

ARTICLES

Peeking Inside the Black Box: Discovery in Fair Housing Act Cases Arising from AI Decision-Making

Jessica A. Mahon Scoles, J.D.^a

Keywords: artificial intelligence, discovery, fair housing act, disparate impact, housing discrimination, civil litigation

Wilmington Law Review

Vol. 1, Issue 1, 2026

In today's real estate market, more housing decisions are being automated. If artificial intelligence renders a biased decision, will the person impacted have any legal recourse? Commentators have expressed concern that the proprietary nature of AI will allow defendants in discrimination lawsuits to avoid discovery, effectively shielding their decision making from judicial scrutiny. This article examines whether AI and algorithm-related information is the discovery black box that some fear it to be.

Addressing the issue in the context of Fair Housing Act (FHA) litigation, the article argues that concerns about access to AI and algorithm-related discovery are overstated. In real litigation, the Federal Rules of Civil Procedure have proven well-equipped to handle the discovery issues raised by novel AI and algorithm technology. Thus, with proper planning, plaintiffs should be able to obtain the AI and algorithm-related information they need to prove their claims. However, Plaintiffs who file FHA claims against AI developers and users do still face significant challenges – and those challenges are likely to grow. In particular, a circuit split regarding whether FHA plaintiffs can challenge policies where the disparate impact arises from existing demographic patterns makes it difficult for plaintiffs in some circuits to survive motions to dismiss. If the Supreme Court takes up this circuit split, disparate impact liability under the FHA is likely to be significantly narrowed.

While there have been many articles that discuss AI and AI bias, few discuss discovery strategies when litigating claims arising from AI bias. *Peeking Inside the Black Box* is the first article to comprehensively review the existing case law addressing the discoverability of AI and algorithm related information. It aims to both contribute to the academic debate about discoverability of AI and provide a practical guide for plaintiffs' attorneys litigating disparate impact claims against AI developers.

^a Associate Professor of Law, Western New England University School of Law. Thank you to my research assistants, Alexander Peña and Brianna McDermott, for their research help on this article; Attorney Christa Douaihy, of Community Legal Aid, for sharing her insights regarding litigating disparate impact claims; Dean Zeld Harris, Associate Dean Nicole Belbin, and Associate Dean Timothy Webster, for their support and encouragement; Professors Harris Freeman and Claudia Quintero, who are always willing to connect me with resources and people to help with my projects; and Dave and Finn Scoles, who kept me motivated to finish this piece.

I. Introduction

Some time ago, while practicing law in Los Angeles, I represented a pro bono client whose landlord had engaged in a pattern of illegal, discriminatory conduct. The landlord's property management company, the decision-maker for much of the illegal conduct, had designated an agent for service of process in Humboldt County. If you are unfamiliar with California geography, Humboldt County is in the heart of California's redwood forests, an eleven-hour drive north of Los Angeles. By designating an agent for service of process in such a remote location, the property manager had made it difficult for low-income tenants in Los Angeles to hand serve the company, a necessity given the short timeframes for many housing cases in California. Effectively, then, the property manager had hidden key evidence in a cabin in the woods.

In today's real estate market, with more housing decisions being automated, some have expressed concern that artificial intelligence (AI) use is the technological equivalent of hiding evidence in a remote cabin. These commentators worry that the proprietary nature of AI will allow defendants to avoid discovery, shielding their decision making from scrutiny.¹ Given AI's well-documented tendency to render biased decisions,² the implications of a litigation regime where plaintiffs cannot access the information about how AI renders decisions are troubling. Accordingly, this article examines whether AI and algorithm-related information is the discovery black box that some fear it to be.

This article approaches the issue through the context of Fair Housing Act³ (FHA) litigation. The FHA, a landmark anti-discrimination law that makes it illegal to discriminate in the sale or rental of housing, is well-suited to target developers and users of AI platforms that make biased decisions.⁴ Like other major anti-discrimination laws, the FHA provides for disparate impact claims, the claim arguably best suited to address the types of unintentional discrimination in which AI is most likely to engage.⁵ And, unlike other major anti-discrimination laws, the FHA does not define a class of defendants who are subject to suit. Instead, it allows suit against anyone who engages in certain proscribed conduct, making its reach broad enough to cover AI developers.⁶

This article argues that access to discovery regarding AI and algorithms is not the obstacle that some commentators fear it might be. The Federal Rules of Civil Procedure have proven to be well-equipped to handle the discovery issues raised by novel AI and algorithm technology. That is not to

¹ See sec. IV.B, *infra*.

² See sec. II.B, *infra*.

³ 42 U.S.C. § 3601, *et seq.*

⁴ See sec. III, *infra*.

⁵ See sec. III.C, *infra*.

⁶ See sec. III.A, *infra*.

say that plaintiffs who file FHA claims against AI developers and users do not face significant challenges.⁷ Litigating discrimination claims has always been difficult. Those challenges have only gotten more acute given the current backlash against diversity, equity, and inclusion.⁸ In particular, a circuit split regarding whether FHA plaintiffs can challenge policies where the disparate impact arises from existing demographic patterns is already making it difficult for plaintiffs in some circuits to survive motions to dismiss.⁹ If the Supreme Court takes up this issue, disparate impact liability under the FHA is likely to be significantly narrowed. Until that happens, the FHA remains one of our best tools for combatting AI bias, and, with proper planning, plaintiffs should be able to obtain the AI and algorithm-related information they need to prove their claims.¹⁰

The body of this article is divided into three sections. Section II provides necessary background on AI. Section II first develops a working definition of AI, then discusses the problem of AI bias, and, finally, explains why litigation presently offers the best option for addressing that bias. Section III provides background on the FHA and discusses the circuit split regarding the three-stage burden-shifting framework for proving disparate impact claims under the FHA. Section IV discusses pleadings and discovery in FHA disparate impact cases involving alleged AI bias. It argues that the most significant hurdle that plaintiffs face is not obtaining discovery about proprietary AI platforms, but rather pre-filing investigation to meet the *Twombly/Iqbal* pleading burden. Section IV also offers some suggestions to help plaintiffs clear pleading and discovery hurdles on the way to trial.

II. A Short Primer on AI and Bias

As AI use has become more visible in recent years, so too have examples of AI bias. Despite the attention that AI bias receives, the problem has proven relatively intractable thus far. To illuminate why this is, this section defines AI, explains the origins of AI bias, and discusses some potential solutions for the problem.

A. What is AI?

Before we can discuss AI, we need to understand what the term “artificial intelligence” means. And defining the term is no easy task. Indeed, even computer scientists do not agree on the definition of AI. To start, “‘artificial intelligence’ is an umbrella term [I]t does not refer to a single technology

⁷ See, e.g., Valerie Schneider, *Locked Out by Big Data: How Big Data Algorithms and Machine Learning May Undermine Housing Justice*, 52 COLUM. HUM. RTS. L. REV. 251, 288 (2020) (quoting a public interest attorney’s statement about the challenges of FHA litigation); Janine S. Hiller, *Fairness in the Eyes of the Beholder: AI, Fairness, and Alternative Credit Scoring*, 123 W. VA. L. REV. 907, 922–23 (2021) (discussing difficulty of proving both disparate impact and disparate treatment claims).

⁸ See sec. III.C, *infra*.

⁹ See sec. III.C, *infra*.

¹⁰ See sec. IV, *infra*.

or function.”¹¹ Technologies that are commonly referred to as AI include “rule bases,” which are databases that store the rules that systems use to make decisions; “large language models” like ChatGPT, which are computational models that use statistical information to generate human language; and “machine learning,”¹² “which enables computers to ‘automatically detect patterns in data, and then use the uncovered patterns to predict future data or to perform decision-making tasks under uncertainty.’”¹³

Moreover, the field of AI is rapidly evolving, so our understanding of what AI is and what it can do must be flexible enough to account for this rapid development. For example, current AI, referred to as “narrow” AI, can only perform discrete tasks.¹⁴ But AI developers hope to create “general AI,” AI that can rival human performance at a full complement of human tasks including self-sustenance.¹⁵ Whether and when that goal will come to pass are open questions. American companies have recently seen the advancement of their AI models slow, with some industry observers questioning whether there is enough data in the world for AI models to continue to improve at historic rates.¹⁶ At the same time, Chinese startup DeepSeek stunned United States tech companies with its development of an AI chatbot that it claims performs as well as ChatGPT while using a fraction of the resources,¹⁷ suggesting that even industry insiders cannot always predict where the AI revolution will take us.

Further complicating things is the tendency to affix the AI label to things that aren’t really AI at all.¹⁸ AI currently is a hot topic, and businesses are keen to jump on the trend. This has led some companies to claim they use AI when they do not. The practice is so prevalent that it has a name: “AI washing.”¹⁹ “There are several types of AI washing. Some companies

11 Paul W. Grimm, Maura R. Grossman & Gordon V. Cormack, *Artificial Intelligence as Evidence*, 19 NW. J. TECH. & INTELL. PROP. 9, 16 (2021).

12 *Id.* Computer scientists disagree as to whether all these technologies actually are AI. For example, during a recent presentation on generative AI, a colleague in the computer science program at WNE explained that many computer scientists do not believe that large language models are AI.

13 Sharona Hoffman & Andy Podgurski, *Artificial Intelligence and Discrimination in Health Care*, 19 YALE J. HEALTH POL’Y L. & ETHICS. 1, 8 (2020) (quoting Kevin P. Murphy, *Machine Learning: A Probabilistic Perspective* 1 (2012)).

14 Grimm et al., *supra* note 11, at 15.

15 *Id.*

16 See Rachel Metz et al., *OpenAI, Google and Anthropic Are Struggling to Build More Advanced AI*, BLOOMBERG, Nov. 13, 2024, <https://www.bloomberg.com/news/articles/2024-11-13/openai-google-and-anthropic-are-struggling-to-build-more-advanced-ai>; Stephanie Palazzolo et al., *OpenAI Shifts Strategy as Rate of ‘GPT’ AI Improvements Slows*, THE INFORMATION, Nov. 9, 2024, <https://www.theinformation.com/articles/openai-shifts-strategy-as-rate-of-gpt-ai-improvements-slows>; see also Cade Metz and Karen Weise, *AI is Getting More Powerful, but Its Hallucinations Are Getting Worse*, NEW YORK TIMES (May 5, 2025).

17 Fan Wang & João da Silva, *DeepSeek: How China’s ‘AI heroes’ overcame US curbs to stun Silicon Valley*, BBC NEWS, Jan. 28, 2025, <https://www.bbc.com/news/articles/czepw096wy4o> [<https://perma.cc/HSA7-4UVR>]; see generally James M. Cooper, *Are “Friends” Electric?: A Comparativist Approach to Guidelines for the Development and Implementation of Artificial Intelligence in the People’s Republic of China and the United States of America*, 42 B.U. INT’L L.J. 119, 131 (2024) (discussing the debate over whether China has overtaken the United States in AI development and regulation).

18 Grimm et al., *supra* note 11, at 28.

19 The term “AI washing,” is derived from the term “green washing” – the practice of making a companies’ products appear more environmentally friendly than they are. Emma Woollacott, *What is ‘AI washing’ and why is it a problem?*, BBC NEWS, Jun. 26, 2024, <https://www.bbc.com/news/articles/c9xx8122893o> [<https://perma.cc/8NH5-F9K5>]

claim to use AI when they're using less-sophisticated computing, while others overstate the efficacy of their AI over existing techniques, or suggest that their AI solutions are fully operational when they are not.”²⁰ Whether innocent or intentional, AI washing further muddies the waters when it comes to defining AI.

Fortunately, resolving the computer science debate about what the term AI really means is outside the scope of this article. For purposes of our discussion, a working definition of AI suffices. In this article, the term “AI” refers to computers learning and solving problems. That includes supervised machine learning—where a computer learns from examples that a trainer provides.²¹ It includes unsupervised machine learning—where a computer is not given the correct answers and instead learns by looking “for patterns, groupings or anomalies that might be of interest.”²² And it includes reinforcement learning—where an AI system learns by facing a “game-like situation” in which it “employs trial and error” to solve problems and the programmer employs “rewards or penalties” to teach the system.²³

Machine learning has the potential to improve the world in many ways, from curing diseases to helping scientists sort through large volumes of data in search of solutions to scientific problems.²⁴ However, it is important to understand the limits of machine learning. The term “machine learning” itself anthropomorphizes AI by “impl[ying] a human-focused process inconsistent with what generative AI does.”²⁵ When we hear the term “machine learning,” we might imagine that an AI model is able to read and understand concepts in the way that humans do.²⁶ What the computer is actually doing is converting data “into binary code and assign[ing] vector numbers of probability” that the AI uses to make predictions.²⁷ In other words, “machine learning does not require a computer to engage in higher-order cognitive skills like reasoning or understanding of abstract concepts. This leaves AI methods vulnerable to pursuing forms of analysis that might be

20 *Id.* AI washing has become so prevalent that it has attracted the attention of the SEC. In several recent enforcement actions, the SEC has gone after registered firms that misrepresented their use of AI. *See, e.g.,* Complaint, *Securities and Exchange Commission v. Raz*, No. 24 Civ. 4466 (S.D.N.Y. Jun. 11, 2024) (alleging that the founder of a tech startup Joonko Diversity, Inc. claimed to use proprietary algorithms and machine learning to find diverse candidates for jobs when they had no such technology).

21 Grimm et al., *supra* note 11, at 12; Kristin Johnson et al., *Artificial Intelligence, Machine Learning, and Bias in Finance: Toward Responsible Innovation*, 88 *FORDHAM L. REV.* 499, 506 (2019).

22 Grimm et al., *supra* note 11, at 26; *see also* Johnson et al., *supra* note 21, at 506.

23 Grimm et al., *supra* note 11, at 12 n.10; (citing Blazej Osiński & Konrad Budek, *What is Reinforcement Learning? The Complete Guide*, *DEEPSense.AI*, July 5, 2018, <https://deepsense.ai/blog/what-is-reinforcement-learning-deepsense-ais-complete-guide/> [<https://perma.cc/3JE4-QTH5>]).

24 *See, e.g.,* Wenlong Li et al., *A Generalizable Low-Data-Dependent and Time-Efficient Foundation Model for Retinal Image Analysis* (2023), <https://ssrn.com/abstract=5319099> (non-peer-reviewed study regarding efficacy of AI model for diagnosing diseases of eye); Henrique de Medeiros Back et al., *Decoding the Evolution of Marine Plastic Litter Research with Large Language Models* (2025), <https://ssrn.com/abstract=5332715> (non-peer-reviewed study using large language model to conduct meta-analysis of prior studies of marine plastic pollution).

25 Stacey M. Lantagne, *Just Like a Human Being: The Anthropomorphism of Generative Artificial Intelligence in Copyright Law Arguments*, 25 *WAKE FOREST J. BUS. & INTELL. PROP. L.* 174, 181 (2025).

26 *Id.* at 182; Johnson et al., *supra* note 21, at 507.

27 Lantagne, *supra* note 25, at 182.

set aside as suspect by [professionals in a field].”²⁸ Thus, delegating to AI tasks that require human judgment and discretion carries risks that may not exist when AI is used to solve problems with clear answers.²⁹ One of the risks of delegating tasks requiring judgment to AI is that AI cannot determine when its own results are the product of faulty logic, such as bias.³⁰

B. AI & Bias

Calling upon AI to make discretionary decisions carries with it an inherent risk that those decisions will be tainted by bias. Indeed, in just the last few years there have been sundry examples that establish that biased AI is—if not inevitable—then, at least, ubiquitous. For example, in 2016, Microsoft was forced to shut down an AI-powered Twitter chatbot when it began spewing racist comments within hours of its release.³¹ In 2014 Amazon scrapped an AI recruiting tool when it became apparent the AI had learned to prefer male applicants over female applicants.³² And AI facial recognition tools are significantly worse at identifying people with darker skin as compared to those with lighter skin, and identifying women as compared to identifying men.³³

While it is theoretically possible to develop AI that is free from bias,³⁴ doing so requires a deliberate effort to avoid bias, and that effort simply was not made when developing current AI systems.³⁵ As scientist and data security expert Matthew Carroll explains, “[t]he average engineer . . . [is] focused on nuances such as ideal network topology, activation functions, training gradients, weight normalization, and data overfitting [They are] not thinking about bias or proper provenance when designing a neural network.”³⁶

28 Johnson et al., *supra* note 21, at 507 (citing Michael L. Rich, *Machine Learning, Automated Suspicion Algorithms, and the Fourth Amendment*, 164 U. PA. L. REV. 871, 880 (2016)). See also Lantagne, *supra* note 25, at 181–87.

29 As an example, if an AI system is trained to determine if tumors are cancerous, then there is little reason to worry about the AI’s lack of reasoning skills. The question whether a tumor is cancerous is binary with clear right and wrong answers. While an AI system may be more or less accurate at assessing tumors, there is little reason to fear that it will abuse its discretion as the task for which it was trained requires none.

30 Lantagne, *supra* note 25, at 182.

31 Matthew Ingram, *Microsoft’s Chat Bot Was Fun for Awhile, Until it Turned into a Racist*, FORTUNE, Mar. 24, 2016, <https://fortune.com/2016/03/24/chat-bot-racism/> [<https://perma.cc/Z3VS-695D>]

32 See Christine Polek & Shastri Sandy, *The Disparate Impact of Artificial Intelligence and Machine Learning*, 21 COLO. TECH. L.J. 85, 92–3 (2023).

33 Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, 81 J. MACH. LEARNING RSCH. 1 (2023). See also <http://gendershades.org/overview.html> [<https://perma.cc/7C6F-7YD1>] (providing an interactive presentation of the study’s results).

34 See Emily Black et al., *Less Discriminatory Algorithms*, 113 GEO. L.J. 53, 67–71 (2024) (discussing computer science research that suggests that, for any problem AI is trained to solve, it is nearly always possible to develop a less discriminatory alternative algorithm that solves the same problem with a similar error rate); Giacomo Capuzzo, *A Comparative Study on Algorithmic Discrimination between Europe and North-America*, 10 COMP. L. REV. 125, 135 (2019) (discussing the potential to address algorithm bias during the training phase).

35 See Sonia K. Katyal, *Private Accountability in the Age of Artificial Intelligence*, 66 UCLA L. REV. 54, 68 (2019) [hereinafter Katyal, *Private Accountability*].

36 *Id.* at 73.

Indeed, in a 2019 survey of machine learning researchers, a majority of respondents said that AI developers should prioritize safety and ethics more than they currently do.³⁷ The same group of researchers reported that they had somewhere between “no trust at all” and “not too much trust” that technology companies—including Google, Facebook, Apple, Microsoft, Amazon, Open AI, DeepMind, Tencent, Baidu and Alibaba—will act in the public’s best interests when developing AI.^{38,39} More recently, computer scientist Geoffrey Hinton, known as the “godfather of AI,” warned that humanity should be taking the risks of AI more seriously.⁴⁰

Moreover, putting trust concerns aside, even the best-intentioned AI engineers are ill-equipped to mitigate the risk of bias in AI. These engineers “do not reflect the diversity of the populations to which the algorithms will be applied[] and have little to no training in ethics or the law”⁴¹ Perhaps most troubling is that it is incredibly difficult to determine when an AI model is relying on protected characteristics or proxies for protected characteristics to render its decisions. As Matthew Carroll explained, “Once a model is trained, engineers quite frequently lack understanding of the model’s actual decision-making process.”⁴² If an engineer cannot understand “why a model made the decision that it did,” then they cannot possibly be sure “it didn’t make a legally questionable decision, like discriminating on the basis of race.”⁴³

Together, AI engineers lack of attention to AI bias, lack of knowledge of anti-bias law, and lack of understanding about how AI systems render decisions mean that there are multiple places in the AI training and development process where biases can enter AI systems unchecked. First, bias may arise when data used to train AI systems is collected and compiled. This is a problem of significant scope because big data routinely fails to collect data from marginalized communities.⁴⁴ If the pool of training data is incomplete or otherwise not representative of the population about which the AI will render decisions, AI is likely to make inaccurate predictions about groups

37 Baobao Zhang et al., *Ethics and Governance of Artificial Intelligence: Evidence from a Survey of Machine Learning Researchers*, 71 J. OF ARTIFICIAL INTEL. RSCH. 591, 593 (2021).

38 *Id.* at 599.

39 That some of the leading technocapitalists have been discretely flirting with racist ideologies like eugenics for some time suggests that machine learning researchers had good reason to question the trustworthiness of the AI industry. See Robert Evans, *Tech Bros Have Built a Cult Around AI: Part One*, BEHIND THE BASTARDS (Jan. 24, 2025), <https://www.iheart.com/podcast/105-behind-the-bastards-29236323/episode/part-one-tech-bros-have-built-146544960/>, and sources cited therein.

40 Steven Bartlett, *Godfather of AI: I Tried to Warn Them, But We’ve Already Lost Control!*, THE DIARY OF A CEO WITH STEVEN BARTLETT (Jun. 16, 2025) (downloaded using Apple Podcasts).

41 Grimm et al., *supra* note 11, at 44–5.

42 Katyal, *Private Accountability*, *supra* note 35, at 73.

43 *Id.* at 73–4.

44 *Id.* at 71–2.

underrepresented in the data.⁴⁵ To illustrate, Professors Stephanie Ben-Ishai and Mandy Bedford offered the example of an AI system that is trained to predict creditworthiness using a data set that includes data about many men but just a small number of women. Because women are underrepresented in the training data, the statistical effect of one of the women in the data set defaulting on a loan will be much higher than when a man defaults. Unless the data is reweighted to account for the underrepresentation of women in the pool, the AI system is likely to learn that women are inherently less credit-worthy than men.⁴⁶

Even where training data is accurate and complete, it can still lead to biased decisions if the data set “reflects or internalizes some forms of structural discrimination.”⁴⁷ Amazon became a real-life example of this problem when it developed an AI bot to assist with hiring. The program was trained using Amazon’s own past hiring data. Because Amazon had historically hired more men than women—apparently because the talent pool in the technology field has historically been disproportionately male—the bot learned that men were “better” job candidates than women and down-graded women’s job applications accordingly.⁴⁸

Second, biases may arise during the training process. In supervised machine learning, humans teach AI systems how to answer questions by providing examples of correct and incorrect answers.⁴⁹ Sometimes this is a simple task. If we want to train AI to identify primates, then pictures of howler monkeys, gorillas, tarsiers, and humans can all confidentially be included in the “primate” group, while pictures of cats, pigeons, hamsters, and squid can be confidently included in the “not primate” group.⁵⁰ But for many decisions that AI will be called upon to make, “it is less obvious what the target variables should be.”⁵¹ Human trainers must therefore make judgment calls about how to categorize data.⁵² For example, if an AI program will be

45 See Stephanie Ben-Ishai & Mandy Bedford, *AI, Consumer Credit, and Discrimination: A Comparative Look at Canada and the United States*, 2 CORP. & BUS. L.J. 271, 280 (2021); Katyal, *Private Accountability*, *supra* note 35, at 68; Capuzzo, *supra* note 34, at 137; Rangita de Silva de Alwis et al., *Equitable Ecosystem: A Two-Pronged Approach to Equity in Artificial Intelligence*, 29 MICH. TECH. L. REV. 165, 172 (2023).

46 Ben-Ishai et al., *supra* note 45, at 280.

47 Katyal, *Private Accountability*, *supra* note 35, at 68.

48 See Polek et al., *supra* note 32, at 93.

49 See Johnson et al., *supra* note 21, at 506 (discussing supervised machine learning).

50 This example is included in honor of my son, Finn, an aspiring primatologist. If you are not familiar with tarsiers, they are a family of tiny prosimians with enormous eyes and kangaroo-like legs that live in Southeast Asia. Taking a moment to Google them is well-worth your time. Tarsiers look very different from other primates, so including them in the training data might initially give the AI system some difficulties. But there are no difficult judgment calls for the human trainer to make because scientists unequivocally classify tarsiers as primates. See, e.g., Peter M. Kappeler, *The Behavioral Ecology of Strepsirrhines and Tarsiers*, in *The Evolution of Primate Societies* (John C. Mitani, et al., eds. 2012).

51 Capuzzo, *supra* note 34, at 136.

52 See *id.*; Katyal, *Private Accountability*, *supra* note 35, at 71 (“[c]ategorization is a powerful semantic and political intervention: what the categories are, what belongs in a category, and who decides how to implement these categories in practice, are all powerful assertions about how things are and are supposed to be.”)(quoting Tarleton Gillespie, *The Relevance of Algorithms*, MEDIA TECHNOLOGIES: ESSAYS ON

used to identify “good tenants,” then a human trainer must decide what the characteristics of a good tenant are. This introduces the potential for human trainers to not only make categorization mistakes by including criteria that are not truly predictive of “good” tenancy, but also to introduce their own biases into the training data.

Unsupervised machine learning would seem to be less prone to this type of bias. After all, the computer trains itself by looking for patterns in the data set.⁵³ However, as illustrated by the Amazon example discussed above, AI can learn biases on its own. In the case of the Amazon bot,⁵⁴ it appears that the AI drew an incorrect inference from a pattern it found in the data.⁵⁵ It saw that Amazon had been more likely to hire men in the past and inferred that Amazon hired the candidates because of their gender.⁵⁶

Some commentators have suggested that where a protected characteristic and an AI system’s objective are correlated, AI has an “inherent tendency” to discriminate.⁵⁷ For example, studies suggest that there is a correlation between race and mortgage loan default rates.⁵⁸ If it is true that AI systems have an inherent tendency to discriminate where discrimination is efficient, then an AI system trained to make mortgage lending decisions will discover this correlation and learn to discriminate against people of color.⁵⁹ What’s more, if the AI system is deprived of information about race, it will optimize its objective by finding a proxy for race and discriminating on that basis.⁶⁰

Finally, biases may arise during the output stage when a human user prompts a system for results. One recent example of this problem involved Facebook’s advertising algorithm, which allows advertisers to exclude certain ethnic groups from seeing their advertising. Because the system was developed without regard for anti-discrimination laws, users were able to prompt it to take the clearly illegal action of targeting housing advertising solely to non-

COMMUNICATION, MATERIALITY, AND SOCIETY 167, 171 (2014)); Black, et al., *supra* note 34, at 67 (“Machine learning systems are built through a series of iterative decisions, often called the machine learning pipeline. These decisions often involve subjective choices for which there is often no correct answer, requiring the developer to weigh competing values and make judgment calls.”).

53 See Johnson et al., *supra* note 21, at 506 (2019) (discussing unsupervised machine learning).

54 See Polek et al., *supra* note 32, at 93.

55 As previously discussed, AI does not think the way that humans do, so “inference” is an imperfect term for how Amazon’s AI hiring bot learned to discriminate against women. However, I use the term inference here because it is a term that captures the essence of what went wrong in a way that will make sense to law-trained human readers.

56 See Ben-Ishai et al., *supra* note 45, at 279 (discussing a hypothetical situation where an AI system learns that people of color are worse drivers because it mistakes correlation, i.e., the presence of more people of color in urban centers where there are more accidents—for causation, i.e., the higher population in urban centers leading to more accidents).

57 Anya E. R. Prince & Daniel Schwarcz, *Proxy Discrimination in the Age of Artificial Intelligence and Big Data*, 105 IOWA L. REV. 1257, 1263–64 (2020).

58 See Kevin Alan Park, *A Comparison of Mortgage Denial and Default Rates by Race, Ethnicity, and Gender* (February 7, 2022), available at <https://ssrn.com/abstract=4030908> or <http://dx.doi.org/10.2139/ssrn.4030908> (comparing default rates on FHA-insured mortgages and concluding that non-Hispanic white borrowers are less likely to default on loans than Black and Hispanic white borrowers, and more likely to default on loans than Asian borrowers).

59 See Prince et al., *supra* note 57, at 1263–64. Humans making lending decisions already engage in this type of discrimination. See Park, *supra* note 58 (concluding that Black, Hispanic white, and Asian borrowers loan denial rates were disproportionate to their risk profile, even taking higher default rates for some of these groups into account).

60 Prince et al., *supra* note 57, at 1263–64.

Hispanic white people.⁶¹ While this is an egregious example, it is easy to imagine how an AI system that lacks controls to prevent discrimination might be prompted to engage in wrongdoing that is harder to spot, such as engaging in proxy discrimination.

C. Solutions to AI Bias

This article focuses on using litigation to combat AI discrimination because it is superior to other options in our current political climate. As discussed in this subsection, both self-regulation by AI developers and national AI regulation are currently unlikely. Those impacted by biased AI in the near term will therefore need to turn to the courts.

1. Developer Self-Regulation Is Unlikely

Unfortunately, there is not much reason to believe that AI developers will successfully remedy bias issues on their own. While extreme incidents of AI bias attract public attention, less visible problems are likely to go unnoticed. One reason is that AI programs may not be collecting the demographic data that would allow AI developers to detect biased decision making. Indeed, in some cases, it is illegal to collect that kind of data.⁶²

Another concern is that most AI models do not have a mechanism to receive and learn from feedback about how their decisions worked in the real world.⁶³ Consider, for example, an AI tenant screening program that rejects a potential tenant. If that tenant goes on to be a model tenant in a building that does not use AI screening software, then the screening software arguably made a mistake in rejecting the tenant. However, there is no way for the software company, or the landlord who rejected the tenant, to learn that the AI program made an incorrect prediction. And, as long as the landlord is able to fill their vacancies despite rejecting some “good” tenants, neither the landlord nor the AI developer has an incentive to investigate the program’s decisions. The AI tenant screening program thus will continue to reject qualified applicants.

Finally, even where AI developers are aware of bias issues, remedying the problem has “complex political ramifications.”⁶⁴ Simply depriving an AI system of information about protected characteristics does not prevent discrimination. As discussed, the AI system instead looks for proxies for those characteristics.⁶⁵ The remedy, then, needs to be race or gender or disability

⁶¹ See de Silva de Alwis et al., *supra* note 45, at 168.

⁶² See Polek et al., *supra* note 32, at 98.

⁶³ Katyal, *Private Accountability*, *supra* note 35, at 90.

⁶⁴ *Id.* at 68.

⁶⁵ See Prince et al., *supra* note 57, at 1276; Shaun Khoo & Zi En Chow, *Algorithmic Fairness: Challenges and Opportunities for Artificial Intelligence Governance*, 34 *SING. ACAD. L.J.* 736, 743 (2022) (discussing failed attempts to debias algorithms by depriving them of information about protected characteristics).

or other protected characteristic conscious.⁶⁶ In other words, the best fix is something that looks an awful lot like affirmative action.⁶⁷ In the current political climate, the legal and public relations risks of implementing that type of solution are a powerful deterrent.⁶⁸ It is therefore likely that AI developers who are concerned about bias will rely on less effective remedies.

2. National Regulation of AI is Unlikely

At least in the short term, we also cannot rely on a legislative solution to AI bias. In the United States, regulation of AI has been slow and piecemeal. AI has operated in what one panelist at a 2022 Critical Tech Talk series described as “kind of a wild west space.”⁶⁹

At the federal level, Congressional efforts to regulate AI have repeatedly stalled.⁷⁰ In fact, the most recent Congressional attempt to legislate around AI was an attempt by House Republicans to include a provision in the so-called “Big, Beautiful Bill” that would have prevented states and municipalities from regulating AI for ten years.⁷¹ In the absence of Congressional leadership on this issue, the Biden administration began to address the issue via executive action. In October 2022, the Biden administration published *The Blueprint for an AI Bill of Rights*, a whitepaper that identified five principles to “guide the design, use, and deployment of automated systems to protect the American public in the age of artificial intelligence.”⁷² Consistent with that blueprint, on October 30, 2023, President Biden issued an executive order that sought to make the United States a leader in addressing concerns about AI safety.⁷³ Many believe that President Biden’s attempts to regulate the tech industry, including attempts to place guardrails on AI, led to a number of prominent tech executives shifting their support to Republicans in the 2024 election.⁷⁴ Those executives received what they hoped for. President

⁶⁶ See Schneider, *supra* note 7, at 290 (suggesting that race-conscious solutions are needed to address AI bias).

⁶⁷ Professor Emily Black, et al. suggest that bias can be rooted out without the need for protected-characteristic conscious solutions by considering alternative methods to solve problems early in the development of AI. Black, et al., *supra* note 34, at 71. But this would require developers who have already brought an AI product to market to go back to the training and development stage. See *id.*

⁶⁸ See generally Nino C. Monea, *Next on the Chopping Block: The Litigation Campaign Against Race-Conscious Policies Beyond Affirmative Action in University Admissions*, 33 B.U. PUB. INT. L.J. 1 (2024); see also *Connecticut v. Teal*, 457 U.S. 440 (1982) (stating that employer may not justify using a promotion process that has a disparate impact on a protected class by also using an affirmative action system to achieve an appropriate racial balance after the fact).

⁶⁹ Michael Simon and Andrew Perry, *Regulation of Artificial Intelligence*, 78 THE BUSINESS LAWYER 317, 317 (Winter 2022–2023) (quoting Kevin Crowley, *AI’s Wild West*, COMMUNITECH (May 17, 2022)), <http://www.communitech.ca/technews/ais-wild-west.html> [<https://perma.cc/WZ3Y-LED3>].

⁷⁰ See, e.g., Algorithmic Accountability Act of 2019, H.R.2231, 116th Congress (2019–2020); Algorithmic Accountability Act of 2023, S.2892, 118th Congress (2023–2024).

⁷¹ Matt Brown & Matt O’Brien, *House Republicans include a 10-year ban on US states regulating AI in ‘big, beautiful’ bill*, Associated Press (May 16, 2025), <https://apnews.com/article/ai-regulation-state-moratorium-congress-39d1c8a0758ffe0242283bb82f66d51a> [<https://perma.cc/ZW6C-4RTH>].

⁷² *The Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People* (October 2022), <https://bidenwhitehouse.archives.gov/ostp/ai-bill-of-rights/> [<https://perma.cc/JUH3-EPLG>].

⁷³ Executive Order 14110, *Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence* (October 30, 2023).

⁷⁴ See, e.g., Gerrit De Vynck, Elizabeth Dwoskin, Nitasha Tiku, *The Silicon Valley realignment leading tech titans to Trump*, THE WASHINGTON POST, Jul. 18, 2024; Lauren Feiner, *What a second Trump presidency means for tech*, THE VERGE, Nov. 8, 2024; Eric Levitz, *Why Big Tech turned right*, Vox, Jan. 31, 2025.

Trump repealed the Biden administration’s executive order on AI almost immediately, replacing it with an order that promises to remove regulatory barriers to AI development.⁷⁵

For the next few years, then, we are unlikely to see any federal regulation of AI. State governments in the United States⁷⁶ and the governments of other nations may step in to fill the regulatory void.⁷⁷ But these efforts will be complicated by the Trump administration’s protectionist stance towards the tech industry.⁷⁸

The failure of legislatures to address the issue of AI bias reflect, in part, legitimate concerns about the difficulty in regulating in an area of rapid change. Lawmakers are reluctant to pass laws that might have the unintended side effect of stifling innovation⁷⁹ and, in the words of legal scholar, Lawrence Lessig, “give[] dinosaurs a veto over the future.”⁸⁰ However, reluctance to regulate also stems from tech industry lobbying that capitalizes upon the commonly held belief that it is unnecessary to regulate against AI bias because AI is not human and therefore inherently neutral and unbiased.⁸¹ As discussed, that belief has repeatedly been shown to be misguided.

Because industry self-regulation and legislation are unlikely to provide a timely or complete solution to the problem of AI bias, litigation is needed to fill the regulatory gaps. And of the existing antidiscrimination laws that provide private rights of action, the Fair Housing Act (FHA) is best suited to addressing bias in AI. It thus provides a natural starting point for attorneys who seek to develop a litigation strategy to tackle AI bias.

75 Executive Order 14179, *Removing Barriers to American Leadership in Artificial Intelligence* (January 23, 2025). The recently released AI Action Plan contains policy suggestions for implementing the White House’s AI deregulation goals. The White House, *Winning the Race: America’s AI Action Plan* (July 2025), available at <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf> [<https://perma.cc/YB5X-55B8>].

76 In May 2024, Colorado became the first state to regulate potential harm to consumers stemming from the use of AI, but the law does not go into effect until 2026. See Tamara Chuang, *Colorado becomes first state with law regulating potential consumer harms of artificial intelligence*, THE COLORADO SUN, May 18, 2024, at <https://coloradosun.com/2024/05/18/colorado-artificial-intelligence-law-signed/> [<https://perma.cc/H6JG-4DM5>].

77 China was the first country to adopt a national regulatory scheme for AI. Its Personal Information Protection Law was enacted in early 2021. Simon et al, *supra* note 69, at 318; David van Boven, et al., *The Legal Implications of Explaining Artificial Intelligence*, 5 J. OF ROBOTICS, A.I. & L. 157, 165 (2022). The EU was next to weigh in, adopting its Artificial Intelligence Act in May 2021. See Council of the EU, Press Release, *Artificial intelligence (AI) act: Council gives final green light to the first worldwide rules on AI* (May 21, 2024), available at <https://www.consilium.europa.eu/en/press/press-releases/2024/05/21/artificial-intelligence-ai-act-council-gives-final-green-light-to-the-first-worldwide-rules-on-ai/pdf/>.

78 See, e.g., Anne Applebaum, *Europe’s Elon Musk Problem*, THE ATLANTIC, January 27, 2025 (suggesting that President Trump’s hostility to the European Union stems from its power to regulate tech companies if the EU countries continue to act as a block); The White House, *Winning the Race: America’s AI Action Plan* (July 2025), available at <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf> (suggesting that federal agencies “consider a state’s AI regulatory climate” when deciding how to award AI-related discretionary funds).

79 See, e.g., Chuang, *supra* note 77. In a letter to the Colorado legislature, the governor noted that he was “concerned about the impact this law may have on an industry that is fueling critical technological advancements across our state for consumers and enterprises alike. Government regulation that is applied at the state level in a patchwork across the country can have the effect to tamper innovation and defer competition in an open market.” *Id.*

80 Lawrence Lessig, *FREE CULTURE: HOW BIG MEDIA USES TECHNOLOGY AND THE LAW TO LOCK DOWN CULTURE AND CONTROL CREATIVITY*, 199 (2004).

81 Katyal, *Private Accountability*, *supra* note 35, at 65 (noting that algorithmic decision making, because it involves “rote analysis of quantitative information,” can appear to be free of the “subjectivity and bias” of human judgment).

III. The FHA

As its name suggests, the FHA aims to advance “the policy of the United States to provide, within constitutional limitations, for fair housing throughout the United States.”⁸² Congress enacted the FHA in 1968, following “a period of considerable social unrest” in the United States.⁸³ In 1967, a period of civil unrest in America’s urban centers fueled a conservative backlash to then President Lyndon Johnson’s Great Society initiatives.⁸⁴ In response, President Johnson created the National Advisory Commission on Civil Disorders, frequently referred to as the “Kerner Commission,” to study what had happened and identify potential solutions.⁸⁵ Ultimately, the Commission “described the riots as the outgrowth of racial inequality and oppression.”⁸⁶ The Commission identified a number of institutional forces, including “residential segregation and unequal housing and economic conditions in the inner cities as significant, underlying causes of the [country’s] social unrest.”⁸⁷

By the 1960s, when the Kerner Commission was writing its report, as many as two-thirds of nonwhite families were concentrated in urban areas that, in the words of the Commission, were “marked by substandard housing and general urban blight.”⁸⁸ This situation did not arise by chance. Although the Supreme Court had declared *de jure* race discrimination in residential housing unconstitutional in the early part of the 20th century, patterns of housing segregation ballooned in the mid-20th century. This increase was due to factors such as suburban development and accompanying white flight; and segregation practices, “sometimes with governmental support,” such as racially restrictive covenants that prevented the conveyance of property to people who were not white, real-estate agents steering “potential buyers

⁸² 42 U.S.C. § 3601 (“Declaration of Policy”); *see also* Stephen Clowney, *Do Fraternities Violate the Fair Housing Act? An Empirical Study of Segregation in the Greek Organizations*, 41 *YALE L. & POL’Y REV.* 152, 191 (2023) (describing the FHA as having “twin goals of integration and nondiscrimination”).

⁸³ *Texas Dep’t of Hous. & Cmty. Affs. v. Inclusive Cmty. Project, Inc.* [hereinafter *Inclusive Cmty. Project, Inc.*], 576 U.S. 519, 529 (2015); *see also* Robert G. Schwemm, *Source-of-Income Discrimination and the Fair Housing Act*, 70 *CAS. W. RES. L. REV.* 573, 578 (2020) (discussing the reasons for the FHA’s enactment).

⁸⁴ *See* The National Advisory Commission on Civil Disorders, *THE KERNER REPORT*, xiv (Princeton University Press 2016). The Great Society was one of the largest social reform movements in modern American history. Its goals included addressing inequality, poverty, and environmental issues.

⁸⁵ *See id.*, at xvi; *see also Inclusive Cmty. Project, Inc.*, 576 U.S. at 530; *see also* Schwemm, *supra* note 84, at 578 (discussing the reasons for the FHA’s enactment).

⁸⁶ The National Advisory Commission on Civil Disorders, *THE KERNER REPORT*, p. xxvii (Princeton University Press 2016).

⁸⁷ *Inclusive Cmty. Project, Inc.*, 576 U.S. at 530; *see also* Schwemm, *supra* note 84, at 578 (discussing the reasons for the FHA’s enactment).

⁸⁸ *Inclusive Cmty. Project, Inc.*, 576 U.S. at 529 (2015); *see also* Schwemm, *supra* note 84, at 578 (discussing the reasons for the FHA’s enactment).

to consider homes in racially homogenous areas,” and redlining, a discriminatory lending practices, that prevented non-white families from purchasing homes in affluent areas.⁸⁹

To address this pattern of discrimination, the Kerner Commission recommended that Congress enact “a comprehensive and enforceable open-occupancy law making it an offense to discriminate in the sale or rental of any housing . . . on the basis of race, creed, color, or national origin.”⁹⁰ The April 1968 assassination of Dr. Martin Luther King, Jr. spurred Congress to adopt the Kerner Commission’s recommendation and pass the Fair Housing Act.⁹¹

This section provides relevant background on the FHA. It first explains why AI developers are subject to suit under the FHA, then explains the claims available under the FHA, and finally discusses the current legal standard for one of those claims: a disparate impact claim.

A. Developers and Users of AI That Makes Housing Decisions Are Subject to Suit under the FHA

Today, the FHA makes it unlawful to discriminate in the sale or rental of housing “because of race, color, religion, sex, familial status, . . . national origin[,]” or disability.⁹² Prohibited practices include 1) “refus[ing] to sell or rent . . . a dwelling;” 2) “discriminat[ing] against any person in the terms, conditions, or privileges of sale or rental of a dwelling[;]” 3) publishing “any notice, statement, or advertisement, with respect to the sale or rental of a dwelling that indicates any preference, limitation, or discrimination based on race, color, religion, sex, handicap, familial status, or national origin[;]” and 4) refusing to permit modifications to units that would allow a person with a disability to use and enjoy the property.^{93,94}

Unlike many other anti-discrimination laws, the FHA does not limit the range of potential defendants by defining a “set of covered entities.”⁹⁵ Instead, anyone who engages in any of the practices that the FHA prohibits can potentially be held liable⁹⁶ in a private lawsuit, Department of Justice enforcement action, or administrative complaint to the Department of

⁸⁹ *Inclusive Cmty. Project, Inc.*, 576 U.S. at 528-29 (2015) (citing *Shelley v. Kraemer*, 334 U.S. 1 (1948); M. Klarman, UNFINISHED BUSINESS: RACIAL EQUALITY IN AMERICAN HISTORY 140–141 (2007); Brief for Housing Scholars as Amici Curiae, *Inclusive Cmty. Project, Inc.*, No. 13–1371, 2014 WL 7405732, at *22-23 (Dec. 23, 2014); K. Clark, DARK GHETTO: DILEMMAS OF SOCIAL POWER 11, 21–26 (1965)); see also Schwemm, *supra* note 84, at 578 (discussing the reasons for the FHA’s enactment).

⁹⁰ *Inclusive Cmty. Project, Inc.*, 576 U.S. at 529–530 (citing KERNER COMMISSION REPORT).

⁹¹ *Id.* at 530.

⁹² 42 U.S.C. § 3604.

⁹³ *Id.*

⁹⁴ In some circumstances, the FHA does not apply to owner-occupied multifamily housing with four or fewer units. 42 U.S.C. § 3603.

⁹⁵ Andrew D. Selbst & Solon Barocas, *Unfair Artificial Intelligence: How FTC Intervention Can Overcome the Limitations of Discrimination Law*, 171 U. PA. L. REV. 1023, 1033 (2023) [hereinafter Selbst, et al., *Unfair AI*].

⁹⁶ See 42 USC § 3604(a)-(b); Statement of Interest of the United States, *Louis v. SafeRent Sols., LLC.*, No. 22-cv-10800, 2023 WL 140377 (D. Mass. Jan. 9, 2023); see also Selbst et al., *Unfair AI*, *supra* note 96, at 1033.

Housing and Urban Development.⁹⁷ The FHA thus is already broad enough to reach misconduct of both AI developers whose products render biased housing decisions and those who rely on AI products to render biased housing decisions.⁹⁸

Examples of AI tools already in use that fall within the purview of the FHA include various tenant screening programs,⁹⁹ algorithms that deliver targeted housing advertisements,¹⁰⁰ and AI-powered mortgage origination and default forecasting models.¹⁰¹ Indeed, there have already been FHA lawsuits against developers of some of these types of programs.¹⁰² For example, in a 2023 putative class action against tenant screening service SafeRent Solutions, LLC and a property management company that used the service, the plaintiffs' FHA claims for race and national origin discrimination survived a motion to dismiss.¹⁰³ The judge rejected SafeRent's argument that it could not be liable under the FHA because "it 'is not a landlord and does not otherwise have authority to make housing decisions'"; the judge held instead that SafeRent does make housing decisions when its program delivers an "accept/decline/conditional decision" on applicants.¹⁰⁴

In a joint statement filed in the SafeRent case, HUD and the DOJ advocated for the broad reading of the FHA that the judge eventually adopted. They reasoned that the FHA's plain language extends to anyone who has "personally committed or contributed to a Fair Housing Act violation" and asked the court to reject the argument that the FHA is limited to those who have final say on "housing decisions."¹⁰⁵ Thus, they concluded, tenant screening companies are subject to liability under the clear language of the FHA.¹⁰⁶

⁹⁷ 42 U.S.C. §§ 3610, 3614 (2012).

⁹⁸ See 42 U.S.C. §§ 3601-3614 (2012); see also Selbst et al., *Unfair AI*, *supra* note 96, at 1033 (arguing that the FHA reaches AI developers); Stevie J. Swanson, *Indignity Perpetuated: Race-Based Housing Post-Reconstruction to the Fair Housing Act's Impact on the Digital Age: Where Do We Go from Here?*, 22 CONN. PUB. INT. L.J. 126, 145-146 (2023) (arguing, based upon similar reasoning, that liability under the FHA extends to anyone who publishes a racially restrictive covenant in an online database).

⁹⁹ See Schneider, *supra* note 7, at 269-270 (discussing tenant screening programs CoreLogic, Tenant Alert, and Smart Move).

¹⁰⁰ Archives, *Justice Department Secures Groundbreaking Settlement Agreement with Meta Platforms, Formerly Known as Facebook, to Resolve Allegations of Discriminatory Advertising*, U.S. DEP'T OF JUSTICE.

(Jun. 21, 2022), <https://www.justice.gov/archives/opa/pr/justice-department-secures-groundbreaking-settlement-agreement-meta-platforms-formerly-known> [<https://perma.cc/EGZ4-3D6T>].

¹⁰¹ Grimm, et al., *supra* note 11, at 33-34 (discussing AI's application in the loan origination space).

¹⁰² See, e.g., *Conn. Fair Hous. Ctr. v. Corelogic Rental Property Sols., LLC*, 369 F. Supp. 3d 362, 371-72 (D. Conn. 2019); Archives, *Justice Department Secures Groundbreaking Settlement Agreement with Meta Platforms, Formerly Known as Facebook, to Resolve Allegations of Discriminatory Advertising*, U.S. DEP'T OF JUSTICE (Jun. 21, 2022), <https://www.justice.gov/archives/opa/pr/justice-department-secures-groundbreaking-settlement-agreement-meta-platforms-formerly-known> [<https://perma.cc/H7C8-Z8YT>].

¹⁰³ Order, *Louis v. SafeRent Solutions, LLC.*, 685 F. Supp. 3d 19, 34-36 (D. Mass. 2023).

¹⁰⁴ *Id.*

¹⁰⁵ Statement of Interest of the United States, *Louis v. SafeRent Solutions, LLC.*, No. 22-cv-10800, 2023 WL 140377 (D. Mass. Jan. 9, 2023) (citing *Sabal Palm Condos. of Pine Island Ridge Ass'n, Inc. v. Fischer*, 6 F. Supp. 3d 1272, 1293 (S.D. Fla 2014)).

¹⁰⁶ Statement of Interest of the United States, *Louis v. SafeRent Solutions, LLC.*, No. 22-cv-10800, 2023 WL 140377 (D. Mass. Jan. 9, 2023).

Based upon the foregoing, the FHA’s language is broad enough to address the modern problem of discrimination by algorithm.¹⁰⁷ Accordingly, a plaintiff who has suffered discrimination in the sale or rental of housing can consider several potential claims against the AI developers and users who played a role in the discriminatory decision.

B. Types of Claims Available Under the FHA

FHA plaintiffs can bring three types of claims. A claim for disparate treatment alleges intentional discrimination.¹⁰⁸ A claim for disparate impact alleges that a policy adopted without the apparent intent to discriminate nonetheless violate the law because it has a disparate impact on a protected class.¹⁰⁹ A failure to accommodate claim alleges failure to provide reasonable accommodations for a person with disabilities.¹¹⁰ This article focuses on litigating disparate impact claims because, as discussed, AI discrimination is generally a byproduct of deficiencies in the AI training process.¹¹¹ As such, defendants are likely to argue that AI discrimination is unintentional, and plaintiffs are therefore well-advised to include in their complaint a claim that can reach unintentional discrimination. However, that does not mean that plaintiffs should ignore the possibility of a disparate treatment claim;¹¹² the discovery concepts discussed *infra* can also be used to uncover evidence of intentional discrimination.

Courts have recognized two types of disparate impact claims under the FHA. The first is a claim for “harm to the community generally by creating, increasing, reinforcing, or perpetuating segregated housing patterns.”¹¹³ The second is a claim for “harm to a particular group of persons by a disparate impact.”¹¹⁴

¹⁰⁷ See 42 USC § 3604(a)-(b); Press Release, Dep’t of Hous. and Urban Dev., *HUD Charges Facebook with Hous. Discrimination over Co.’s Targeted Advert. Practices* (Mar. 28, 2019), <https://archives.hud.gov/news/2019/pr19-035.cfm> [<https://perma.cc/4H35-G55D>] (“Even as we confront new technologies, the fair housing laws enacted over half a century ago remain clear—discrimination in housing-related advertising is against the law.”); Statement of Interest of the United States, *Louis v. SafeRent Solutions, LLC.*, No. 22-cv-10800, 2023 WL 140377 (D. Mass. Jan. 9, 2023).

¹⁰⁸ See *Astralis Condo. Ass’n v. Sec’y*, *U.S. Dep’t of Hous. & Urb. Dev.*, 620 F.3d 62, 66 (1st Cir. 2010); *Smith & Lee Assocs. v. City of Taylor*, 102 F.3d 781, 790 (6th Cir. 1996).

¹⁰⁹ Statement of Interest of the United States, *Louis v. SafeRent Solutions, LLC.*, No. 22-cv-10800, 2023 WL 140377 (D. Mass. Jan. 9, 2023).

¹¹⁰ See *Astralis Condo. Ass’n v. Sec’y*, 620 F.3d at 66; *Smith & Lee Assocs.*, 102 F.3d at 790.

¹¹¹ See sec. II.B, *supra*.

¹¹² Professor Stephanie Bornstein has persuasively argued that at least some AI training decisions are best understood as intentional discrimination because the decision parameters and training data sets reflect stereotype bias. To illustrate, she offers the example of a human programmer who decides that good leaders are “aggressive” and gives an AI hiring program a training set that associates aggression with masculinity. Stephanie Bornstein, *Antidiscriminatory Algorithms*, 70 ALA. L. REV. 519, 564 (2018).

¹¹³ Implementation of the Fair Housing Act’s Discriminatory Effects Standard, 78 Fed. Reg. at 11460, 11469 (Feb. 15, 2013).

¹¹⁴ *Id.*

C. *Inclusive Communities* and the Current Legal Standard for Disparate Impact Claims under the FHA

Since its enactment, the FHA has been understood to protect against housing policies with a disparate impact on protected classes.¹¹⁵ Indeed, in the nearly 50 years between the FHA's enactment and the Supreme Court's decision in *Inclusive Communities*,¹¹⁶ every circuit court that considered the issue held that disparate impact claims are cognizable under the FHA.¹¹⁷ In 2015, the Supreme Court decided *Texas Dept of Housing and Community Affairs v. Inclusive Communities Project*, which confirmed that the FHA provides for disparate impact claims, and adopted a three-step, burden-shifting framework for determining disparate impact liability.¹¹⁸

Given this history, it would be reasonable to characterize the law on this issue as settled. That has not prevented conservative groups' persistent efforts to undermine the disparate impact doctrine. Those who would eliminate disparate impact liability argue that it is inconsistent with the Equal Protection Clause because addressing the disparate impact of a policy requires evaluating how the policy impacts protected classes.¹¹⁹ It follows, they argue, that some combination of fear and frugality will compel businesses to abandon necessary policies in favor of adopting race, gender, and other protected class quotas that disadvantage majority groups.¹²⁰ Although this argument is charitably described as speculative,¹²¹ it provides the ideological underpinnings for the Trump administration's attacks on disparate impact liability.¹²²

115 See Schneider, *supra* note 7, at 262 (discussing the history of disparate impact claims under the FHA).

116 576 U.S. 519.

117 See *Huntington Branch, NAACP v. Town of Huntington*, 844 F.2d 926, 935-36 (2d Cir. 1988); *Arthur v. City of Toledo*, 782 F.2d 565, 574-75 (6th Cir. 1986); *Hanson v. Veterans Admin.*, 800 F.2d 1381, 1386 (5th Cir. 1986); *United States v. Marengo Cty. Comm'n*, 731 F.2d 1546, 1559, n.20 (11th Cir. 1984); *Smith v. Town of Clarkton*, 682 F.2d 1055, 1065 (4th Cir. 1982); *Halet v. Wend Inv. Co.*, 672 F.2d 1305, 1311 (9th Cir. 1982); *Resident Advisory Bd. v. Rizzo*, 564 F.2d 126, 146 (3d Cir. 1977); *Metro. Hous. Dev. Corp. v. Vill. of Arlington Heights*, 558 F.2d 1283, 1290 (7th Cir. 1977); *United States v. City of Black Jack*, 508 F.2d 1179, 1184-85 (8th Cir. 1974).

118 576 U.S. 519.

119 See, e.g., Alison Somin, *Disparate Impact as a Non-Delegation Violation and Major Question*, 2024 HARV. J.L. & PUB. POL'Y PER CURIAM 18 (2024); Kenneth L. Marcus, *The War Between Disparate Impact and Equal Protection*, CATO SUP. CT. REV., Vol. 2008-2009 (August 26, 2009).

120 See, e.g., Somin, *supra* note 118, at 18 (contending that companies may be afraid to adopt any policy with a disparate impact—even one supported by business necessity—out of fear of litigation); Marcus, *supra* note 118 (contending that disparate impact forces employers to adopt racial quotas because doing so is less expensive than developing hiring policies that actually identify the most qualified candidates); see also *Ricci v. DeStefano*, 557 U.S. 557, 594 (2009) (Scalia, J., concurring) (“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.”).

121 If remedying disparate impact required (or strongly encouraged) the use of quotas, one would expect to see evidence of the wide adoption of these types of quotas reflected in statistics related to hiring, housing, and other areas covered by civil rights law. For example, 50 plus years of race-based hiring quotas would be expected to produce employment statistics that show that BIPOC people are represented in the work force at a rate roughly proportionate to their representation in the general population. Actual employment statistics reflect quite the opposite. For example, in the legal profession, employment gains by attorneys of color have been slow, and BIPOC attorneys are still significantly underrepresented in leadership roles in the profession. See Jessica A. Mahon Scoles, *Can I Bring Myself to Court?: Teaching Presentation Style in the Trial Advocacy Courtroom*, 49 VERMONT L.R. 1, 11 (2024). A full analysis of the alleged conflict between disparate impact liability and the Equal Protection Clause is outside the scope of this article.

122 Executive Order 14281, *Restoring Equality of Opportunity and Meritocracy* (Apr. 23, 2025), <https://www.whitehouse.gov/presidential-actions/2025/04/restoring-equality-of-opportunity-and-meritocracy/> [<https://perma.cc/D59M-2L85>].

President Trump’s hostility to diversity, equity and inclusion makes it reasonable to assume that his administration will continue its efforts to eliminate disparate impact liability. While it is not possible to predict precisely what shape those efforts will take, the President’s past behavior suggests that his administration will focus on executive action. For example, in his first term, President Trump attempted to revise the regulations governing the FHA in a manner that effectively eliminated disparate impact liability for AI users by creating a defense to liability for the disparate impact of AI decision making if: 1) the algorithms inputs were not “substitutes of close proxies” for protected characteristics; and 2) a neutral third party either developed or certified the fairness of the system.¹²³ More recently, on April 23, 2025, President Trump signed an executive order with the purported aim of eliminating disparate impact liability. Among other things, the order directs administrative agencies to deprioritize enforcement of disparate impact claims and directs the Attorney General “to repeal or amend the implementing regulations for Title VI of the Civil Rights Act of 1964 for all agencies to the extent they contemplate disparate-impact liability.”¹²⁴

As troubling as these actions are, they do not impact private plaintiff’s ability to bring disparate impact claims under the FHA. Even if the administration successfully amends the regulations implementing the FHA, such an amendment is not likely to change the federal courts’ approach in disparate impact cases. That is because the existing circuit split on the question of how a plaintiff may prove disparate impact liability derives from competing interpretations of the Supreme Court’s *Inclusive Communities* decision, and not from the courts’ interpretation of the relevant regulations.¹²⁵ Thus, unless the Supreme Court overrules *Inclusive Communities*, courts are likely to continue to decide disparate impact claims along the same ideological faults that existed before President Trump took office.

Inclusive Communities adopted a three-step burden shifting framework for litigating disparate impact claims under the FHA.¹²⁶ Under the first step in the *Inclusive Communities* framework, the plaintiff must demonstrate a robust causal connection between the defendant’s challenged policy and the disparate impact on the protected class.¹²⁷ Under the second step, the defendant has the burden to “state and explain the valid interest served by their policies.” The defendant’s step two burden is analogous to Title VII

123 HUD’s Implementation of the Fair Housing Act’s Disparate Impact Standard, 85 Fed. Reg. 60288 (Sept. 24, 2020).

124 Executive Order 14281, *Restoring Equality of Opportunity and Meritocracy* (Apr. 23, 2025), <https://www.whitehouse.gov/presidential-actions/2025/04/restoring-equality-of-opportunity-and-meritocracy/> [<https://perma.cc/2MNG-NSM4>].

125 *See, e.g., Reyes v. Waples Mobile Home Park L.P.*, 903 F.3d 415, 424 at n.4 (“Without deciding whether there are meaningful differences between the frameworks, we note that the standard announced in *Inclusive Communities* rather than the HUD regulation controls our inquiry.”); *see also* sec. III.C, *infra*.

126 *Inclusive Cmty. Project, Inc.*, 576 U.S. at 540–42.

127 *Id.* at 542.

of the Civil Rights Act’s business necessity standard,¹²⁸ which requires proof of a “relationship between the employment policy and overall job success.”¹²⁹ Under the third step, the burden shifts back to the plaintiff to prove that the defendant’s asserted interests “could be served by another practice that has a less discriminatory effect.”¹³⁰

Courts are split on how this framework is to be applied. As discussed in more detail, *infra*, the majority apply the *Inclusive Communities* framework in a manner consistent with 24 C.F.R. § 100.500, a regulation that the Department of Housing and Urban Development (HUD) promulgated in 2013. The minority reads *Inclusive Communities* as establishing a more stringent standard than the 2013 regulation with regard to pleading, proof, or both. The challenges that an FHA plaintiff faces when litigating a disparate impact claim thus vary depending upon the circuit in which they file.

1. Plaintiff’s Stage One Burden

The circuit split regarding the *Inclusive Communities* framework mainly impacts the plaintiff’s burden at stage one. This section first discusses the majority approach that treats *Inclusive Communities* as implicitly endorsing the 2013 regulation. It then discusses minority approaches that impose tougher pleading or proof standards at stage one.

a. Approach One: Treat the Inclusive Communities Standard as Consistent with the 2013 Regulation.

Several circuits and a handful of district courts have interpreted *Inclusive Communities* as an endorsement of the 2013 regulation’s three-stage burden shifting scheme. The 2013 regulation describes the burden-shifting in FHA disparate impact cases as follows:

- (1) The charging party ..., or the plaintiff..., has the burden of proving that a challenged practice caused or predictably will cause a discriminatory effect. (2) Once the charging party or plaintiff satisfies the burden of proof set forth in paragraph (c)(1) of this section, the respondent or defendant has the burden of proving that the challenged practice is necessary to achieve one or more substantial, legitimate, nondiscriminatory interests of the respondent or defendant. (3) If the respondent or defendant satisfies the burden of proof set forth in paragraph (c)(2) of this section, the charging party or plaintiff may still

¹²⁸ *Id.*

¹²⁹ Garrett A. R. Yursza Warfield & David J. Rini, *New EEOC Guidance: Implications for Ex-Offender Reentry and Employment, or “It Is Hard to Articulate the Minimum Qualification for Posing A Low Risk of Attacking Someone?”*, 95 MASS. L. REV. 195, 206 (2013); see also *Watson v. Fort Worth Bank & Tr.*, 487 U.S. 977, 1005 (1988) (Blackmun, J. concurring in part) (stating in interpretation of a superseded version of Title II, that to qualify as a business necessity, “[t]he criterion must directly relate to a prospective employee’s ability to perform the job effectively”).

¹³⁰ *Inclusive Cmty. Project, Inc.*, 576 U.S. at 527 (quoting 24 C.F.R. § 100.500(c)(3)).

prevail upon proving that the substantial, legitimate, nondiscriminatory interests supporting the challenged practice could be served by another practice that has a less discriminatory effect.”¹³¹

An approach consistent with this regulation has two key features. First, the plaintiff’s pleading burden is limited to stage one in the burden shifting scheme. In other words, the plaintiff must plead facts that plausibly show that the defendant has a practice that has caused or predictably will cause a discriminatory effect.¹³² The defendant’s stage two burden is treated like an affirmative defense that the plaintiff does not need to plead around.¹³³ It follows that the plaintiff does not need to plead facts that show that the defendant’s policy is illegitimate or unnecessary.¹³⁴

Second, plaintiffs can meet their burden of establishing that the defendant’s policy caused a discriminatory effect by pointing to “sufficiently substantial” statistical disparities in how the policy impacts protected class members versus those who are not in the class.¹³⁵ It is not enough to show that a policy denies housing to a large number of people within a protected class. Instead, the plaintiff must provide statistical evidence that compares the policy’s impact on protected class members to the policy’s impact on those outside the protected class.¹³⁶ The comparative data ensures that the plaintiff

¹³¹ 24 C.F.R. § 100.500.

¹³² See *Reyes v. Waples Mobile Home Park L.P.*, 903 F.3d 415, 425 (4th Cir. 2018); *Oviedo Town Ctr. II, L.L.L.P. v. City of Oviedo*, 759 F. App’x 828, 835 (11th Cir. 2018); see also *Huskey v. State Farm Fire & Cas. Co.*, No. 22C7014, 2023 WL 5848164 (N.D. Ill. Sept. 11, 2023); *Louis v. SafeRent Solutions, LLC*, 685 F. Supp. 3d 19, 26 (D. Mass. 2023); *Jeter v. City of Cincinnati*, No. 1:20-CV-581, 2021 WL 4504247, at *9 (S.D. Ohio Sept. 30, 2021) (rejecting argument that plaintiff’s complaint must allege that defendant’s policy is the sole factor contributing to causation and requiring only facts giving rise to a plausible inference of causation); *Conn. Fair Hous. Ctr.*, 369 F. Supp. 3d at 371-72.

¹³³ See, e.g., *Huskey*, 2023 WL 5848164, at *8 (noting the difference in pleading and proof requirements for disparate impact claims).

¹³⁴ See *id.*

¹³⁵ *Reyes*, 903 F.3d at 425; see also *Oviedo Town Ctr. II, L.L.L.P.*, 759 F. App’x at 835 (plaintiff must provide comparative data that shows how the challenged policy impacts the protected class versus the nonprotected class); *Oxford House, Inc. v. Twp. of N. Bergen*, 158 F.4th 486, 495 (3d Cir. 2025) (stating that plaintiffs complaint must “y suggests an actual statistical disparity for members of the protected class and raises a reasonable expectation that discovery will reveal methodologically sound statistical evidence at trial”).

¹³⁶ Compare *Oviedo Town Ctr. II, L.L.L.P.*, 759 F. App’x at 835; *Louis*, 685 F. Supp. 3d at 26 (denying motion to dismiss disparate impact claim against landlord and tenant screening company where plaintiffs alleged that screening algorithm was more likely to reject Black and Hispanic applicants because those applicants were more likely to have low credit scores as compared to applicants outside the protected class); *Jeter v. City of Cincinnati*, No. 1:20-CV-581, 2021 WL 4504247, at *9 (S.D. Ohio Sept. 30, 2021) (denying motion to dismiss where plaintiff’s alleged that city’s tax abatement policy accelerated racial segregation because abatements went disproportionately to white neighborhoods); *Crossroads Residents Organized for Stable & Secure Residencies v. MSP Crossroads Apts. LLC*, No. CV 16-233 ADM/KMM, 2016 WL 3661146, at *7 (D. Minn. Jul. 5, 2016) (stating that disparate impact is commonly shown in cases challenging a landlord’s rental policy by “alleg[ing] that tenants at the complex harmed by Defendants’ actions are disproportionately protected class members”); *Fussell v. Sun Vista RV Resort LLC*, No. CV-24-03414-PHX-SHD, 2025 WL 2161491, at *6 (D. Ariz. July 29, 2025) (allegations that the defendants policy of requiring a “700+ credit score,” impact[ed] recent immigrants because immigrants lack an ‘extensive U.S. credit history.’” sufficient to meet stage 1 pleading burden); and *Boykin v. Fenty*, 650 Fed. App’x. 42, 44 (D.C. Cir. 2016) (per curiam) (affirming dismissal of plaintiff’s disparate impact claim arising from city’s closure of two homeless shelters because complaint “failed to allege facts suggesting that the closure affected a greater proportion of disabled individuals than non-disabled”); *Burbank Apartments Tenant Ass’n v. Kargman*, 474 Mass. 107 (2016) (affirming dismissal of disparate impact suit against landlord who switched from participating in project based subsidy program for low income tenants to accepting Section 8 vouchers because allegation that project based subsidy recipients were more likely to be members of a protected class than the general public did not establish disparate impact); *Binns v. City of Marietta Ga., (Hous. Choice Voucher Program)*, 704 F. App’x 797, 802 (11th Cir. 2017) (affirming dismissal of FHA disparate impact claim because the plaintiff “provided no comparative data, statistical or otherwise, to show that elderly and disabled participants are disproportionately impacted by the City’s policy”); *Comfy Cottage, Inc. v. Pasco Cnty.*, No. 8:24-CV-991-MSS-TGW, 2025 WL 2443712, at *8 (M.D. Fla. Aug. 25, 2025) (granting defendant’s motion for summary judgment of disparate impact claim because while

has identified a true disparity in how different groups in a neighborhood are treated, rather than merely identifying a policy that exists in a neighborhood that includes many members of a protected class.¹³⁷

The Fourth Circuit’s decision in *Reyes v. Waples Mobile Home Park L.P.* provides a good illustration of this approach. In that case, the court held that the plaintiffs stated a disparate impact claim against a mobile home park based upon the park’s policy of requiring that all individuals living in the park document their immigration status. Statistical evidence that established that “Latinos were ten times more likely than non-Latinos to be adversely affected by the Policy” was sufficient to plead substantial disparity.¹³⁸

In comparison, in *Oviedo Town Center II, LLLP v. City of Oviedo*, the Eleventh Circuit applied the same test to reach an opposite result. In that case, the plaintiff, which operated a multi-unit affordable housing community, alleged that the city’s policy of setting utility rates based upon the number of units in the building had a disparate impact on racial minorities. To establish disparate impact, the plaintiff alleged that the percentage of “racial minority heads of households” in its building was greater than the percentage of “racial minority households” in the city.¹³⁹ The Eleventh Circuit held that this was insufficient to state a claim. The court explained: “To establish a prima facie case of disparate impact, the appellants would have needed to draw a comparison between (a) the percentage of racial minorities occupying multifamily properties impacted by the 2012 Policy throughout the City of Oviedo and (b) the percentage of non-minorities living in such properties, again throughout the City of Oviedo.”¹⁴⁰

At the pleading stage, statistical information is generally enough, by itself, to establish causation; as the case progresses, if statistical disparities are not enough to make causation self-evident, then the plaintiff may need additional evidence to establish the link between the defendant’s policy and the disparate impact.¹⁴¹

“[e]vidence that shows a significant discriminatory effect is enough to establish a prima facie violation of the FHA,” the plaintiff had no evidence, statistical or otherwise, to show that an ordinance requiring sprinklers in group housing had a disparate impact on people with disabilities.”); *Farhan v. 2715 NMA LLC*, 161 F.4th 475, 485-86 (7th Cir. 2025) (affirming dismissal of disparate impact claim based upon landlord’s ban of pro-Palestine symbols because, while the court was “concern[ed] over any slide towards heightened pleading standards,” the complaint did not “plead that the policy negatively impacts Palestinians more by proportion than non-Palestinians”).

¹³⁷ See generally cases cited in note 136.

¹³⁸ *Reyes*, 903 F.3d at 421.

¹³⁹ *Oviedo Town Ctr. II, L.L.L.P.*, 759 F. App’x at 833.

¹⁴⁰ *Id.* at 835. Cf. *Nat’l Fair Hous. All. v. Travelers Indem. Co.*, 261 F. Supp. 3d 20, 34 (D.D.C. 2017) (plaintiff’s complaint challenging insurers’ refusal to insure homes rented to Section 8 voucher recipients stated a claim where it pleaded facts to show that “because of the different composition of the affected population (voucher recipients) as compared to the District’s population as a whole, members of a protected class are more likely to be harmed by [insurer’s] policy than are other individuals”), and *Boykin*, 650 Fed. App’x. at 44 (plaintiff could have pled disparate impact on disabled homeless individuals by “for instance, includ[ing] an allegation that disabled homeless individuals are more likely to rely on low-barrier shelters than non-disabled homeless individuals”); *Burbank Apartments Tenant Ass’n*, 474 Mass. 107 (to establish that switch from low income subsidy to low income voucher program had a disparate impact on a protected class, plaintiff should have presented evidence that showed that subsidy recipients were more likely to belong to a protected class than voucher recipients).

¹⁴¹ See *Reyes*, 903 F.3d at 428 (large statistical anomalies may make “disparate impact of a policy ‘self-evident’”); *Oliver v. Navy Fed. Credit Union*, No. 1:23-CV-1731 (LMB/WEF), 2024 WL 2786905, at *8 (E.D. Va. May 30, 2024) (indicating willingness to revisit causation argument if plaintiffs did not uncover via discovery evidence linking underwriting policies to disparate impact)

Putting it all together, in jurisdictions that follow an approach consistent with the 2013 regulation, the plaintiff's stage one burden is to (1) identify a specific policy of the defendant, and (2) show that the policy caused (3) a substantial disparity between the treatment of protected and non-protected class members. The burden at stage two is on the defendant, who can defend a disparate impact claim by establishing that the policy at issue serves a legitimate non-discriminatory interest.

b. Approach Two: Treat the Inclusive Communities Standard as Requiring Heightened Pleading or Causation Requirements.

Other courts interpret the “robust causality” language in *Inclusive Communities* as imposing a “more demanding test” than that of the 2013 regulation.¹⁴² In courts that follow this approach, plaintiffs must still identify a policy and show that it has a disparate impact on a protected class. However, these courts that apply the “more demanding test” impose an additional pleading requirement, an additional proof requirement for causation, or both. These requirements limit plaintiffs’ ability to bring disparate impact claims. Indeed, when the first Trump administration sought to limit disparate impact liability for AI developers, they looked to cases that imposed this more demanding test for guidance.¹⁴³

First, some courts require that the plaintiff’s complaint “alleged facts plausibly demonstrating that the housing-code standards were arbitrary and unnecessary under the FHA.”¹⁴⁴ This approach effectively flips the burden at stage two of the analysis. Instead of allowing the defendant to avoid liability by proving that its policies serve legitimate ends, the plaintiff has the burden of pleading that the defendant’s policies do not serve legitimate ends because they are arbitrary and unnecessary.¹⁴⁵

Second, some courts require that the plaintiff not only prove that the defendant’s policy resulted in a disparate impact, but also that the policy caused the underlying disparities that led to the disparate impact.¹⁴⁶ In

¹⁴² See *Inclusive Communities Project, Inc. v. Lincoln Prop. Co.*, 920 F.3d 890, 902 (5th Cir. 2019) [hereinafter *Lincoln Prop.*]

¹⁴³ According to the rule amendment proposed in 2020, a disparate impact plaintiff must plead: “(1) That the challenged policy or practice is arbitrary, artificial, and unnecessary to achieve a valid interest or legitimate objective such as a practical business, profit, policy consideration, or requirement of law; (2) That the challenged policy or practice has a disproportionately adverse effect on members of a protected class; (3) That there is a robust causal link between the challenged policy or practice and the adverse effect on members of a protected class, meaning that the specific policy or practice is the direct cause of the discriminatory effect; (4) That the alleged disparity caused by the policy or practice is significant; and (5) That there is a direct relation between the injury asserted and the injurious conduct alleged.” *Mass. Fair Hous. Ctr. v. United States Dep’t of Hous. & Urb. Dev.*, 496 F. Supp. 3d 600, 605 (D. Mass. 2020) (enjoining the 2020 rule before implementation). As will be shown, *infra*, these concepts derive from 5th and 8th Circuit case law.

¹⁴⁴ *Ellis v. City of Minneapolis*, 860 F.3d 1106, 1112 (8th Cir. 2017); see also *Cobb Cnty. v. Bank of Am. Corp.*, No. 1:15-CV-04081-LMM, 2020 WL 13200158, at *4 (N.D. Ga. Sept. 18, 2020) (adopting *Ellis*’s approach to pleading disparate impact claims); *Smith v. New Neighborhoods, Inc.*, No. 3:24-CV-00327 (KAD), 2025 WL 2615052, at *7 (D. Conn. Sept. 10, 2025) (citing *Moody v. Related Companies, L.P.*, 620 F. Supp. 3d 51, 57 (S.D.N.Y. 2022) for the proposition that plaintiffs must plead, inter alia, “facts from which the Court can plausibly infer that (1) ‘the challenged policy or practice is arbitrary, artificial, and unnecessary to achieve a valid interest or legitimate objective such as a practical business, profit, policy consideration, or requirement of law’”).

¹⁴⁵ *Ellis*, 860 F.3d at 1112; *Cobb Cnty.*, No. 2020 WL 13200158, at *4.

¹⁴⁶ See *Lincoln Prop.*, 920 F.3d at 922; *Inclusive Cmty. Project, Inc. v. Heartland Cmty. Ass’n, Inc.*, 824 F. App’x 210, 213 (5th Cir. 2020) [hereinafter *Heartland*]; *Treeco v. Perrier Condo. Owners Ass’n, Inc.*, 519 F. Supp. 3d 342 (E.D. La.) (modified, 569 F. Supp. 3d 347 (E.D. La. 2021)).

jurisdictions that follow this approach, a plaintiff cannot prove a disparate impact claim if the disparate impact of the defendant’s policy arises from preexisting demographic patterns. Instead, the plaintiff must prove that the policy caused the relevant protected class to be the dominant group that the policy affects.¹⁴⁷

In jurisdictions that follow this approach, a defendant’s decision to adopt policies that aggravate existing patterns of inequality appears to be beyond the reach of a disparate impact claim. For example, the 5th Circuit has dismissed disparate impact claims against landlords who refuse to rent to Section 8 voucher holders because the plaintiffs could not show that the landlord’s policy of refusing to rent to voucher holders caused the percentage of voucher holders in the area to be disproportionately members of protected classes.¹⁴⁸ Similarly, a federal district court in Louisiana granted a defendant’s motion for summary judgment of a disparate impact claim against a condominium association that enacted a rule requiring a certain square footage per occupant in each unit.¹⁴⁹ Although the requirement had a disproportionate impact on families with children, the court reasoned that the plaintiffs had not offered “evidence to support an inference [the association’s] policies caused households with children to make up a larger percentage of households of six than households without children.”¹⁵⁰

2. Defendant’s Stage Two Burden

If the plaintiffs meet their stage one burden, the burden shifts to the defendant to provide a “legally sufficient justification.”¹⁵¹ The defendant must offer evidence to show “that the policy ‘serves in a significant way,’ its legitimate interests.”¹⁵² The “defendant need not demonstrate that the challenged policy is ‘essential’ or ‘indispensable’ to its business.”¹⁵³ “A legally sufficient justification must be supported by evidence and may not be hypothetical or speculative.”¹⁵⁴

In practice, the defendant’s stage two burden is not onerous.¹⁵⁵ Defendants must explain the business or governmental interest that the challenged policy advances, and they must present evidence that the concerns underlying those

¹⁴⁷ See *Lincoln Prop. Co.*, 920 F.3d at 922; *Heartland*, 824 F. App’x at 213; *Trecece*, 519 F. Supp. 3d 342.

¹⁴⁸ *Lincoln Prop. Co.*, 920 F.3d 890; *Heartland*, 824 F. App’x 210.

¹⁴⁹ *Trecece v. Perrier Condo. Owners Ass’n, Inc.*, 519 F. Supp. 3d 342, 363 (E.D. La.), *modified*, 569 F. Supp. 3d 347 (E.D. La. 2021).

¹⁵⁰ *Trecece*, 519 F. Supp. 3d at 362. The plaintiff’s argument as to one of the occupancy rules also failed because the plaintiffs could not show that the rule was being enforced. *Id.* at 362-63.

¹⁵¹ 24 C.F.R. § 100.500(b)(2); see also *Inclusive Cmty. Project, Inc.*, 576 U.S. at 540-41.

¹⁵² *Sw. Fair Hous. Council, Inc. v. Maricopa Domestic Water Improvement Dist.*, 17 F.4th 950, 967 (9th Cir. 2021).

¹⁵³ *Id.*

¹⁵⁴ 24 C.F.R. § 100.500(b)(2); see also *Fair Hous. Ctr. of Wash. v. Breier-Scheetz Props., LLC*, 743 F. App’x 116, 118 (9th Cir. 2018) (holding that the defendant’s support for policy was uncorroborated and too hypothetical to carry the burden of proof).

¹⁵⁵ See *Trecece*, 519 F. Supp. 3d at 354 (citing *Lincoln Prop.*, 920 F.3d at 902) (stating “leniency must be given to defendants to explain the valid interest served by their policy”).

interests are grounded in fact, not speculation.¹⁵⁶ For example, in *Graul*, a District of Rhode Island judge found that a landlord did not carry its burden at stage two of the framework where it pointed to the state housing code to justify a policy that had a disparate impact on families with children.¹⁵⁷ While complying with state law can be a legitimate interest in the abstract, the court found that the landlord's concerns were speculative on the facts of the case. Because the landlord's interpretation of the housing code was incorrect, it was not at risk of violating the housing code.¹⁵⁸

3. Plaintiff's Stage Three Burden

If the defendant introduces evidence sufficient to meet its stage two burden, then the burden shifts back to the plaintiff¹⁵⁹ to show that the defendant's asserted interests "could be served by another practice that has a less discriminatory effect."¹⁶⁰ While some courts have applied a standard requiring that the proposed less discriminatory alternative be as effective as the challenged policy, regulatory guidance interpreting the FHA suggests that the less discriminatory alternative that the plaintiff proposes need not be as effective as the challenged policy and that it is acceptable to impose some administrative cost on the defendants to develop a new alternative.¹⁶¹

Understanding what is causing the disparate impact of an AI system's decision making, whether the use of AI serves a legitimate interest, and whether there are less discriminatory alternatives, will often require access to details about the AI system at issue. For example: how was the AI trained; what criteria does it weigh when rendering decisions; and are humans involved in the decision-making process at any point? Because AI decision making is something of a black box, even to those who design AI, plaintiffs may need access to the AI system itself to answer some of these questions. In short, FHA plaintiffs who believe that AI has discriminated against them will need discovery about AI systems. This raises the question whether these plaintiffs are likely to get that discovery.

¹⁵⁶ Compare *Sv. Fair Hous. Council, Inc.*, 17 F.4th at 967; *Mbany Mgmt. v. Cty. of Nassau*, 819 F.3d 581, 617 (2d Cir. 2016) (holding that defendants carried their burden) and *R.I. Comm'n for Human Rights v. Graul*, 120 F. Supp. 3d 110 (D.R.I. 2015); *Reyes*, 91 F.4th at 277 (holding that defendants did not carry their burden).

¹⁵⁷ *Graul*, 120 F. Supp. 3d at 127.

¹⁵⁸ *Id.* at 128-29.

¹⁵⁹ Some courts have also required the defendant to address the availability of less discriminatory alternatives as part of its stage 2 burden. See *Rizzo*, 564 F.2d at 149 ("the defendant must show that no alternative course of action could be adopted that would enable that interest to be served with less discriminatory impact"); *Huntington Branch, NAACP*, 844 F.2d at 939; *Salute v. Stratford Greens Garden Apartments*, 136 F.3d 293, 302 (2d Cir. 1998). However, this line of cases predates *Inclusive Communities*.

¹⁶⁰ *Inclusive Cmty. Project, Inc.*, 576 U.S. at 527-28 (quoting 24 C.F.R. § 100.500(c)(3)); see also Robert G. Schwemm & Calvin Bradford, *Proving Disparate Impact in Fair Housing Cases after Inclusive Communities*, 19 N.Y.U. J. LEGIS. & PUB. POL'Y 685, 697 (2016).

¹⁶¹ Compare *Nat'l Fair Hous. All. v. Bank of Am., N.A.*, No. CV SAG-18-1919, 2025 WL 2030225, at *20 (D. Md. July 21, 2025) (relying on an employment discrimination case for the proposition that a plaintiff must show that its proposed alternative is "equally effective" at achieving the defendant's articulated objectives' and will result in "less disparate impact") (quoting *Hardie v. Nat'l Collegiate Athletic Ass'n*, 876 F.3d 312, 320 & n.9 (9th Cir. 2017)) with Federal Housing Finance Agency, Advisory Bulletin AB 2021-04 (Dec. 20, 2021).

IV. Pretrial Investigation and Discovery in FHA Cases Arising from AI Bias

While some commentators have expressed concern that discovery limitations will make it difficult for plaintiffs to obtain the evidence needed to prove FHA claims involving AI bias,¹⁶² existing case law suggests that these concerns are overstated. In cases where FHA plaintiffs clear the hurdle of pleading facts sufficient to withstand a motion to dismiss for failure to state a claim, courts are quite willing to allow discovery regarding algorithms and artificial intelligence provided that the discovery is relevant to the case. It is clearing the pleading hurdle that poses more of a challenge. As discussed, *supra*, courts, particularly those in conservative leaning circuits, are skeptical of disparate impact claims. Thus, a prudent plaintiff will want to have some understanding of how a potential defendant is using AI before filing their case. And, even with that understanding, pleading certain disparate impact claims may be nearly impossible in circuits that treat *Inclusive Communities* as imposing heightened pleading or causation requirements.

This section first discusses some considerations for pleading a successful FHA disparate impact claim arising from biased AI decision making. It then discusses how courts have handled discovery in cases involving algorithms and artificial intelligence, concluding that the discovery procedures outlined in the Federal Rules of Civil Procedure provide the tools necessary for a plaintiff to obtain the evidence needed to prove their case.

A. Pleading a Viable FHA Disparate Impact Claim Arising from AI Decision-Making

A complaint alleging a disparate impact claim must satisfy the familiar *Twombly/Iqbal* standard, which requires sufficient factual allegations to “plausibly suggest an entitlement to relief.”¹⁶³ Although plaintiffs do not need to label the elements of their claim or otherwise plead any magic words,¹⁶⁴ the safe practice is nonetheless to consider the elements of the claim and plead facts that plausibly suggest that each element is met.¹⁶⁵

Considering the circuit split discussed in section III.C, *supra*, the facts that are likely to satisfy a plaintiff’s pleading obligation can vary significantly depending upon the circuit in which the case is filed. Plaintiffs should thus consider forum selection, to the extent possible. Particularly in class actions involving AI housing products in use throughout the country, filing the

¹⁶² See sec. IV.B., *infra*.

¹⁶³ *Ashecroft v. Iqbal*, 556 U.S. 662, 678 (2009); see also *Bell Atlantic Corp. v. Twombly*, 550 U.S. 544, 555 (2007) (“plaintiff’s obligation to provide the ‘grounds’ of his ‘entitle[ment] to relief’ requires more than labels and conclusions, and a formulaic recitation of the elements of a cause of action will not do”).

¹⁶⁴ See *Johnson v. City of Shelby*, 574 U.S. 10, 11 (2014).

¹⁶⁵ Jessica Marks & Marcus Gadson, *Text, Fairness, and Efficiency: The Case Against the Plausibility Standard for Affirmative Defenses*, 17 FED. CTS. L. REV. 1, 5-6 (2025) (noting that the lack of clarity in the *Twombly/Iqbal* standard has led to inconsistent interpretations by trial courts and that “[s]ome courts have suggested that plaintiffs must plead facts corresponding to each element of a claim”).

complaint in a circuit that follows the majority approach and allows plaintiffs to establish disparate impact via statistical evidence may be the difference between dismissal and proceeding to discovery.¹⁶⁶

1. Pleadings in circuits that follow the majority approach

In circuits that follow the majority approach, a plaintiff's complaint must identify a specific policy and allege that the policy has a statistically significant impact on those in the protected class. The specificity with which the defendant's policy must be identified in cases involving AI discrimination is unclear. Arguably, pointing to the defendant's use of AI tools should be sufficient to identify the policy at issue.¹⁶⁷ However, providing some detail about how the AI functions, such as the inputs that it uses, will strengthen the complaint.¹⁶⁸ It is here that the concerns about defendants' ability to hide behind the proprietary nature of AI tools discussed in section IV.B, *infra*, are most acute because options for pretrial investigation into how AI functions may be limited. However, carefully combing through publicly available information about the AI product at issue may allow plaintiffs to flesh out their allegations beyond a bare-bones reference to AI.

To plead a disparity, most plaintiffs will rely on statistical evidence that compares the policy's impact on protected class members to the policy's impact on those outside the protected class.¹⁶⁹ At this stage, a statistically significant disparity is enough to plead causation.¹⁷⁰ However, if the plaintiff has other facts that suggest causation, it would be wise to include them in the complaint. Notably, in one case alleging discrimination by an insurance claims processing algorithm, the plaintiffs bolstered their causation allegations with an explanation regarding how algorithms are prone to bias.¹⁷¹

¹⁶⁶ See sec. III.C, *supra*. I recognize that individual plaintiffs may not have a choice of forum.

¹⁶⁷ See, e.g., *Huskey*, 2023 WL 5848164, at *8 (allegation that insurer used “algorithmic decision-making tools to automate claims-processing” sufficient to identify challenged policy at pleading stage).

¹⁶⁸ See, e.g., *Oliver*, 2024 WL 2786905, at *7 (allegation that insurer used a “semi-automated” underwriting process that relied upon a “proprietary algorithm that consider[ed] a host of variables that can be proxies for race” sufficient to identify policy); *Louis*, 685 F. Supp. 3d at 26 (policy identified as the tenant-screening algorithm's use of credit scores). In both of these cases, the plaintiffs apparently used publicly available information about the variables considered by the defendant's algorithm to more specifically identify the policy at issue.

¹⁶⁹ Compare *Oviedo Town Ctr. II, L.L.L.P.*, 759 F. App'x at 835; *Louis*, 685 F. Supp. 3d at 26 (denying motion to dismiss disparate impact claim against landlord and tenant screening company where plaintiffs alleged that screening algorithm was more likely to reject Black and Hispanic applicants because those applicants were more likely to have low credit scores as compared to applicants outside the protected class); *Jetter*, 2021 WL 4504247, at *9 (denying motion to dismiss where plaintiff's alleged that city's tax abatement policy accelerated racial segregation because abatements went disproportionately to white neighborhoods); *Crossroads Residents Organized for Stable & Secure Residencies*, 2016 WL 3661146, at *23 (stating that disparate impact is commonly shown in cases challenging a landlord's rental policy by “alleg[ing] that tenants at the complex harmed by Defendants' actions are disproportionately protected class members”); and *Boykin*, 650 Fed. Appx. at 44 (affirming dismissal of plaintiff's disparate impact claim arising from city's closure of two homeless shelters because complaint “failed to allege facts suggesting that the closure affected a greater proportion of disabled individuals than non-disabled”); *Burbank Apartments Tenant Ass'n*, 474 Mass. 107 (affirming dismissal of disparate impact suit against landlord who switched from participating in project based subsidy program for low income tenants to accepting Section 8 vouchers because allegation that project based subsidy recipients were more likely to be members of a protected class than the general public did not establish disparate impact); *Binns*, 704 F. App'x at 797 (affirming dismissal of FHA disparate impact claim because the plaintiff “provided no comparative data, statistical or otherwise, to show that elderly and disabled participants are disproportionately impacted by the City's policy”).

¹⁷⁰ See sec. III.C.1.a, *supra*. For a detailed discussion of how to make use of statistical evidence in disparate impact cases, see Schwemm & Bradford, *supra* note 159; see also *Polek et al.*, *supra* note 32, at 89-91 (discussing mathematical methods for identifying disparate impact).

¹⁷¹ *Huskey*, 2023 WL 5848164, at *9.

2. Pleadings in circuits that follow the minority approach

In circuits that follow the minority approach, the plaintiff's complaint must plead facts that identify a policy of the defendant and demonstrate that the policy creates a substantial statistical disparity between members of the protected class and others.¹⁷² However, those allegations are not sufficient by themselves to state a claim. The plaintiff will also need to plead facts that establish that the policy—and not preexisting demographic patterns—caused the disparity.¹⁷³ It is not clear precisely what these allegations should look like as it does not appear any plaintiffs have successfully pled “robust causation” under the minority rule.

Because some courts also require that plaintiffs plead facts that demonstrate the defendant's policy is not supported by a sufficient business interest, it is good practice to plead facts that tend to show that the challenged policy is arbitrary or unnecessary.

B. Obtaining the Discovery Needed to Try an FHA Case Involving AI Discrimination

Some commentators who have written about AI discrimination have expressed concern that defendants will be able to hide behind the proprietary nature of AI algorithms, escaping liability by refusing to disclose how those algorithms render decisions.¹⁷⁴ As discussed in more detail, *infra*, this concern has been overstated. To be sure, “[d]isparate impact claims are difficult to prove—particularly because plaintiff's will often have to ‘prove that an alternative policy could serve the business purpose with a less discriminatory effect.’”¹⁷⁵ And legal scholars are correct to conclude that “the use of data-driven decision making and artificial intelligence likely makes both the potential harms and proof of discrimination by disparate treatment harder yet.”¹⁷⁶ However, these difficulties of proof do not arise from defendants' ability to shield their algorithms from discovery; they arise from the inherent difficulty in understanding and explaining how AI models render decisions.

¹⁷² See, e.g., *Trecee*, 519 F. Supp. 3d 342 (concluding that the plaintiffs had sufficiently identified a policy that created a statistically significant disparate impact, but dismissing complaint for failure to plead “robust causation”).

¹⁷³ See sec. III.C.1.b, *supra*.

¹⁷⁴ See, e.g., Emile Loza de Siles, *Artificial Intelligence Bias and Discrimination: Will We Pull the Arc of Moral Universe towards Justice?*, 2022 REVISTA FORUMUL JUDECATORILOR 40, 61 (2022) (Actions against AI developers “are deeply challenging, including, for instance, during discovery conflicts where plaintiffs' access to potential sources of AI bias and discrimination, for example training datasets, statistical models, weights assigned to data features, are stymied by trade secret arguments, despite courts' long-standing practice to protect such disclosures under seal”); Grimm, et al., *supra* note 11, at 63 (“[W]hen it comes to litigation, AI providers generally assert that information concerning the data and the algorithms are proprietary trade secrets and refuse to disclose them, thereby impeding the ability to challenge their scientific validity and reliability, and to address the many other questions they raise.”); Katyal, *Private Accountability*, *supra* note 35, at 85-86 (discussing trade secret protections as a general impediment to public inquiry into how algorithms work); Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 *Wash. L. Rev.* 1, 14 (2014) (stating that algorithmic credit scoring means that “[c]redit bureaus may be laundering discrimination into black-boxed scores, which are immune from scrutiny,” but also noting elsewhere the availability of protective orders to access credit score information in litigation); see also Andrea Roth, *Trial by Machine*, 104 *Geo. L.J.* 1245, 1275 (2016) (discussing the concern in the context of criminal cases).

¹⁷⁵ Hiller, *supra* note 7, at 922 (citing Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 *CAL. L. REV.* 671, 706-09 (2016)).

¹⁷⁶ *Id.*

Under Federal Rule of Civil Procedure 26, “[u]nless otherwise limited by court order . . . [p]arties may obtain discovery regarding any nonprivileged matter that is relevant to any party’s claim or defense and proportional to the needs of the case”¹⁷⁷ As is relevant here, in addition to limiting discovery based upon proportionality, courts may limit discovery of inaccessible electronically stored information,¹⁷⁸ enter a protective order requiring that “a trade secret or other confidential research, development, or commercial information not be revealed or be revealed only in a specified way,”¹⁷⁹ or enter an order shifting the cost of discovery to the requesting party.¹⁸⁰

Courts generally treat Rule 26 as establishing a burden shifting standard. The requesting party has the burden of explaining the relevance of the discovery sought.¹⁸¹ If they can do so, then the resisting party has the burden¹⁸² of explaining why the court should deny the requested discovery. At this stage, the resisting party can argue that the discovery is disproportionate to the needs of the case “considering the importance of the issues at stake in the action, the amount in controversy, the parties’ relative access to relevant information, the parties’ resources, the importance of the discovery in resolving the issues, and whether the burden or expense of the proposed discovery outweighs its likely benefit.”¹⁸³ The resisting party can also raise concerns about trade secrets, inaccessibility of electronically stored information, and the expense of producing the requested information.¹⁸⁴ If the resisting party has a sound basis to oppose discovery, then the court must balance the requesting party’s interest in receiving the requested information against the burden on the resisting party in producing it.¹⁸⁵

Case law shows that judges treat discovery regarding AI and algorithms no differently than any other discovery. There is no presumption either for or against production of AI or algorithm-related evidence. Judges have been quite willing to order production of such evidence, including production of the algorithm itself, where the requesting party has articulated a concrete basis for seeking AI or algorithm-related discovery and has met and conferred

¹⁷⁷ Fed. R. Civ. P. 26(b)(1).

¹⁷⁸ Fed. R. Civ. P. 26(b)(2)(B).

¹⁷⁹ Fed. R. Civ. P. 26(c)(1)(G).

¹⁸⁰ *Id.* at (c)(3) (citing Fed. R. Civ. P. 37(a)(5)).

¹⁸¹ *See, e.g., O’Gara v. Equifax Info. Servs., LLC*, No. 1:16-CV-01237-TWP-MPB, 2018 WL 513535, at *2 (S.D. Ind. Jan. 23, 2018); *Xinuos, Inc. v. IBM Corp.*, No. 22-CV-09777-CS-VR, 2024 WL 5010717, at *2 (S.D.N.Y. Nov. 13, 2024); *AmTote Int’l Inc. v. Kentucky Downs, LLC*, No. 1:15-CV-0047-GNS, 2017 WL 6329874, at *4 (W.D. Ky. Dec. 11, 2017); *Armstrong Pump, Inc. v. Hartman*, No. 10-CV-446S, 2016 WL 7208753, at *2 (W.D.N.Y. Dec. 13, 2016).

¹⁸² *See, e.g., O’Gara*, 2018 WL 513535, at *1; *Panduit Corp. v. Corning Inc.*, No. 5:18-CV-00229-FL, 2021 WL 6201203, at *1 (E.D.N.C. June 30, 2021); *In re Subpoena to Third Party Senticon, Inc.*, No. 22-MC-80281-VKD, 2022 WL 17477092, at *2 (N.D. Cal. Dec. 6, 2022); *Armstrong Pump, Inc.*, 2016 WL 7208753, at *2. *But see Shared P’ship v. Meta Platforms, Inc.*, No. 22-CV-02366-RS (RMI), 2023 WL 2526645, at *3 (N.D. Cal. Mar. 14, 2023) (placing burden on the requesting party to establish both relevance and proportionality).

¹⁸³ Fed. R. Civ. P. 26(b)(1).

¹⁸⁴ Fed. R. Civ. P. 26 (c).

¹⁸⁵ *Chen-Oster v. Goldman, Sachs & Co.*, 293 F.R.D. 557, 565 (S.D.N.Y. 2013) (“Under Rule 26(b), [the requesting party’s interest in production] must be balanced against the burden of production.”).

in good faith with the responding party.¹⁸⁶ Indeed, in the majority of cases where a court denied a motion to compel AI or algorithm-related discovery, it is because the requesting party failed to explain why their motion met the relevant standards.¹⁸⁷ Attorneys who are mindful of the relevant discovery rules and strategic in their approach to discovery should be able to obtain the information they need in most cases.

This section addresses some of the considerations that plaintiffs should consider when developing their discovery strategy. First, any discovery objections based upon trade secrets or the proprietary nature of an AI algorithm are easily addressed with protective orders; thus, plaintiffs should be proactive in proposing an appropriate protective order early in litigation. Second, plaintiffs must be able to articulate with specificity the relevance of the defendant's AI algorithm to the claims and defenses in the action. Third, plaintiffs should consider whether narrowing the scope of discovery sought or offering to pay some of the processing costs makes sense in a particular case; courts rarely use Rule 26's proportionality requirement to deny discovery regarding AI altogether, but they have shown willingness to limit the scope of requests, order cost shifting, or require further meet and confer efforts. Finally, the Federal Rules of Civil Procedure give plaintiffs multiple tools to access information on how a defendant's AI algorithms function; using interrogatories, requests for admission, depositions, and expert disclosures to gain an understanding of the defendant's AI may strengthen a plaintiff's argument for the production of the AI algorithm and documents related to its development.

1. Addressing Trade Secret Concerns

Multiple commentators have expressed concern that algorithmic decision-making could take place in an impenetrable black box, with trade secret protections thwarting discovery into potential sources of bias and discrimination.¹⁸⁸ It is certainly true that trade secret law gives artificial intelligence companies “a strong incentive . . . to avoid transparency.”¹⁸⁹ That is because artificial intelligence companies will risk losing trade secret protection for their products if their algorithms, training data, and other

¹⁸⁶ See sec. IV.B., *infra*.

¹⁸⁷ See secs. IV.B.1 and II.B.2, *infra*.

¹⁸⁸ See, e.g., Loza de Siles, *supra* note 174, at 61; Grimm, et al., *supra* note 11, at 63; Katyal, *Private Accountability*, *supra* note 35, at 85-86; Keats Citron, et al., *supra* note 174, at 14; Roth, *supra* note 174, at 1275.

¹⁸⁹ de Silva de Alwis et al, *supra* note 45, at 184; see also Peter K. Yu, *The Algorithmic Divide and Equality in the Age of Artificial Intelligence*, 72 FLA. L. REV. 331, 375 (2020) (noting “the lack of incentives for technology developers to reveal the algorithms used or to make algorithmic outcomes available for public scrutiny” and the risk to intellectual property if such disclosures were made).

details of their systems are made public.¹⁹⁰ Accordingly, artificial intelligence companies have resisted calls to make more information about their systems available in the name of transparency.¹⁹¹

But it would be a mistake to look at how AI companies have handled journalistic inquiries and extrapolate to FHA cases. The Federal Rules of Civil Procedure give litigation plaintiffs access to discovery tools that simply are not available to the public and provide mechanisms for protecting trade secret information that should obviate most trade secrets objections to producing information about AI.¹⁹² Thus, although there have not yet been many cases analyzing the discoverability of algorithmic decision-making materials, there is no reason to believe that trade secret protections will be a significant barrier to discovery in cases where such materials are relevant to the plaintiff's claims.

“[T]here is no absolute privilege [against disclosure] for trade secrets and similar confidential information.”¹⁹³ Instead, when faced with a discovery dispute involving purported trade secrets, courts weigh the requesting party's need for disclosure against the responding party's confidentiality concerns. The Tenth Circuit's decision in *In re Cooper Tire & Rubber Co.*,¹⁹⁴ provides a good example of the burden shifting test that courts apply in these situations. The party resisting discovery has the initial burden. It “must first establish that the information sought is a trade secret and then demonstrate that its disclosure might be harmful.”¹⁹⁵ If the resisting party meets its initial burden, “the burden shifts to the party seeking discovery to establish that the disclosure of trade secrets is relevant and necessary to the action.”¹⁹⁶ If both parties meet their burden, the court then balances the interests of the parties. “It is within the sound discretion of the trial court to decide whether trade secrets are relevant and whether the need outweighs the harm of disclosure.”¹⁹⁷

¹⁹⁰ de Silva de Alwis et al, *supra* note 45, at 184.

¹⁹¹ For example, several commentators have expressed concern that Northpointe, developer of COMPAS, an AI tool that is used to predict criminal recidivism, refused to produce information about its tool in response to a ProPublica investigation. *See, e.g.,* Katyal, *Private Accountability*, *supra* note 35, at 85-86; Grimm, et al., *supra* note 11, at 63-64. ProPublica describes itself as “an independent, nonprofit newsroom that produces investigative journalism with moral force.” ProPublica, About Us, <https://www.propublica.org/about/> [<https://perma.cc/F63G-Z3AA>].

¹⁹² *See* Fed. R. Civ. P. 26-37 (disclosures and discovery).

¹⁹³ *Federal Open Mkt. Comm. of Fed. Reserve Sys. v. Merrill*, 443 U.S. 340, 362 (1979) (quoting 8 Charles Alan Wright & Arthur R. Miller, *Federal Practice and Procedure* § 2043, p. 300 (1st ed. 1970)).

¹⁹⁴ 568 F.3d 1180, 1190 (10th Cir. 2009).

¹⁹⁵ *Id.* (quoting *Centurion Indus., Inc. v. Warren Steurer & Assocs.*, 665 F.2d 323, 325 (10th Cir. 1981)).

¹⁹⁶ *Id.*

¹⁹⁷ *Id.* (quoting *Centurion Indus., Inc.*, 665 F.2d at 326).

Where a requesting party has met its burden of establishing a need for access to the responding party's trade secrets, the usual solution is to order production pursuant to the terms of an appropriate protective order.¹⁹⁸ The availability of attorney's eyes only protective orders, which limit disclosure of sensitive information to specified members of a litigation team, ensures that even where parties are direct competitors, relevant trade secrets can be produced.¹⁹⁹ "As Judge Lee H. Rosenthal and Professor Steven S. Gensler have observed, [p]rotective orders have been used to shield private information in part because they are wonderfully flexible."²⁰⁰ Not only can they prevent the discovery of certain otherwise discoverable information, they can also "reduce the intrusion [if discovery is allowed] by restricting how the information is accessed, used, or disseminated."²⁰¹ Moreover, both parties and nonparties can seek protective orders.²⁰² The availability of protective orders is generally enough to rebut arguments about potential misuse of discovery because courts assume that parties will comply with their obligations under a protective order.²⁰³

Courts that have considered discovery disputes involving trade secret claims approach their analysis the same whether the requested discovery relates to algorithms or other types of sensitive information. Accordingly, in the vast majority of cases involving motions for protective orders or motions to compel production of discovery related to proprietary algorithms, courts

¹⁹⁸ See, e.g., *Gardner v. Boone*, No. 19-CV-5572, 2023 WL 1927992, at *10 (E.D. Pa. Feb. 9, 2023) (ordering production of algorithm related information pursuant to terms of protective order); *Flores v. Stanford*, No. 18-CV-02468 (VB)(JCM), 2021 WL 4441614, at *11-12 (S.D.N.Y. Sept. 28, 2021) (ordering that expert be given access to material already produced pursuant to protective order); *Sorenson v. Riffio*, No. 2:06-CV-749 TS, 2008 WL 2966785, at *6 (D. Utah July 30, 2008) (ordering production of algorithm pursuant to terms of protective order); *ADA-ES, Inc. v. Big Rivers Elec. Corp.*, No. 4:18-CV-00047-HBB, 2019 WL 1988661, at *2, *5 (W.D. Ky. May 6, 2019) (ordering production of documents that "reflect and contain...algorithms" pursuant to terms of protective order); *Clayton v. Tri City Acceptance, Inc.*, No. 3:18-CV-00308-DJH-LLK, 2019 WL 5458002, at *3 (W.D. Ky. Oct. 24, 2019) (ordering production of credit reporting algorithm pursuant to protective order); *Johnson v. Allstate Ins. Co.*, No. 07-CV-0781-MJR, 2008 WL 4279992, at *4 (S.D. Ill. Sept. 17, 2008) (ordering production of data used by coverage algorithm pursuant to terms of protective order); *Pictoso, Inc. v. Republic Servs., Inc.*, No. 4:19-CV-00397-JAR, 2023 WL 3198465, at *2 (E.D. Mo. Mar. 28, 2023) (ordering production of algorithms that set price increases subject to a protective order); *Marshall v. Georgetown Mem'l Hosp.*, No. 2:21-CV-2733-RMG, 2025 WL 411661, at *5 (D.S.C. Feb. 6, 2025) (ordering that non-party software provider produce algorithm for physical fitness assessment program pursuant to protective order). Interestingly, the District of Idaho has a standing protective order that parties may use to produce confidential information in any case in the district. See *Anne A. v. United HealthCare Ins. Co.*, No. 2:20-CV-00814, 2023 WL 197301, at *2 (D. Utah Jan. 17, 2023).

¹⁹⁹ *Structural Pres. Sys., LLC v. Andrews*, No. 12-CV-1850-MJG, 2014 WL 12738909, at *3, *6 (D. Md. Feb. 3, 2014) (resolving dispute about which documents to designate as attorneys' eyes only pursuant to terms of stipulated protective order); *BEC Cap., LLC v. Bistrovic*, 177 A.D.3d 438, 438-39 (2019) (ordering that production of source code be made pursuant to protective order's attorneys' eyes only provision); *Util. Sys. Servs., Inc. v. Careguard Warranty Servs., Inc.*, No. 16-CV-505-SDD-EWD, 2017 WL 3713472, at *8, *10 (M.D. La. Aug. 14, 2017) (resolving dispute about which documents to designate as attorneys' eyes only pursuant to terms of stipulated protective order); *Inovision Software Sols., Inc. v. Autis Ingenieros, S.L.U.*, No. 19-CV-12580, 2021 WL 1422777, at *4 (E.D. Mich. Apr. 14, 2021) (ordering production of information about algorithm pursuant to protective order with attorneys' eyes only designation because parties are direct competitors); *Modern Font Applications v. Alaska Airlines*, No. 2:19-CV-00561-DBB-CMR, 2021 WL 364189, at *4, *6 (D. Utah Feb. 3, 2021) (denying request to modify attorneys' eyes only designation that would have allowed in-house counsel to access trade secret materials of opposing party).

²⁰⁰ Hon. James C. Francis IV (Ret.), *Good Intentions Gone Awry: Privacy As Proportionality Under Rule 26(b)(1)*, 59 SAN DIEGO L. REV. 397, 409 (2022) (quoting Lee H. Rosenthal & Steven S. Gensler, *The Privacy-Protection Hook in the Federal Rules*, 105 JUDICATURE 77, 78 (2021)).

²⁰¹ *Id.*

²⁰² *Id.*

²⁰³ See *Flores*, 2021 WL 4441614.

have ordered production pursuant to the terms of a protective order.²⁰⁴ Recognizing this reality, holders of trade secrets often stipulate to a protective order that meets their business needs, rather than wasting time and money resisting discovery altogether.²⁰⁵

The most significant hurdle that the requesting party must clear to obtain discovery of trade secret materials is establishing relevance. This requires clearly articulating why those materials are necessary to the case. This is a significant task in the context of obtaining discovery regarding artificial intelligence. Because AI is relatively novel, lawyers cannot assume that judges will be familiar with what AI is, how it works, or its potential to render biased decisions. Lawyers must be prepared to clearly articulate how particular trade secret information relates to the claims or defenses in the case. General averments that specific evidence would be “nice to have” are unlikely to tip the balance in favor of production. Indeed, in most cases where courts have blocked production of algorithmic information, the decisions rested on the moving party’s failure to explain the relevance of the requested discovery.²⁰⁶

Some courts have also demonstrated reluctance to order production of source code — particularly where a responding party has produced information sufficient to show how an algorithm works.²⁰⁷ Arguably, the reluctance of those courts is misguided. There is no “source code exception” to the Federal Rules of Civil Procedure.²⁰⁸ “As summarized by one scholar in the context of discovery in both patent infringement and trade secret litigation, “while source code production can be maddening, time consuming, and costly, it is by now relatively common in software cases.”²⁰⁹ Because source code production is now common, litigators “have ready-made tools at their disposal to address the merit of software-related disputes while ensuring that the source code remains protected and yet disclosed in a litigation dispute.”²¹⁰ A plaintiff who seeks access to the AI system itself should thus be entitled to that access based upon the same showing of necessity needed for other trade secret information. With that in mind,

²⁰⁴ See, e.g., *id.*; *Gardner*, 2023 WL 1927992, at *10; *Sorenson*, 2008 WL 2966785, at *6; *ADA-ES, Inc.*, 2019 WL 1988661, at *2, *5; *Clayton*, 2019 WL 5458002, at *3; *Allstate Ins. Co.*, 2008 WL 4279992, at *4; *Pietoso, Inc.*, 2023 WL 3198465, at *2; *Marshall*, 2025 WL 411661, at *5; *Anne A.*, 2023 WL 197301; *Structural Pres. Sys., LLC*, 2014 WL 12738909, at *3; *BEC Cap., LLC*, 177 A.D.3d at 438–39; *Util. Sys. Servs., Inc.*, 2017 WL 3713472, at *8, *10; *Inovision Software Sols., Inc.*, 2021 WL 1422777, at *4; *Modern Font Applications*, 2021 WL 364189, at *4, *6.

²⁰⁵ See, e.g., *Structural Pres. Sys.*, 2014 WL 12738909, at *3 (involving a stipulated protective order).

²⁰⁶ See, e.g., *Viacom Int'l Inc. v. YouTube Inc.*, 253 F.R.D. 256, 260-61 (S.D.N.Y. 2008) (granting motion for protective order preventing production of source code for YouTube algorithm where requesting party had not made sufficient showing of need); *DIRECTV, Inc. v. Trone*, 209 F.R.D. 455, 460 (C.D. Cal. 2002) (declining to order production of algorithm and other trade secrets where requesting party did not make sufficient showing of need); *Congoo, LLC v. Revcontent LLC*, No. 16-CV-401 (MAS), 2017 WL 3584205, at *1, *4 (D.N.J. Aug. 10, 2017) (declining to order production of advertising algorithm where harm from disclosure of confidential material outweighed requesting party’s need for same).

²⁰⁷ See, e.g., *Congoo, LLC*, 2017 WL 3584205, at *1 (declining to order production of advertising algorithm where resisting party was willing to produce information sufficient to show how algorithm worked).

²⁰⁸ *Thales Visionix, Inc. v. United States*, 149 Fed. Cl. 38, 48 (2020).

²⁰⁹ *Id.* (quoting Sonia K. Katyal, *The Paradox of Source Code Secrecy*, 104 CORNELL L. REV. 1183, 1275 (2019)).

²¹⁰ *Id.* (quoting Katyal, *supra* note 210 at 1275-76).

a plaintiff who seeks production of source code should be prepared to explain why the production of documents and information about how an AI algorithm works (such as information on inputs, training, and decision making) is not a good substitute for the source code itself in a particular case.

2. *Relevance and Proportionality Concerns*

Defendants in disparate impact cases are unlikely to avoid discovery regarding the workings of their AI programs based upon trade secret protections. But that does not mean that plaintiffs will automatically be entitled to this information. The ability to clearly and specifically articulate how discovery about AI is relevant to the claims and defenses in an action is essential both to meet the plaintiff's burden of establishing relevance and to rebut arguments that algorithm discovery is disproportionate to the needs of the case.

Under Federal Rule of Civil Procedure 26(b)(1), nonprivileged information is discoverable only if it is “relevant to any party’s claim or defense.”²¹¹ Despite amendments designed to reduce the expense of discovery,²¹² courts continue to interpret Rule 26(b) liberally. Generally, courts will treat discovery requests as relevant “if there is *any* possibility that the information sought may be relevant to the claim or defense of any party.”²¹³ “Information sought only fails the relevance test if it is clear that it could have ‘no possible bearing on the claim.’”²¹⁴ Although relevance is interpreted broadly, it is not unlimited. To give the relevance requirement teeth, most courts place the burden on the requesting party to make a showing of relevance with “precision or particularity.”²¹⁵

Additionally, as previously discussed, the scope of discovery is limited to that which is “proportional to the needs of the case, considering the importance of the issues at stake in the action, the amount in controversy, the parties’ relative access to relevant information, the parties’ resources, the importance of the discovery in resolving the issues, and whether the burden

211 Fed. R. Civ. P. 26(b)(1).

212 Notes of Advisory Committee on Rules - 2015, Fed. R. Civ. P. 26.

213 *Folkner v. Ford Motor Credit Co.*, LLC, Nos. 18-50214-RLJ-11, 20-05005, 2021 Bankr. LEXIS 3004, at *3 (Bankr. N.D. Tex. Oct. 29, 2021)(emphasis in original) (quoting *Ries v. Ardinger*, 555 B.R. 579, 589 (Bankr. N.D. Tex. 2016)); see also *Dotson v. Edmonson*, No. 16-15371, 2017 U.S. Dist. LEXIS 229685, at *1 (E.D. La. Nov. 17, 2017)[hereinafter *Edmonson*]; *Cottonwood Acres, LLC v. First Am. Title Ins. Co.*, No. 2:22-cv-00680-TC-DBP, 2024 U.S. Dist. LEXIS 104835, at *3 (D. Utah June 12, 2024); *Wellemeier v. Trans Union, LLC*, No. 3:20-CV-00814-DJH-LLK, 2022 U.S. Dist. LEXIS 92889, at *8 (W.D. Ky. May 23, 2022); *VeroBlue Farms USA, Inc. v. Cassels Brock & Blackwell LLP*, Nos. 18-01297, 19-09015, 2021 Bankr. LEXIS 1083, at *10 (Bankr. N.D. Iowa Apr. 22, 2021); *Hagbajegbi v. Guess?, Inc.*, 168 F. Supp. 3d 1277, 1280 (S.D. Cal. 2016); *Jonathan R. v. Jim Justice*, No. 3:19-CV-00710, 2024 U.S. Dist. LEXIS 24186, at *5 (S.D. W. Va. Feb. 12, 2024); *South Capitol Bridgebuilders v. Lexington Ins. Co.*, No. 1:21-CV-1436-RCL, 2023 U.S. Dist. LEXIS 196709, at *6 (D.D.C. Mar. 24, 2023); *Dass v. City Univ. of N.Y.*, No. 18-CV-11325 (VSB)(OTW), 2022 U.S. Dist. LEXIS 202699, at *6 (S.D.N.Y. Nov. 7, 2022); *Hicks v. Deepwater Glob. Distrib.*, No. 6:17-CV-1472-Orl-41TBS, 2018 U.S. Dist. LEXIS 117635, at *4 (M.D. Fla. July 16, 2018); *Kegerise v. Susquehanna Twp. Sch. Dist.*, No. 1:14-CV-0747, 2016 U.S. Dist. LEXIS 62124, at *3 (M.D. Pa. May 11, 2016); *Gamon Plus, Inc. v. Campbell Soup Co.*, No. 1:15-CV-08940, 2022 U.S. Dist. LEXIS 236014, at *3 (N.D. Ill. May 26, 2022).

214 *Folkner*, 2021 Bankr. LEXIS 3004, at *3 (quoting *Ries*, 555 B.R. at 589).

215 *Johnson v. Sherman*, No. 1:21-CV-01313-DAD-EPG-HC, 2022 WL 3567257, at *7 (E.D. Cal. Aug. 18, 2022), report and recommendation adopted, No. 1:21-CV-01313-ADA-EPG-HC, 2022 WL 16701964 (E.D. Cal. Nov. 3, 2022) (denying motion to compel production of source code where requesting party did not provide a “declaration or showing, with any precision or particularity, how a review of the ... source code” would help the requesting party examine the opposing party’s expert).

or expense of the proposed discovery outweighs its likely benefit.”²¹⁶ This language is the product of a 2015 amendment to Rule 26 that was intended to address long-standing concerns about discovery abuse.²¹⁷ Although the concept of proportionality has been part of Rule 26 since 1983,²¹⁸ the Advisory Committee on Rules hoped that incorporating proportionality into the definition of discoverable information would stress “the obligation of the parties to consider [proportionality] factors in making discovery requests, responses, or objections.”²¹⁹ The Advisory Committee on Rules’ goals were to discourage “overdiscovery” and encourage judges to be more proactive about limiting burdensome discovery.²²⁰

Both supporters and opponents of the 2015 amendment to Rule 26 viewed the amendment as “limiting plaintiffs’ access to discovery, and ‘predicted that the impact would be profound and immediate.’”²²¹ Passage of the amendment was seen as a coup for the defense bar that created “a new blanket objection to all discovery requests and impose[d] a new burden on the requesting party to justify each and every request.”²²² Opponents of the amendment were particularly concerned about the subjectivity of the proportionality factors. Without measurable standards, the availability of discovery in a case might be subject to the whims of the assigned judge, making it difficult for parties to predict what would be “discoverable” in any given case.²²³

Given these concerns, it seems prudent to ask whether the predictions about the 2015 amendment to Rule 26 have borne out. Is proportionality a significant barrier to discovery? And has the subjectivity of the proportionality factors in the Rule left parties without guidance as to the scope of discovery in a given case? Fortunately, it seems that the answer to these questions is no. First, there is no strong evidence to suggest that the 2015 amendment led to a general curtailment of discovery. Second, cases

216 Fed. R. Civ. P. 26(b)(1).

217 Notes of Advisory Committee on Rules - 2015, Fed. R. Civ. P. 26; see also Allyson Haynes Stuart, *The Privacy Paradox in Discovery*, 26 VAND. J. ENT. & TECH. L. 615, 633; Sean Farhang, *Supreme Court Oversight of the Federal Rules: A Principal-Agent Problem?*, 72 DEPAUL L. REV. 363, 371 (2022) (citing Robert H. Klonoff, *Application of the New “Proportionality” Discovery Rule in Class Actions: Much Ado About Nothing*, 71 VAND. L. REV. 1949, 1949 (2018)); Andrew M. Pardieck, *The Shifting Sands of Cost Shifting*, 69 CLEV. ST. L. REV. 349, 353 (2021). The Advisory Committee had made several previous efforts to use the concept of proportionality to limit discovery. The 1983 amendments introduced the concept of proportionality in Rule 26. Stuart, *supra*, at 629-30; Stephen L. Rispoli et al., *When to Leave The Stones Unturned: Using Proportionality to Navigate Discovery Efficiently, Effectively, and Ethically*, 107 MARQ. L. REV. 487, 496 (2023). Francis, *supra* note 201, at 410. Because of concerns that judges were not actually considering proportionality when issuing discovery rulings, the 1993 and 2000 amendments included further tweaks designed to encourage more consideration of whether discovery was proportional to the needs of a case. Rispoli et al., *supra* note 218, at 497; Francis, *supra* note 201, at 412.

218 Stuart, *supra* note 218, at 629-30; Rispoli et al., *supra* note 218, at 496. Hon. James C. Francis IV (Ret.), *Good Intentions Gone Awry: Privacy As Proportionality Under Rule 26(b)(1)*, 59 SAN DIEGO L. REV. 397, 410 (2022).

219 Fed. R. Civ. P. 26(b)(1) advisory committee’s note to 2015 amendment, at 19.

220 *Id.* at 17, 18, 21.

221 Farhang, *supra* note 218, at 371 (citing Klonoff, *supra* note 218, at 1949).

222 Pardieck, *supra* note 218, at 358.

223 See Seth Katsuya Endo, *Discovery Dark Matter*, 101 TEX. L. REV. 1021, 1051-52 (discussing the concerns of opponents to the 2015 amendments). *But see* Farhang, *supra* note 218, at 367 (explaining how the subjectivity of Rule 26 also makes it difficult for conservatives to use the 2015 amendment to advance their goal of curtailing discovery).

analyzing the discoverability of information about AI and algorithms suggest that judges rarely use proportionality as a basis to deny discovery entirely. These cases also offer a fairly consistent set of principles that plaintiffs can use to argue for the discoverability of AI and algorithm information.

Although it has been nearly ten years since proportionality became part of the scope of discovery, there is not much empirical data about the amendment's impact.²²⁴ Moreover, the articles that have attempted to assess the amendment's impact have reached inconsistent conclusions.

In 2017, Professor Steven Baicker-McKee, reviewed every federal discovery decision over the prior two years in an effort to determine whether—after the 2015 amendment—courts were more frequently using proportionality to limit discovery.²²⁵ He concluded that courts were, finding that courts cited proportionality as a basis to restrict discovery roughly twice as often after the 2015 amendment as compared to the pre-amendment comparison period.²²⁶ However, Professor Baicker-McKee acknowledged that it was not possible to determine whether the increased reliance on proportionality in judicial decisions actually translated to less discovery. For example, it is possible that a judge who might have denied discovery as “unduly burdensome” pre-2015 now denies similar discovery as “not proportional,” thus increasing the number of decisions that use the term “proportional” without actually decreasing the overall volume of discovery allowed.²²⁷

A year later, Professor Robert H. Klonoff reached the opposite conclusion about the impact of the 2015 amendment.²²⁸ Professor Klonoff examined the discovery disputes decided in class action cases following the 2015 amendment and concluded that the amendments had not changed the way that judges approached discovery disputes. As he explained, “First, numerous courts emphasize that the basic overarching principles existed prior to the amendment. Second, a number of courts continue to rely heavily on pre-amendment case law in resolving discovery disputes. Third, contrary to the fears of many plaintiffs’ attorneys and academics, the amendment has not shifted the burden of proof to the party seeking discovery to prove proportionality.”²²⁹

224 See Edith Beardsen, *Discovery Culture*, 57 GA. L. REV. 981, 987 (2023) (citing Kevin E. Davis & Helen Hershkoff, *Contracting for Procedure*, 53 WM. & MARY L. REV. 507, 516-17 (2011)).

225 Steven Baicker-McKee, *Mountain or Molehill?*, 55 DUQ. L. REV. 307, 309 (2017).

226 *Id.* at 313. Specifically, Professor Baicker-McKee found that “Three hundred thirty-five cases have applied the new proportionality provision in the first year of amended Rule 26(b)(1). Of those cases, in 192 (57%) the court restricted discovery in whole or in part based on proportionality. By comparison, courts applied proportionality 79 times and restricted discovery in 46 cases (58%) during the pre-amendment comparison period. These numbers suggest that parties and courts are applying proportionality more than four times more frequently than before the amendments, and that courts are narrowing discovery on proportionality grounds more than four times more frequently post-amendment.” *Id.*

227 *Id.* at 314.

228 Klonoff, *supra* note 218.

229 *Id.* at 1972.

The contrasting outcomes in these two articles is—to a certain extent—the product of the authors’ differing methodologies, but it also highlights that any change in discovery outcomes following the 2015 amendment is not so stark as to preclude differing interpretations. In other words, it is not clear that the 2015 amendment has had any measurable impact on discovery writ large.²³⁰ A prudent plaintiff’s attorney should certainly consider the proportionality of discovery they are seeking, but proportionality concerns should not deter that attorney from seeking discovery that is important to their case.

When it comes to cases that specifically address the discovery of information about algorithms and AI, the trend is similarly encouraging with the cases discussing the discoverability of information regarding AI and algorithms applying consistent analysis and often allowing the requested discovery.²³¹ To determine whether courts are inclined to limit discovery regarding AI based upon Rule 26’s proportionality requirement, I reviewed cases decided after the 2015 amendment where parties sought production of information concerning AI or algorithms.²³²

This research demonstrated that proportionality has rarely been used to deny discovery regarding AI or algorithms outright. The research results included just three cases where a court denied a motion to compel discovery on proportionality grounds. The first involved a request to produce documents going back twenty years.²³³ In the second, the court denied a motion to compel without prejudice to give the requesting party time to explore whether there was a less burdensome means of obtaining the information it needed.²³⁴ In the third, the court denied a motion to compel production of Experian’s credit matching algorithm where Experian had already produced documents sufficient to establish how it had incorrectly matched adverse credit information to the plaintiff.²³⁵

²³⁰ Reaching any conclusion on this question is further complicated by the fact that there does not appear to be more recent scholarship on this question. See Miyoko T. Pettit-Toledo, *The Politics of Proportionality in State Civil Rulemaking*, 101 DENV. L. REV. 641, 663 (2024) (noting the lack of newer scholarship on the impact of the 2015 amendments); David Freeman Engstrom & Jonah B. Gelbach, *Legal Tech, Civil Procedure, and the Future of Adversarialism*, 169 U. PA. L. REV. 1001, 1048–49 (2021) (“Despite their centrality in legal and policy debates, discovery costs have drawn little careful rigorous empirical inquiry, particularly recently. Indeed, much of the best empirical evidence dates back to the 1970s and 1980s...”).

²³¹ See notes 233–242, *infra*.

²³² Specifically, I searched Westlaw for all federal cases starting from January 1, 2015, that included the terms “artificial intelligence” or “algorithm” and also included the terms “proportional to the needs of the case” and “26.” I reviewed all the cases in the search result list where either party sought discovery about AI or an algorithm or sought production of the AI tool or algorithm itself. This significantly narrowed the pool of search results because many cases discussing AI and discovery deal with technology assisted review, the use of AI to review documents. To ensure that my focus on proportionality did not screen out any particularly relevant cases on AI, I ran a broader search for cases that discussed “artificial intelligence” and “motions to compel,” “motions to quash,” or “motions for protective orders.” This research was last updated on September 15, 2025.

²³³ *Prudhomme v. Geico Ins. Co.*, No. 15-CV-00098, 2019 WL 5096085, at *7 (W.D. La. Oct. 10, 2019).

²³⁴ *Xinuos, Inc.*, 2024 WL 5010717, at *2–3.

²³⁵ *Dunmore v. Experian Infa. Sols. Inc.*, No. 2:24-CV-225-PPS-JEM, 2025 WL 2331036, at *2 (N.D. Ind. Aug. 13, 2025).

Narrowing discovery requests due to proportionality concerns is somewhat more common; in nine cases judges granted-in-part the requesting party's motion to compel discovery regarding AI or algorithms.²³⁶ The most common reasons for narrowing a request were where the moving party had requested "all" documents on a topic and where the parties failed to fully meet and confer before the requesting party filed its motion.²³⁷

Establishing relevance proved to be the bigger hurdle. Of the eighteen cases in which courts outright denied motions to compel discovery regarding AI or algorithms,²³⁸ eleven of the opinions rested on the moving party's failure to establish the relevance of the requested information.²³⁹ These opinions reveal a few common themes. First, the moving party should explain the relevance of the information sought as specifically as it can, given the stage of the proceedings. Using boilerplate language or cutting and pasting from motions in other cases will almost certainly result in a court denying the motion to compel. Second, before requesting discovery regarding AI or algorithms, it is helpful to have a basic understanding of what the AI or algorithm does and does not do. If the decision or practice at issue in the case is one where humans have the final say, then courts are less likely to allow discovery regarding AI or algorithms. Third, relevance is limited to claims and defenses that are actually pled; courts are unlikely to

- 236 *Vance v. Int'l Bus. Machines Corp.*, No. 20-CV-577, 2021 WL 12305778 (N.D. Ill. Oct. 26, 2021), objections sustained in part and overruled in part, No. 20-CV-577, 2021 WL 12305777 (N.D. Ill. Nov. 23, 2021); *O'Gara*, 2018 WL 513535; *CoStar Grp., Inc. v. Restb.AI, LLC.*, No. 22-MC-00228-CBM(ASX), 2023 WL 3325201, at *5 (C.D. Cal. Mar. 10, 2023); *Monet v. Tesla, Inc.*, No. 2:24-CV-00107-JPH-MKK, 2025 WL 1298224, at *4 (S.D. Ind. Apr. 16, 2025); *Rumble, Inc. v. Google LLC*, No. 21-CV-00229-HSG(LJC), 2023 WL 3751797, at *8 (N.D. Cal. May 31, 2023); *Muff ex. rel. Huff v. Wells Fargo Bank N.A.*, No. 21-CV-4003-CJW, 2022 WL 601051, at *5 (N.D. Iowa Jan. 28, 2022); *Thales Visionix, Inc.*, 149 Fed. Cl. 38 at 41; *Cleary v. CoreLogic Rental Prop. Sols., LLC*, No. 2:17-CV-14158-BAF-DRG, 2019 WL 13193362 (E.D. Mich. Jan. 4, 2019); *Armstrong Pump, Inc.*, 2016 WL 7208753, at *2. In four additional cases, the court came to other types of compromise resolutions. *Frost v. BNSF Ry. Co.*, 218 F. Supp. 3d 1122, 1135 (D. Mont. 2016) (ordering party who resisted production of documents to instead produce witness for deposition); *Concord Music Grp., Inc. v. X Corp.*, No. 3:23-CV-00606, 2025 WL 209184, at *8 (M.D. Tenn. Jan. 15, 2025) (ordering further meet and confer efforts); *In re Soc. Media Adolescent Addiction/Pers. Inj. Prods. Liab. Litig.*, No. 22-MD-03047-YGR (PHK), 2024 WL 3488659, at *4–6 (N.D. Cal. June 20, 2024) (same); *Craigville Tel. Co. v. T-Mobile USA, Inc.*, No. 19-CV-7190, 2022 WL 1499908, at *5 (N.D. Ill. May 12, 2022) (same).
- 237 *Craigville Tel. Co.*, 2022 WL 1499908, at *5 (ordering parties to meet and confer to narrow scope of requests seeking "all documents that refer or relate to" certain issues with defendant's algorithm); *CoStar Grp., Inc.*, 2023 WL 3325201, at *5 (narrowing requests for production of "all" documents regarding AI that looks for and crops out watermarks in copyrighted images); *Cleary*, 2019 WL 13193362 (limiting request for "related to" matching algorithms to information that establishes or directly addresses the procedures).
- 238 *Edmar Fin. Co., LLC v. Currenex, Inc.*, 347 F.R.D. 641, 643–44 (S.D.N.Y. 2024) (denying plaintiff's motion to compel and granting defendant's in same opinion); *Shared P'ship*, 2023 WL 2526645, at *2; *In re Domestic Airline Travel Antitrust Litig.*, No. 15-MC-1404, 2020 WL 5847779, at *9, *10 (D.D.C. Feb. 10, 2020), report and recommendation adopted, No. 15—MC-1404 (CKK), 2020 WL 5834892 (D.D.C. Feb. 11, 2020); *New York Times Co. v. Microsoft Corp.*, 757 F. Supp. 3d 594, 596 (S.D.N.Y. 2024); *Bright v. Brookdale Senior Living, Inc.*, No. 3:19-CV-0374, 2022 WL 18639035, at *7 (M.D. Tenn. Feb. 14, 2022), review granted, decision aff'd, No. 3:19-CV-00374, 2022 WL 18639048 (M.D. Tenn. Aug. 25, 2022); *Runton by & through Adult Advoc. & Representation v. Brookdale Senior Living, Inc.*, No. 17-CV-60664, 2018 WL 1083493, at *5 (S.D. Fla. Feb. 27, 2018); *Conn. Fair Hous. Cr. v. CoreLogic Rental Prop. Sols., LLC*, No. 3:18-CV-705 (VLB), 2020 WL 401776, at *5 (D. Conn. Jan. 24, 2020); *Xinuos, Inc.*, 2024 WL 5010717, at *2–3; *Royal Mile Co., Inc. v. UPMC & Highmark Inc.*, No. 2:10-CV-01609-JFC, 2016 WL 6915978, at *17 (W.D. Pa. June 24, 2016); *Los Amigos Mkt., Inc. v. United States*, No. 21-CV-3746, 2022 WL 2899271, at *1 (E.D. Pa. May 25, 2022); *Dotson v. Experian Info Sols., Inc.*, No. 17_CV-575-D, 2019 WL 440588, at *6 (W.D. Okla. Feb. 4, 2019); *Landmark Infrastructure Holding Co., LLC v. R.E.D. Invs., LLC*, No. 15-CV-4064-C-MJW, 2016 WL 11339515, at *3 (W.D. Mo. Nov. 30, 2016); *Doc v. Fullstory, Inc.*, No. 23-CV-00059-WHO (VKD), 2025 WL 474907, at *3 (N.D. Cal. Feb. 12, 2025); *Freeman Inv. Mgmt. Co., LLC v. Frank Russell Co.*, No. 13-CV-2856-JLS(RBB), 2016 WL 11620828, at *5–8 (S.D. Cal. Feb. 29, 2016); *Prudhomme*, 2019 WL 5096085, at *7; *Fed. Trade Comm'n v. Marcus*, No. 0:17-CV-60907, 2020 WL 1482250, at *1 (S.D. Fla. Mar. 27, 2020); *Trio v. Amazon Web Servs., Inc.*, No. 23-CV-1389, 2025 U.S. Dist. LEXIS 45330, at *3 (N.D. Ill. Feb. 28, 2025).
- 239 *Edmar Fin. Co.*, 347 F.R.D. at 647–49; *In re Domestic Airline Travel Antitrust Litig.*, 2020 WL 5847779, at *9; *New York Times Co.*, 757 F. Supp. 3d at 596; *Runton*, 2018 WL 1083493, at *5; *Conn. Fair Hous. Cr.*, 2020 WL 401776, at *5; *Royal Mile Co.*, 2016 WL 6915978, at *8; *Los Amigos Mkt.*, 2022 WL 2899271, at *1; *Dotson*, 2019 WL 440588, at *6; *Landmark Infrastructure Holding Co., LLC*, 2016 WL 11339515, at *3; *Fullstory, Inc.*, 2025 WL 474907, at *3; *Trio*, 2025 U.S. Dist. LEXIS 45330, at *4.

compel discovery regarding alternate theories that are outside the scope of the pleadings. Finally, in bifurcated cases, premature attempts at merits discovery are unlikely to succeed.

In summary, attorneys who are strategic in framing discovery requests, attentive to the meet and confer requirements in their jurisdiction, and thorough in explaining the relevance of the information sought are likely to obtain the AI and algorithm discovery they need to litigate their case. Indeed, in the majority of cases discussing the discoverability of information regarding AI or algorithms, the court granted discovery either in whole or in part.²⁴⁰ And, in the majority of cases where courts granted this type of discovery, the motion was granted in full.²⁴¹

Flores v. Stanford,²⁴² a class action arising from the New York State Board of Parole's use of an AI decision-making tool, provides a good example of the degree of precision that courts expect from a moving party's relevance explanation. In *Flores*, the plaintiffs' expert sought access to third-party discovery concerning COMPAS, an AI-powered tool that purports to predict the risk that parole candidates will reoffend.²⁴³ The discovery requested "(1) the normative dataset used to create and normalize COMPAS []; and (2) the regression models for two COMPAS 'scales': (a) the General Recidivism Risk

²⁴⁰ Research results included 35 cases where the court granted motions to compel in whole or in part, see notes 235, *supra*, and 242, *infra*, one where further meet and confer was ordered (without granting any part of the motion), and 18 cases where the court denied motions to compel. See cases cited in note 239, *supra*.

²⁴¹ In 21 cases, the court granted the motion to compel in full. *RVassets Ltd. v. Marex Cap. Markets Inc.*, No. 1:23-CV-14192, 2025 WL 1928056, at *4 (N.D. Ill. July 14, 2025) (ordering production of discovery regarding algorithm in trade secrets case); *Renaissance Specialty Ins., LLC v. Landrum*, No. 5:24-CV-00392-WGY-PRL, 2025 WL 2106691, at *3 (M.D. Fla. July 28, 2025) (granting motion to compel production of pricing algorithm, among other things); *Flores*, 2021 WL 4441614, at *1; *Edmar Fin. Co., LLC*, 347 F.R.D. 641, 643-44 (S.D.N.Y. 2024) (granting defendant's motion to compel production of algorithm source code in same opinion denying plaintiff's motion); *Fed. Trade Comm'n v. DIRECTV, Inc.*, No. 15-CV-01129-HSG-MEJ, 2016 WL 4154851, at *5 (N.D. Cal. Aug. 5, 2016) (ordering production of text mining algorithm); *ADA-ES, Inc.*, 2019 WL 1988661, at *5 (denying motion for protective order precluding discovery of algorithm); *Panduit Corp.*, 2021 WL 6201203, at *1 (granting motion to compel production of documents describing algorithm-driven fiber selection process); *In re Facebook, Inc. Consumer Priv. User Profile Litig.*, No. 18-MD-2843 VC (JSC), 2021 WL 10282215, at *3 (N.D. Cal. Sept. 29, 2021) (ordering production of Mark Zuckerberg's notebooks discussing ideas for how Facebook's algorithms should work); *Vesta Corp. v. Amdocs Mgmt. Ltd.*, No. 3:14-CV-01142-HZ, 2017 WL 714354, at *1 (D. Or. Feb. 21, 2017) (ordering production of source code); *IQVIA, Inc. v. Veeva Sys., Inc.*, No. 2:17-CV-00177-CCC-MF, 2018 WL 11693781, at *8 (D.N.J. Nov. 30, 2018), modified, No. 17-CV-00177-JXN-JSA, 2022 WL 17990836 (D.N.J. Dec. 29, 2022) (ordering production of source code); *Thuman v. Dembski*, No. 13-CV-01087A(F), 2022 WL 1197551, at *5 (W.D.N.Y. Apr. 22, 2022) (ordering production of documents regarding development of hedge fund trading algorithm); *BlackBerry Ltd. v. Facebook, Inc.*, No. 18-CV-1844-GW (KSX), 2019 WL 13447227, at *3, *7 (C.D. Cal. Sept. 30, 2019) (granting motion to compel production of documents about algorithm alleged to infringe patent); *Williams v. Corelogic Rental Prop. Sols., LLC*, No. CV PX 16-58, 2016 WL 6277675, at *5, *7 (D. Md. Oct. 26, 2016) (ordering production of non-privileged documents about tenant screening algorithm); *MedImpact Healthcare Sys., Inc. v. IQVIA Inc.*, No. 19-CV-1865-GPC (DEB), 2021 WL 5605209, at *3 (S.D. Cal. June 29, 2021) (ordering production of source code); *In re Subpoena to Third Party Senticon, Inc.*, 2022 WL 17477092, at *2 (ordering production of source code); *Ingham Reg'l Med. Ctr. v. United States*, 169 Fed. Cl. 12, 26-27 (2024) (reopening discovery to allow requesting party to seek information about algorithm driven reimbursement system); *Nicholes v. Combined Ins. Co. of Am.*, No. 5:16-CV-10203, 2018 WL 1098246, at *9, *10 (S.D. W. Va. Feb. 28, 2018) (granting motion to compel production of documents related to algorithms for computer systems that price, sell, retain, cancel, and handle claims for insurance policies); *AmTote Int'l Inc.*, 2017 WL 6329874, at *4 (ordering production of information regarding development and operation of a horse racing odds program); *Stagger v. Experian Info. Sols., Inc.*, No. 21-CV-2001, 2021 WL 5299791, at *2-3 (N.D. Ill. Nov. 15, 2021) (ordering production of documents related to policy for scrubbing debt records, which appears to involve algorithm); *In re Clearview AI, Inc., Consumer Priv. Litig.*, No. 21-CV-135, 2022 WL 22903326, at *1-4 (N.D. Ill. Aug. 18, 2022) (ordering production of documents related to defendant's training of its AI); *Jenkins v. Regal Cinemas*, No. 20-CV-3782, 2025 WL 81341, at *5-7 (N.D. Ill. Jan. 13, 2025) (granting motion to compel production of information regarding defendant's storage of employee fingerprints, which are converted into algorithmic information before storage); *Software Automation Holdings, LLC v. Ins. Toolkits, LLC*, No. 5:23-CV-140-D, 2025 U.S. Dist. LEXIS 110938, at *7-8 (E.D.N.C. June 11, 2025) (compelling inspection of source code).

²⁴² *Flores*, 2021 WL 4441614.

²⁴³ *Id.* at *1.

Scale, and (b) the Violent Recidivism Risk Scale”²⁴⁴ In their opposition to third-party Northpointe’s motion for protective order, the plaintiffs clearly demonstrated the relevance of the requests at both the class certification and merits stages of the discovery.²⁴⁵ Specifically, the plaintiffs understood from prior discovery that the COMPAS algorithm used age as a datapoint, and argued that their expert needed access to the discovery materials so the plaintiffs could determine whether the algorithm improperly weighted this factor.²⁴⁶ Further, they argued that the requested information was relevant to whether the alleged illegal practices were common to all class members.²⁴⁷

Certainly, it is true that the discretion afforded to trial judges can lead to different outcomes depending upon the judge to whom a case is assigned.²⁴⁸ Nonetheless, it appears that the Advisory Committee on Rules was largely successful in crafting rules for the production of electronically stored information that are “broad enough to cover all current types of computer-based information, and flexible enough to encompass future changes and developments.”²⁴⁹ Faced with the new technology of algorithms and AI, courts have successfully analyzed discovery requests using the same standards that they apply to other types of discovery requests.²⁵⁰

3. Cost Shifting and Ways to Reduce Discovery Costs

While proportionality has not yet proven to be a major obstacle to AI and algorithm discovery, there are options to consider if the party resisting discovery substantiates a proportionality objection. Where discovery “is not reasonably accessible because of undue burden or cost” the court “may specify conditions for the discovery” that make the production less burdensome.²⁵¹ Cost shifting—placing some or all of the cost of reviewing and producing the requested information on the requesting party—is one specifically enumerated option.²⁵²

²⁴⁴ *Id.*

²⁴⁵ *Id.* at *4-5.

²⁴⁶ *Id.* at *5-6.

²⁴⁷ *Id.* at *7.

²⁴⁸ Farhang, *supra* note 219, at 387.

²⁴⁹ Fed. R. Civ. P. 34 advisory committee’s notes to 2006 amendment, at 28.

²⁵⁰ See the preceding discussion in this section. Because this article is about AI, I thought it would be worthwhile to see how AI analyzed the questions raised in this section. In June 2025, I asked Lexis+AI: “Are courts granting motions to compel production of information related to artificial intelligence, including artificial intelligence algorithms, training, and functionality?” While Lexis cited to some cases that were off topic, I largely agree with its ultimate conclusion: “Courts have addressed motions to compel production of information related to artificial intelligence (AI), including AI algorithms, training, and functionality, with varying outcomes depending on the relevance, proportionality, and burden of the requested discovery . . . [case citations omitted]. These cases illustrate that courts carefully evaluate motions to compel AI-related information, balancing the requesting party’s need for the information against the burden and proportionality considerations under the Federal Rules of Civil Procedure.”

²⁵¹ Fed. R. Civ. P. 26(b)(2)(B).

²⁵² Fed. R. Civ. P. 26(c)(1)(B).

In a case where the expense of producing algorithm and AI related information is high, for example due to the volume of information to be reviewed or because of the format in which information is stored, plaintiffs should consider whether cost shifting makes sense. However, Rule 26's list of options for mitigating the burden of discovery is non-exhaustive.²⁵³ Other options—either in lieu of or in addition to cost-shifting—include using technology assisted review tools to reduce the expense of document review and arranging for the plaintiff's expert to access the AI system at the defendant's place of business subject to the provisions of an appropriate protective order.

B. The Need for Comprehensive Discovery Planning

The case law discussing the discoverability of information regarding algorithms and AI also points to the need for comprehensive discovery planning. In several of the cases where discovery was denied, it appears that the requesting party assumed they would be entitled to discovery as a matter of course and therefore made little effort to justify their request.²⁵⁴ Requesting parties who gained an understanding of the algo²⁵⁵ Therefore, in cases where a discovery request is likely to draw an objection, plaintiffs should consider making use of the full range of discovery devices and disclosures that the Federal Rules of Civil Procedure provide to assess the relevance of algorithmic information.

For example, plaintiffs can use interrogatories and depositions to gain a basic understanding of how an AI product works before requesting production of the AI itself.²⁵⁶ This will not only help the plaintiff understand the relevance of the information they are requesting, but also preempt any court concerns that the plaintiff explore other options before requesting access to sensitive information.²⁵⁷ Involving technology experts in crafting discovery requests, reviewing discovery that has been produced, and responding to meet and confer requests can also help plaintiffs to tailor their discovery and demonstrate relevance with the required degree of specificity.

Comprehensive discovery planning has another benefit. In many FHA cases involving AI discrimination, plaintiffs will be able to meet their stage one burden with publicly available statistics. The algorithm's decision-making process will be more relevant to the questions whether (1) use of AI serves a legitimate defense interest; and (2) whether there are less discriminatory

253 Fed. R. Civ. P. 26(c) (“The court may, for good cause, issue an order to protect a party or person from annoyance, embarrassment, oppression, or undue burden or expense, *including* one or more of the following”) (emphasis added).

254 *See, e.g., New York Times Co.*, 757 F. Supp. 3d at 596.

255 *See, e.g., Flores*, 2021 WL 4441614.

256 *See In re Domestic Airline Travel Antitrust Litig.*, 2020 WL 5847779, at *9 (during the course of litigation, parties should use discovery to “refine[] their understanding of the issues [so that they can] offer an explanation as to the importance of the requested discovery to their case”).

257 *Xinuos, Inc.*, 2024 WL 5010717, at *3 (denying motion to compel where plaintiff did not explore “other discovery tools like interrogatories and depositions” before requesting production of electronically stored information).

alternatives to use of the AI system as designed.^{258,259} The defense always has the burden on the first question and may have some burden on the second.²⁶⁰ It thus is often in the defendant's interest to produce at least some information about how their AI system works.²⁶¹ Proactively proposing an appropriate protective order and meeting and conferring with the defendants can help to convince them that it would ultimately be counterproductive to object to algorithm and AI-related discovery.

V. Conclusion

AI-driven tools are quickly becoming ubiquitous in the real estate section. Given the biases embedded in many AI models, AI decisions that improperly taken into account protected characteristics are a near certainty. Under current law, the FHA provides one of our best tools for addressing this type of bias. As discussed, Plaintiffs who are strategic in their prefling investigation and pretrial discovery should be able to clear the *Twombly/Iqbal* pleading bar and obtain the discovery about AI and algorithms that they need to litigate their case.

A review of the existing case law regarding AI and algorithm-related discovery requests demonstrates that judges have adapted to the challenges posed by AI's rapid development, applying familiar principles from the Federal Rules of Civil Procedure to balance plaintiffs' need for algorithm-related discovery against trade secret and proportionality concerns. While the ability to obtain discovery is not an insurmountable obstacle for FHA plaintiffs, future developments in disparate impact case law may be. Accordingly, there is some urgency to addressing the problem of biased AI now, while litigators still have a full range of tools at their disposal.

Submitted: August 12, 2025 EDT. Accepted: August 15, 2025 EDT.



This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CCBY-4.0). View this license's legal deed at <http://creativecommons.org/licenses/by/4.0> and legal code at <http://creativecommons.org/licenses/by/4.0/legalcode> for more information.

258 Interestingly, for AI models with an error rate, there are almost always alternative ways to build the model without increasing the error rate. This suggests that less discriminatory alternatives could be developed in most case involving AI bias. Black, et al., *supra* note 34, at 67–71. Adding a “human in the loop” to review AI decisions for possible discrimination is another potentially less discriminatory alternative. *See generally*, Ethan Mollick, *Co-Intelligence: Living and Working with AI* 52–54 (2024).

259 Because human decision making is not bias-free, it is possible that some AI platforms may render decisions that – though they have a disparate impact – are less discriminatory than the average human-made decision. Whether and how this might constitute a defense to a disparate impact claim is an interesting question, but one that is ultimately beyond the scope of this article.

260 *See* Section III.C, *supra*.

261 In some cases, transparency about how an AI system operates may also serve the defendant's business interests by creating trust in their product. *See* John Villasenor, *Artificial Intelligence, Trade Secrets, and the Challenge of Transparency*, 25 N.C. J. L. & TECH. 495, 526–29 (2024).