* note 127, at 173.

¹⁵¹ In many instances, identifying any explanatory variables may not be possible. See, e.g., Dave Bergmann, *What Is Deep Learning?*, IBM (Sep. 15, 2025), <https://www.ibm.com/think/topics/deep-learning#763338456> ("While [a deep learning] framework is very powerful and versatile, it's comes [sic] at the expense of *interpretability*. There's often little, if any, intuitive explanation—beyond a raw mathematical one—for how the values of individual model parameters learned by a neural network reflect real-world characteristics of data."); *supra* notes 18–19 and accompanying text. Commentators have pointed out that predictive AI models do not necessarily purport to explain the reasons behind the statistical relationships they uncover, and their developers may not start from any hypothesis about what inputs will prove useful for predicting certain outcomes. See, e.g., Prince & Schwarcz, *supra* note 11, at 1274; Gillis & Spiess, *supra* note 19, at 477; Bruckner, *supra* note 16, at 16; Kim, *supra* note 5, at 865, 874–75. Professors Anya Prince and Daniel Schwarcz conclude, "For this reason, the ultimate statistical models that AIs derive are often nearly impossible to explain intuitively; the models work, but no one—including the programmer, the firm that relies on it, or the AI itself—can explain *why* or *how* it does so." Prince & Schwarcz, *supra* note 11, at 1274; see also, e.g., GAO, PROPERTY TECHNOLOGY, *supra* note 108, at 22 ("According to HUD officials, the lack of transparency in AI models may hinder consumers' and regulators' ability to assess or challenge discriminatory practices.").

antidiscrimination laws, “rational” discrimination is illegal. Even if a protected trait like race, sex, or age is correlated with a legitimate target variable—such as likelihood of graduating from a local police academy—acting on that correlation may be illegal.¹⁵² In many cases, however, the protected trait is correlated both with the target variable and with other variables that are legitimate to consider. To take a real-world example, ECOA allows financial institutions to consider income when setting the terms of credit transactions, even if income is correlated with sex, race, or other protected traits.¹⁵³ Just because an AI system’s output is correlated with a protected trait does not mean that the protected trait is a hidden cause of the output.

Complicating matters further, antidiscrimination laws generally do not bar irrational or arbitrary classifications (if not based on protected traits), and this is especially true of laws that do not prohibit disparate impact discrimination.¹⁵⁴ If a financial institution decides to deny credit to people born on odd days of the year, for instance, federal antidiscrimination law generally would not prohibit it.

If a human decisionmaker claimed that they denied a consumer credit because of seemingly arbitrary criteria, a court might scrutinize whether those criteria masked discriminatory intent, especially if the arbitrary criteria were correlated with a protected trait. Courts and juries may infer that seemingly irrational decisionmaking—like selecting police recruits based on music taste—is likely a pretext for prohibited discrimination.¹⁵⁵ In cases involving human decisionmakers, other evidence of their motive may also be available.¹⁵⁶ In cases involving AI models, without an assumption of human motivation behind the decision and potentially without any explanation for why a model incorporated a particular variable, courts may have a more difficult time disentangling true proxies for prohibited traits from proxies for other omitted, permissible variables, whether those variables appear rational or not.

Additionally, proxy discrimination is already a legal gray area. As discussed above, the law is not clear as to whether some traits are so closely related to protected characteristics as to be inherent proxies,¹⁵⁷ and assuming they can be, there is no clear standard for how close a proxy must be to a protected characteristic to count. Potential proxy discrimination in the context of an AI system that assigns weights to a constellation of traits that together may stand in for a protected characteristic is a new development for courts to confront in an area of existing uncertainty.¹⁵⁸

The more a plaintiff must know about how a complex algorithm works, the more difficult the case may be to prove. In cases under laws that do not allow for motivating-factor claims, like some ADEA matters, plaintiffs could have to prove both that the algorithm considered variables that served only to proxy for protected characteristics and that those variables were a but-for cause of

¹⁵² See *supra* “Mixed Motives and “Rational” Discrimination.”

¹⁵³ See 15 U.S.C. § 1691(b)(2); 12 C.F.R. § 202.6(b)(5) (2026). ECOA does prohibit discrimination based on *source* of income, as well as discrimination “on the basis of race, color, religion, national origin, sex or marital status, or age (provided the applicant has the capacity to contract).” 15 U.S.C. § 1691(a).

¹⁵⁴ See *Simms v. First Gibraltar Bank*, 83 F.3d 1546, 1556 (5th Cir. 1996) (stating that bank’s decision “may have been unsound, unfair, or even unlawful, yet not have been violative of the FHA if there is no evidence from which a jury reasonably could infer that race was a significant factor”); *Davis v. Corelogic Platinum Valuation Servs., LLC*, No. 21-CV-01372, 2023 WL 2652477, at *13 (D. Minn. Mar. 27, 2023) (finding no age discrimination under state law where plaintiff received fewer work assignments than younger employee due to “clerical error”). Irrational classifications may run afoul of other laws, such as those prohibiting “unfair or deceptive acts or practices.” *E.g.*, 15 U.S.C. § 45.

¹⁵⁵ See discussion *supra* notes 38–39 and accompanying text.

¹⁵⁶ Scherer, King & Mrkonich, *supra* note 16, at 492–93.

¹⁵⁷ *Supra* “Intentional Classifications.”

¹⁵⁸ See Scherer, King & Mrkonich, *supra* note 16, at 491–92.

an adverse action against them.¹⁵⁹ That is, they would have to prove that the algorithm would have produced a different output if it had not considered the proxy data. In many cases, plaintiffs may not even be able to show what variables an AI model incorporated into its decision rules. With many state-of-the-art AI models, it may be impossible to understand any reason for the model's output.¹⁶⁰ Plaintiffs may never be able to prove disparate treatment in such cases.

Disparate Impact

How Algorithms May Cause Disparate Impacts

Algorithmic decisionmaking could result in particular groups experiencing disproportionately negative outcomes in several ways. As a non-exhaustive list of examples, disparate impacts could result when algorithms rely on proxies for protected characteristics; when a model uses a target variable or decision rules that disfavor certain groups; when a model relies on criteria that are less predictive for particular groups; or when models are trained on biased data.

An algorithm's consideration of proxies for protected characteristics, even if not considered or proven to be disparate treatment, could give rise to disparate impact claims.¹⁶¹ Take a hypothetical AI system used to evaluate employees for "fit" with an organization's culture that disfavors applicants who use the email sign-off "Have a blessed day." Even if one could not prove the system was simply disfavoring religious applicants, such a system could result in an adverse impact on people of certain faiths, requiring the employer to present a legitimate reason for using it.

As with human decisionmaking, an algorithmic decisionmaker could also focus on target variables or decision rules that favor certain groups above others, even if most people would not consider those criteria closely related enough to a particular group to be proxy variables. For example, a landlord who designed an algorithm to select for predicted renter tenure could end up disproportionately rejecting military families if they tend to move more often than others.¹⁶² As another example, one lawsuit alleged that a tenant-screening tool had an unjustifiable disparate impact on racial minorities by considering arrests without convictions, because, the plaintiffs contended, minorities "are significantly more likely to be arrested, charged, and indicted."¹⁶³

A model could be built or trained in a way that results in the model emphasizing attributes that are less predictive for some groups than others or that results in the model failing to consider factors that are particularly predictive for certain groups.¹⁶⁴ For example, one study showed that an algorithm that used health care spending as a proxy for health care need understated the needs of Black patients. Black consumers, the study found, spend less on health care than comparably sick White consumers.¹⁶⁵ Separately, one lawsuit alleged that a tenant-screening tool disproportionately screened Black and Hispanic renters out of housing in part because it failed to

¹⁵⁹ See *supra* note 52 and accompanying text.

¹⁶⁰ See discussion *supra* note 151.

¹⁶¹ See, e.g., BIPARTISAN HOUSE TASK FORCE REPORT, *supra* note 2, at 83; Barocas & Selbst, *supra* note 11, at 691. Plaintiffs have filed lawsuits making allegations using this theory of disparate impact. *E.g.*, Knight-Bell v. Assurant, Inc., No. 26-CV-670, 2026 WL 2958723, at *2 (N.D. Ala. Oct. 1, 2026); Harper v. Sirius XM Radio, LLC, No. 25-CV-12403, 2026 WL 2942711, at *1 (E.D. Mich. Sept. 30, 2026).

¹⁶² See Barocas & Selbst, *supra* note 11, at 680, 726–27. The FHA does not cover discrimination based on military status, but some state and local laws do. *E.g.*, N.Y. EXEC. LAW § 296(2-a) (McKinney 2026).

¹⁶³ Conn. Fair Hous. Ctr. v. Corelogic Rental Prop. Sols., LLC, 369 F. Supp. 3d 362, 378 (D. Conn. 2019).

¹⁶⁴ See Gillis & Spiess, *supra* note 19, at 471–72; Scherer, King & Mrkonich, *supra* note 16, at 489–90.

¹⁶⁵ Obermeyer et al., *supra* note 124, at 449–50.

factor in housing vouchers when considering an applicant's likelihood of making payments.¹⁶⁶ Medical researchers have cautioned against training AI models on datasets using measurements from “wearable technologies that perform sub-optimally or lack validation in people with darker skin.”¹⁶⁷

The data used to train an AI model could be biased in a number of ways as well.¹⁶⁸ The training data could itself be affected by discrimination. For example, a school may train an AI model to spot the kinds of applicants it wants to admit using data from past students, and those past students may have been admitted using discriminatory criteria.¹⁶⁹ An insurance company could flag claims for additional scrutiny based on past data reflecting the bias of human claims adjudicators. In a lawsuit against a home insurance company, Black policyholders alleged that the company's claims processing algorithms “use inputs that correlate with race or learn from ‘historically biased housing and claims data.’”¹⁷⁰ An image generator trained on pictures scraped from the internet may generate images that reflect widespread social stereotypes. One study found that one such model “exhibits a bias towards the sexualization of women.”¹⁷¹ In circumstances involving the use of biased training data, commentators have asserted that an AI model “will necessarily infer rules that exhibit the same prejudice.”¹⁷²

Training data could also be biased in the sense that they are unrepresentative, and algorithms may “systematically disadvantage those who are under- or overrepresented in the dataset.”¹⁷³ For example, if certain demographic groups are underrepresented in the training data for a facial recognition model, that model may misidentify members of those groups at higher rates.¹⁷⁴ An AI model trained to estimate creditworthiness based on consumer banking data may lack adequate information about groups that tend to be unbanked.¹⁷⁵ The model may therefore predict the

¹⁶⁶ *Louis v. Saferent Sols., LLC*, 685 F. Supp. 3d 19, 26 (D. Mass. 2023).

¹⁶⁷ Ariana Mihan, Ambaris Pandey & Harriette GC Van Spall, *Mitigating the Risk of Artificial Intelligence Bias in Cardiovascular Care*, 6 LANCET DIGIT. HEALTH e749, e751 (2024), <https://www.thelancet.com/action/showPdf?pii=S2589-7500%2824%2900155-9> [<https://perma.cc/Q78G-9VDD>].

¹⁶⁸ BIPARTISAN HOUSE TASK FORCE REPORT, *supra* note 2, at 81. For examples of disparate impacts caused by algorithms using biased training data, see Obermeyer et al., *supra* note 124, at 453.

¹⁶⁹ See Stella Lowry & Gordon Macpherson, *A Blot on the Profession*, 296 BRIT. MED. J. 657, 657 (1988).

¹⁷⁰ *Huskey v. State Farm Fire & Cas. Co.*, No. 22-CV-7014, 2023 WL 5848164, at *3 (N.D. Ill. Sep. 11, 2023) (quoting First Amended Complaint at 18, *Huskey*, No. 22-CV-7014).

¹⁷¹ Ryan Steed & Aylin Caliskan, Image Representations Learned with Unsupervised Pre-Training Contain Human-Like Biases, in FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency 701, 702, 709 (Mar. 1, 2021), <https://dl.acm.org/doi/epdf/10.1145/3442188.3445932> [<https://perma.cc/4FVJ-5H88>].

¹⁷² Barocas & Selbst, *supra* note 11, at 680–82; see, e.g., FTC, BIG DATA, *supra* note 54, at 8.

¹⁷³ Barocas & Selbst, *supra* note 11, at 680–81; see *id.* at 684–87; Price II, *supra* note 108 (“For instance, if the data available for AI are principally gathered in academic medical centers, the resulting AI systems will know less about—and therefore will treat less effectively—patients from populations that do not typically frequent academic medical centers. Similarly, if speech-recognition AI systems are used to transcribe encounter notes, such AI may perform worse when the provider is of a race or gender underrepresented in training data.”).

¹⁷⁴ See CRS Report R48935, *Facial Recognition Technology: Definitions, Applications, and Policy Considerations for Congress*, by Dominique T. Greene-Sanders (2026).

¹⁷⁵ This situation has arisen with algorithms that have long been in use, such as credit scores. In 2015, the Consumer Financial Protection Bureau (CFPB) estimated that approximately 45 million American consumers could not receive credit scores and therefore had greatly reduced access to credit because they lacked the necessary credit history to be scored. These consumers were disproportionately Black and Hispanic. CFPB OFF. OF RSCH., DATA POINT: CREDIT INVISIBLES 6, 16–18 (2015), https://files.consumerfinance.gov/f/201505_cfpb_data-point-credit-invisibles.pdf [<https://perma.cc/VSZ7-JYWZ>]; see also Kate Crawford, *Think Again: Big Data*, FOREIGN POL’Y (May 10, 2013), <https://foreignpolicy.com/2013/05/10/think-again-big-data> [<https://perma.cc/A94A-EBJG>] (discussing “‘signal problems’ in big-data sets—dark zones or shadows where some citizens and communities are overlooked or underrepresented”).

behavior of members of that group less reliably.¹⁷⁶ Some scholars predict that these sorts of errors will tend to affect “historically disadvantaged groups at higher rates” because their “lives are less ‘datafied’ than the general population’s.”¹⁷⁷ Models may also extrapolate inaccurate conclusions about groups that are overrepresented, such as a predictive policing model that concludes members of one race are more likely to commit crimes because it is trained on a dataset that overrepresents that group in the criminal population.¹⁷⁸

Open Questions on Disparate Impact and Algorithms

Under antidiscrimination laws that allow for disparate impact claims—which, as explained above, is not all of them—using disparate impact doctrine would relieve plaintiffs in algorithmic discrimination cases of the burden of proving intent or showing what variables an algorithm based its output on. Nevertheless, these cases would likely not be straightforward. There are a number of unanswered questions about how disparate impact doctrine might apply in this context.

What Is the Specific Practice Causing the Disparity?

As discussed above, disparate impact plaintiffs must identify the specific practice that causes a disparity. In algorithmic discrimination cases, one question is whether using an algorithm or AI system is a specific enough practice, or whether plaintiffs will need to prove something more detailed—how the algorithm works, bias in the training data, or specific problematic inputs, for example.¹⁷⁹ Title VII allows plaintiffs to analyze the entire decisionmaking process as “one employment practice” when the “elements” of the process “are not capable of separation for analysis.”¹⁸⁰ While that provision seems to point to allowing plaintiffs to challenge an algorithm or AI system as a whole, whether courts will follow that interpretation in practice, and whether they will apply that rule outside of Title VII, remains to be seen.¹⁸¹

¹⁷⁶ See BIPARTISAN HOUSE TASK FORCE REPORT, *supra* note 2, at 82 (acknowledging “biases [that] arise when algorithms are trained on one type of data and cannot accurately extrapolate beyond that”).

¹⁷⁷ Barocas & Selbst, *supra* note 11, at 684–85 (quoting Jonas Lerman, *Big Data and Its Exclusions*, 66 STAN. L. REV. ONLINE 55, 57 (2013)); see Kim, *supra* note 5, at 885–86 (“[E]vidence suggests that errors are more likely to occur for members of subgroups that are farther from the mainstream.”). For example, one article examining potential racial bias in AI models used in medical settings reported that less data were collected about minority groups during the COVID-19 pandemic, because “African Americans were less likely to be offered COVID-19 testing, less likely to have testing clinics available in their neighborhood, and less likely to be able to follow public health guidelines with their living situation.” Atin Jindal, *Misguided Artificial Intelligence: How Racial Bias Is Built Into Clinical Models*, 2 J. OF BROWN HOSP. MEDICINE 1, 2 (2023), <https://doi.org/10.56305/001c.38021> (discussing other ways minority groups may be vulnerable to underrepresentation in the data used to train clinical AI models); see also, e.g., Mihan, Pandey & Van Spall, *supra* note 167, at e750 (“[F]emale patients are under-referred for intracardiac devices and cardiac synchronization therapy and would be under-represented in related digital datasets.” (footnote omitted)).

¹⁷⁸ See, e.g., Rashida Richardson, Jason M. Schultz & Kate Crawford, *Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice*, 94 N.Y.U. L. REV. ONLINE 192, 199, 201 (2019); Kristian Lum & William Isaac, *To Predict and Serve?*, 13 SIGNIFICANCE 14, 15–16 (2016), <https://academic.oup.com/jrssig/article/13/5/14/7029190> [<https://perma.cc/8KDE-JTL9>].

¹⁷⁹ Gillis & Spiess, *supra* note 19, at 479 n.57 (stating that courts “rarely provide guidance” on this question).

¹⁸⁰ 42 U.S.C. § 2000e-2(k)(1)(B).

¹⁸¹ Courts may not take an all-or-nothing approach. At the initial stage of a lawsuit, they may allow plaintiffs to proceed with much less information about how an algorithm works, while demanding more at later stages once more information has been exchanged between the parties. One court, for example, allowed plaintiffs to plead a disparate impact claim against their home insurance company by pointing to the company’s use of algorithmic claims-processing software. “Plaintiffs need not pin down the policy with greater precision *at this stage*,” the court stated. *Huskey v. State Farm Fire & Cas. Co.*, No. 22-CV-7014, 2023 WL 5848164, at *8 (N.D. Ill. Sep. 11, 2023) (emphasis added). Discovery in that case is ongoing; the court might require more at a later point. In contrast, a district court dismissed an (continued...)

Relatedly, disparate impact claims are often brought as class actions, which require the plaintiffs to show a common cause for their injuries.¹⁸² At least one court has refused to certify a class alleging discriminatory mortgage lending when the plaintiffs pled only that the lender's algorithm led to discriminatory results. The court held that the plaintiffs had not identified a specific enough common reason for their injuries.¹⁸³ The data necessary to prove commonality, or other parts of a disparate impact claim, may be difficult to obtain. For example, the plaintiffs in *Mobley v. Workday*, the case discussed above alleging that the defendant's resume-screening software causes disparate impacts, have struggled to obtain data about applicants to third-party employers who used the defendant's software. According to the plaintiffs, several of the defendant's customers responded to subpoenas for the data by directing the plaintiffs back to the defendant.¹⁸⁴ The court has refused to order the defendant to produce the data, agreeing with the defendant that the data belong to its individual customers and are not in the defendant's legal "possession" because the information "resides on segregated, password-protected customer tenants that Workday is contractually prohibited from freely accessing."¹⁸⁵ Commentators observe that this restriction could inhibit the plaintiffs' ability to certify a class and potentially to prove a disparate impact at all.¹⁸⁶

It is also unclear whether plaintiffs could successfully challenge an element of algorithmic decisionmaking when the final output does not create a disparate impact. On the one hand, under Title VII, the Supreme Court has held that plaintiffs may challenge one part of a decisionmaking process even when the process as a whole does not cause a disparate impact. In *Connecticut v. Teal*, the Court allowed plaintiffs to challenge their employer's use of a written examination in its promotion process.¹⁸⁷ Although the examination disproportionately screened out Black employees, the employer nevertheless ultimately promoted Black employees at a higher rate than White ones considering other criteria.¹⁸⁸ A court might read *Connecticut v. Teal* to allow plaintiffs to challenge elements of an algorithmic system even if the system overall does not cause a disparate impact. A broad reading of *Teal* might even permit plaintiffs to challenge a single input among thousands.

On the other hand, *Teal* did not address the viability of a disparate impact challenge to, for instance, a single question on the exam at issue, and the Supreme Court has never held that plaintiffs could challenge such subcomponents of a decisionmaking process. In a different context, considering an FHA challenge to a federal grant program, the D.C. Circuit observed that

ADEA disparate impact claim based on an employer's "use of modern technology and artificial intelligence" because the plaintiff "fail[ed] to specifically identify which features of the technology have an adverse impact on older workers." *Robinson v. Piedmont Healthcare, Inc.*, No. 23-CV-04099, 2024 WL 7162870, at *3 n.3 (N.D. Ga. Sept. 17, 2024).

¹⁸² See CRS In Focus IF12763, *Class Action Lawsuits: An Introduction*, by Bryan L. Adkins (2024).

¹⁸³ *In re Wells Fargo Mortg. Discrimination Litig.*, 794 F. Supp. 3d 662, 670–71 (N.D. Cal. 2025). Over a dissent in a different case, the Fourth Circuit held that plaintiffs could satisfy the commonality requirement by pleading that the defendant "uses a *single* algorithm to evaluate *every* applicant." *Oliver v. Navy Fed. Credit Union*, 167 F.4th 106, 116 (4th Cir. 2026). *But see id.* at 127–28 (Richardson, J., concurring in the judgment in part and dissenting in part) (arguing that the plaintiffs had not satisfied the commonality requirement in part because the complaint lacked information about how the challenged algorithm worked and how the defendant used it).

¹⁸⁴ Plaintiff's Motion to Compel Production of Requested Documents and Discovery at 17 n.6, *Mobley v. Workday, Inc.*, 23-CV-00770 (N.D. Cal. filed Apr. 20, 2026), Dkt. No. 310.

¹⁸⁵ Discovery Order at 2, *Mobley*, No. 23-CV-00770 (N.D. Cal. July 12, 2026), Dkt. No. 394 (discovery order of magistrate judge), *aff'd*, 2026 WL 2211616 (N.D. Cal. July 28, 2026).

¹⁸⁶ Chris Marr, *Workday's Discovery Wins Make It Harder to Prove AI Bias Case*, BLOOMBERG LAW (Aug. 21, 2026, at 09:14 ET), <https://news.bgov.com/daily-labor-report/workdays-discovery-wins-make-it-harder-to-prove-ai-bias-case>.

¹⁸⁷ 457 U.S. 440, 442 (1982).

¹⁸⁸ *Id.* at 443–44.

focusing on each element of a grant formula, instead of the formula as a whole, would make it nearly impossible for an agency to use any formula “without *some* element having a disproportionate impact on a protected group.”¹⁸⁹ As a practical matter, potential plaintiffs may be unlikely to know enough about an algorithmic system to bring such tailored challenges,¹⁹⁰ but courts may have to decide which parts of an algorithmic process, if any, are separable enough for plaintiffs to contest.

What Is a Legally Adequate Justification?

After a plaintiff proves a specific practice causing a disparate impact, the defendant must show that the practice serves a legitimate, nondiscriminatory purpose. As explained above, there is general uncertainty regarding when a defendant’s justification is legally adequate.¹⁹¹ When defendants use AI systems to make decisions, it raises the question of whether they need to show more than that the system is predictive of some legitimate consideration, such as a renter’s tenure or a worker’s productivity. In other words, do defendants need to show a logical connection between the inputs the system considers and the defendant’s valid goal, or only that the system is predictively accurate, even if no one knows why?¹⁹² Commentators have pointed out that if a defendant needs to show only that an AI system has some predictive value for the defendant’s target variable, most defendants will pass the test.¹⁹³ Many AI models are designed to predict outcomes based on historical data and inputs. As long as the target variable appears legitimate, one legal scholar has opined with respect to Title VII that “to ask whether the model is ‘job related’ in the sense of ‘statistically correlated’ is tautological.”¹⁹⁴

Models trained on biased, incomplete, or erroneous data may prove less predictive of reality, and an AI model could be so badly designed or trained that a court would not consider its use to be related to a legitimate purpose.¹⁹⁵ Courts may need to decide how far to require defendants to go to make sure a model is predictive in the environment in which it is used, and just how predictive it needs to be. As a practical matter, however, putative plaintiffs might have difficulty rebutting a

¹⁸⁹ Greater New Orleans Fair Hous. Action Ctr. v. HUD, 639 F.3d 1078, 1086 (D.C. Cir. 2011).

¹⁹⁰ See Knight-Bell v. Assurant, Inc., No. 26-CV-670, 2026 WL 2958723, at *8 (N.D. Ala. Oct. 1, 2026) (dismissing claim that an employer’s automated resume-screening tool used keyword searches that disparately screened out Black applicants in part because the plaintiff did not identify any of the keywords allegedly used); Saas v. Major, Lindsey & Afr., LLC, No. 23-CV-02102, 2024 WL 2113654, at *5 (D. Md. May 10, 2024) (dismissing case alleging that employer used an algorithm with a disparate impact on mothers because the allegations were based “solely on supposition that Defendants use artificial intelligence tools”), *aff’d*, No. 24-1527, 2024 WL 4850774 (4th Cir. Nov. 21, 2024); Roy Maurer, *New York City AI Law Is a Bust*, SOC’Y FOR HUM. RES. MGMT. (Feb. 18, 2024), <https://www.shrm.org/topics-tools/news/technology/new-york-city-ai-law> [<https://perma.cc/X7S4-R8FC>].

¹⁹¹ See *supra* “Legally Adequate Justification.”

¹⁹² See Scherer, King & Mrkonich, *supra* note 16, at 490–91 (“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 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.”).

¹⁹³ E.g., Prince & Schwarcz, *supra* note 11, at 1305; Bruckner, *supra* note 16, at 35–36; Kim, *supra* note 5, at 908; Barocas & Selbst, *supra* note 11, at 707, 709.

¹⁹⁴ Kim, *supra* note 5, at 866–67.

¹⁹⁵ *Id.* at 887; Barocas & Selbst, *supra* note 11, at 707–08; see also Dothard v. Rawlinson, 433 U.S. 321, 331 (1977) (holding that defendants had not identified a legitimate interest in imposing certain height and weight requirements on correctional officers when they “produced no evidence correlating the height and weight requirements with the requisite amount of strength thought essential to good job performance”). But see Nat’l Fair Hous. All. v. Bank of Am., N.A., No. 18-CV-01919, 2025 WL 2030225, at *18 (D. Md. July 21, 2025) (stating that courts do not consider “the successfulness (or lack thereof) of Defendants’ policies in achieving their professed business objectives” at this point in the analysis).

defendant's argument that a model is predictive due to lack of access to the model and its training data, difficulty reproducing the model's decision rules, or expense in designing and executing a study, among other reasons.

Rather than focus on whether a model is truly predictive of the user's target variable, courts might instead focus on whether defendants can articulate some logical connection between an algorithm's decision rules and the defendant's legitimate end.¹⁹⁶ For example, if a financial model evaluating creditworthiness downranks people who attend historically Black colleges and universities, or institutions with a religious affiliation, a court could require the lender to prove, or at least posit, a connection between those variables and creditworthiness divorced from borrowers' race or religion.¹⁹⁷ This approach would require the users of an AI model to know how it arrived at its output, which may be impossible.¹⁹⁸ It may also require some limits on which factors a defendant needs to explain, given the impracticality of addressing thousands, millions, or even billions of parameters governing a model's output.¹⁹⁹

Courts may also consider how much a defendant's interests in cost savings or efficiency can count in and of themselves as legitimate justifications for using an algorithm, even when a model's predictions are flawed. Cost savings may not be "job related," but they could be "business necessities," underscoring the lack of clarity in case law as to whether these are the same standard or different standards that must be met separately.²⁰⁰ Accepting time and money savings as legitimate justifications could lighten defendants' burden.

How Could Plaintiffs Identify a Less Discriminatory Alternative?

Plaintiffs may struggle to suggest equally effective, less discriminatory alternatives to allegedly discriminatory algorithms "given that AIs are indeed uniquely effective at optimizing their programmed objective."²⁰¹ Some scholars have suggested that plaintiffs may be able to identify less discriminatory algorithms, for example, by using "techniques" to "'repair' data sets" to "cleanse . . . data of its proxies for protected classes."²⁰² Others are skeptical that disparate

¹⁹⁶ See, e.g., Grimmelmann & Westreich, *supra* note 127, at 170 ("We believe that where a plaintiff has identified a disparate impact, the defendant's burden to show a business necessity requires it to show not just that its model's scores are not just *correlated* with job performance but *explain* it."); see also Scherer, King & Mrkonich, *supra* note 16, at 486–88 (discussing the difference between traditional employment tests, where developers selected questions based on their predicted connection to a work-related skill, and algorithmic tools, which "do not consider theoretical or logical relationships between variables" and are more likely to "capitalize on chance correlations").

¹⁹⁷ See STUDENT BORROWER PROTECTION CENTER, EDUCATIONAL REDLINING 15–19 (2020), <https://protectborrowers.org/wp-content/uploads/2020/02/Education-Redlining-Report.pdf> [<https://perma.cc/2MD7-XEAQ>].

¹⁹⁸ See discussion *supra* note 160. In addition to the opacity inherent in black box algorithms, many developers of algorithmic decisionmaking systems may consider their inventions proprietary and resist releasing detailed information about them. See, e.g., CHRISTIAN, *supra* note 146, at 7; Obermeyer et al., *supra* note 124, at 447.

¹⁹⁹ See *What Are Model Parameters?*, IBM (May 5, 2025), <https://www.ibm.com/think/topics/model-parameters> [<https://perma.cc/X68Y-DU84>] ("Model parameters are the learned values within a machine learning model that determine how it maps input data to outputs, such as generated text or a predicted classification. . . . Large models, such as those used in generative AI, can have billions of parameters."); Prince & Schwarcz, *supra* note 11, at 1317; Scherer, King & Mrkonich, *supra* note 16, at 490–91.

²⁰⁰ See *supra* "Legally Adequate Justification."

²⁰¹ Prince & Schwarcz, *supra* note 11, at 1305.

²⁰² Bruckner, *supra* note 16, at 36–37. In 2023, EEOC issued guidance anticipating that employers developing "algorithmic decision-making tools" may "produce a variety of comparably effective alternative algorithms," and that "[f]ailure to adopt a less discriminatory algorithm that was considered during the development process therefore may give rise to liability." EEOC, EEOC-NVTA-2023-2, SELECT ISSUES: ASSESSING ADVERSE IMPACT IN SOFTWARE, (continued...)

impacts can be eliminated this way.²⁰³ In any case, plaintiffs may lack the technical expertise, access to the algorithm, appropriate training data, or other necessary resources to suggest these changes or to develop independent alternatives.²⁰⁴

Plaintiffs may especially struggle to prove that any alternative model was available to the defendant.²⁰⁵ As reviewed above, to be available, courts have held, an alternative must be concrete rather than theoretical, and it must be an option the defendant should have known about.²⁰⁶ Defendants in lawsuits involving algorithmic discrimination may not have developed the challenged software themselves and may not have had the power to affect a model's design.²⁰⁷ Even if plaintiffs prove that theoretically, someone could have developed an equally valid but less discriminatory algorithm, if that program was not commercially available, there is no obvious legal avenue for claiming that an end user should have known about it or created it themselves.

It may also be difficult to obtain the necessary proof that an alternative would work just as well. Some courts have held plaintiffs to a high evidentiary bar to make this showing. In *Allen v. City of Chicago*, for example, the Seventh Circuit acknowledged that the city had previously successfully used the plaintiffs' proposed promotions process without generating a disparate impact but concluded nevertheless that prior successful use "merely predicts, but does not establish" future success.²⁰⁸ In an area of emerging technology, such a standard of proof may circumscribe the kinds of solutions available to plaintiffs. Plaintiffs seeking to force defendants to deploy a differently designed AI model, for example, may not have even past results to rely on.

If plaintiffs cannot rely on less discriminatory algorithms, they would need to come up with other solutions that serve defendants' needs. This, too, raises the question of whether or how a defendant in a disparate impact suit may rely on cost control and efficiency—that is, in the context of the less discriminatory alternative inquiry, whether and at what threshold a defendant could argue that an alternative is less effective because of its higher cost or other burdens.²⁰⁹

Considerations for Congress

Thus far, courts have considered few cases alleging algorithmic discrimination. How federal antidiscrimination law may apply to such claims is therefore largely undeveloped. In this fast-

ALGORITHMS, AND ARTIFICIAL INTELLIGENCE USED IN EMPLOYMENT SELECTION PROCEDURES UNDER TITLE VII OF THE CIVIL RIGHTS ACT OF 1964 (2023) [hereinafter EEOC, SELECT ISSUES], <https://www.eeoc.gov/laws/guidance/select-issues-assessing-adverse-impact-software-algorithms-and-artificial> [<https://perma.cc/WWX9-TLXY>]. That guidance is no longer available on EEOC's website.

²⁰³ See Gillis & Spiess, *supra* note 19, at 468–70.

²⁰⁴ See Barocas & Selbst, *supra* note 11, at 710 (explaining that "[f]ixing the models . . . is not a trivial task" and elaborating on hurdles plaintiffs may face).

²⁰⁵ *Id.*

²⁰⁶ See *supra* "Less Discriminatory Alternative."

²⁰⁷ See Griffith, *supra* note 107, at 375 ("Even if a financial institution can get in-depth access to see how a third-party AI works, it is unlikely to have the expertise to adequately interpret the algorithm or evaluate its compliance with fair lending laws.").

²⁰⁸ 351 F.3d 306, 315 (7th Cir. 2003); *accord* Johnson v. City of Memphis, 770 F.3d 464, 477 (6th Cir. 2014).

²⁰⁹ See *supra* notes 92–93 and accompanying text; *Resetting Antidiscrimination Law*, *supra* note 129, at 1567 ("[A]s the use of AI becomes ubiquitous, it may be difficult to argue that a non-AI system could also adequately meet an employer's legitimate business needs." (footnote omitted)). OLC has opined that plaintiffs cannot succeed if their proposed alternative "is 'less effective, more burdensome, or more costly,'" or "if the cost or effectiveness of an alternative is reasonably debatable." OLC Disparate Impact Op., *supra* note 59, at 17–18.

changing area of technological development and litigation, Congress could wait to see what fact patterns emerge and how courts handle them before considering legislative interventions.

If Congress did seek to legislate on how federal antidiscrimination law applies to algorithmic decisionmaking, it would not be writing on a blank slate. Congress could look to various approaches proposed by states and federal agencies. Some states and localities have passed laws or issued regulations to prevent and/or remedy algorithmic “bias,”²¹⁰ some of which are described further below. Federal agencies over the past two decades have also issued regulations and guidance on algorithmic discrimination. Some of these rules persist, while others have been significantly revised or rescinded.²¹¹ Several federal agency documents discuss mitigating the risk of algorithmic discrimination in documents addressing AI regulation more generally. Such documents include the Government Accountability Office’s (GAO’s) “accountability framework” for the federal development and deployment of AI;²¹² the National Institute of Standards and Technology’s (NIST’s) Artificial Intelligence Risk Management Framework;²¹³ the Office of Management and Budget’s (OMB’s) memorandum on “Accelerating Federal Use of AI through Innovation, Governance, and Public Trust,”²¹⁴ and the National Telecommunications and Information Administration’s Artificial Intelligence Accountability Policy Report,²¹⁵ to name a few.²¹⁶ Congress could delve more deeply into the approaches attempted by other regulators as potential models for federal legislation.

The following is a discussion of some of the questions and options Congress might consider were it to seek to address potential algorithmic discrimination through legislation.²¹⁷

²¹⁰ E.g., CAL. HEALTH & SAFETY CODE § 1367.01(k) (West 2026) (regulating the use of algorithms in health insurance); CAL. INS. CODE § 10123.135(j) (West 2026) (regulating the use of algorithms in disability insurance); TEX. BUS. & COM. CODE § 552.056 (West 2026) (prohibiting the development or deployment of AI with an intent to discriminate); N.J. ADMIN. CODE § 13:16-3.2(c) (West 2026) (explaining how the use of “automated employment decisions tools” may result in unlawful disparate impacts).

²¹¹ Examples of agency regulations and guidance currently in force include 42 C.F.R. § 92.210 (2025) (prohibiting federally funded health entities from discriminating “through the use of patient care decision support tools”) and DOJ, *Algorithms, Artificial Intelligence, and Disability Discrimination in Hiring*, ADA.GOV (May 12, 2022), <https://www.ada.gov/resources/ai-guidance> [<https://perma.cc/245U-KG28>]. Examples of rescinded guidance include U.S. DEP’T OF EDUC., *AVOIDING THE DISCRIMINATORY USE OF ARTIFICIAL INTELLIGENCE* (2025), <https://www.ed.gov/media/document/avoiding-discriminatory-use-of-artificial-intelligence-112274.pdf> [<https://perma.cc/Q48D-Q9S9>]; CFPB, CIRCULAR 2023-03, *ADVERSE ACTION NOTIFICATION REQUIREMENTS AND THE PROPER USE OF THE CFPB’S SAMPLE FORMS PROVIDED IN REGULATION B* (2023) [hereinafter CFPB CIRCULAR 2023-03], https://files.consumerfinance.gov/f/documents/cfpb_adverse_action_notice_circular_2023-09.pdf [<https://perma.cc/A495-7MZ9>]; and CFPB, CIRCULAR 2022-03, *ADVERSE ACTION NOTIFICATION REQUIREMENTS IN CONNECTION WITH CREDIT DECISIONS BASED ON COMPLEX ALGORITHMS* (2022) [hereinafter CFPB CIRCULAR 2022-03], https://files.consumerfinance.gov/f/documents/cfpb_2022-03_circular_2022-05.pdf [<https://perma.cc/D5YU-EDRX>].

²¹² GAO, GAO-21-519SP, *ARTIFICIAL INTELLIGENCE: AN ACCOUNTABILITY FRAMEWORK FOR FEDERAL AGENCIES AND OTHER ENTITIES* (2021) [hereinafter GAO AI ACCOUNTABILITY FRAMEWORK], <https://www.gao.gov/assets/gao-21-519sp.pdf> [<https://perma.cc/XM3A-V4BA>].

²¹³ NIST AIRMF 1.0, *supra* note 4.

²¹⁴ Memorandum from Russell T. Vought, Dir., Off. of Mgmt. and Budget, on Accelerating Federal Use of AI Through Innovation, Governance, and Public Trust, M-25-21 (2025), <https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf> [<https://perma.cc/HV3T-C5A3>].

²¹⁵ NAT’L TELECOMM. & INFO. ADMIN., *ARTIFICIAL INTELLIGENCE ACCOUNTABILITY POLICY REPORT* (2024), <https://www.ntia.gov/sites/default/files/ntia-ai-report-final.pdf> [<https://perma.cc/AR82-WJTV>].

²¹⁶ Other federal frameworks for regulating AI are discussed in CRS Report R48555, *Regulating Artificial Intelligence: U.S. and International Approaches and Considerations for Congress*, by Laurie Harris (2025).

²¹⁷ CRS also has other products that review approaches to the broader regulation of AI, as well as products addressing (continued...)

Regulating Algorithmic Decisionmaking Under General Antidiscrimination Laws or Using Standalone Antidiscrimination Rules

Congress may consider whether to regulate algorithmic decisionmaking under general antidiscrimination authorities or whether to pass new antidiscrimination laws specifically targeting algorithms or AI.²¹⁸ Congress could amend existing antidiscrimination laws to clarify areas that may generate confusion in the context of algorithmic decisionmaking. Congress could, for example, clarify the intent standard for disparate treatment; the standard to show an adequate interest justifying a policy or practice that causes disparate impacts; or the standards for less discriminatory alternatives. Such changes could, depending on how they were framed, impact both algorithmic discrimination claims and other kinds of claims.²¹⁹

Alternatively, rather than relying on general antidiscrimination principles to reach algorithmic decisionmaking, Congress could decide to specify what uses of algorithms are or are not discriminatory. Some of the ways algorithms may be alleged to be “unfair” are difficult to classify as disparate impact or disparate treatment, or as discrimination under current law at all, leaving unclear what standards apply. For example, it is not clear if using an algorithm trained on unrepresentative data violates antidiscrimination law, and assuming it does, whether the relevant antidiscrimination framework would be disparate treatment or disparate impact. Congress could legislate specific rules governing when certain uses of algorithmic models constitute discrimination. In so doing, it could try to adhere to current legal frameworks or write new ones specific to this context.

Adopting Broad Versus Tailored Requirements

Congress could seek to apply antidiscrimination standards to algorithmic decisionmaking across a variety of domains or choose to focus on particular contexts. Some state laws apply the same or similar standards across many contexts in which algorithms may be used. Colorado, for example, passed a bill with transparency requirements for when “automated decision-making technology” is used to “materially influence” a “consequential decision,” including access to education, employment, housing, financial services, insurance, health care, or “essential government services and public benefits.”²²⁰ Other laws are more narrowly targeted. Montana and New Hampshire, for example, have passed laws regulating the states’ own use of potentially discriminatory AI.²²¹

In addition to determining whether to target algorithmic decisionmaking broadly or in specific areas, Congress may consider how narrow or broad a range of technology it wishes to regulate.

AI regulation in specific domains. See CRS Report R48909, *Artificial Intelligence: CRS Products*, by Laurie Harris and Rachael D. Roan (2025); CRS Report R47644, *Artificial Intelligence: Overview, Recent Advances, and Considerations for the 118th Congress*, by Laurie Harris (2023); CRS Report R46795, *Artificial Intelligence: Background, Selected Issues, and Policy Considerations*, by Laurie Harris (2021).

²¹⁸ *Resetting Antidiscrimination Law*, *supra* note 129, at 1563 (framing the debate over whether new technology can be regulated using old frameworks or whether it requires new ones).

²¹⁹ *Cf. id.* at 1563–64 (observing that attempts to regulate AI decisionmaking may generate new ideas about how to regulate human decisionmaking).

²²⁰ 2026 Colo. Sess. Laws 569 (to be codified at COLO. REV. STAT. §§ 6-1-1701 to -1709).

²²¹ MONT. CODE ANN. §§ 2-21-101 to -104 (2025); N.H. REV. STAT. ANN. §§ 5-D:1 to D:6 (2025).

Some laws regulate a wide range of automated decisionmaking tools,²²² while others apply only to AI²²³ or to a particular subset of automated tools, like facial recognition technology.²²⁴

Broad and tailored approaches to regulation each involve trade-offs. For example, broad regulations can tackle many instantiations of similar problems and apply consistent standards across domains. They also risk treating conduct, entities, or problems that are different similarly, potentially creating confusion or resulting in under- or over-regulation. Tailored regulations can take into account the specificities of different sectors or technologies and adjust accordingly, but they also risk leaving manifestations of a problem unaddressed.

Deciding Who Should Be Regulated

Congress may wish to consider whether federal antidiscrimination laws target the right defendants when assigning liability for discriminatory algorithms. Antidiscrimination laws often, but not always, cover specific actors. The federal employment discrimination laws, for example, forbid discrimination by an “employer,” “employment agency,” “labor organization,” or “joint labor-management committee.”²²⁵ ECOA forbids discrimination by a “creditor.”²²⁶ Several provisions of the FHA, in contrast, provide only that “it shall be unlawful” to engage in certain discriminatory conduct, without specifying a particular kind of person or entity to whom the law applies.²²⁷ The parties responsible for developing allegedly discriminatory algorithms may not be the same parties who deploy them, and the law governing a claim may or may not cover an algorithm’s developer.²²⁸ Courts have held that some laws can render software developers liable for discriminatory algorithms, depending on the facts.²²⁹

²²² E.g., 2026 Colo. Sess. Laws 569, at 570–71 (defining “automated decision-making technology”) (to be codified at COLO. REV. STAT. § 6-1-1701(2)); CAL. CODE REGS. tit. 2, § 11008.1(a) (2026) (defining “automated-decision system”).

²²³ E.g., CONN. GEN. STAT. § 4a-2e (2026) (requiring inventories and assessments of “systems that employ artificial intelligence” in state government).

²²⁴ For a discussion of different definitions of facial recognition technology and related regulations and policy proposals, see CRS Report R48935, *Facial Recognition Technology: Definitions, Applications, and Policy Considerations for Congress*, by Dominique T. Greene-Sanders (2026).

²²⁵ 29 U.S.C. §§ 623(a)–(c), 630(d) (ADEA); 42 U.S.C. § 2000e-2(a)–(d) (Title VII); *id.* §§ 2000ff-1 to 2000ff-4 (GINA); *id.* § 12111(2) (Title I of the ADA).

²²⁶ 15 U.S.C. § 1691.

²²⁷ 42 U.S.C. § 3604. The Second Circuit explained, “The FHA does not exclude any class of defendants from its scope. Whether any defendant is subject to liability under the FHA turns on whether the plaintiff has established a *prima facie* case of discrimination—including whether the defendant proximately caused the alleged harm.” *Conn. Fair Hous. Ctr. v. CoreLogic Rental Prop. Sols., LLC.*, 167 F.4th 605, 620 (2d Cir. 2026). Some provisions of the FHA do apply to specific categories of people or entities, for example, “person[s] in the business of selling or renting dwellings,” 42 U.S.C. § 3603(b)–(c), or “person[s] or other entit[ies] whose business includes engaging in residential real-estate transactions,” *id.* § 3605.

²²⁸ More than one entity may be involved in an algorithm’s development, a factor Congress may also consider. *See, e.g.*, Katherine Lee, A. Feder Cooper & James Grimmelmann, *Talkin’ ‘Bout AI Generation: Copyright and the Generative-AI Supply Chain*, 72 J. COPYRIGHT SOC’Y 251, 291–302 (2025) (describing the stages of generative AI development and the different actors involved in each).

²²⁹ *See Mobley v. Workday, Inc.*, 740 F. Supp. 3d 796, 805–08 (N.D. Cal. 2024); *Louis v. Saferent Sols., LLC*, 685 F. Supp. 3d 19, 34–35 (D. Mass. 2023); *Conn. Fair Hous. Ctr. v. CoreLogic Rental Prop. Sols., LLC*, 369 F. Supp. 3d 362, 372 (D. Conn. 2019). The district court in *CoreLogic* ruled on the motion to dismiss that the developer of tenant-screening software could be liable under the FHA. However, at trial the court held that the developer was not liable because the plaintiffs’ evidence did not show that the developer “makes housing unavailable.” The landlord company, the court determined, made all the relevant decisions, and the developer’s software did not “suppress” the information those decisions were based on. *Conn. Fair Hous. Ctr. v. CoreLogic Rental Prop. Sols., LLC*, No. 18-CV-705, 2023 WL 4669482, at *15–20 (D. Conn. July 20, 2023), *aff’d in relevant part, vacated in part, rev’d in part*, 167 F.4th 605 (2d Cir. 2026).

Given this, Congress could determine that current antidiscrimination law does not cover all of the actors potentially responsible for discriminatory algorithms. Conversely, it could determine that the law sweeps in parties too remote from discriminatory decisionmaking. Congress could amend antidiscrimination laws to specify which entities are to be held liable for algorithmic discrimination and under what facts.

Considering Requirements Beyond Nondiscrimination

Beyond establishing new antidiscrimination requirements in federal law, or amending existing ones, Congress could consider other frameworks for addressing algorithmic decisionmaking exercising different levels of oversight, such as transparency, audit, or design requirements. Other requirements could be added to the current requirements of federal antidiscrimination law or substituted for them. Congress might determine, for example, that requiring litigating parties to explain how AI systems work is prohibitively difficult. It could allow AI developers or users to defend against discrimination claims by showing that their systems meet federal requirements for transparency, testing, design, or other features. Similarly, it could allow plaintiffs to prove claims by showing that an AI system did not comply with such requirements. Congress could also impose requirements on AI systems without providing for a private right of action to enforce them, leaving enforcement to federal agencies or the market.

Transparency

In some cases, plaintiffs have struggled to challenge algorithmic decisionmaking because they did not know whether algorithmic tools were being used or how they worked.²³⁰ Regulators and commentators commonly suggest that algorithmic decisionmaking should be transparent.²³¹ Transparency can mean a number of things, from reporting the use of an algorithm, to sharing information about how the algorithm was trained, to disclosing how it functions, to allowing individuals to correct the data on which an algorithm relies.²³² Transparency requirements can direct reporting to regulators, or to the public, or both. As an example of a transparency-focus law, Illinois requires employers to disclose to job applicants when they are using AI to analyze video interviews and to “explain[] how the artificial intelligence works and what general types of characteristics it uses to evaluate applicants.”²³³ Colorado law requires a “deployer” of certain automated decisionmaking systems to give consumers the opportunity to correct “factually inaccurate or materially inaccurate personal data” used by such systems in some circumstances.²³⁴

Congress has already imposed transparency requirements to combat discrimination in some circumstances. Under ECOA, creditors must provide written statements telling applicants the

²³⁰ See *Knight-Bell v. Assurant*, No. 26-CV-670, 2026 WL 2958723, at *8 (N.D. Ala. Oct. 1, 2026) (dismissing claim that an employer’s automated resume-screening tool used keyword searches that disparately screened out Black applicants in part because the plaintiff did not identify any of the keywords allegedly used); *Saas v. Major, Lindsey & Afr., LLC*, No. 23-CV-02102, 2024 WL 2113654, at *5 (D. Md. May 10, 2024) (dismissing case alleging that employer used an algorithm with a disparate impact on mothers because the allegations were based “solely on supposition that Defendants use artificial intelligence tools”), *aff’d*, No. 24-1527, 2024 WL 4850774 (4th Cir. Nov. 21, 2024).

²³¹ See, e.g., Chris Marr, *AI Hiring Bias Laws Limited by Lack of Transparency in Tools*, DAILY LABOR REPORT (Nov. 12, 2024, at 05:30 ET), <https://news.bloomberglaw.com/daily-labor-report/ai-hiring-bias-laws-limited-by-lack-of-transparency-in-tools>; Prince & Schwarcz, *supra* note 11, at 1311; Kim, *supra* note 5, at 879–80.

²³² See NIST AI RMF 1.0, *supra* note 4, at 15–16.

²³³ 820 ILL. COMP. STAT. 42/5 (2026). The law gives applicants the right to opt out of interviews analyzed by AI. *Id.*

²³⁴ 2026 Colo. Sess. Laws 569, 578 (to be codified at COLO. REV. STAT. § 6-1-1705).

“specific reasons” for an “adverse action.”²³⁵ As several circuit courts have explained, the notice requirement both “discourages discrimination” and “educates consumers” as to how they can improve their creditworthiness.²³⁶ In 2022 and 2023, the Consumer Financial Protection Bureau (CFPB) issued ECOA guidance clarifying that creditors could not “use complex algorithms when doing so means they cannot provide the specific and accurate reasons for adverse actions.”²³⁷ The CFPB explained that creditors could not simply explain that an algorithm dictated the decision.²³⁸ This guidance was withdrawn in 2025.²³⁹

Congress could take several actions around transparency in algorithmic decisionmaking. It could, for example, amend antidiscrimination laws to include transparency requirements; pass stand-alone transparency laws about algorithmic decisionmaking; or direct agencies to issue guidance interpreting existing transparency requirements in the context of algorithms. Congress could also legislate to require less transparency around algorithmic systems, if it were concerned that transparency requirements stifle innovation or expose trade secrets, for example.²⁴⁰ Congress could amend ECOA to allow creditors to report the reasons for adverse actions at a higher level of generality.

Any approach Congress takes could face legal challenges and constitutional scrutiny. Transparency requirements in state laws have been challenged on various constitutional (and other) grounds, including the Takings Clause²⁴¹ and First Amendment.²⁴² One AI developer, for example, argued that California’s Artificial Intelligence Training Data Transparency law constituted a taking because, the developer argued, the law required disclosure of protected trade secrets.²⁴³ The developer also argued that the law’s compelled disclosures violated the First Amendment.²⁴⁴ The district court accepted that “hypothetically, datasets and details about them could be trade secrets,” but it denied a preliminary injunction, reasoning that the developer had not pled the existence of trade secrets with enough specificity.²⁴⁵ On the First Amendment claim, the district court held that the law regulated commercial speech, subjecting it to a lower level of scrutiny than most content-based regulations of noncommercial speech receive, and denied the

²³⁵ 15 U.S.C. § 1691(d). The reason may be given verbally in limited circumstances. *Id.* § 1691(2)(B), (5). Regulations implementing ECOA elaborate that “[t]he statement of reasons for adverse action . . . must be specific and indicate the principal reason(s) for the adverse action. Statements that the adverse action was based on the creditor’s internal standards or policies or that the applicant, joint applicant, or similar party failed to achieve a qualifying score on the creditor’s credit scoring system are insufficient.” 12 C.F.R. § 202.9(b)(2).

²³⁶ *Treadway v. Gateway Chevrolet Oldsmobile Inc.*, 362 F.3d 971, 977 (7th Cir. 2004); *accord* *Barat v. Navy Fed. Credit Union*, 127 F.4th 833, 836 (11th Cir. 2025); *Fischl v. Gen. Motors Acceptance Corp.*, 708 F.2d 143, 146 (5th Cir. 1983) (citing S. REP. NO. 94-589, 94th Cong. (1976)).

²³⁷ CFPB CIRCULAR 2022-03, *supra* note 211, at 1; *see* CFPB CIRCULAR 2023-03, *supra* note 211.

²³⁸ *See* CFPB CIRCULAR 2023-03, *supra* note 211, at 5.

²³⁹ Interpretive Rules, Policy Statements, and Advisory Opinions; Withdrawal, 90 Fed. Reg. 20084, 20086 (May 12, 2025).

²⁴⁰ *See* sources cited *infra* notes 241–247 and accompanying text.

²⁴¹ U.S. CONST. amend. V, cl. 5 (“[N]or shall private property be taken for public use, without just compensation.”).

²⁴² *Id.* amend I.

²⁴³ *X.AI LLC v. Bonta*, No. 25-CV-12295, 2026 WL 626926, at *5 (C.D. Cal. Mar. 4, 2026), *appeal filed* No. 26-1591 (9th Cir. Mar. 17, 2026).

²⁴⁴ *Id.* at *5.

²⁴⁵ *Id.*

request for a preliminary injunction.²⁴⁶ The case is in a preliminary posture, and the developer has appealed.²⁴⁷

Audits

Some commentators have expressed skepticism that greater transparency, at least alone, would combat algorithmic “bias.”²⁴⁸ Some regulators and commentators have gone further and required or suggested that developers and/or users of algorithmic systems conduct some form of audit or testing to discern whether their programs unfairly or illegally impact certain groups.²⁴⁹

Commentators have observed that one of the advantages of algorithmic tools is that, unlike some human decisionmaking, the rules can (sometimes) be knowable, tested, and changed before being implemented.²⁵⁰ As one example of a regulation requiring auditing, New York City has implemented a law regulating “automated employment decision tools,” which prohibits employers from using certain algorithmic tools unless the tool has undergone a “bias audit” within the past year.²⁵¹ Laws can also incentivize bias testing or auditing without requiring it. California regulations, for example, state that “evidence, or the lack of evidence, of anti-bias testing or similar proactive efforts to avoid unlawful discrimination” is “[r]elevant to any . . . claim or available defense” in employment discrimination matters.²⁵²

²⁴⁶ *Id.* at *5–8.

²⁴⁷ *X.AI LLC v. Bonta*, No. 26-1591 (9th Cir. filed Mar. 17, 2026). For more information on Takings Clause claims regarding disclosure of trade secrets, see CRS Legal Sidebar LSB11143, *Vanda Pharmaceuticals, Inc. v. United States: Fifth Amendment Takings Claims for Alleged Disclosures of Trade Secrets and Confidential Information*, by Dorothy C. Kafka and Christina L. Shifton (2024). For more information on the First Amendment and government disclosure requirements, see CRS In Focus IF12388, *First Amendment Limitations on Disclosure Requirements*, by Valerie C. Brannon et al. (2023); and CRS Report R45700, *Assessing Commercial Disclosure Requirements under the First Amendment*, by Valerie C. Brannon (2019).

²⁴⁸ *E.g.*, Dwork & Mulligan, *supra* note 146, at 37; Maurer, *supra* note 190; see *Resetting Antidiscrimination Law*, *supra* note 129, at 1573; Griffith, *supra* note 107, at 375; Kim, *supra* note 5, at 899–900; see also NIST AI RMF 1.0, *supra* note 4, at 16 (“A transparent system is not necessarily an accurate, privacy-enhanced, secure, or fair system.”).

²⁴⁹ See, e.g., CONN. GEN. STAT. § 4a-2e(c) (2026); GAO AI ACCOUNTABILITY FRAMEWORK, *supra* note 212, at 6, 47, 52–53, 58, 61; EVA WOLKOWITZ & SARAH PARKER, CTR. FOR FIN. SERVS. INNOVATION, BIG DATA, BIG POTENTIAL: HARNESSING DATA TECHNOLOGY FOR THE UNDERSERVED MARKET 24 (2015), <https://s3.amazonaws.com/cfsi-innovation-files/wp-content/uploads/2017/02/13062352/Big-Data-Big-Potential-Harnessing-Data-Technology-for-the-Underserved-Market.pdf> [<https://perma.cc/9WY8-WZMG>].

²⁵⁰ See, e.g., Chouldechova & Hemel, *supra* note 146, at 10; Gillis & Spiess, *supra* note 19, at 479–80; see also EEOC, SELECT ISSUES, *supra* note 202 (“One advantage of algorithmic decision-making tools is that the process of developing the tool may itself produce a variety of comparably effective alternative algorithms.”); Mark MacCarthy, *Fairness in Algorithmic Decision-Making*, BROOKINGS (Dec. 6, 2019), <https://www.brookings.edu/articles/fairness-in-algorithmic-decision-making> [<https://perma.cc/359H-JCPX>] (similar).

²⁵¹ N.Y.C. Local Law No. 144 (2021) (codified at N.Y.C. Admin. Code §§ 20-870 to § 20-874 (2021)). A local rule clarifies the requirements for a “bias audit.” N.Y.C. DEP’T OF CONSUMER & WORKER PROT., NOTICE OF ADOPTION OF FINAL RULE (2023), <https://rules.cityofnewyork.us/wp-content/uploads/2023/04/DCWP-NOA-for-Use-of-Automated-Employment-Decisionmaking-Tools-2.pdf> [<https://perma.cc/MZ8J-FAPZ>].

²⁵² CAL. CODE REGS. tit. 2, § 11009 (2026). Connecticut law similarly states that adjudicators of discrimination claims involving “automated employment-related decision processes” “may consider evidence of anti-bias testing or similar proactive efforts to avoid the discriminatory practice, including, but not limited to, the quality, efficacy, recency and scope of such testing or efforts, the results of such testing or efforts and the response thereto.” CONN. PUB. ACT No. 26-15, § 13(b)(1) (2026).

Auditing requirements could take many forms. In addition to determining what an audit would consist of (or directing agencies to do so), Congress could consider which parties would be responsible for conducting an audit, how often audits would occur, and the consequences following audit results, among other things. Bills have been introduced that would require some kind of evaluation or audit of certain algorithms for bias or discriminatory potential, including the Algorithmic Accountability Act of 2025²⁵³ and the Artificial Intelligence Civil Rights Act of 2025.²⁵⁴ The latter also includes disclosure requirements.²⁵⁵

As with transparency requirements, state-level audit requirements have been challenged on constitutional grounds. In April 2026, DOJ sought to intervene in a lawsuit by an AI developer challenging Colorado's then-existing statute requiring developers of certain AI systems to, among other things, "use reasonable care to protect consumers from any known or reasonably foreseeable risks of algorithmic discrimination arising from" their products.²⁵⁶ The law also required "deployer[s]" of certain AI systems to adopt a "risk management policy and program" and engage in "impact of assessment[s]" to analyze and mitigate the risk of algorithmic discrimination.²⁵⁷ DOJ argued that these requirements violated the Equal Protection Clause, contending that they forced developers and AI users to engage in disparate treatment by adjusting their systems to avoid disparate outcomes on the basis of protected characteristics.²⁵⁸ DOJ described this process as "racial (or other protected-class-based) engineering."²⁵⁹ The parties agreed to a stay,²⁶⁰ and Colorado subsequently repealed and replaced the challenged law, removing the "reasonable care," "risk management," and "impact assessment" elements.²⁶¹ Further discussion of potential constitutional issues that may arise when legislating to avoid disparate impacts is below.²⁶²

Design Requirements

Congress might consider laws that would directly control the development of algorithmic systems. Congress could enact policies regulating AI training data or the criteria on which AI models can base outputs, for example. Under Maryland law, for example, software tools used for health insurance utilization review must rely on an individual's "clinical" data and may not rely "solely on a group dataset."²⁶³ Congress could also require developers to try to minimize discrimination risk or maximize transparency in algorithmic design.²⁶⁴ The Department of Health and Human Services, for example, requires federally funded health programs to "make reasonable

²⁵³ H.R. 5511, 119th Cong. (2025); S. 2164, 119th Cong. (2025).

²⁵⁴ H.R. 6356, 119th Cong. (2025); S. 3308, 119th Cong. (2025); *see also, e.g.*, Artificial Intelligence Education and Workforce Readiness Act of 2026, H.R. 10515, 119th Cong. §§ 6(b)(1)(D)(v), 504(c) (2026) (requiring bias testing for some automated systems).

²⁵⁵ H.R. 6356, § 301; S. 3308, § 301.

²⁵⁶ United States Complaint in Intervention ¶ 24, *X.AI LLC v. Weiser*, No. 26-CV-01515 (D. Colo. filed Apr. 24, 2026), Dkt. No. 17 (quoting COLO. REV. STAT. § 6-1-1702(1) (2026) (repealed by S.B. 26-189 (May 14, 2026))).

²⁵⁷ *Id.* ¶¶ 37–38 (quoting COLO. REV. STAT. §§ 6-1-1703(2)(a), 6-1-1703(3)(a)–(b)(II) (2026) (repealed by S.B. 26-189 (May 14, 2026))).

²⁵⁸ *Id.* ¶¶ 40, 46–53.

²⁵⁹ *Id.* ¶ 50.

²⁶⁰ Minute Order, *X.AI LLC v. Weiser*, No. 26-CV-01515 (D. Colo. Apr. 27, 2026), Dkt. No. 24.

²⁶¹ 2026 Colo. Sess. Laws 569 (to be codified at COLO. REV. STAT. §§ 6-1-1701 to -1709).

²⁶² *See infra* "Examining Disparate Impact."

²⁶³ MD. CODE ANN., INS. § 15-10B-05.1(c)(1)–(2) (West 2026).

²⁶⁴ For a discussion of the development of "explainable AI," *see* CRS Report R46795, *Artificial Intelligence: Background, Selected Issues, and Policy Considerations*, by Laurie Harris (2021).

efforts to mitigate the risk of discrimination” from any “patient care decision support tools” that “employ input variables or factors that measure race, color, national origin, sex, age, or disability.”²⁶⁵ Some sources assert that there are technical ways to ensure algorithms do not discriminate on the bases of protected characteristics or their proxies;²⁶⁶ or to select among available algorithms those that are substantially equally effective at achieving a legitimate goal but produce less discriminatory outputs.²⁶⁷ Congress could consider requiring such methods to be explored when algorithmic decisionmaking tools are developed or implemented.

Role of Human Decisionmakers

Congress could consider if, when, and how humans are to participate in certain decisions. A bipartisan task force in the House of Representatives recommended that “[h]av[ing] humans in the loop to actively identify and remedy potential flaws when AI is used in highly consequential decision-making” is “necessary to safeguard the constitutional rights and civil liberties of Americans affected by AI systems, whether those systems are used by government or private entities.”²⁶⁸ Keeping humans involved can mean anything from forbidding the use of automated tools in certain domains to providing for human review when decisions are challenged. At the federal level, for example, one bill introduced in 2026 would prohibit the federal government and state and local governments receiving certain federal funds from using facial recognition technology or other “biometric surveillance systems” except under specific circumstances.²⁶⁹ A bill introduced in 2025 would require health care entities to adopt policies ensuring that algorithmic “outputs are not substituted for the independent judgment of a health care professional.”²⁷⁰ The GAO observes that despite the widespread use of automated valuation tools in the mortgage market, lenders still engage in “traditional appraisals” because they are required to do so by certain federal loan programs.²⁷¹ At the state level, for example, Montana and New Hampshire passed state laws requiring human review of certain outputs from state government AI systems.²⁷² Colorado law currently requires that a “consumer” be given “an opportunity for meaningful human review” of certain adverse decisions when “commercially reasonable,” if those decisions were “materially influence[d]” by automated tools.²⁷³ Congress could legislate to clarify when entities may and may not allow for fully autonomous decisionmaking by algorithmic systems and what, if any, human participation is required.

²⁶⁵ 45 C.F.R. § 92.210(b) (2026).

²⁶⁶ See, e.g., Prince & Schwarcz, *supra* note 11, at 1313–14; Scherer, King & Mrkonich, *supra* note 16, at 507–09; Pope & Sydnor, *supra* note 149, at 207.

²⁶⁷ See EEOC, SELECT ISSUES, *supra* note 202; cf. Louisiana v. Callais, 608 U.S. 85, 118 (2026) (observing that with “[modern computer] technology, if it is possible to identify an alternative map that fully achieves all the State’s legitimate goals while producing ‘greater racial balance,’ then a [Voting Rights Act] plaintiff can easily do so” (emphasis in original) (quoting Alexander v. S.C. State Conf. of the NAACP, 602 U.S. 1, 34–35 (2024))).

²⁶⁸ BIPARTISAN HOUSE TASK FORCE REPORT, *supra* note 2, at 86.

²⁶⁹ Facial Recognition and Biometric Technology Moratorium Act of 2026, H.R. 10563, 119th Cong. (2026).

²⁷⁰ Right to Override Act, S. 2997, 119th Cong. § 101(a)(1)(A) (2025).

²⁷¹ GAO, PROPERTY TECHNOLOGY, *supra* note 108, at 10.

²⁷² MONT. CODE ANN. § 2-21-104 (2025); N.H. REV. STAT. ANN. § 5-D:4 (2025).

²⁷³ 2026 Colo. Session Laws 569, at 578 (to be codified at COLO. REV. STAT. § 6-1-1705).

Examining Disparate Impact

The viability of disparate impact doctrine generally is in dispute.²⁷⁴ Federal agencies have historically assumed that algorithmic decisionmaking could result in disparate treatment or (where statutes allow it), disparate impact violations.²⁷⁵ In 2025, President Trump issued an executive order that described disparate impact doctrine as “wholly inconsistent with the Constitution,” stating that it “imperils the effectiveness of civil rights laws by mandating, rather than proscribing, discrimination.”²⁷⁶ The order declared that “[i]t is the policy of the United States to eliminate the use of disparate-impact liability in all contexts to the maximum degree possible” and instructed agencies to “deprioritize” enforcement of disparate impact liability and take steps to “amend[] or repeal, as appropriate under applicable law,” agency rules and guidance imposing disparate impact liability.²⁷⁷ This order, and agency actions implementing it,²⁷⁸ may constrain federal agencies in how they approach the discriminatory potential of algorithmic decisionmaking.

If Congress intends for agencies to pursue disparate impact cases, it could strengthen their statutory mandates to do so. It may, for example, add disparate impact liability to statutes that lack it or require agencies to devote resources to disparate impact enforcement. Alternatively, if Congress endorses the Trump Administration’s position on disparate impact, it could repeal disparate impact authorities or specify that disparate impact causes of action are not available.²⁷⁹ Congress could also alter or clarify the standards by which claimants prove disparate impact claims, generally or in cases involving algorithmic decisionmaking. Congress may take different approaches to different statutes or sectors.

²⁷⁴ See *supra* “Tensions Between Disparate Treatment and Disparate Impact.”

²⁷⁵ See, e.g., EEOC, EEOC-NVTA-2022-2, THE AMERICANS WITH DISABILITIES ACT AND THE USE OF SOFTWARE, ALGORITHMS, AND ARTIFICIAL INTELLIGENCE TO ASSESS JOB APPLICANTS AND EMPLOYEES (2022), <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence> [<https://web.archive.org/web/20250119202615/https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence>]; EEOC, SELECT ISSUES, *supra* note 202; HUD, GUIDANCE ON THE APPLICATION OF THE FAIR HOUSING ACT TO THE SCREENING OF APPLICANTS FOR RENTAL HOUSING 9–10 (2024), https://www.hud.gov/sites/dfiles/FHEO/documents/FHEO_Guidance_on_Screening_of_Applicants_for_Rental_Housing.pdf [https://web.archive.org/web/20250117202458/https://www.hud.gov/sites/dfiles/FHEO/documents/FHEO_Guidance_on_Screening_of_Applicants_for_Rental_Housing.pdf] These guidance documents are no longer available on EEOC or HUD’s websites.

²⁷⁶ Exec. Order No. 14281, 90 Fed. Reg. 17537, 17537 (Apr. 23, 2025).

²⁷⁷ *Id.* at 17537–38.

²⁷⁸ See, e.g., Equal Credit Opportunity Act (Regulation B), 91 Fed. Reg. 21620, 21620, 21624 (Apr. 22, 2026) (codified at 12 C.F.R. pt. 1002) (rescinding ECOA’s disparate impact regulations and citing Executive Order 14281); Rescinding Portions of Department of Justice Title VI Regulations to Conform More Closely with the Statutory Text and to Implement Executive Order 14281, 90 Fed. Reg. 57141 (Dec. 10, 2025) (codified at 28 C.F.R. pt. 42).

²⁷⁹ GINA, for example, expressly states that “‘disparate impact’ . . . does not establish a cause of action.” 42 U.S.C. § 2000ff-7.

The Supreme Court held in *Ricci* that some attempts to eliminate disparate impacts result in disparate treatment.²⁸⁰ Additionally, some Justices and other federal judges have concluded that majority-group plaintiffs also must be able to bring disparate impact claims if the doctrine is to survive constitutional scrutiny.²⁸¹ These perspectives could significantly limit how far Congress can go to require algorithms to avoid disparate impacts, at least when regulating discrimination on the basis of constitutionally suspect traits, such as race, national origin, and sex.²⁸² Courts may find, for example, that laws that impose liability without giving defendants enough leeway to proffer nondiscriminatory justifications for their actions cross the line to mandating demographic balancing.²⁸³ Similar results could occur with laws that allow plaintiffs to challenge isolated elements of a system, especially given that members of both majority and minority groups can bring disparate impact cases, regardless of the ultimate impact of the system as a whole. The more individual parts of a system can be challenged, the more it exposes a user to liability, and the more it may incentivize users to balance outcomes to avoid the pain of litigation.

Congress does not face these constitutional constraints when addressing discrimination on the basis of traits like disability, age, and family status, which are not constitutionally suspect.²⁸⁴ To the extent disparate impact and disparate treatment frameworks are in tension with regard to discrimination on the basis of such characteristics, Congress has significant leeway to determine how to resolve those tensions.

Balancing the Risks of Discrimination with the Risks of Regulation

Algorithmic decisionmaking is a rapidly changing field. Many commentators, including federal regulators, have emphasized the promise in addition to the risks of automated decisionmaking systems.²⁸⁵ These potential upsides include greater efficiency, cost-effectiveness, and accuracy in decisionmaking.²⁸⁶ The upsides also include potential progress in combating discrimination and creating equal opportunity.²⁸⁷ In 2023, for example, a job-search website agreed in a resolution

²⁸⁰ See *supra* “Tensions Between Disparate Treatment and Disparate Impact.”

²⁸¹ *Tex. Dep’t of Hous. & Cmty. Affs. v. Inclusive Cmty. Project, Inc.*, 576 U.S. 519, 555 (2015) (Thomas, J., dissenting); *Greater New Orleans Fair Hous. Action Ctr. v. HUD*, 639 F.3d 1078, 1088 (D.C. Cir. 2011); see OLC Disparate Impact Op., *supra* note 59, at 10–11 (agreeing with this perspective).

²⁸² Congress’s power to regulate on the basis of certain traits, most prominently race, national origin, and sex, is significantly constrained by equal protection principles. See *Bolling v. Sharpe*, 347 U.S. 497, 498–500 (1954) (holding that the Fifth Amendment’s Due Process Clause incorporates the same equal protection principles as the Fourteenth Amendment’s Equal Protection Clause); *Libr. of Cong., Amdt5.7.3 Equal Protection*, CONSTITUTION ANNOTATED, https://constitution.congress.gov/browse/essay/amdt5-7-3/ALDE_00013730 (last visited Sep. 24, 2026). Laws that classify on the basis of race and national origin receive strict scrutiny, while laws that classify on the basis of sex receive heightened scrutiny. *United States v. Skrametti*, 605 U.S. 495, 510 (2025). Congress has more flexibility to regulate on the basis of other traits, such as age and disability. Laws classifying on the basis of age, disability, and most other traits receive rational basis review. *City of Cleburne v. Cleburne Living Ctr.*, 473 U.S. 432, 441–42 (1985).

²⁸³ See *Inclusive Cmty. Project, Inc.*, 576 U.S. at 542.

²⁸⁴ See discussion *supra* note 282.

²⁸⁵ See, e.g., NIST AI RMF 1.0, *supra* note 4, at 4.

²⁸⁶ See, e.g., CRS Report R47644, *Artificial Intelligence: Overview, Recent Advances, and Considerations for the 118th Congress*, by Laurie Harris (2023) (“Broadly, AI technologies can accelerate and provide insights into data processing, augment human decisionmaking, optimize performance for complex tasks and systems, and improve safety for people in dangerous occupations.”). For example, the FTC reported that when one company started using an AI tool to assess job applicants, it discovered that one of the criteria it previously prioritized “turns out to have no bearing on either productivity or retention.” FTC, *BIG DATA*, *supra* note 54, at 5–6. GAO reported that automated underwriting systems “may improve default risk predictions by identifying patterns not captured through traditional underwriting.” GAO, *PROPERTY TECHNOLOGY*, *supra* note 108, at 24.

²⁸⁷ E.g., FTC, *BIG DATA*, *supra* note 54, at 5–6; Gillis & Spiess, *supra* note 19, at 460, 479–87.

with EEOC to implement an AI-powered tool to “scrape” its website for job postings that discriminated on the basis of national origin.²⁸⁸ In 2019, the CFPB reported findings from a company using “alternative data and machine learning” for credit decisions showing increased access to credit and lower rates “across all tested race, ethnicity, and sex segments” with “no disparities that require further fair lending analysis.”²⁸⁹

If Congress finds that the regulation of AI or algorithmic decisionmaking has inhibited useful innovation, it could try to scale back or eliminate regulations to promote AI use and development, for example, by preempting state laws.²⁹⁰ Congress could also take a nonregulatory approach to incentivizing its preferred developments. Congress could, for example, support further private research into and development of solutions to algorithmic discrimination. It could fund research, incentivize, or compensate parties for auditing algorithms, or use other methods. It could also continue to support the development of technical expertise and standards within federal agencies.²⁹¹ Examples of prior laws directing federal resources toward the development and deployment of AI are included in other CRS products.²⁹²

As commentators have observed, modern algorithmic decisionmaking may be biased, but human decisionmaking and older algorithmic methods can be as well.²⁹³ In many instances, the adoption of policies like those described above regarding transparency or auditing would impose more restrictions on algorithmic decisionmaking than currently exist on human decisionmaking.²⁹⁴ If regulators decide to enact barriers to the deployment of algorithmic decisionmaking models to eliminate bias, they may wish to consider whether alternatives to algorithms—like human decisionmaking—are more or less effective at addressing discrimination.

²⁸⁸ Press Release, EEOC, DHI Group, Inc., Conciliates EEOC National Origin Discrimination Finding (Mar. 20, 2023), <https://www.eeoc.gov/newsroom/dhi-group-inc-conciliates-eeoc-national-origin-discrimination-finding> [https://perma.cc/P8U7-G64X]; see also, e.g., *Breaking New Ground in Generative AI: Zillow’s Fair Housing Classifier*, ZILLOW (May 21, 2024), <https://www.zillow.com/news/zillows-fair-housing-classifier> [https://perma.cc/D8UP-5NGG] (describing an “AI innovation” that detects potentially discriminatory language in user questions and Zillow’s chatbot responses and alerts developers to formulate a response).

²⁸⁹ Patrice Alexander Ficklin & Paul Watkins, *An Update on Credit Access and the Bureau’s First No-Action Letter*, CFPB (Aug. 6, 2019), <https://www.consumerfinance.gov/archive/blog/update-credit-access-and-no-action-letter> [https://perma.cc/XW2S-7FQT]; see, e.g., GAO, PROPERTY TECHNOLOGY, *supra* note 108, at 24; STAFF OF H. COMM. ON FIN. SERVS., BIPARTISAN WORKING GRP. ON AI, 118TH CONG., REPORT ON AI INNOVATION EXPLORED: INSIGHTS INTO AI APPLICATIONS IN FINANCIAL SERVICES AND HOUSING 12, 18 (2024), https://financialservices.house.gov/uploadedfiles/bipartisan_working_group_on_ai_staff_report.pdf [https://perma.cc/EL5E-V3SN]; Griffith, *supra* note 107, at 361; Bruckner, *supra* note 16, at 27.

²⁹⁰ See Exec. Order No. 14365, 90 Fed. Reg. 58499 (Dec. 11, 2025) (directing the establishment of an “AI Litigation Task Force” to “challenge State AI laws inconsistent with the policy set forth” in the executive order).

²⁹¹ See BIPARTISAN HOUSE TASK FORCE REPORT, *supra* note 2, at 86–87.

²⁹² See, e.g., CRS Report R47644, *Artificial Intelligence: Overview, Recent Advances, and Considerations for the 118th Congress*, by Laurie Harris (2023).

²⁹³ See, e.g., GAO, PROPERTY TECHNOLOGY, *supra* note 108, at 19–20; Alex P. Miller, *Want Less-Biased Decisions? Use Algorithms.*, HARV. BUS. REV. (July 26, 2018), <https://hbr.org/2018/07/want-less-biased-decisions-use-algorithms> [https://perma.cc/RX6F-6QDB]; Bruckner, *supra* note 16, at 27; Kim, *supra* note 5, at 869; Barocas & Selbst, *supra* note 11, at 730; Price II, *supra* note 108; cf. *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 448 (1975) (Blackmun, J., concurring in the judgment) (extolling the potential virtues of well-designed “pre-employment testing” to “measure all applicants objectively on a standardized basis”).

²⁹⁴ Cf. Chouldechova & Hemel, *supra* note 146, at 6 (arguing that humans cannot be made completely “race blind” “insofar as all of us have learned to draw inferences in an informational environment that includes race as a relevant variable” and that requiring algorithms to be stripped of all “race consciousness” even at the training stage would hold them to a different standard than applied to humans).

Author Information

Abigail A. Graber
Legislative Attorney

Disclaimer

This document was prepared by the Congressional Research Service (CRS). CRS serves as nonpartisan shared staff to congressional committees and Members of Congress. It operates solely at the behest of and under the direction of Congress. Information in a CRS Report should not be relied upon for purposes other than public understanding of information that has been provided by CRS to Members of Congress in connection with CRS's institutional role. CRS Reports, as a work of the United States Government, are not subject to copyright protection in the United States. Any CRS Report may be reproduced and distributed in its entirety without permission from CRS. However, as a CRS Report may include copyrighted images or material from a third party, you may need to obtain the permission of the copyright holder if you wish to copy or otherwise use copyrighted material.