William & Mary
Law Review Online

VOLUME 63 NO. 5, 2022

A TITLE VII DEAD END? MACHINE LEARNING AND EMPLOYEE MONITORING

KAYLA BURRIS

TABLE OF CONTENTS

I. INTRODUCTION

Advancements in technology have made it easier than ever for employers to monitor their employees.[1] These new capabilities, combined with a growing number of employees working remotely, have encouraged companies to adopt monitoring technologies at an alarming rate.[2] According to an American Management Association survey, "80 percent of major companies monitor the internet usage, phone and email of their employees."[3]

Although employees generally oppose monitoring,[4] companies have little incentive to eliminate their programs. Insider threats—employees who misuse their company's network access—pose a substantial threat to companies' networks and cost them a significant amount of money every year.[5] A study from the Ponemon Institute—a research center that focuses on privacy, data protection, and information security—found that cybersecurity incidents caused by employees have increased by 47 percent since 2018, and the costs of such incidents have also risen by 31 percent.[6] However, employers

---

1. *See* Tam Harbert, *Watching the Workers*, SOC'Y FOR HUM. RES. MGMT. (Mar. 16, 2019), https://www.shrm.org/hr-today/news/all-things-work/pages/watching-the-workers.aspx [https://perma.cc/UYG4-UDLW].

2. *See* Bobby Allyn, *Your Boss is Watching You: Work-From-Home Boom Leads to More Surveillance*, NPR (May 13, 2020, 5:00 AM), https://www.npr.org/2020/05/13/854014403/your-boss-is-watching-you-work-from-home-boom-leads-to-more-surveillance [https://perma.cc/Y5BZ-L2Z3]; *see also* Will Douglas Heaven, *This Startup is Using AI to Give Workers a "Productivity Score"*, MIT TECH. REV. (June 4, 2020), https://www.technologyreview.com/2020/06/04/1002671/startup-ai-workers-productivity-score-bias-machine-learning-business-covid/ [https://perma.cc/6BRS-Q6T7].

3. Hannah George, *How Much Employee Monitoring Is Too Much?*, A.B.A. (Jan. 2018), https://www.americanbar.org/news/abanews/publications/youraba/2018/january-2018/how-much-employee-monitoring-is-too-much-/ [https://perma.cc/8EE5-RDXU].

4. *See, e.g.*, Emma Woollacott, *Should You be Monitoring Your Staff with AI?*, RACONTEUR (May 14, 2019), https://www.raconteur.net/technology/artificial-intelligence/ai-workplace-surveillance/ [https://perma.cc/NQ8G-RZ79]; Robert Booth, *UK Businesses Using Artificial Intelligence to Monitor Staff Activity*, THE GUARDIAN (Apr. 7, 2019, 7:38 AM), https://www.theguardian.com/technology/2019/apr/07/uk-businesses-using-artifical-intelligence-to-monitor-staff-activity [https://perma.cc/CP2Y-29LX].

5. *See What is an Insider Threat?*, PROOFPOINT, https://www.proofpoint.com/us/threat-reference/insider-threat [https://perma.cc/GLS9-9KT4].

6. PONEMON INST., 2020 COST OF INSIDER THREATS GLOBAL REPORT 3 (2020), https://cdw-prod.adobecqms.net/content/dam/cdw/on-domain-cdw/brands/proofpoint/ponemon-global-cost-of-insider-threats-2020-report.pdf [https://perma.cc/8HVY-3EQT].

are not just dealing with cybersecurity concerns. They also have to worry about employees wasting time online,[7] employees sending company information outside the network,[8] and other instances of inappropriate computer use.[9]

As employers ingest ever-growing quantities of employee network data, they are starting to turn to machine learning to help them analyze and make sense of this information.[10] Although machine learning tools make dealing with large amounts of data easier than before, they can come with significant drawbacks, namely algorithmic discrimination.[11] In an employee monitoring context, algorithmic discrimination could result in the algorithm flagging people in one group as potentially more likely to commit violations than others, which could thus result in the employer taking more adverse actions against that group.[12]

As a preliminary matter, it is important to first address the looming question of why discrimination in employee monitoring matters. If the programs are still catching people who do bad things, such as sending out proprietary information or making fantasy football spreadsheets on company time, should anyone care if more of the people who are caught belong to a particular group? This is a legitimate question, particularly given the assumption that many companies are unlikely to intentionally use these tools to discriminate against their own employees (it is in their self-interest to avoid such practices which, if discovered, could lead to lawsuits).

---

7. *See, e.g.*, Cheryl Conner, *Wasting Time at Work: The Epidemic Continues*, FORBES (July 15, 2015, 5:39 PM), https://www.forbes.com/sites/cherylsnappconner/2015/07/31/wasting-time-at-work-the-epidemic-continues/?sh=430551011d94 [https://perma.cc/SQQ2-YYN7].

8. *See, e.g.*, Zack Whittaker, *Amazon Fires Employees for Leaking Customer Email Addresses and Phone Numbers*, TECHCRUNCH (Jan. 10, 2020, 5:22 PM), https://techcrunch.com/2020/01/10/amazon-employees-email-address/ [https://perma.cc/GGB2-M8EJ].

9. *See The Latest on Workplace Monitoring and Surveillance*, AM. MGMT. ASSOC. (Apr. 8, 2019), https://www.amanet.org/articles/the-latest-on-workplace-monitoring-and-surveillance/ [https://perma.cc/LAY6-R5ZT].

10. *See* Rick Bales, *Artificial Intelligence in the Workplace*, OHIO STATE BAR ASS'N (July 20, 2020), https://www.ohiobar.org/member-tools-benefits/practice-resources/practice-library-search/practice-library/section-newsletters/2020/artificial-intelligence-in-the-workplace/ [https://perma.cc/CG9L-3A45].

11. *See* Jenifer Winter, *Algorithmic Discrimination: Big Data Analytics and the Future of the Internet*, *in* 17 PUB. ADMIN. & INFO. TECH. 125, 131-32 (Jennifer Winter & Ryota Ono eds., 2015).

12. *See id.*

To illustrate why discriminatory employee monitoring is concerning, consider the police practice of patrolling African-American neighborhoods more than neighboring white neighborhoods.[13] Discriminatory employee monitoring functions in a similar fashion. In both practices, people are caught committing acts that are wrong, whether they be crimes or violations of company policy. The more an employer targets or watches particular employees, the more likely it is to find that those people did something wrong.[14] This wrong can be large, such as sending company earnings reports to competitors, or small, such as sending home a company newsletter that includes some proprietary information.[15]

This Note will argue that Title VII, as courts currently apply the law, does not adequately protect employees from algorithmic discrimination when companies use machine learning to monitor their employees' computers. Part I will provide an introduction to how employee monitoring tools work, how employers are using machine learning in their monitoring programs, and how these programs can discriminate. Because scholars have already done significant work in this area, this Note will not try to replicate this research but will provide an overview of how this discrimination can occur. Parts II and III will then analyze how an employee might prove a Title VII claim. Part II will analyze an employee's claim under the disparate treatment theory of discrimination and ultimately conclude that an employee is unlikely to succeed under this theory of discrimination. Part III then analyzes a potential claim under the disparate impact theory of discrimination, analyzing each of the three prongs of the disparate impact test. This Note ultimately concludes that, although

13. *See, e.g.*, Ronald Weitzer & Rod K. Brunson, *Policing Different Racial Groups in the United States*, 35 CAHIERS POLITIESTUDIES 129, 135-40 (2015); Will Douglas Heaven, *Predictive Policing Algorithms Are Racist. They Need to Be Dismantled.*, MIT TECH. REV. (July 17, 2020), https://www.technologyreview.com/2020/07/17/1005396/predictive-policing-algorithms-racist-dismantled-machine-learning-bias-criminal-justice/ [https://perma.cc/5MHQ-WH5X].

14. *Cf.* Weitzer & Brunson, *supra* note 13; Heaven, *supra* note 13.

15. *See* Joshua Stowers, *7 Ways Your Work Tech is Betraying Your Privacy*, BUS. NEWS DAILY (Dec. 21, 2021), https://www.businessnewsdaily.com/7928-work-computer-employee-monitoring.html [https://perma.cc/H8YK-HTV4]; *see also* Christopher M. Sullivan & Zachary P. O'Keeffe, *Does More Policing Lead to Less Crime—or Just More Racial Resentment?*, WASH. POST (July 25, 2016), https://www.washingtonpost.com/news/monkey-cage/wp/2016/07/25/does-more-policing-lead-to-less-crime-or-just-more-racial-resentment/ [https://perma.cc/NKX2-5HC6].

disparate impact appears better suited to address algorithmic discrimination in employee monitoring, an employee is still unlikely to succeed under this theory. Part IV discusses potential ways to address the issue of algorithmic discrimination in employee monitoring and ultimately concludes that a negligent use of technology standard would best suit the interests of both employers and employees.

I. EMPLOYEE MONITORING, MACHINE LEARNING, AND ALGORITHMIC DISCRIMINATION

Generally, employers try to answer four main questions with their employee monitoring programs: (1) which employees may pose liability risks?; (2) who threatens to share its data without authorization?; (3) who is trying to harm its network?; and (4) which employees are more, or less, productive?[16] These questions are usually incorporated into an employer's computer usage policy.[17] The usage policy typically outlines what employees are and are not allowed to do on company computers and warns that any violation of the policies can lead to employment consequences.[18] Examples of computer policy provisions include prohibitions on personal email use, harassing or explicit content, downloading unknown software, and providing unauthorized access to company systems.[19]

To answer these questions and to identify such actions, employers install monitoring programs on employees' computers.[20] The programs themselves vary in sophistication and degree of analysis.[21] Some provide only basic information to the employer, such as keystrokes and email contents, leaving the employer to read and

16. *See* Shuchih Ernest Chang, Anne Yenching Liu & Sungmin Lin, *Exploring Privacy and Trust for Employee Monitoring*, 115 INDUST. MGMT. & DATA SYS. 88, 89 (2015); *Employee Monitoring Software: Productivity, Security & Compliance Made Simple*, VERIATO, https://www.veriato.com/solutions/use-cases/employee-monitoring-software [https://perma.cc/2SP9-U2V6] [hereinafter *Employee Monitoring Software*].

17. *See, e.g.*, *Computer, E-mail and Internet Usage Policy*, SHRM, https://www.shrm.org/ [https://perma.cc/6C4J-JF7M] (access by searching "SHRM Computer Use Policy" in web browser search engine).

18. *E.g.*, *id.*

19. *Id.*

20. *See* Andrew Milam Jones, *Employee Monitoring*, 83 TEX. BAR J. 98, 98 (2020).

21. *See id.*

interpret the data.[22] Other tools perform more analysis on the data to help the employer see patterns and focus on particular employees.[23]

The most advanced analysis tool to date is machine learning. Machine learning is the process of teaching an algorithm to make decisions on its own.[24] It works by giving a computer program "large amounts of data with [target] output variables."[25] The computer program then searches for "useful patterns" between the data and the target variables and self-adjusts its algorithms until it determines the best result possible with the training data provided.[26] At this point, the program is capable of ingesting live data and using the algorithms it made previously to make a decision based on the new data.[27]

One popular example of machine learning is Google Maps.[28] Google Maps attempts to predict the speed of traffic flow, and thus, how long it will take a user of the app to get from one point to another.[29] In order to calculate this, it ingests large amounts of data, including "historical traffic data, information like speed limits and construction sites from local governments, and also factors like the

---

22. *See id.*

23. *See id.*; *Employee Monitoring Software*, *supra* note 16; *Best Practices for Detecting Insider Security Threats*, INTERGUARD (Jan. 21, 2019), https://www.interguardsoftware.com/best-practices-for-detecting-insider-security-threats-in-2019/ [https://perma.cc/Y2NY-NKEK] [hereinafter *Best Practices*].

24. *See* Ignacio N. Cofone, *Algorithmic Discrimination is an Information Problem*, 70 HASTINGS L.J. 1389, 1395 (2019); Anya E.R. Prince & Daniel Schwarcz, *Proxy Discrimination in the Age of Artificial Intelligence and Big Data*, 105 IOWA L. REV. 1257, 1273 (2020); Bernard Marr, *What is the Difference Between Artificial Intelligence and Machine Learning?*, FORBES (Dec. 6, 2016, 2:24 AM), https://www.forbes.com/sites/bernardmarr/2016/12/06/what-is-the-difference-between-artificial-intelligence-and-machine-learning/#5bba6c892742 [https://perma.cc/UL35-KCNR].

25. Cofone, *supra* note 24, at 1395; *see also* Cliff Kuang, *Can A.I. be Taught to Explain Itself?*, N.Y. TIMES MAG. (Nov. 21, 2017), https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-to-explain-itself.html [https://perma.cc/N55W-8X4V].

26. Prince & Schwarcz, *supra* note 24, at 1273 (quoting Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CALIF. L. REV. 671, 677 (2016)). This is a basic description of machine learning.

27. *See id.*

28. *See* James Vincent, *How Google Maps Uses DeepMind's AI Tools to Predict Your Arrival Time*, VERGE (Sept. 3, 2020, 10:00 AM), https://www.theverge.com/2020/9/3/21419632/how-google-maps-predicts-traffic-eta-ai-machine-learning-deepmind [https://perma.cc/5AFR-ZQUC].

29. *See id.*

quality, size, and direction of any given road."[30] It uses this information to build an algorithmic model, that is, a formula that gives more weight to certain features over others in order to predict how the traffic flow will change when certain variables are added or removed.[31] So for example, based on historical data, the algorithm may decide that it is more likely to determine the correct historical traffic flow when it weighs speed limits more than road quality in a particular area.[32] The app then starts ingesting live data—such as information from peoples' phones and updated construction information—and applies this into the existing algorithm to determine the traffic flow and offer a prediction for how long the drive will take.[33] This model also constantly incorporates live data into its algorithm to continue to improve its predictions.[34]

Machine learning in monitoring technologies works in a similar fashion. In an employment context, an employer could train the algorithm by giving it information about a large number of employees, some who committed violations on the network and some who did not.[35] That way, when the algorithm ingests live employee data, the algorithm can use the past data to inform its analysis of the live data and predict who may be more likely to commit a violation in the future.[36] Once an employer identifies particular employees as risks, it can then monitor these employees more closely in order to determine whether they have in fact committed a violation.[37]

There has been substantial scholarship describing how algorithms discriminate.[38] This Note will focus on how this discrimination can occur in the context of employee monitoring.

One way algorithms can discriminate is by learning from discriminatory inputs. As discussed, an algorithm learns and refines itself

---

30. *Id.*
31. *Id.*
32. *See id.*
33. *See id.*
34. *See id.*
35. *See* Kuang, *supra* note 25 (providing a similar example about using machine learning to evaluate loan decisions).
36. *See id.*
37. *See Best Practices*, *supra* note 23.
38. *See generally* Talia B. Gillis & Jann L. Spiess, *Big Data and Discrimination*, 86 U. CHI. L. REV. 459 (2019); Pauline T. Kim, *Auditing Algorithms for Discrimination*, 166 U. PA. L. REV. ONLINE 189 (2017).

by analyzing past data and adjusting its algorithm so that it comes to the correct result based on past instances.[39] Thus, if discrimination existed in the past data, the machine would incorporate this bias in future results in order to reach the desired output that it was trained to reach.[40]

Amazon experienced this phenomenon when it tried to implement a machine learning hiring tool in 2014.[41] This tool analyzed resumes and gave applicants a "score[]" based on how good of an employee the algorithm thought the applicant would be.[42] It was not until a year later that Amazon realized the tool had learned to bias itself against women.[43] This bias occurred because Amazon trained its models off of its hiring data over a ten-year period, and this data reflected the current "male dominance across the tech industry."[44] The model used this information and essentially "learned" that it should prefer men over women in its algorithm.[45] The same phenomenon can occur in monitoring—if a monitoring system were to notice based on company records of employee violations that a certain group created those violations, it could "learn" that the group was more likely to create such a violation in the future.[46]

This process of incorporating biased inputs into the algorithm can occur whether or not the protected class is included in the data set through the use of proxies.[47] Algorithms use proxies—that is, other features in the data set—to infer a characteristic about an individual.[48] For example, the resumes Amazon used likely did not

---

39. *See* Sandra G. Mayson, *Bias In, Bias Out*, 128 YALE L.J. 2218, 2218 (2019); Fredric I. Lederer, *Here There Be Dragons: The Likely Interaction of Judges with the Artificial Intelligence Ecosystem*, 59 JUDGES' J. 12, 12 (2020) ("The accuracy of an AI depends on its original programming, the quality of its training, and the quantity and quality of data it uses.")

40. *See* Gillis & Spiess, *supra* note 38, at 467.

41. *See* Jeffrey Dastin, *Amazon Scraps Secret AI Recruiting Tool that Showed Bias Against Women*, REUTERS (Oct. 10, 2018, 7:04 PM), https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G [https://perma.cc/93ZP-FFPE].

42. *Id.*

43. *Id.*

44. *Id.*

45. *Id.*

46. *Cf. id.*

47. *See* Barocas & Selbst, *supra* note 26, at 691-92.

48. *See id* at 691.

include the applicant's gender; however, the algorithm may have learned to infer gender based on the person's name.[49] It then could have used the person's name as a proxy for gender in future scenarios, thus scoring people higher or lower based on their "name."[50]

It is also possible for data creators to use masking techniques, such as the intentional use of biased training data or feature selection, as a way to insert discriminatory bias into the algorithm.[51] For example, it is now well-known that due to our fractured society, zip codes can stand as a proxy for race in certain instances.[52] Programmers with this knowledge could thus intentionally include zip codes in the algorithm's data set in order to discriminate against certain races.[53]

These instances of discrimination are concerning, particularly as employers begin to use machine learning in their monitoring programs to detect possible violations among employees.

II. ESTABLISHING A TITLE VII CLAIM: DISPARATE TREATMENT

According to Title VII, employers are not allowed "to discharge ... or otherwise to discriminate against" employees on the grounds of a protected characteristic—"race, color, religion, sex, or national origin."[54] The law also prohibits employers from limiting or classifying employees on the basis of a protected characteristic "in any way which would ... adversely affect" their employment status.[55] To allege a Title VII claim, an employee may assert either a disparate impact or a disparate treatment theory of discrimination.[56]

---

49. *See* Barocas & Selbst, *supra* note 26, at 691-92; Dastin, *supra* note 41.

50. *See* Barocas & Selbst, *supra* note 26, at 691-92; Dastin, *supra* note 41.

51. *See* Barocas & Selbst, *supra* note 26, at 692.

52. *See, e.g.*, Alexandra George, *Thwarting Bias in AI Systems*, CARNEGIE MELLON U., https://engineering.cmu.edu/news-events/news/2018/12/11-datta-proxies.html [https://perma.cc/4WFV-N9VW]; Nicol Turner Lee, Paul Resnick & Genie Barton, *Algorithmic Bias Detection and Mitigation: Best Practices and Policies to Reduce Consumer Harms*, BROOKINGS (May 22, 2019), https://www.brookings.edu/research/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/ [https://perma.cc/D3VB-F3MC].

53. Barocas & Selbst, *supra* note 26, at 692.

54. 42 U.S.C. § 2000e-2(a)(1).

55. *Id.* § 2000e-2(a)(2).

56. *See* Ricci v. DeStefano, 557 U.S. 557, 577-78 (2009); Joseph A. Seiner, *Disentangling Disparate Impact and Disparate Treatment: Adapting the Canadian Approach*, 25 YALE L. & POL'Y REV. 95, 96-97 (2006).

This Part will first explore how an employee may go about alleging a Title VII claim under the disparate treatment theory of discrimination. Part III will then analyze how an employee may allege a discrimination claim under the disparate impact theory of discrimination. Ultimately, this Part concludes that an employee is unlikely to succeed in a disparate treatment claim because they will likely not have enough information on the employer's monitoring practices. Regardless, even if the employee uses this theory of discrimination, the claim will likely not succeed because the employer can dispute the claims of pretext by proffering other valid motivations for the action.

It is important to note that this analysis assumes that the only evidence of discrimination the employee possesses are the facts surrounding the employee's computer usage and the employer's adverse actions against the employee. Any additional evidence demonstrating possible discriminatory intent would likely bolster an employee's Title VII claim. However, because this Note focuses on machine learning and employee monitoring in particular, it will not analyze how additional evidence may affect potential cases.

Disparate treatment involves claims of intentional discrimination based on a protected characteristic.[57] In order to establish a disparate-treatment claim, a plaintiff must show "that the defendant had a discriminatory intent or motive" in terminating the employee or taking some other form of adverse action against the employee.[58] The treatment cannot consist of isolated or sporadic occurrences of discrimination.[59]

Under a disparate-treatment theory of discrimination, an employee would allege that the employer set up its monitoring program in a way that intentionally discriminated against a particular indvidual or group based on a protected characteristic.[60] The employee could allege this discrimination in two ways.[61] First, an employee could allege that an employer used a protected class in its algorithm

57. Int'l Brotherhood of Teamsters v. United States, 431 U.S. 324, 335 n.15 (1977).

58. *Ricci*, 557 U.S. at 577 (quoting Watson v. Fort Worth Bank & Trust, 487 U.S. 977, 986 (1988)); *see* 42 U.S.C. § 2000e-2(a)(1)-(2).

59. *See Int'l Brotherhood of Teamsters*, 431 U.S. at 336.

60. *See Ricci*, 557 U.S. at 577; *Int'l Brotherhood of Teamsters*, 431 U.S. at 335 n.15; 42 U.S.C. § 2000e-2(a)(1)-(2).

61. *See* Barocas & Selbst, *supra* note 26, at 699.

in order to discriminate against those employees.[62] For example, an employee could allege that an employer intentionally coded its monitoring tool to find members of a certain religious group as more likely to take company data.[63] This alone would constitute a violation of Title VII; however, it may be difficult for the employee to prove such a classification occurred.

Proving such a classification may be difficult for several reasons. First, employees would need to know or have a sense that the employer was using such a classification system—rather than unintentionally discriminating against, for example, a particular race, which would be disparate impact.[64] Although many employers notify employees that they may be monitored,[65] they generally do not advertise their exact algorithms.[66] Second, the algorithms are often proprietary and thus are "resistant to discovery and scrutiny" by employees.[67] Thus without other indications of bias or knowledge of how the system actually works, an employee may have a difficult time alleging a specific claim against an employer.

Second, an employee could allege that an employer intentionally discriminated against her *because* of a protected characteristic.[68] An employee can do this under either the *McDonnell Douglas* test or through the mixed motive regime.[69] In *McDonnell Douglas Corporation v. Green*, the Supreme Court articulated that in order to establish a Title VII claim, an employee must first establish "a prima facie case of [protected status] discrimination."[70]

---

62. *Id.*

63. *Id.*

64. *See infra* Part III.

65. *See, e.g.*, *Managing Workplace Monitoring and Surveillance*, SHRM (Mar. 13, 2019), https://www.shrm.org/resourcesandtools/tools-and-samples/toolkits/pages/workplaceprivacy.aspx [https://perma.cc/Z8DG-NGUB].

66. Many of these systems use proprietary algorithms and thus cannot be shared with employees. *See* Ulrich Leicht-Deobald, Thorsten Busch, Christoph Schank, Antoinette Weibel, Simon Schafheitle, Isabelle Wildhaber & Gabriel Kasper, *The Challenges of Algorithm-Based HR Decision-Making for Personal Integrity*, 160 J. BUS. ETHICS 377, 381 (2019). Furthermore, doing so may enable employees to alter their activities in order to get around the system.

67. Sonia K. Katyal, *Private Accountability in the Age of Artificial Intelligence*, 66 UCLA L. REV. 54, 59 (2019).

68. *See* Barocas & Selbst, *supra* note 26, at 699.

69. *See id.* at 696.

70. 411 U.S. 792, 802 (1973).

McConnell Douglas involved an allegation of discriminatory hiring.[71] There, the defendant, McDonnell Douglas Corporation, originally laid off Green as part of a reduction in work force.[72] Believing the termination was racially motivated, Green protested by blocking entrances to the company and was arrested in the process.[73] Soon thereafter, McDonnell Douglas advertised for the same position and turned down Green's application due to his illegal conduct.[74] Green countered, alleging that McDonnell Douglas refused to hire him because of his race and his involvement in legitimate civil rights activities.[75] In analyzing the case, the Court found that Green met his initial prima facie burden, which in the hiring context involved a showing

> (i) that he belongs to a racial minority; (ii) that he applied and was qualified for a job for which the employer was seeking applicants; (iii) that, despite his qualifications, he was rejected; and (iv) that, after his rejection, the position remained open and the employer continued to seek applicants from persons of complainant's qualifications.[76]

However, the Court noted that the evidentiary burden then shifts to the employer "to articulate some legitimate, nondiscriminatory reason for the employee's rejection," and ultimately held that unlawful conduct met this second step.[77]

In a case for termination, an employee could present a prima facie case by showing that she was a member of a protected class, she suffered an adverse employment action, her performance was satisfactory, and "the circumstances give rise to an inference of discrimination."[78] Although the prima facie burden is supposed to be

---

71. *Id.* at 796.
72. *Id.* at 794.
73. *Id.* at 794-95.
74. *Id.* at 796.
75. *Id.* at 796.
76. *Id.* at 802.
77. *Id.* at 802-03.
78. Pye v. Nu Aire, Inc., 641 F.3d 1011, 1019 (8th Cir. 2011) (citing Wierman v. Casey's Gen. Stores, 638 F.3d 984, 993 (8th Cir. 2011)); *see also* Holland v. Washington Homes, Inc., 487 F.3d 208, 214 (4th Cir. 2007); 42 U.S.C. § 2000e-2(a)(1)-(2); *cf. McDonnell Douglas*, 411 U.S. at 802.

a low bar,[79] it may prove difficult for employees to pass in a monitoring situation.[80] On its own, a violation of company policy oftentimes defeats an employee's prima facie burden of discrimination because she is no longer able to point to satisfactory performance to raise an inference of discrimination.[81] Instead, the employee will need some additional information to establish "an inference of discrimination."[82]

In other cases when plaintiffs violate a company policy and are then terminated, they typically establish an inference of discrimination by showing that similarly situated employees also violated the policy but did not face the same consequences or reprimands.[83] This would "present evidence that the stated reason [for termination] is a 'pretext.'"[84] In this scenario, an employee would need to find evidence that other similarly situated employees not part of the same protected class committed similar violations, such as sending emails outside the network or personal use of a work computer.[85] Given the nature of computer violations—that employees generally do not use their work computers in front of an audience, especially for questionable actions—finding sufficient proof could in itself prove daunting for many employees.

An employee can also try a disparate treatment case under the mixed-motive framework, in which "a plaintiff need only present sufficient evidence for a reasonable jury to conclude, ... that [a protected class] 'was a motivating factor'" for the adverse action.[86] As Barocas and Selbst note, practically this means that the employee "must show that the same action would not have been taken absent the discriminatory motive."[87] Under this framework, an employee would virtually never succeed under most employee monitoring scenarios because an employer would always have another motivation

79. *See* Stephanie Bornstein, *Antidiscriminatory Algorithms*, 70 ALA. L. REV. 519, 559 (2018) ("The prima facie burden is not a hard one to meet.").

80. *See supr*a Part I.

81. *See, e.g.*, *Pye*, 641 F.3d at 1019.

82. *See id.*

83. *See* Coleman v. Donahoe, 667 F.3d 835, 841 (7th Cir. 2012); Jackson v. VHS Detroit Receiving Hosp., Inc., 814 F.3d 769, 776-77 (6th Cir. 2016).

84. *Coleman*, 667 F.3d at 845 (quoting McDonnell Douglas Corp. v. Green, 411 U.S. 792, 804 (1973)).

85. *See supra* Part I.

86. Desert Palace, Inc. v. Costa, 539 U.S. 90, 101 (2003) (quoting 42 U.S.C. § 2000e-2(m)).

87. Barocas & Selbst, *supra* note 26, at 697.

for the action. In an insider threat scenario or a scenario of violating company policy, an employer would be able to cite the violation as the motivating factor—not the discrimination—even if it did in fact occur.[88] Similarly, in a productivity tool scenario, the employer could point to instances of the employee violating computer usage policies or simply to the employee's efficiency itself as a sufficient motivating factor for termination.[89]

In summary, absent some clear discriminatory policy on the part of the employer, the disparate treatment theory of discrimination appears to be of little use to most employees who face adverse action due to monitoring.

III. ESTABLISHING A TITLE VII CLAIM: DISPARATE IMPACT

On its face, disparate impact theory appears better-suited to deal with claims of discriminatory machine-learning algorithms in employee monitoring.[90] Disparate impact theory refers to facially neutral employment practices that disproportionally affect one group more than another, regardless of the employer's actual intent in implementing the practice.[91] Accordingly, in a case alleging discriminatory monitoring practice, an employee could point to the *impact* of the monitoring on a particular class.[92] To give a hypothetical example, an employee could allege that although the employer claimed to have a neutral machine-learning monitoring system, African Americans appeared to receive a disproportionate number of violations as compared to other persons in the office.[93]

On its face, this theory appears particularly apt to solving the problem of algorithmic discrimination. However, because of the way the law is structured as well as the particular nature of employee

88. *See supra* Part I.

89. *See supra* Part I.

90. *See* Barocas & Selbst, *supra* note 26, at 701; Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 WM. & MARY L. REV. 857, 887 (2017) ("When statistical bias coincides with systematic disadvantage to protected classes, it causes discriminatory harm.").

91. *See* Griggs v. Duke Power Co. 401 U.S. 424, 430 (1971) ("Under the Act, practices, procedures, or tests neutral on their face, and even neutral in terms of intent, cannot be maintained if they operate to 'freeze' the status quo of prior discriminatory employment practices.").

92. *See id.*

93. *See id.*

monitoring, employees will likely have a difficult time mounting a successful discrimination claim under this theory as well.

A. Establishing a Claim

In order to establish a claim of disparate-impact liability, an employee must first establish a prima facie basis for the suit.[94] An employee can do this by showing that there was in fact "a significant statistical disparity" between different groups of employees.[95]

The Equal Employment Opportunity Commission (EEOC) has defined the required statistical disparity for hiring in the "four-fifths rule," which says that "[a] selection rate for any [protected class] which is less than four-fifths (4/5) (or eighty percent) of the rate for the group with the highest rate will generally be regarded ... as evidence of adverse impact."[96] Although the rule's language focuses on hiring, courts have sometimes applied these guidelines to termination cases, though many courts have abandoned the "rule" (which is in fact a guideline and not legally binding)[97] altogether in favor of analyzing statistical significance.[98] Needless to say, although the measure and degree of disparity may differ between circuits, employees will need to prove that a particular class was significantly more affected by the monitoring policies[99]—whether that be in the form of disciplinary measures, terminations, or some other form of adverse treatment.[100]

Apart from the usual troubles a plaintiff may face in establishing statistical significance,[101] an employee in an employee monitoring case does not appear to face any unique challenges in meeting the prima facie burden. However, that is not the end of a disparate impact case.

---

94. *See* Ricci v. DeStefano, 557 U.S. 557, 557 (2009).

95. *Id.* at 587 (citing Connecticut v. Teal, 457, U.S. 440, 446 (1982)).

96. 29 C.F.R. § 1607.4(D).

97. Clady v. County of Los Angeles, 770 F.2d 1421, 1428 (9th Cir. 1985).

98. *See id.* ("[T]he 80 percent rule has been sharply criticized by courts and commentators."); Jones v. City of Boston, 752 F.3d 38, 43 (1st Cir. 2014). For more information on the different tests, see Katie Eissenstat, *Lies, Damned Lies, and Statistics: The Case to Require "Practical Significance" to Establish a Prima Facie Case of Disparate Impact Discrimination*, 68 OKLA. L. REV. 641 (2016).

99. *See* Eissenstat, *supra* note 98, at 642-43.

100. *See* 42 U.S.C. § 2000e-2(a)(1)-(2).

101. *See* Eissenstat, *supra* note 98, at 642-43.

B. Business Necessity and Alternative Employment Practices

After establishing a prima facie case under the disparate impact theory, an employer may then challenge the claim by raising a business necessity defense.[102] The Supreme Court has called the business necessity defense the "touchstone" of a disparate impact case.[103] The Supreme Court initially defined the defense in *Griggs v. Duke Power Co.* as a showing that the practice in question was "related to job performance."[104]

Congress directly dealt with business necessity when it passed the Civil Rights Act of 1991, which says that "a complaining party demonstrates that a respondent uses a particular practice that causes a disparate impact on the basis of [a protected class] and the respondent fails to demonstrate that the challenged practice is job related for the position in question and consistent with business necessity."[105]

Since the 1991 Act, courts have dealt with business necessity in different ways, some requiring a "manifest relationship" to employment and others requiring the actions be "significantly correlated" to employment.[106] In the Supreme Court's most recent discussion of disparate impact—which took place in the context of a Fair Housing Act (FHA) dispute—it described an even more lenient standard.[107] The Court described business necessity as "mandat[ing] the 'removal of artificial, arbitrary, and unnecessary barriers.'"[108]

So far, there are no cases answering the question of whether some of these employee monitoring programs would qualify as a business necessity under Title VII. However, federal law allows employers to monitor employees' computers,[109] and federal courts have also taken

102. *See* 42 U.S.C. § 2000e-2(k)(1)(A)(i); Lewis v. City of Chicago, 560 U.S. 205, 213 (2010).

103. Griggs v. Duke Power Co., 401 U.S. 424, 431 (1971) ("The touchstone is business necessity.").

104. *Id.*

105. 42 U.S.C. § 2000e-2(k)(1)(A)(i); *see also* Barocas & Selbst, *supra* note 26, at 703.

106. Barocas & Selbst, *supra* note 26, at 704 (first quoting Gallagher v. Magner, 619 F.3d 823, 834 (8th Cir. 2010), then quoting Gulino v. N.Y. State Educ. Dep't, 460 F.3d 361 (2d Cir. 2006)).

107. *See id.* at 704 n.161; Tex. Dep't of Hous. & Cmty. Affs. v. Inclusive Cmtys. Project, Inc., 576 U.S. 519, 540 (2015).

108. *Tex. Dep't of Hous.*, 576 U.S. at 540 (quoting *Griggs*, 401 U.S. at 431).

109. *See* Corey A. Ciocchetti, *The Eavesdropping Employer: A Twenty-First Century Framework for Employee Monitoring*, 48 AM. BUS. L.J. 285, 293 (2011).

a generally pro-employer stance on issues of employee monitoring so long as employers provide sufficient notice of the monitoring.[110] Furthermore, unlike in the hiring context where courts are determining whether particular traits or qualities are job-related,[111] in an employer monitoring context, employers are arguably using machine learning for more critical goals.[112] In terms of a business necessity argument, companies would likely succeed in arguing that monitoring for potential insider threats, violations, liabilities, and inefficient workers[113] is critical to most modern companies.[114]

C. Alternative Employment Options

If an employer is able to meet the business necessity burden, then the burden will shift back to the employee for the final prong of disparate impact analysis: the alternative employment options test.[115] Under this prong, an employee may still succeed in a disparate impact claim by showing "that the employer refuses to adopt an available alternative employment practice that has less disparate impact and serves the employer's legitimate needs."[116] On its face, this would appear to be a perfect solution to the problem of discriminatory machine learning algorithms: if an employee is able to point to less discriminatory options for monitoring—perhaps by using a different monitoring system or tweaking the code—then the employee should theoretically be able to defeat a showing of business necessity.[117] However, deeper scrutiny into how courts have applied this analysis reveals that employees will likely have an extremely difficult time in overcoming this hurdle.

---

110. *See* Janna Fischer, *Big Boss is Watching: Circumstances Under Which Employees Waive the Attorney-Client Privilege by Using E-Mail at Work*, 12 COLO TECH. L.J. 365, 378-85 (2014).

111. *See, e.g.*, Barocas & Selbst, *supra* note 26, at 703-05.

112. *See supra* Part I.

113. *See supra* Part I.

114. *See, e.g.*, *2020 Cost of Insider Threats Global Report*, *supra* note 6, at 3; Kiely Kuligowski, *Distracted Workers Are Costing You Money*, BUS. NEWS DAILY (Nov. 23, 2020), https://www.businessnewsdaily.com/267-distracted-workforce-costs-businesses-billions.html [perma.cc/RPN7-GZWG].

115. *See* 42 U.S.C. § 2000e-2(k)(1)(A)(ii); Ricci v. DeStefano, 557 U.S. 557, 578 (2009).

116. *Ricci*, 557 U.S. at 578.

117. *See id.*; *see also* Barocas & Selbst, *supra* note 26, at 709.

In order to meet this burden, employees must show that (1) there existed an alternative employment practice that was equally valid in meeting the employer's needs; (2) that practice would have resulted in less of a disparate impact; and (3) the employer refused to adopt the practice.[118] To grasp the difficulties employees may face in asserting this claim, this Note will walk through each of the required elements.

1. Existence of an Alternative Employment Practice

First, an employee would need to establish an alternative monitoring practice that is *equally* valid in meeting the employer's needs.[119] As courts have applied it, this is a strict test that would be difficult to apply in the context of machine learning software, no matter which alternative employees choose.

If the court has ruled that the monitoring is a business necessity, a proposed alternative of simply cancelling the program is unlikely to succeed because it would not address the employer's legitimate business needs.[120] As discussed earlier, employers generally use machine learning tools in their monitoring programs to answer important questions related to corporate liability, unauthorized disclosure, cybersecurity, and employee productivity.[121] Thus, courts are unlikely to find doing away with monitoring an equally viable alternative because it does not actually help the employer address the goals of the program.[122]

An employee could also suggest as an alternative that the employer modify its algorithm to reduce disparate impact; however, this would also prove difficult because of the nature of machine learning and the difficulties in addressing the disparate impact in the algorithm.[123]

---

118. *See* §§ 2000e-2(k)(1)(A)(ii); *Ricci*, 557 U.S. at 589-91; Jones v. City of Boston, 845 F.3d 28, 34 (1st Cir. 2016); United States v. Brennan, 650 F.3d 65, 109 (2d Cir. 2011).

119. *See Ricci*, 557 U.S. at 589-91; *Jones*, 845 F.3d at 34; *Brennan*, 650 F.3d at 109.

120. *See, e.g.*, Davis v. Ashcroft, 355 F. Supp. 2d 330, 345 (D.D.C. 2005) (holding that simply doing away with an employment requirement was unjustified because plaintiff did not establish that it served no purpose for the employer).

121. *See supra* Part I.

122. *See, e.g.*, *Ricci*, 557 U.S. at 589-90; *Davis*, 355 F. Supp. 2d at 345.

123. *See,* Barocas & Selbst, *supra* note 26, at 709-10.

The nature of machine learning makes this difficult because unlike a regular algorithm where one can tweak and adjust the inputs, in instances of machine learning, the *algorithm*—not the programmer—determines how to weigh different factors in its decision making.[124] This makes it difficult—if not impossible with current technology—to determine how exactly a machine learning tool is reaching a particular conclusion, a problem scholars refer to as the "black box" problem.[125] Thus, it would be difficult—and in some instances impossible—for an employee to generate *specific* alternative modifications for an employer to implement.[126] The alternative would need to be specific in order to meet all three prongs of the test. For example, in order to determine whether an alternative is equally effective, a court must have real data to analyze, not simply a request that the employer make a change.[127] Specificity is also required in order to evaluate the other two prongs—whether the alternative would produce less of a disparate impact and whether the employer knew about and refused the alternative.[128]

An employer could also deal with a disparate impact in results by adding the protected class into its inputs and then training the algorithm to produce a less biased result.[129] However, this alternative is barred by Title VII, which prohibits using protected classifications in any form, including to correct for disparities in data.[130]

---

124. *See* Ariel Bleicher, *Demystifying the Black Box that is AI*, SCI. AM. (Aug. 9, 2017), https://www.scientificamerican.com/article/demystifying-the-black-box-that-is-ai/[https://perma.cc/F3MN-BU5D].

125. Professor Fredric Lederer explained this issue in a recent article:

> The difficulty in determining how the "black box" AI reached the result it implemented ... may make it impossible to determine causation.... [T]he very nature of AI is problematic as the number of possible causes and the identity of the data points involved and data owners may be so large as to create qualitatively different problems than in the past.

Fredric I. Lederer, *Here There Be Dragons*, 59 JUDGES' J. 12, 13 (2020); *see also* Cade Metz, *Google Researchers are Learning How Machines Learn*, N.Y. TIMES (Mar. 6, 2018), https://www.nytimes.com/2018/03/06/technology/google-artificial-intelligence.html [https://perma.cc/3MJC-N7WN]; Charles McLellan, *Inside the Black Box: Understanding AI Decision-Making*, ZDNET (Dec. 1, 2016, 16:24 GMT), https://www.zdnet.com/article/inside-the-black-box-understanding-ai-decision-making/ [https://perma.cc/FW58-X7AX].

126. *See, e.g.*, Jones v. City of Boston, 845 F.3d 28, 34 n.2 (1st Cir. 2016).

127. *See id.*

128. *See id* at 34.

129. *See* Gillis & Spiess, *supra* note 38, at 471-72.

130. *See, e.g.*, Ricci v. DeStefano, 557 U.S. 557, 590 (2009).

Finally, an employee could suggest an alternative monitoring system, one that does not produce such disparate results.[131] However, once again, this will prove challenging because of the strict "equally valid" standard courts impose.[132] In determining this standard, courts consider factors such as ability to meet employer's needs, cost, and burden on the employer.[133] In order to equally meet an employer's needs in a monitoring context, the alternative monitoring program would have to obtain substantially similar results to the original program[134]—this could be number of threats detected, ability to determine unproductive employees, or other goals by the company.[135] The proposed alternative would also need to not be prohibitively costly to purchase or to implement.[136] Finally, the alternative could not impose other burdens on the employer, meaning, in the context of employment monitoring tools, network compatibility or other technical factors that make the tool difficult to operate.[137] For these reasons, an employee would likely have a difficult time in coming up with an equally viable monitoring practice.

2. Demonstration of a Less Adverse Impact

Should an employee come up with a viable alternative monitoring practice, the employee would next need to demonstrate that this practice would actually result in a less adverse impact.[138] Again, this would be difficult to establish, because as the Court said in *Ricci v. DeStefano*, "isolated statements" by experts claiming a practice would be less discriminatory are not enough to meet this requirement.[139] There has been no case law on alternative employment practices and employment monitoring so it is difficult to posit how an employee would successfully meet this prong; perhaps if the

131. *See id.* at 578.

132. *See id.* at 589.

133. *See* Watson v. Fort Worth Bank & Trust, 487 U.S. 977, 998 (1988).

134. *See id.*; *see also* Ass'n of Mexican-Am. Educators v. California, 937 F. Supp. 1397, 1427-28 (N.D. Cal. 1996) (holding alternative competency requirements, such as GPA and coursework, were not an adequate substitution for a hiring exam).

135. *See supra* Part I.

136. *See* Ernest F. Lidge III, *Financial Costs as a Defense to an Employment Discrimination Claim*, 58 ARK. L. REV. 1, 32-38 (2005).

137. *See Watson*, 487 U.S. at 998.

138. *See* Ricci v. DeStefano, 557 U.S. 557, 591 (2009).

139. *Id.* at 591-92.

company had tried different tools and recorded results from those tests, then that would be enough. However, the employee would have to perform a significant amount of analysis in order to prove to the court that the alternative would have actually resulted in less discrimination.

3. Employer Refused to Adopt the Practice

Finally, an employee would need to demonstrate that the employer in fact knew about the alternative and refused to adopt it anyway.[140] Courts have not yet firmly established what the law means by "refuses"; however, based on the clear meaning of the text and decisions thus far, it would appear to mean that the employer was at least aware of the alternative to some extent.[141] Therefore, the employee would need to establish that the employer knew, at least to some extent, about the alternative practice.

Given courts' strict interpretation of the alternative employment practices test, an employee will likely be unable to meet this burden even after meeting the initial prima facie burden for disparate impact discrimination. Although, as this analysis shows, it is not *impossible* to meet this burden, the extreme difficulty an employee would likely face in meeting this burden should give pause. If Title VII does not do a sufficient job in preventing discrimination by way of algorithms, are companies free to do what they please with their monitoring programs regardless of the discriminatory effect? How should courts handle Title VII cases when it appears machine learning tools are doing the discrimination, particularly in light of the challenges highlighted in employee monitoring cases?

IV. ADOPTING A NEGLIGENCE STANDARD

Given the difficulty for employees in making successful Title VII claims, this Note argues that Congress should consider amending Title VII to give employees a better chance of pursuing an effective discrimination claim. This Part will discuss three proposed options:

140. *See id.* at 578.

141. § 2000e-2(k)(1)(A)(ii); *see* Jones v. City of Boston, 845 F.3d 28, 36-38 (1st Cir. 2016).

anti-discriminatory algorithms, regulations around machine learning, and a negligent use of technology standard, ultimately concluding that a negligent use of technology standard is the best option for both employers and employees.

A. Alternative Solutions

Some scholars have proposed anti-discriminatory algorithms as a potential solution to the problem of discriminatory algorithms in employment.[142] Under this solution, employers would not only prevent algorithms from discriminating,[143] but would use them "in a way that actually improves upon current human decision-making—to make them affirmatively *anti*discriminatory."[144] These algorithms would be trained to identify both stereotypes and biases, as well as to suppress potential biased inputs in the data set.[145] So for example, say an algorithm learned that visiting a particular site—perhaps something innocuous like a sports site—made it more likely that an employee would pose a threat to a company and thus weighted that information in its algorithm. In this scenario, an antidiscriminatory algorithm may be able to figure out that the monitoring algorithm was actually using the sports site as a proxy for men and discriminating against a protected class inadvertently: the algorithm would react by suppressing that feature in the monitoring program in order to prevent a discriminatory outcome.[146] With these anti-discriminatory algorithms in place, scholars suggest framing algorithmic discrimination as a disparate treatment violation.[147]

While this would appear helpful in curing discriminatory issues, it appears to fall short in two respects. First, this approach would necessitate using the protected classes in at least some fashion in order to train the data in the first place. Although in effect this use

142. *See, e.g.*, Bornstein, *supra* note 79, at 550-551.

143. *See, e.g.*, JON KLEINBERG, JENS LUDWIG, SENDHIL MULLAINATHAN & CASS R. SUNSTEIN, ALGORITHMS AS DISCRIMINATION DETECTORS 1 (Proc. Nat'l Acad. Scis. U.S. 2020), https://www.pnas.org/content/pnas/early/2020/07/27/1912790117.full.pdf [https://perma.cc/YS7G-GS2E].

144. Bornstein, *supra* note 79, at 550.

145. *See id.* at 552.

146. *See id.*; Barocas & Selbst, *supra* note 26, at 712.

147. *See* Bornstein, *supra* note 79, at 520.

would appear to lessen discrimination,[148] the Supreme Court has rejected similar remedial remedies in *Ricci*, where it held that the employer's refusal to certify examination results due to its belief in disparate impact was in itself a Title VII violation.[149] While this holding alone does not necessarily preclude using anti-discriminatory algorithms, it is important to note that commentators have viewed this holding as representative of a potential shift in the Court toward an anti-classification view of discrimination.[150] Under this view, the best way to reduce differential treatment is to stop using identity classifications in the first place.[151] This contrasts with the anti-subordination view, which instead focuses on undermining racial hierarchies and the effects of such discrimination.[152] Thus, even if anti-discriminatory algorithms are well-intentioned, the use of protected classes—even to combat discrimination—may well be considered suspect by the Court in the future.

Furthermore, while commentators have discussed using anti-classification principles in the context of hiring,[153] it is unclear whether it would be equally as beneficial in monitoring. In a perfect scenario, the only function of anti-discriminatory algorithms would be to reduce discrimination; however, this is not such a simple task in employee monitoring.[154] In the real world, particularly on a company level, there is likely to be some naturally occurring disparate impact in monitoring results.[155] That is, it is likely that in an individual company a particular group may have a higher violation rate than other groups.[156] In fact, it would be rather odd, statistically-speaking, if the number of violations was perfectly even among groups.[157] Thus, by trying not to discriminate, these anti-discriminatory algorithms may have the potential to suppress real trends and indicators, thereby making the results less accurate. In the hiring

148. *See id.* at 550-51.

149. 557 U.S. 557, 590-93 (2009).

150. *See* Helen Norton, *The Supreme Court's Post-Racial Turn Towards a Zero-Sum Understanding of Equality*, 52 WM. & MARY L. REV. 197, 206-07, 224-28 (2010).

151. *See id.* at 207.

152. *Id.* at 206.

153. *See* Bornstein, *supra* note 79, at 520.

154. *See id.* at 550.

155. *See, e.g.*, Dastin, *supra* note 41.

156. *See* Bornstein, *supra* note 79, at 551-52.

157. *See* Barocas & Selbst, *supra* note 26, at 673-74.

context, this slight inaccuracy, when considering what traits make a "good" employee, is acceptable because it leads to a desired side effect—diversity in the workplace.[158] In employee monitoring there is no such secondary benefit to this inaccuracy; in fact, such a suppression of potentially useful indicators may lead to significant harm.

Scholars have also suggested ways in which employers may improve the algorithms themselves.[159] However, relying on employers to police themselves in this endeavor may prove futile without external financial pressure.[160] Moreover, implementing these technical computer science recommendations into workable regulations is likely to prove challenging both for regulators and for employers who would be forced to strictly comply.[161] Imposing such costs on a new and burgeoning technology appears unlikely to encourage better and more accurate algorithms.

B. Negligent Use of Technology Standard

That leads to the final standard—a negligent use of technology standard imposed on Title VII. While not a perfect solution in terms of fixing algorithmic discrimination, it provides a workable alternative to the current statutory scheme and would likely not be as cumbersome for regulators to adopt.

David Oppenheimer originally articulated a similar theory for Title VII, though he wrote his article before concerns about machine learning and big data came to the fore.[162] Oppenheimer's approach proposed that employers should be liable under Title VII "when the

158. *See* Bornstein, *supra* note 79, at 551.

159. *See* Barocas & Selbst, *supra* note 26, at 731-32.

160. Competition and efficiency concerns are unlikely to eliminate discriminatory models:

> [D]ata models are more likely to exhibit bias, and market competition will not reliably eliminate them. First, biased data models may be accurate *enough* to persist in a competitive market, even though they are biased against certain groups. Second, feedback effects may appear to confirm the accuracy of biased data models, entrenching their use. And finally, biased data models may be efficient precisely *because* they are discriminatory, and therefore pressures toward efficiency will not eliminate them.

Kim, *supra* note 90, at 894.

161. *See* Barocas & Selbst, *supra* note 26, at 714-22.

162. *See generally* David Benjamin Oppenheimer, *Negligent Discrimination*, 141 U. PA. L. REV. 899 (1993).

employer fails to take all reasonable steps to prevent discrimination that it knows or should know is occurring, or that it expects or should expect to occur."[163] However, this Note proposes that this standard—while useful—should be modified slightly in order to better address problems with machine learning and monitoring in particular.

Rather than imposing a traditional negligence standard related to the employer's knowledge, which may tend to blur the lines between disparate impact and treatment entirely,[164] this Note instead proposes a negligent use of technology standard. This standard would fall under disparate impact analysis and make employers liable when they negligently use machine learning and contribute to a disparate impact on a particular class of employees.[165]

To fit this test into a workable framework for the courts, it should be situated within the alternative employment practices prong of Title VII.[166] This solution would give employees more options in overcoming this prong: they could either raise an alternative employment practice or a negligent use claim contending that proper use and implementation of the machine learning technology likely would have prevented such disparate impact, and thus the employer violated Title VII, regardless of the company's business necessity in having the monitoring program.[167]

A negligence standard would likely not be as precise in reducing disparate impact as anti-discriminatory algorithms and a strict rule against a showing of disparate impact.[168] However, it would provide a more workable standard for companies because rather than worrying and adjusting their algorithms in order to ensure *no* disparate impact, they could instead focus on following industry-best practices and modifying their algorithms as technology progressed.[169] Furthermore, because the standard does not require using identity—either in the algorithms or in the testing and evaluation process—it would be less likely to face opposition from anti-classification

163. *Id.* at 900.

164. *See supra* Part II (discussing disparate treatment).

165. *See supra* Part III (describing disparate impact analysis).

166. *See* 42 U.S.C. § 2000e-2(k)(1)(A)(ii); *supra* Part III.C.

167. *See supra* Part III.C.

168. *See supra* Part IV.A.

169. *See supra* Part IV.A. For a summary of potential algorithmic changes an employer could consider, see generally Barocas & Selbst, *supra* note 26.

proponents in the courts.[170] The negligence standard also gives more flexibility to courts and regulators than per se regulations prohibiting and requiring certain conduct, allowing regulators and courts to adapt to rapid shifts in technology.[171]

Regulators such as the EEOC and the Federal Trade Commission (FTC) could issue guidelines for employers and courts to use in determining what negligent conduct looks like.[172] The EEOC and the FTC have already started broadly discussing machine learning and employer best practices,[173] indicating that more specific technical guidance may not be too far off.

On its face, this may appear a seemingly simple solution. However, putting courts in charge of determining what constitutes "negligence" in machine learning software and implementation would have to be an ongoing, iterative process as technology progresses. Current regulatory guidance agrees on some basic practices, such as regular model testing and protecting the algorithm from misuse.[174] However, when it comes to more advanced steps, it may be difficult for courts to determine what counts as "negligent" in the industry, particularly in a fast-changing field where companies oftentimes refuse to share their proprietary algorithms.[175] While these practices pose challenges to developing industry norms, some companies, such as IBM with its AI Fairness 360 toolkit, are open-sourcing their best practices and code.[176] Similar guidance and tools from major players

---

170. *See* Norton, *supra* note 150, at 206-07, 224-28.

171. *See* Barocas & Selbst, *supra* note 26, at 714-22; Darrell M. West & John R. Allen, *How Artificial Intelligence Is Transforming the World*, BROOKINGS (Apr. 24, 2018), https://www.brookings.edu/research/how-artificial-intelligence-is-transforming-the-world/ [https://perma.cc/VA3D-CSDA].

172. The EEOC has already issued guidelines for Title VII, which are non-binding on the courts but do provide a helpful guide. *See, e.g.*, Clady v. County of Los Angeles, 770 F.2d 1421, 1428 (9th Cir. 1985).

173. *See* Andrew Smith, *Using Artificial Intelligence and Algorithms*, FED. TRADE COMM'N (Apr. 8, 2020, 9:58 AM), https://www.ftc.gov/news-events/blogs/business-blog/2020/04/using-artificial-intelligence-algorithms [https://perma.cc/E2MA-JUAA]; EQUAL EMP. OPPORTUNITY COMM'N, MEETING OF OCTOBER 13, 2016—BIG DATA IN THE WORKPLACE: EXAMINING IMPLICATIONS FOR EQUAL EMPLOYMENT OPPORTUNITY LAW 4-5 (2016), https://www.eeoc.gov/meetings/24068/transcript [https://perma.cc/7LHT-T9FL].

174. *See* Smith, *supra* note 173.

175. *See* Andrew D. Selbst, *Negligence and AI's Human Users*, 100 B.U. L. Rev. 1315, 1363-65 (2020).

176. Kush R. Varshney, *Introducing AI Fairness 360*, IBM RSCH. BLOG (Sept. 19, 2018), https://www.ibm.com/blogs/research/2018/09/ai-fairness-360/ [https://perma.cc/472P-P2B9].

in the industry could soon become workable norms for courts to consider when evaluating whether an employer negligently employed its monitoring tools.

Although this change is imperfect, it would be a step in the right direction and would give employees another avenue to challenge negligent and unfair practices as employers adopt new monitoring technologies in the future.[177]

CONCLUSION

As demonstrated in this Note, Title VII, as courts currently apply the law, appears unlikely to be able to protect employees from discriminatory monitoring practices when employers use machine learning tools in monitoring programs. This significant gap in the current anti-discrimination framework should concern both employees and employers alike.[178] Implementing a negligence standard within the Title VII framework would give employees another opportunity to recover against employers who use discriminatory monitoring practices.

177. *See* Christina Pazzanese, *Great Promise but Potential for Peril*, HARV. GAZETTE (Oct. 26, 2020), https://news.harvard.edu/gazette/story/2020/10/ethical-concerns-mount-as-ai-takes-bigger-decision-making-role/ [https://perma.cc/4CEW-UNCN].

178. *See* Charlotte S. Alexander & Elizabeth Tippett, *The Hacking of Employment Law*, 82 MO. L. REV. 973, 973 (2017) ("[W]hen employers use software to avoid the employer-employee relationship entirely, employment law is weakened as more workers operate in spaces beyond the law's reach.").