

MANIPULATING OPPORTUNITY

*Pauline T. Kim**

Concerns about online manipulation have centered on fears about undermining the autonomy of consumers and citizens. What has been overlooked is the risk that the same techniques of personalizing information online can also threaten equality. When predictive algorithms are used to allocate information about opportunities like employment, housing, and credit, they can reproduce past patterns of discrimination and exclusion in these markets. This Article explores these issues by focusing on the labor market, which is increasingly dominated by tech intermediaries. These platforms rely on predictive algorithms to distribute information about job openings, match job seekers with hiring firms, or recruit passive candidates. Because algorithms are built by analyzing data about past behavior, their predictions about who will make a good match for which jobs will likely reflect existing occupational segregation and inequality. When tech intermediaries cause discriminatory effects, they may be liable under Title VII, and Section 230 of the Communications Decency Act should not bar such actions. However, because of the practical challenges that litigants face in identifying and proving liability retrospectively, a more effective approach to preventing discriminatory effects should focus on regulatory oversight to ensure the fairness of algorithmic systems.

I. INTRODUCTION..... 868
 II. PLATFORMS AS LABOR MARKET INTERMEDIARIES..... 875
 A. Recruitment Practices 875
 B. Discriminatory Choices..... 881
 C. Predictive Analytics and Discriminatory Effects 885
 D. Manipulating Opportunity..... 892

* Daniel Noyes Kirby Professor of Law, Washington University School of Law, St. Louis, Missouri. I am grateful to Victoria Schwarz, Miranda Bogen, Aaron Riecke, Greg Magarian, Neil Richards, Peggie Smith, Dan Epps, John Inazu, Danielle Citron, Ryan Calo, Andrew Selbst, Margot Kaminski, and Felix Wu for helpful comments on earlier drafts of this Article. I also benefited from feedback from participants at the 2019 Privacy Law Scholar’s Conference, Washington University School of Law’s faculty workshop, and Texas A&M School of Law’s Faculty Speaker Series. Many thanks to Adam Hall, Theanne Liu, Joseph Tomchak, and Samuel Levy for outstanding research assistance.

<i>E. Possible Objections</i>	898
III. INTERMEDIARIES AND ANTI-DISCRIMINATION LAW	905
<i>A. Who is an Applicant?</i>	907
<i>B. Platforms as Employment Agencies</i>	911
<i>C. Interference with Employment Opportunities</i>	917
IV. PLATFORM RESPONSIBILITY	919
<i>A. Avoiding the CDA Section 230 Defense</i>	919
<i>B. Practical Obstacles and Alternative Strategies</i>	929
V. CONCLUSION	934

I. INTRODUCTION

Our online experiences are increasingly personalized. Facebook and Google micro-target advertisements aimed to meet our immediate needs. Amazon, Netflix, and Spotify offer up books, movies, and music tailored to match our tastes. Our news feeds are populated with stories intended to appeal to our particular interests and biases. This drive toward increasing personalization is powered by complex machine learning algorithms built to discern our preferences and anticipate our behavior. Personalization offers benefits because companies can efficiently offer consumers the precise products and services they desire.

Online personalization, however, has come under considerable criticism lately. Shoshana Zuboff assails our current economic system, which is built on companies amassing and exploiting ever more detailed personal information.¹ Ryan Calo and Tal Zarsky explain that firms are applying the insights of behavioral science to manipulate consumers by exploiting their psychological or emotional vulnerabilities.² Daniel Susser, Beate Roessler, and Helen Nissenbaum describe how information technology is enabling manipulative practices on a massive scale.³ Julie Cohen similarly argues that “[p]latform-based, massively-intermediated processes of search and social networking are inherently processes of

¹ Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* 8–11 (2019).

² See Ryan Calo, *Digital Market Manipulation*, 82 *Geo. Wash. L. Rev.* 995, 996, 999 (2014); Tal Z. Zarsky, *Privacy and Manipulation in the Digital Age*, 20 *Theoretical Inquiries L.* 157, 158, 160–61 (2019).

³ Daniel Susser, Beate Roessler & Helen Nissenbaum, *Online Manipulation: Hidden Influences in a Digital World*, 4 *Geo. L. Tech. Rev.* 1, 2, 10 (2019).

market manipulation.”⁴ In the political sphere as well, concerns have been raised about manipulation, with warnings that news personalization is creating “filter bubble[s]” and increasing polarization.⁵ These issues were highlighted by revelations that Cambridge Analytica sent personalized ads based on psychological profiles of eighty-seven million Facebook users in an effort to influence the 2016 presidential election.⁶ The extensive criticism of personalization is driven by concerns that online manipulation undermines personal autonomy and compromises rational decision making.

Largely overlooked in these discussions is the possibility that online manipulation also threatens equality. Online platforms increasingly operate as key intermediaries in the markets for employment, housing, and financial services—what I refer to as opportunity markets. Predictive algorithms are also used in these markets to segment the audience and determine precisely what information will be delivered to which users. The risk is that in doing so, these intermediaries will direct opportunities in ways that reproduce or reinforce historical forms of discrimination. Predictive algorithms are built by observing past patterns of behavior, and one of the enduring patterns in American economic life is the unequal

⁴ Julie E. Cohen, *Law for the Platform Economy*, 51 U.C. Davis L. Rev. 133, 165 (2017); see also Julie E. Cohen, *Between Truth and Power: The Legal Constructions of Informational Capitalism* 75–77, 83–89, 96 (2019) (describing how techniques for behavioral surveillance and micro-targeting contribute to social harms such as polarization and extremism).

⁵ See, e.g., Eli Pariser, *The Filter Bubble: What the Internet Is Hiding from You* 13–14 (2011); Michael J. Abramowitz, *Stop the Manipulation of Democracy Online*, N.Y. Times (Dec. 11, 2017), <https://www.nytimes.com/2017/12/11/opinion/fake-news-russia-kenya.html> [<https://perma.cc/9YWF-PED7>]; James Doubek, *How Disinformation and Distortions on Social Media Affected Elections Worldwide*, NPR (Nov. 16, 2017, 2:28 PM), <https://www.npr.org/sections/alltechconsidered/2017/11/16/564542100/how-disinformation-and-distortions-on-social-media-affected-elections-worldwide> [<https://perma.cc/ZJ97-GQSZ>]; Jon Keegan, *Blue Feed, Red Feed: See Liberal Facebook and Conservative Facebook, Side by Side*, Wall St. J. (Aug. 19, 2019), <http://graphics.wsj.com/blue-feed-red-feed/> [<https://perma.cc/GJA8-4U9W>].

⁶ Carole Cadwalladr & Emma Graham-Harrison, *Revealed: 50 Million Facebook Profiles Harvested for Cambridge Analytica in Major Data Breach*, Guardian (Mar. 17, 2018, 6:03 PM), <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election> [<https://perma.cc/72CR-9Y8K>]; Alex Hern, *Cambridge Analytica: How Did It Turn Clicks into Votes?*, Guardian (May 6, 2018, 3:00 AM), <https://www.theguardian.com/news/2018/may/06/cambridge-analytica-how-turn-clicks-into-votes-christopher-wylie> [<https://perma.cc/AD8H-PF3M>]; Matthew Rosenberg, Nicholas Confessore & Carole Cadwalladr, *How Trump Consultants Exploited the Facebook Data of Millions*, N.Y. Times (Mar. 17, 2018), <https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html> [<https://perma.cc/3WYQ-3YKP>].

distribution of opportunities along the lines of race, gender, and other personal characteristics. As a result, these systems are likely to distribute information about future opportunities in ways that reflect existing inequalities and may reinforce historical patterns of disadvantage.

The way in which information about opportunities is distributed matters, because these markets provide access to resources that are critical for human flourishing and well-being. In that sense, access to them is foundational. People need jobs and housing before they can act as consumers or voters. They need access to financial services in order to function in the modern economy. Of course, many other factors contribute to inequality, such as unequal educational resources, lack of access to health care, and over-policing in certain communities. Decisions by landlords, employers, or banks can also contribute to inequality. Tech intermediaries are thus just one part of a much larger picture. Nevertheless, they will be an increasingly important part as more and more transactions are mediated online.⁷ Because they control access to information about opportunities, they have the potential to significantly impact how these markets operate.

Online intermediaries have unprecedented potential to finely calibrate the distribution of information. In the past, traditional print or broadcast media might aim at a particular audience, but they could not prevent any particular individual from accessing information that they published. And if an advertiser tried to signal its interest in only a particular group—as has happened with real estate ads that used code words or featured only white models—the attempts at exclusion were plainly visible. In contrast, online intermediaries have the ability to precisely target an audience, selecting some users to receive information and others to be excluded in ways that are not at all transparent.

The issue is illustrated by Facebook's ad-targeting tools. Several lawsuits alleged that employers or landlords could use the company's tools to exclude users on the basis of race, gender, or age from their audience.⁸ To a large extent, these concerns were resolved by a recent settlement in which Facebook agreed to bar the use of sensitive

⁷ See, e.g., Miranda Bogen & Aaron Rieke, *Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias* 5–6 (2018) (describing the role of platforms in the hiring process); Geoff Boeing, *Online Rental Housing Market Representation and the Digital Reproduction of Urban Inequality*, 52 *Env't & Plan. A* 449, 450 (2019) (documenting the growing impact of Internet platforms in shaping the rental housing market).

⁸ See *infra* Section II.B.

demographic variables to target employment, housing, and credit advertisements.⁹ However, the settlement failed to address another potential source of bias—Facebook’s ad-delivery algorithm, which determines which users within a targeted audience actually receive an ad. As explained below, *even if* an advertiser uses neutral targeting criteria and intends to reach a diverse audience, an ad-targeting algorithm may distribute information about opportunities in a biased way.¹⁰ This is an example of a much broader concern—namely, that when predictive algorithms are used to allocate access to opportunities, there is a significant risk that they will do so in a way that reproduces existing patterns of inequality and disadvantage.

Concerns about the distributive effects of predictive algorithms are relevant to all kinds of opportunity markets, including for housing, employment, and basic financial services. Each of these markets operates somewhat differently and is regulated under different laws. They deserve separate attention and more detailed consideration than can be provided here. This Article focuses on the labor market and the relevant laws regulating it; however, the issues it raises likely plague other opportunity markets as well.

Examining employment practices reveals dramatic change. Just a couple of decades ago, employers had a handful of available strategies for recruiting new workers, such as advertising in newspapers or hiring through an employment agency. Today, firms increasingly rely on tech intermediaries to fill job openings.¹¹ Recent surveys suggest that somewhere from 84% to 93% of job recruiters use online strategies to find potential employees.¹² Employers distribute information about positions

⁹ See Galen Sherwin & Esha Bhandari, Facebook Settles Civil Rights Cases by Making Sweeping Changes to Its Online Ad Platform, ACLU (Mar. 19, 2019, 2:00 PM), <https://www.aclu.org/blog/womens-rights/womens-rights-workplace/facebook-settles-civil-rights-cases-making-sweeping> [<https://perma.cc/H6D6-UMJ4>].

¹⁰ See *infra* Section II.C.

¹¹ See Bogen & Rieke, *supra* note 7, at 5–6.

¹² Soc’y for Human Res. Mgmt., SHRM Survey Findings: Using Social Media for Talent Acquisition—Recruitment and Screening 3 (Jan. 7, 2016), <https://www.shrm.org/hr-today/trends-and-forecasting/research-and-surveys/Documents/SHRM-Social-Media-Recruiting-Screening-2015.pdf> [<https://perma.cc/L6NT-N4KL>]. The Society for Human Resource Management conducts biennial surveys of job recruiters. The surveys demonstrated an increase in the use of online recruiting by employers, rising from fifty-six percent in 2011 to seventy-seven percent in 2013 to eighty-four percent in 2015. *Id.*; Soc’y for Human Res. Mgmt., SHRM Survey Findings: Social Networking Websites and Recruiting/Selection 2 (Apr. 11, 2013), <https://www.shrm.org/hr-today/trends-and-forecasting/research-and-surveys/Pages/shrm-social-networking-websites-recruiting-job-candidates.aspx>

through social media. They also rely on specialized job platforms like ZipRecruiter, LinkedIn, and Monster to recruit applicants and recommend the strongest candidates.¹³ In addition, passive recruiting—using data to identify workers who are not actively looking for another position—is a growing strategy for recruiting new talent.¹⁴

The use of algorithms and artificial intelligence in the hiring process has not gone unnoticed. Numerous commenters and scholars have described how employers are using automated decision systems and have raised concerns that these developments may cause discrimination or threaten employee privacy.¹⁵ However, previous work has focused on whether *employers* can or should be held liable when they use predictive algorithms or other artificial intelligence tools to make personnel decisions. What is missing from this literature is close scrutiny of how *tech intermediaries* are shaping labor markets and the implications for equality.

This Article undertakes that analysis, arguing that the use of predictive algorithms by labor market intermediaries risks reinforcing or even worsening existing patterns of inequality and that these intermediaries

[<https://perma.cc/U4HN-E7U7>]; see also Jobvite's New 2015 Recruiter Nation Survey Reveals Talent Crunch, Jobvite (Sept. 22, 2015), https://www.jobvite.com/news_item/-jobvites-new-2015-recruiter-nation-survey-reveals-talent-crunch-95-recruiters-anticipate-similar-increased-competition-skilled-workers-coming-year-86-expect-exp/ [<https://perma.cc/H66S-8E5Z>] (stating that 92% of recruiters use social media to discover or evaluate candidates).

¹³ See Bogen & Rieke, *supra* note 7, at 5, 19–20, 24.

¹⁴ *Id.* at 22.

¹⁵ See, e.g., Ifeoma Ajunwa, Kate Crawford & Jason Schultz, *Limitless Worker Surveillance*, 105 *Calif. L. Rev.* 735, 738–39 (2017); Ifeoma Ajunwa, *The Paradox of Automation as Anti-Bias Intervention*, 41 *Cardozo L. Rev.* (forthcoming 2020) (manuscript at 14) (on file with author); Richard A. Bales & Katherine V.W. Stone, *The Invisible Web of Work: Artificial Intelligence and Electronic Surveillance in the Workplace*, 41 *Berkeley J. Lab. & Emp. L.* (forthcoming 2020) (manuscript at 3) (on file with author); Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 *Calif. L. Rev.* 671, 673–75 (2016); Matthew T. Bodie, Miriam A. Cherry, Marcia L. McCormick & Jintong Tang, *The Law and Policy of People Analytics*, 88 *U. Colo. L. Rev.* 961, 989–92 (2017); James Grimmelmann & Daniel Westreich, *Incomprehensible Discrimination*, 7 *Calif. L. Rev. Online* 164, 170–72, 176–77 (2017); Jeffrey M. Hirsch, *Future Work*, 2020 *U. Ill. L. Rev.* (forthcoming 2020) (manuscript at 3) (on file with author); Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 *Wm. & Mary L. Rev.* 857, 860–61 (2017) [hereinafter Kim, *Data-Driven Discrimination at Work*]; Pauline T. Kim, *Data Mining and the Challenges of Protecting Employee Privacy Under U.S. Law*, 40 *Comp. Lab. L. & Pol'y J.* 405, 406 (2019); Pauline T. Kim & Erika Hanson, *People Analytics and the Regulation of Information Under the Fair Credit Reporting Act*, 61 *St. Louis U. L.J.* 17, 18–19 (2016); Charles A. Sullivan, *Employing AI*, 63 *Vill. L. Rev.* 395, 396 (2018).

should be accountable for those effects. A number of studies have documented instances of biased delivery of employment ads.¹⁶ Although the exact mechanism is unclear, it should not be surprising that predictive algorithms distribute information about job opportunities in biased ways. These algorithms are built by analyzing existing data, and one of the most persistent facts of the U.S. labor market is ongoing occupational segregation along the lines of race and gender.¹⁷ If predictions are based solely on observations about past behavior—without regard to what social forces shaped that behavior—then they are likely to reproduce those patterns.

Tech intermediaries may not intend to cause discriminatory effects, but they are nevertheless responsible for them.¹⁸ They make choices when designing the algorithms that distribute information about job opportunities or suggest the best matches for job seekers and hiring firms. In doing so, they decide what goals to optimize—typically revenue—and those choices influence how information is channeled, making some opportunities visible and obscuring others. Thus, these technologies shape how the market participants—both workers and employers—perceive their available options and thereby also influence their behavior.¹⁹ When these intermediaries structure access to opportunities in ways that reflect

¹⁶ See *infra* Section II.C.

¹⁷ See *infra* Section II.D.

¹⁸ Building predictive models involves numerous choices, many of them implicating value judgments. See, e.g., Barocas & Selbst, *supra* note 15, at 674; Margot E. Kaminski, Binary Governance: Lessons from the GDPR's Approach to Algorithmic Accountability, 92 S. Cal. L. Rev. 1529, 1539 (2019); David Lehr & Paul Ohm, Playing with the Data: What Legal Scholars Should Learn About Machine Learning, 51 U.C. Davis L. Rev. 653, 703–04 (2017); Andrew D. Selbst & Solon Barocas, The Intuitive Appeal of Explainable Machines, 87 Fordham L. Rev. 1085, 1130–31 (2018).

¹⁹ Karen Levy and Solon Barocas have explored how the design choices made by platforms “can both mitigate and aggravate bias.” Karen Levy & Solon Barocas, Designing Against Discrimination in Online Markets, 32 Berkeley Tech. L.J. 1183, 1185 (2017). The focus of their analysis is on *user* bias in online markets like ride matching, consumer-to-consumer sales, short-term rentals, and dating. *Id.* at 1189–90. Because the design choices platforms make will structure users’ interactions with one another, these choices influence behavior, affecting whether or to what extent users can act on explicit or implicit biases. Levy and Barocas review multiple platforms across domains and develop a taxonomy of policy and design elements that have been used to address the risks of bias. Although the focus of this Article is on the impact of predictive algorithms rather than user bias, the issues are obviously interrelated. Past bias by users can cause predictive algorithms to discriminate. Conversely, algorithmic outputs in the form of recommendations or rankings can activate or exacerbate implicit user biases. To that extent, some, but not all, of the strategies they identify may be relevant to addressing bias in online opportunity markets.

historical patterns of discrimination and exclusion, they pose a threat to workplace equality. Even if the discriminatory effects are unintentional, the harm to workers can be real. Employment discrimination law has long targeted discriminatory effects, not just invidious motivation.²⁰

The risk that tech intermediaries will contribute to workplace inequality poses a number of challenges for the law. Discrimination law has largely focused on employers, examining their decisions and practices for discriminatory intent or impact. However, if bias affects how potential applicants are screened out before they even interact with a hiring firm, then focusing on employer behavior will be inadequate to dismantle patterns of occupational segregation. Holding tech intermediaries directly responsible for their effects on labor markets, however, will raise a different set of challenges. Some of these are legal, such as whether existing law reaches these types of intermediaries,²¹ and whether they can avoid liability by relying on Section 230 of the Communications Decency Act (CDA),²² which gives websites a defense to some types of liability. Other obstacles are more practical in nature, which suggests that preventing discriminatory effects may require alternative strategies.²³

This Article proceeds as follows. Part II first explores the role that tech intermediaries play in the labor market and how targeting tools can be misused for discriminatory purposes. It next explains that *even if* employers are no longer permitted to use discriminatory targeting criteria, a significant risk remains that platforms' predictive algorithms will distribute access to opportunities in ways that reproduce existing patterns of inequality. Because tech intermediaries have a great deal of power to influence labor market interactions, and may do so in ways that are not transparent, I argue in Part II that they should bear responsibility when they cause discriminatory effects.

Parts III and IV consider the relevant legal landscape. Part III discusses how the growing importance of tech intermediaries in the labor market poses challenges for existing anti-discrimination law. It first shows how the question "who is an applicant?"—an issue critical for finding employer liability—is complicated as platforms increasingly mediate job seekers' interactions with firms. It then explores the possibilities for holding these intermediaries directly liable under existing employment

²⁰ See *Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971).

²¹ See *infra* Part III.

²² 47 U.S.C. § 230 (2012).

²³ See *infra* Section IV.B.

discrimination law, either as employment agencies or for interfering with third party employment relationships. Part IV considers some obstacles to holding tech intermediaries liable for their discriminatory labor market effects. Section IV.A examines and rejects the argument that Section 230 of the Communications Decency Act would automatically bar such claims. Section IV.B explains that significant practical obstacles remain, suggesting that a post hoc liability regime may not be the best way to prevent discriminatory harms. Thus, Section IV.B also argues that we should look to regulatory models in order to minimize the risks of discrimination from the use of predictive algorithms.

II. PLATFORMS AS LABOR MARKET INTERMEDIARIES

A. Recruitment Practices

The process by which employers recruit and hire new workers is undergoing a transformational shift. A few decades ago, employers had just a handful of strategies for making new hires.²⁴ They could utilize informal methods, such as relying on their current employees to spread the word about job openings. Alternatively, they could engage more formal mechanisms, such as publicizing opportunities at job fairs or in help wanted advertisements in newspapers. Apart from newspapers and other general media, the only significant intermediaries in the hiring process were unions, joint labor-management apprenticeship programs, which trained and referred skilled craft workers, and employment agencies, which referred workers to employers.

Depending upon the circumstances, each of these methods might alleviate or contribute to race- and sex-segregated workplaces. In a homogeneous workplace, relying on word-of-mouth hiring effectively excluded workers from different racial or ethnic backgrounds. Similarly, intermediaries like labor unions or employment agencies sometimes operated in ways that reinforced existing occupational segregation. And before the Civil Rights Act of 1964, newspaper advertisements often announced employers' racial, ethnic, or sex-based preferences. Once discriminatory advertising ceased, more formal methods of advertising job openings through broadly distributed media had the potential for reducing occupational segregation and opening opportunities for women

²⁴ See F. Carson Mencken & Idee Winfield, *Employer Recruiting and the Gender Composition of Jobs*, 32 *Soc. Focus* 201, 202 (1999).

and racial minorities. By publishing job openings, information was available to a much broader audience, and members of previously excluded groups could learn about those opportunities without being part of the social networks of current employees. Studies confirmed that more formal recruitment methods tended to be correlated with reduced occupational segregation.²⁵

The labor market today is increasingly mediated by technology platforms. Job seekers go online not only to apply for jobs but also to first find out what jobs exist. Employers advertise available positions through social media and other online platforms like Google and Facebook. Both job seekers and employers utilize platforms to help them efficiently find the best opportunities. These technology-mediated practices are not merely updates of traditional forms of recruitment and hiring.²⁶ Rather, they are transforming labor markets in ways that are poorly understood and may have profound implications for workplace equality.

Tech-mediated recruiting might appear to be much like the old help wanted ads in a newspaper, except with the possibility of reaching a much larger audience. In fact, however, the operation of these tools is profoundly different in ways that likely have distributive consequences. Unlike traditional media, which framed job ads generically and made them visible to all readers, tech intermediaries distribute information about job opportunities in a highly structured way—targeting certain recipients and excluding others. Just like ads for consumer goods, job postings are targeted at specific users based on inferences made about them through predictive algorithms.²⁷ Similarly, job-matching and recruiting tools leverage large amounts of personal data to identify and recruit the most promising candidates, rather than broadly disseminating

²⁵ See, e.g., *id.* at 211 (finding that jobs that were recruited through advertising rather than informal networks were more likely to have more women).

²⁶ See Ofer Sharone, *LinkedIn or LinkedOut? How Social Networking Sites Are Reshaping the Labor Market*, in 30 *Research in the Sociology of Work: Emerging Conceptions of Work, Management and the Labor Market* 1, 3 (Steven Vallas ed., 2017) (examining the social effects of job seekers' use of social networking websites to find employment). See generally Ifeoma Ajunwa & Daniel Greene, *Platforms at Work: Automated Hiring Platforms and Other New Intermediaries in the Organization of Work*, in 33 *Research in the Sociology of Work: Work and Labor in the Digital Age* 61, 81 (Steven P. Vallas & Anne Kovalainen eds., 2019) (analyzing how automated hiring platforms impact the terms on which workers enter, live within, or exit workplaces, worsen information asymmetries, and alter power dynamics, and suggesting important directions for future empirical work).

²⁷ Pauline T. Kim & Sharion Scott, *Discrimination in Online Employment Recruiting*, 63 *St. Louis U. L.J.* 93, 97–101 (2018).

information about jobs. As Tarleton Gillespie has observed, social media platforms “do not just mediate public discourse: they constitute it.”²⁸ Similarly, online job-matching and recruiting platforms do not merely pass along job information. By sorting, screening, and curating information, they actively constitute the labor market.

Miranda Bogen and Aaron Rieke have cataloged the ways in which predictive analytic tools are currently being used throughout the hiring process.²⁹ They describe the process as a funnel—a series of decisions that progressively narrows the pool of potential hires and ultimately culminates in a job offer or a rejection.³⁰ They identify several stages—sourcing, screening, interviewing, and selection—and describe how the predictive tools used at each stage effectively remove some candidates from the pool.³¹ The focus of this Article is on the sourcing stage—the step in which candidates learn about and apply for available job opportunities. Of course, the combination of these stages determines actual hiring outcomes, but the sourcing stage is particularly important because it is the first step in winnowing the available pool of workers. The sourcing stage also shapes the perceptions of labor market participants—employers and employees alike—about the availability of jobs and the types of workers who can fill them. These perceptions, in turn, likely affect their behavior in the later stages of the hiring process.

Bogen and Rieke explain in detail the ways in which predictive algorithms are used during the sourcing stage: to advertise job postings to selected audiences, to match applicants with prospective employers through recommender algorithms, and to identify passive applicants.³²

Online Advertising Platforms. Targeted advertising occurs on all kinds of platforms, including search engines, social media sites, and websites that provide content. Facebook’s advertising platform provides an illustrative example of the available tools. By leveraging the enormous amount of data it collects about its users,³³ Facebook offers advertisers

²⁸ Tarleton Gillespie, *Platforms Are Not Intermediaries*, 2 *Geo. L. Tech. Rev.* 198, 199 (2018).

²⁹ Bogen & Rieke, *supra* note 7.

³⁰ *Id.* at 13.

³¹ *Id.* at 13–42.

³² *Id.* at 14–25.

³³ Facebook collects extensive data about user behavior on the site. It purchases additional data about users’ offline behaviors such as purchasing history and combines it with information gleaned from its own site. See Julia Angwin, Surya Mattu & Terry Parris, Jr., *Facebook Doesn’t Tell Users Everything It Really Knows About Them*, ProPublica (Dec. 27,

the opportunity to display “the right ad, to the right person, at the right time.”³⁴ The data are analyzed to assign attributes to individual users, and then these attributes—hundreds of thousands of them—are available for advertisers to use to define their target audience. Some of these attributes capture basic demographic information such as gender, age, and geographic location.³⁵ Others are focused on interests (e.g., cat lovers), affiliations (e.g., veterans), or behaviors (e.g., ultramarathoners). Facebook encourages advertisers to narrow their audience by using the available attributes to include or exclude certain types of users from its target audience.³⁶

In the past, all these attributes were available for narrowing the target audience, regardless of the type of ad. After several lawsuits alleged these tools could be used to discriminate, Facebook agreed in March 2019 to a settlement restricting the types of attributes that can be used to select an audience for employment, housing, and credit advertisements.³⁷ However, these restrictions will not necessarily prevent biased distribution of ads. Although advertisers make the initial decisions about how to define the target audience, not all users who match those criteria will see the advertisement. From the advertiser-identified group of eligible users, the platform selects which users will *actually* see a given ad. As explained below, the platform does so through a complex process that involves analyzing users’ data to predict who is most likely to interact with the ad by reading it, clicking on a link, or actually engaging with the advertiser—for example, buying a product, submitting an application, etc.

2016, 9:00 AM), <https://www.propublica.org/article/facebook-doesnt-tell-users-everything-it-really-knows-about-them> [<https://perma.cc/CG76-EYU4>]; Dylan Curran, Are You Ready? Here Is All the Data Facebook and Google Have on You, *Guardian* (Mar. 30, 2018, 3:17 PM), <https://www.theguardian.com/commentisfree/2018/mar/28/all-the-data-facebook-google-has-on-you-privacy> [<https://perma.cc/PS2P-VCP6>]; Hanna Kozłowska, “Why Am I Seeing This Ad” Explanations on Facebook Are Incomplete and Misleading, a Study Says, *Quartz* (Apr. 6, 2018), <https://qz.com/1245941/why-am-i-seeing-this-ad-explanations-on-facebook-are-incomplete-and-misleading-a-study-says/> [<https://perma.cc/KNQ4-9US5>].

³⁴ People-Based Marketing: Thinking People-First Planning and Measurement, Facebook for Business (Sept. 13, 2017), <https://www.facebook.com/business/news/insights/the-future-of-marketing-people-based-planning-and-measurement> [<https://perma.cc/ZL2Y-NQYW>].

³⁵ About Reaching New Audiences, Facebook for Business, <https://www.facebook.com/business/help/717368264947302> [<https://perma.cc/LDA4-S78N>] (last updated Dec. 26, 2019).

³⁶ Ad Targeting: Help Your Ads Find the People Who Will Love Your Business, Facebook for Business, <https://www.facebook.com/business/ads/ad-targeting> [<https://perma.cc/EE8H-EAGA>] (last visited Apr. 2, 2020).

³⁷ Sherwin & Bhandari, *supra* note 9.

Another option offered by Facebook is the “Lookalike Audience” tool.³⁸ To use this feature, an advertiser provides Facebook with information about an existing group—the source audience—that represents its target audience. In the recruitment context, an employer might identify its current workforce or a group of successful recent hires. Facebook then uses its extensive data to analyze the attributes of users in the source audience and to identify other Facebook users who have similar attributes. Although the employer initially defines the targeted group by choosing the original source audience, Facebook’s algorithm determines the relevant attributes that define this audience and which other users will *actually* receive the advertisement.³⁹ The March 2019 settlement also prohibited Facebook from using characteristics like age and gender in creating Lookalike Audiences. Once again, however, these changes do not ensure a nonbiased audience, as explained in Section II.C.

Job-Matching Platforms. Workers who are actively looking for employment may register on platforms designed to assist with the job search process, such as ZipRecruiter and LinkedIn. Employers seeking to hire also register, providing information about the available opportunity and the skills or experience required. Rather than passively posting this information, these platforms actively facilitate the job-matching process by nudging job seekers to apply for certain jobs and recommending particular applicants to employers. These platforms use recommender systems—the same type of technology familiar to consumers who receive recommendations of books or movies on Amazon or Netflix.

Recommender systems rely on both user-provided data as well as data collected by the platform to make predictions. They typically use two methods to develop recommendations: content-based filtering and collaborative filtering.⁴⁰ Content-based filtering relies on users’ express or demonstrated preferences—for example, movies viewed or books purchased in the past—to recommend similar items.⁴¹ Collaborative filtering draws inferences about a user’s preferences or attributes by

³⁸ About Lookalike Audiences, Facebook for Business, [https://www.facebook.com/business/help/164749007013531?helpref=faq_content# \[https://perma.cc/H398-UDBS\]](https://www.facebook.com/business/help/164749007013531?helpref=faq_content# [https://perma.cc/H398-UDBS]) (last updated Dec. 24, 2019).

³⁹ *Id.*

⁴⁰ Bogen & Rieke, *supra* note 7, at 20; see generally Sirui Yao & Bert Huang, Beyond Parity: Fairness Objectives for Collaborative Filtering, 31st Conf. Neural Info. Processing Sys. (2017) (describing the “collaborative filtering” approach and explaining how various types of unfairness can occur with collaborative filtering).

⁴¹ Bogen & Rieke, *supra* note 7, at 20.

comparing her with similar users.⁴² When used by job-matching platforms, the latter method might, for example, recommend a job seeker to a particular employer because similar applicants received positive responses from that employer in the past.

ZipRecruiter offers an example of how predictive tools are used in multiple stages of the job-matching process.⁴³ After an employer posts a job opening, ZipRecruiter's algorithm scans tens of thousands of resumes available to it, looking for candidates that match the employer's specified requirements, and then nudges those candidates to apply for the job. The algorithm also predicts which candidates the employer will be most interested in, labeling them a "Great Match" and listing them at the top of the queue visible to employers. Bogen and Rieke further explain:

[O]n ZipRecruiter, employers can opt to give incoming applicants a "thumbs up." As ZipRecruiter collects these positive signals, it uses a machine learning algorithm to identify other jobseekers in its system with similar characteristics to those who have already been given a "thumbs up"—who have not yet applied for that role—and automatically prompts them to apply. The details of the matching process make up ZipRecruiter's special sauce, which considers not only basic demographic and skills information from resumes and other information added by jobseekers, but also insights gleaned from their behavior on the website.⁴⁴

Thus, ZipRecruiter's algorithms intervene in the job-matching process in several ways—by nudging job seekers to apply for certain jobs, by highlighting some applicants as "Great Matches," and by searching out additional candidates based on a favorable employer response.

With each such intervention, a recommender system necessarily promotes some opportunities while obscuring others. Job seekers have some choices about what information to provide about themselves,⁴⁵ but the platform ultimately determines which factors are used to predict likely matches and therefore which opportunities they are more likely to see. Similarly, the platform controls who is encouraged to apply and who is

⁴² *Id.*

⁴³ ZipRecruiter, <https://www.ziprecruiter.com/> [<https://perma.cc/6UBC-PKB6>] (last visited Apr. 2, 2020).

⁴⁴ Bogen & Rieke, *supra* note 7, at 20.

⁴⁵ See Sharone, *supra* note 26, at 20–21 (describing how LinkedIn constrains job seekers' ability to determine their own self-presentation on the job market).

promoted as a good match, thereby influencing which candidates an employer sees. Thus, the platform ultimately controls the visibility and salience of potential opportunities for both employee and employer.

Passive Recruiting Systems. Passive recruiting refers to the strategy of identifying top job candidates who are not currently on the market.⁴⁶ Employers try to discover potential hires who have the skills they need and then cultivate relationships that may ultimately lead them to leave their current firm. A handful of technology platforms are also working in this space, using machine learning and predictive analytics to analyze hundreds of thousands of available profiles in order to help companies identify workers with a needed skill set. Much like the recommender systems used by matching platforms, these search firms analyze available data about potential candidates to infer targeted characteristics or to predict who is most likely to change jobs. A recruiting platform exercises a great deal of control because employers will only be aware of the people the platform identifies as promising hires, and workers may only learn of opportunities if an employer actually tries to recruit them.

B. Discriminatory Choices

Depending upon how they are constructed, the sourcing tools described above may allow employers to deliberately discriminate against protected groups. In choosing how to target its job advertisements, a discriminatory employer might use demographic attributes to define the audience along protected class lines—for example, showing the ad only to men or only to women, or only to users under a certain age. Alternatively, an employer might rely on interest or behavioral attributes that are closely correlated with a protected status in order to include or exclude certain groups. Similarly, by utilizing a biased target sample to train recommender systems or passive recruiting tools, employers could reproduce existing biases in their workforces.

Speicher et al. show that attributes available on Facebook, like “BlackNews.com” and “Nuestro Diario,” are highly correlated with specific racial or ethnic groups and could be used to effectively skew a targeted audience.⁴⁷ They also show how the Lookalike Audience tool

⁴⁶ Sam Holzman, 5 Tried-and-True Methods for Sourcing Passive Candidates, Glassdoor for Employers (Nov. 28, 2018), <https://www.glassdoor.com/employers/blog/5-tried-and-true-methods-for-sourcing-passive-candidates/> [<https://perma.cc/4CWC-KE4U>].

⁴⁷ Till Speicher, Muhammad Ali, Giridhari Venkatadri, Filipe Nunes Riberio, George Arvanitakis, Fabrício Benevenuto, Krishna P. Gummadi, Patrick Loiseau & Alan Mislove,

could be used for biased targeting. By beginning with a highly biased source audience, an advertiser could use the tool to scale up those biases, reproducing them in the larger targeted audience.⁴⁸ Similarly, because of patterns of racial and ethnic residential segregation, Facebook’s highly precise location-targeting tools⁴⁹ could be used to exclude members of disfavored groups.

This risk that *advertisers* will discriminate by using these micro-targeting tools has received the most popular and legal attention. In October 2016, ProPublica reported that Facebook allowed advertisers, including employers and landlords, to exclude users using attributes for “Ethnic Affinities” like “African-American,” “Asian-American,” and “Hispanic.”⁵⁰ Around the same time, a New York Times and ProPublica investigation revealed that large employers such as Amazon and Verizon had placed multiple recruitment ads on Facebook that targeted particular age groups—most often younger workers.⁵¹ Other advertising platforms

Potential for Discrimination in Online Targeted Advertising, 81 Proc. Machine Learning Res. 1, 9 (2018).

⁴⁸ Id. at 11–12.

⁴⁹ Irfan Faizullahoy & Aleksandra Korolova, Facebook’s Advertising Platform: New Attack Vectors and the Need for Interventions, Workshop on Tech. & Consumer Protection (2018), <https://www.ieee-security.org/TC/SPW2018/ConPro/papers/faizullahoy-conpro-18.pdf> [<https://perma.cc/5BQZ-DL35>].

⁵⁰ Julia Angwin & Terry Parris, Jr., Facebook Lets Advertisers Exclude Users by Race, ProPublica (Oct. 28, 2016, 1:00 PM), <https://www.propublica.org/article/facebook-lets-advertisers-exclude-users-by-race> [<https://perma.cc/TJJ5-NU8U>]. According to a spokesperson, Facebook does not ask users to identify their race or ethnic origin but assigns these “ethnic affinity” attributes based on users’ posts, likes, and other behaviors on the site. In response to the ProPublica report and the criticism that followed, see, e.g., Jessica Guynn, Facebook Must Stop Ads That Exclude Races: Lawmakers, USA Today (Nov. 1, 2016), <https://www.usatoday.com/story/tech/news/2016/11/01/congressional-black-caucus-asks-facebook-to-stop-letting-advertisers-exclude-racial-ethnic-groups-in-housing-ads/93147048> [<https://perma.cc/YA3D-H7KC>], Facebook announced that it would prohibit the use of ethnic affinity attributes to target ads for housing, employment, and credit. A year later, however, it appeared that it was still possible for these types of ads to be targeted on the basis of ethnic affinity. See Julia Angwin, Facebook Says It Will Stop Allowing Some Advertisers To Exclude Users by Race, ProPublica (Nov. 11, 2016, 10:00 AM), <https://www.propublica.org/article/facebook-to-stop-allowing-some-advertisers-to-exclude-users-by-race> [<https://perma.cc/V36H-G6JG>]; Julia Angwin, Ariana Tobin & Madeleine Varner, Facebook (Still) Letting Housing Advertisers Exclude Users by Race, ProPublica (Nov. 21, 2017, 1:23 PM), <https://www.propublica.org/article/facebook-advertising-discrimination-housing-race-sex-national-origin> [<https://perma.cc/K46X-35VQ>].

⁵¹ Julia Angwin, Noam Scheiber & Ariana Tobin, Dozens of Companies Are Using Facebook To Exclude Older Workers from Job Ads, ProPublica (Dec. 20, 2017, 5:45 PM), <https://www.propublica.org/article/facebook-ads-age-discrimination-targeting> [<https://perma.cc/3KRG-HFEH>]; Julia Angwin, Noam Scheiber & Ariana Tobin, Facebook Job Ads Raise

like Google and LinkedIn similarly make it possible for employers to select their target audience by age.⁵² Separate investigations found examples of job ads on Facebook that were targeted by gender.⁵³

Following these media reports, several lawsuits were filed against Facebook alleging that its audience selection tools permitted advertisers to discriminate in the delivery of employment, housing, and credit ads.⁵⁴ Facebook's response to these allegations was to place responsibility for any discrimination on the advertisers, arguing that others were misusing the "neutral" tools it provided. For example, in 2016, Steve Satterfield, privacy and public policy manager at Facebook, said that "[w]e take a strong stand against advertisers misusing our platform: [o]ur policies prohibit using our targeting options to discriminate, and they require

Concerns About Age Discrimination, N.Y. Times (Dec. 20, 2017), <https://www.nytimes.com/2017/12/20/business/facebook-job-ads.html> [<https://perma.cc/5NH7-5MXX>].

⁵² Jennifer Valentino-DeVries, AARP and Key Senators Urge Companies To End Age Bias in Recruiting on Facebook, ProPublica (Jan. 8, 2018, 8:00 AM), <https://www.propublica.org/article/aarp-and-key-senators-urge-companies-to-end-age-bias-in-recruiting-on-facebook> [<https://perma.cc/SM7J-TSFE>].

⁵³ Facebook EEOC Complaints, ACLU (last updated Sept. 25, 2019), <https://www.aclu.org/cases/facebook-eec-complaints> [<https://perma.cc/A8ZN-K54H>]; Ariana Tobin & Jeremy B. Merrill, Facebook Is Letting Job Advertisers Target Only Men, ProPublica (Sept. 18, 2018, 6:39 PM), <https://www.propublica.org/article/facebook-is-letting-job-advertisers-target-only-men> [<https://perma.cc/8L74-WPFT>].

⁵⁴ In November 2016, plaintiffs in *Onuoha v. Facebook* brought a class action alleging that the platform's tools "enable and encourage" advertisers to discriminate against African-Americans, Latinos, and Asian Americans by excluding them from receiving ads for employment, housing, and credit. First Amended Complaint at 1, *Onuoha v. Facebook, Inc.*, No. 5:16-cv-06440-EJD, 2017 BL 115835 (N.D. Cal. Feb. 13, 2017). In another class action, *Bradley v. T-Mobile*, plaintiffs sued under the Age Discrimination in Employment Act, alleging that a class of employers had used Facebook's advertising tools to exclude older users from receiving job advertisements. First Amended Class and Collective Action Complaint at 2, *Bradley v. T-Mobile US, Inc.*, No. 17-cv-07232-BLF, 2019 WL 2358972 (N.D. Cal. May 29, 2018). Several fair housing groups brought a separate suit under the Fair Housing Act, arguing that the availability of targeting attributes such as "English as a second language," "parents with toddlers," "moms of grade school kids," and "Disabled American Veterans" permitted discrimination against racial minorities, women, families with children, and people with disabilities. Complaint at 18–20, *Nat'l Fair Hous. All. v. Facebook, Inc.*, No. 18-cv-02689 (S.D.N.Y. Mar. 27, 2018); see also Third Amended Class Action Complaint at 1–2, *Riddick v. Facebook, Inc.*, No. 3:18-cv-04529-LB, 2018 WL 8786810 (N.D. Cal. Nov. 19, 2018) (alleging unlawful discrimination under California state laws). Charges were also filed with the Equal Employment Opportunity Commission alleging that Facebook enabled employers to exclude women from seeing certain job advertisements. ACLU and Workers Take on Facebook for Gender Discrimination in Job Ads, ACLU (Sept. 18, 2018), <https://www.aclu.org/press-releases/aclu-and-workers-take-facebook-gender-discrimination-job-ads> [<https://perma.cc/H585-GPF9>]; Tobin & Merrill, *supra* note 53.

compliance with the law.”⁵⁵ Similarly, in moving to dismiss the lawsuits against it, Facebook characterized its targeting tools as “neutral” and disclaimed responsibility for any misuse by others.⁵⁶ It is *advertisers*, Facebook emphasized, who “are responsible for . . . what targeting criteria to use, if any.”⁵⁷

In March 2019, Facebook announced a “historic” agreement⁵⁸ to resolve the pending lawsuits by creating a separate portal for housing, employment, and credit ads (the HEC portal) with restricted targeting options.⁵⁹ In this portal, advertisers will be unable to use gender, age, or zip code to target audiences. Location-based targeting will only be permitted by designating a fifteen-mile minimum radius around the center of a city or town, address, or pin-drop.⁶⁰ In addition, the portal will not permit the use of attributes that are “direct descriptors of, or semantically or conceptually related to,” protected classes,⁶¹ presumably making unavailable attributes correlated with race or ethnic affinity, or attributes such as “soccer moms” or “young and hip.” The agreement further specifies that the Lookalike Audience tool will no longer use characteristics like age, gender, and religious views to create an audience that matches a source audience.⁶²

⁵⁵ Angwin & Parris, Jr., *supra* note 50.

⁵⁶ See, e.g., Defendant’s Notice of Motion and Motion To Dismiss First Amended Complaint at 2, *Onuoha*, 2017 BL 115835.

⁵⁷ *Id.* (“Advertisers, not Facebook, are responsible for . . . what targeting criteria to use, if any. Facebook’s provision of these neutral tools to advertise falls squarely within the scope of CDA immunity.”).

⁵⁸ Sheryl Sandberg, *Doing More To Protect Against Discrimination in Housing, Employment and Credit Advertising*, Facebook (Mar. 19, 2019), <https://newsroom.fb.com/news/2019/03/protecting-against-discrimination-in-ads/> [<https://perma.cc/3XDK-8VG4>].

⁵⁹ Settlement Agreement and Release at 3, *Nat’l Fair Hous. All. v. Facebook, Inc.*, No. 1:18-cv-02689-JGK (S.D.N.Y. Mar. 19, 2019). Those posting housing, employment, or credit advertisements on Facebook will be required to use the special portal. *Id.* The company also agreed to develop classifiers to detect HEC advertisements that are not submitted through the special portal and to sanction advertisers who repeatedly fail to comply. *Id.* at 2.

⁶⁰ About Audiences for Special Ad Categories, Facebook for Business, <https://www.facebook.com/business/help/2220749868045706> [<https://perma.cc/7UFT-GFMA>] (last updated Oct. 14, 2019).

⁶¹ Settlement Agreement and Release, *supra* note 59, at 3.

⁶² *Id.* ex. A, at 4. Facebook also agreed to create a tool that would allow users to search all current housing listings. *Id.* The agreement includes a statement that Facebook will conduct an algorithmic fairness review. *Id.* at 5. While the agreement states that Facebook will “engage” with various experts and advocates “to study the potential for unintended bias in algorithmic modeling,” it does not specify in any detail what such engagement will look like or what level of access these groups will have to relevant data or the results of any internal studies. *Id.*

While the civil rights organizations that brought the suits welcomed the agreement,⁶³ the changes are largely consistent with the stance that Facebook has taken all along. In her statement announcing the settlement, Facebook's Chief Operating Officer Sheryl Sandberg stated, "Our job is to make sure . . . that our ads tools aren't *misused*."⁶⁴ Facebook, she explained, had engaged a civil rights firm to "help us understand what more we could do to guard against misuse."⁶⁵ In other words, Facebook regards the problem as one of malicious advertisers misusing its "neutral" tools. By focusing on advertisers' discriminatory choices, the company is implicitly disclaiming any responsibility for discriminatory ad targeting.⁶⁶ However, as explained in the next Section, the reforms agreed to by Facebook in March 2019 are unlikely to prevent biased ad delivery.

C. Predictive Analytics and Discriminatory Effects

The possibility that landlords or employers can deliberately discriminate by targeting advertisements based on race, gender, or age is both intuitively obvious and easy to condemn. However, the advertising, matching, and recruiting technologies described in Section II.A above can produce discriminatory patterns even when advertisers do not expressly rely on protected characteristics or their close proxies. Merely prohibiting the use of these characteristics cannot guarantee that platforms will distribute information about key opportunities in an unbiased way.

Consider online advertising platforms like Facebook and Google. After an employer or landlord has selected attributes to define its target audience, its ads are not simply displayed to everyone who fits the selected criteria, or even to a random sample of those users. Instead, the

⁶³ Facebook Agrees to Sweeping Reforms To Curb Discriminatory Ad Targeting Practices, ACLU (Mar. 19, 2019), <https://www.aclu.org/news/facebook-agrees-sweeping-reforms-curb-discriminatory-ad-targeting-practices> [<https://perma.cc/8FCZ-YN9T>] (stating that "[t]he settlement encompasses sweeping changes that the tech giant will make to its paid advertising platform to prevent discrimination in employment, housing, and credit advertising").

⁶⁴ See Sandberg, *supra* note 58 (emphasis added).

⁶⁵ *Id.*

⁶⁶ Shortly after the settlement between Facebook and the civil rights organizations, the U.S. Department of Housing and Urban Development filed a charge against Facebook alleging that the company's algorithms caused discrimination in violation of the Fair Housing Act. See Charge of Discrimination at 1, Facebook, Inc., FHEO No. 01-18-0323-8 (H.U.D. Mar. 28, 2019); Katie Benner, Glenn Thrush & Mike Isaac, Facebook Engages in Housing Discrimination with Its Ad Practices, U.S. Says, N.Y. Times (Mar. 28, 2019), <https://www.nytimes.com/2019/03/28/us/politics/facebook-housing-discrimination.html> [<https://perma.cc/5VXK-FGJQ>].

platform determines who in the eligible audience will *actually* see a particular ad. Exactly how it does so is not easy to untangle because the process involves complex interactions among different actors. The online ad ecosystem is a “vast, distributed, and decentralized system” with multiple parties.⁶⁷ Where a particular ad appears is influenced by the advertiser (who specifies its target audience), other advertisers (who are competing for advertising space), users themselves (who choose whether or not to click on particular ads), and the platform that coordinates these preferences.

The interaction between all these actors is mediated by the platforms’ proprietary algorithms, rendering the entire process opaque. Nevertheless, certain things are known. Advertisers not only provide the ad content, they also select targeting criteria and indicate how much they are willing to pay.⁶⁸ When advertising space becomes available—in a user’s news feed or alongside search results—an online auction automatically occurs among advertisers.⁶⁹ Because platforms seek to optimize revenue, their algorithms try to predict which ads will be most relevant to which users. These predictions are based not only on the known interests and behaviors of that particular individual but also on what is inferred about her from the behavior of similar users.⁷⁰ Precisely which ads an individual will see is ultimately determined through an algorithmic process controlled by the platform.

Because multiple factors influence online ad delivery, the targeting criteria chosen by the advertiser do not determine the demographic distribution of the ad’s recipients. An employer may try to avoid a demographic skew by not using characteristics like sex or age as targeting variables. If it truly wishes to reach a diverse audience, it may also avoid obvious proxies for those characteristics. Nevertheless, even if it has chosen neutral targeting criteria, the actual audience receiving the ad may be skewed along the lines of race, sex, or other protected characteristics due to the platform’s targeting algorithm.

⁶⁷ Amit Datta, Anupam Datta, Jael Makagon, Deirdre K. Mulligan & Michael Carl Tschantz, *Discrimination in Online Advertising: A Multidisciplinary Inquiry*, 81 *Proc. Machine Learning Res.* 20, 22 (2018).

⁶⁸ *Id.* at 23.

⁶⁹ *Id.* at 25.

⁷⁰ *Id.* app. A, at 2 (explaining that Google predicts how likely a given consumer is to click on an ad if shown based upon data about the behavior of millions of other users).

Studies have documented instances of employment-related ads being delivered in a biased way.⁷¹ Amit Datta et al. showed that simulated computer users identified as female received fewer advertisements for a “career coaching agency that promised large salaries,” even though their web search histories were identical to the simulated users identified as male.⁷² The researchers could not determine with certainty what caused the gender skew in ad delivery. While deliberate discrimination by either the advertiser or the platform (Google, in the case of their study) was possible, the skew might also occur if other advertisers’ targeting choices and bidding rates crowded out the career coaching ad for female computer users.⁷³ Alternatively, it could be the result of different behaviors of men and women. Even though the simulated users in the study behaved identically, if men overall click on this ad (or ads like it) more often than women, then Google’s algorithm would recognize that pattern and deliver the ad more often to men.⁷⁴ An algorithm will not necessarily respond in this way, but a system designed to maximize interaction will inevitably reflect past behavioral patterns.

In another study, Lambrecht and Tucker field-tested an ad that promoted careers in STEM (Science, Technology, Engineering and Math).⁷⁵ Although they targeted the ad on a gender-neutral basis, it “was shown to more than 20% more men than women.”⁷⁶ Because they had chosen neutral targeting criteria, it was clear that the discriminatory

⁷¹ Latanya Sweeney, *Discrimination in Online Ad Delivery*, 56 *Comm. ACM* 44, 46–47. Latanya Sweeney documented that ads suggesting an arrest record (e.g., “Latanya Sweeney, arrested?” with a link to a site offering background checks) are more likely to appear next to Google searches for people with black-identified names (e.g., Latanya, Imani, Trevor) than white-identified names (e.g., Brad, Jill, Katelyn). These types of ads are not technically employment ads, but they likely have an indirect effect on job opportunities. If employers conducting online searches about prospective employees are prompted by ads suggesting a criminal record more often for black applicants than white applicants, it may alter their perceptions of applicants based on race. Thus, although this is not an example of discriminatory ad delivery to job seekers, it provides another example of how discriminatory online advertising could affect labor markets.

⁷² Amit Datta, Michael Carl Tschantz & Anupam Datta, *Automated Experiments on Ad Privacy Settings: A Tale of Opacity, Choice, and Discrimination*, 2015 *Proc. Privacy Enhancing Technologies* 92, 93.

⁷³ Datta et al., *supra* note 67, at 25.

⁷⁴ *Id.* at 26.

⁷⁵ Anja Lambrecht & Catherine Tucker, *Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads*, 65 *Mgmt. Sci.* 2966, 2966 (2019).

⁷⁶ *Id.* at 2966.

delivery pattern was not intentional. Nor did it reflect different levels of interest: women who received the ad clicked on it at a higher rate than men.⁷⁷ Instead, the researchers hypothesized that the gender skew resulted from the market effects of other advertisers' choices. Citing evidence that women are a highly prized demographic for advertisers,⁷⁸ Lambrecht and Tucker explained that spillover effects could cause the discriminatory delivery of the STEM ad because other advertisers vying for the attention of female users might have outbid the STEM ad.

In a recent paper, Muhammad Ali et al. describe a series of experiments designed to tease out the factors that might cause discriminatory ad delivery on Facebook.⁷⁹ The researchers ran numerous advertisements using race- and gender-neutral targeting criteria, and they repeatedly found significant race and gender skews in the populations actually receiving the ads. By running parallel ads while changing only one parameter, they were able to test various hypotheses about what causes discriminatory ad delivery. They found evidence that both market effects and Facebook's predictive algorithms were causing the observed skews.

Ali et al. first tested Lambrecht and Tucker's hypothesis that spillover effects from the bidding process could create a gender skew in ad delivery. They found that the lower the budget for distributing an ad, the lower the proportion of women in the actual audience. This finding is consistent with the explanation that women are a more highly valued audience and therefore more expensive to reach. If advertisers compete to capture their attention, then market effects alone could cause a gender skew with no discriminatory intent on the part of the advertiser. Although these market effects reflect advertisers' collective behavior, the platforms themselves decide how to structure the auction. Platforms do not rely solely on bid prices to determine which ad is shown, but they consider other factors as well, such as the predicted response of the user.⁸⁰ In other

⁷⁷ Id. at 2972.

⁷⁸ Alex Cocotas, *Young Women Are the Most Valuable Mobile Ad Demographic*, *Business Insider* (Feb. 29, 2012, 2:55 PM), <https://www.businessinsider.com/young-women-are-most-valuable-mobile-ad-demographic-2012-2> [<https://perma.cc/Z4T2-63NL>].

⁷⁹ Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove & Aaron Rieke, *Discrimination Through Optimization: How Facebook's Ad Delivery Can Lead to Biased Outcomes*, 3 *Proc. ACM on Hum.-Computer Interaction* 1 (2019).

⁸⁰ Id. at 2–3; see also Datta et al., *supra* note 67, app. A, at 2 (explaining that the most important factor is the “expected click-through rate,” which predicts the likelihood the user will click on the ad (emphasis omitted)).

words, *the platform* decides how to weigh different factors and thus ultimately controls what ads to display to which users.

Ali et al. also found that the text and image comprising an advertisement—what is referred to as the “ad creative”—can cause a demographically skewed audience. As they describe their findings:

[A]ds targeting the same audience but that include a creative that would stereotypically be of the most interest to men (e.g., bodybuilding) can deliver to over 80% men, and those that include a creative that would stereotypically be of the most interest to women (e.g., cosmetics) can deliver to over 90% women. Similarly, ads referring to cultural content stereotypically of most interest to Black users (e.g., hip-hop) can deliver to over 85% Black users, and those referring to content stereotypically of interest to white users (e.g., country music) can deliver to over 80% white users, even when targeted identically by the advertiser. Thus, despite placing the same bid on the same audience, the advertiser’s ad delivery can be heavily skewed based on the ad creative alone.⁸¹

These skewed audiences appear from the very beginning of the ad-delivery period, indicating that they result from Facebook’s predictions rather than users’ actual responses to the ad. When an ad is first placed, Facebook’s algorithm apparently analyzes the elements of the ad creative and predicts who is most likely to be interested based on that content.⁸²

Ali et al. also explored whether the same dynamics can affect the delivery of employment and housing ads. They found that for some types of jobs, delivery of the ads was markedly skewed on race and/or gender lines, even when neutral text and images were selected. For example, recipients of an ad for lumberjack positions were over 90% male and over 70% white, while those receiving ads for a janitor position were over 65% women and 75% black.⁸³ The researchers similarly tested housing ads, varying their content between types of property (rental vs. purchase), implied cost, and the type of image used (white family, black family, or

⁸¹ Ali et al., *supra* note 79, at 3.

⁸² Ali et al. found that the image used in the ad creative had a very strong effect on the demographic skew in delivery. *Id.* at 18–19.

⁸³ *Id.* at 20. These figures represent aggregate results across all ad types used in the experiment. Looking only at the results of ads that used a neutral stock photo with no person in the image, the skews in the audience receiving them are still quite pronounced. *Id.* at 20–22.

no people pictured). They found significant differences in ad delivery to black and white users depending upon the content of the ads.⁸⁴

The recent Facebook settlement is not likely to eliminate discriminatory ad delivery. That agreement was intended to remove the ability of housing, employment, and credit advertisers to intentionally discriminate,⁸⁵ but it left untouched the role of the targeting algorithm. Even if an employer or landlord uses neutral targeting criteria and wishes to reach a broad, inclusive audience, Facebook's algorithm can deliver the ad to a demographically skewed audience.

In a follow-up study, researchers tested the effects of the changes implemented as a result of the settlement.⁸⁶ They found that merely removing demographic inputs like gender and age from the targeting criteria failed to prevent the biased distribution of employment ads. Ads targeted at a gender- and age-balanced audience were delivered to demographically skewed audiences depending upon the type of job advertised. For example, ads for supermarket jobs delivered to a 72% female (and mostly older) audience; ads for AI jobs delivered to a 60% male (and mostly younger) audience.⁸⁷ An investigation by ProPublica uncovered similar effects.⁸⁸ It documented the case of a construction and

⁸⁴ *Id.* at 20–22.

⁸⁵ According to the ProPublica investigation, as of November 2019, it was still possible to place housing and employment ads without going through the special portal as required by the settlement. Ava Kofman & Ariana Tobin, Facebook Ads Can Still Discriminate Against Women and Older Workers, Despite a Civil Rights Settlement, ProPublica (Dec. 13, 2019, 5:00 AM), <https://www.propublica.org/article/facebook-ads-can-still-discriminate-against-women-and-older-workers-despite-a-civil-rights-settlement> [<https://perma.cc/F3YU-HCTQ>]. Facebook has asserted that since that time, it has more closely monitored compliance by housing and employment advertisers. *Id.*

⁸⁶ Piotr Sapiezynski, Avijit Ghosh, Levi Kaplan, Alan Mislove & Aaron Rieke, Algorithms That “Don’t See Color”: Comparing Biases in Lookalike and Special Ad Audiences (arXiv, Working Paper No.1912.07579, 2019), <https://arxiv.org/pdf/1912.07579.pdf> [<https://perma.cc/HG3E-NLCY>].

⁸⁷ *Id.* at 8, fig.9. The researchers also specifically tested whether prohibiting Facebook from using demographic characteristics to build a Lookalike Audience would eliminate biased ad delivery. Following the settlement, Facebook created an alternative to the Lookalike Audience called “Special Ad Audiences.” Like the Lookalike Audience tool, the Special Ad Audience targets a larger audience that is similar to a source audience identified by the advertiser, but it does so without relying on gender, age, and other demographic characteristics to identify similar users in the larger audience. *Id.* at 1–2. The researchers found that if an advertiser started with a source audience biased along gender or age lines, the resulting Special Ad Audience largely replicated the bias. *Id.* at 4–5. Their study offers an experimental demonstration that merely removing demographic inputs will not necessarily prevent biased outputs from algorithmic systems that have access to large amounts of data.

⁸⁸ Kofman & Tobin, *supra* note 85.

supply company that advertised for truck drivers on Facebook using the special portal for employment ads. Although it was impossible to target by gender, and the company intended to reach a gender-neutral audience, the ad was in fact delivered to an audience that was 87% male.⁸⁹

Similar to targeted advertising, job-matching platforms and passive recruitment systems are powered by predictive algorithms that can cause demographic skews even if the employer wishes to recruit a diverse pool of candidates. These systems typically base their predictions on the observed behavior of both employers and applicants. If express or implicit biases shape behavior, those biases will be reproduced in the resulting predictions and recommendations.⁹⁰ Amazon encountered this problem when it tried to create an automated tool for hiring software developers that would “crawl the web and spot candidates worth recruiting.”⁹¹ Its model was trained using the resumes of past hires—mostly men—and as a result, it learned to systematically prefer men over women, downgrading resumes that included phrases like “women’s chess club captain” or listed certain women’s colleges.⁹²

Not only can recommender systems reproduce past patterns that reflect bias, but they could also reinforce or worsen them. Bogen and Rieke explain how this might happen in the context of ZipRecruiter’s model. If two job seekers are identified as similar and an employer gives one of them a “thumbs up,” the other will be encouraged to apply for the same job. If the second person does so, the system will highlight that applicant as a “‘great match[,]’ . . . essentially reinforcing the employer’s initial screening decision.”⁹³ And if the employer also gives the second candidate a “thumbs up,” the prediction that this type of candidate will be successful is further reinforced. If the employer’s reaction to the first

⁸⁹ Other examples reported include an ad for a union apprenticeship program that was delivered to an audience that was two-thirds male and an ad for an internship program for college students that was delivered to a 73% male audience. *Id.*

⁹⁰ Levy & Barocas, *supra* note 19, at 1209 (“[I]f the platform relies on historical patterns of successful user interaction to guide its future recommendations, these suggestions might reproduce or even exacerbate the prejudices and biases that influenced previous users’ decisions to interact with others . . .”).

⁹¹ Jeffrey Dastin, Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women, *Reuters* (Oct. 9, 2018, 11:12 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/T6NZ-T4ZW>].

⁹² *Id.*

⁹³ Bogen & Rieke, *supra* note 7, at 20.

candidate was influenced by bias, the predicted success of certain favored types will become a self-fulfilling prophecy.

Because recommendation and passive recruiting systems rely on collaborative filtering, they also “risk[] stereotyping users because of the actions of others like them.”⁹⁴ Predictions are based on the behavior of similar users, thereby instantiating group generalizations. To the extent that a model organically perceives similarities along race, gender, or other demographic bases, recommendations made to job seekers will tend to match stereotyped expectations. For example, if women as a group are more likely to apply for jobs in female-dominated occupations like nursing or child care, that pattern will influence predictions about individual women job seekers. As a result, a woman who is open to non-traditional jobs in manufacturing or construction may not see those opportunities because she is not predicted to have any interest in them.⁹⁵

As predictive algorithms are increasingly incorporated into the recruitment process, they risk distributing information about employment opportunities in a biased way. A number of studies have documented instances of employment advertisements being delivered to audiences skewed along gender or racial lines. Data on how job-matching platforms and passive recruiting systems operate are not as readily available, but because the same dynamics are in play, those systems can produce biased outcomes as well. Importantly, these effects can occur solely or primarily through the operation of predictive algorithms, independent of or even despite the intentions of employers to act in a nondiscriminatory way. Thus, the recent focus on preventing advertisers from using targeting criteria to discriminate is far too limited. The platforms themselves play an important role in how job opportunities are distributed because they design the targeting or matching algorithms that control information flows. For that reason, the effects of their choices when building these systems warrant closer scrutiny.

D. Manipulating Opportunity

Because predictive algorithms determine the flow of information about housing, employment, and credit opportunities to particular users, these tools have the potential for manipulating opportunity markets. This possibility is related to concerns about manipulation in consumer and

⁹⁴ Id. at 21.

⁹⁵ See id.

political markets but differs in important ways. Most significantly, while manipulation in those markets may undermine autonomy and rational self-determination, manipulation of opportunity markets raises concerns about discrimination and inequality.

Criticism of manipulative marketing practices is not new,⁹⁶ but concerns have intensified as consumer transactions have increasingly moved online. Scholars like Ryan Calo and Tal Zarsky have argued that the immense amounts of personal data available increase the risk that firms will manipulate consumers.⁹⁷ In a marketplace mediated by technology, firms can leverage that data to personalize advertising pitches in ways that exploit consumers' cognitive biases and individual emotional vulnerabilities. Worries about manipulation have spread to the political realm as well.⁹⁸ These fears crystallized following revelations that a political consulting firm used personalized ads in an attempt to trigger emotional responses and influence voters in the 2016 presidential election.⁹⁹

Considerable debate exists over how exactly to define manipulation,¹⁰⁰ but there is agreement on two critical elements. First, manipulation entails some kind of influence over another.¹⁰¹ And second, that influence is hidden or covert.¹⁰² Scholars concerned with manipulation in consumer

⁹⁶ See, e.g., Jon D. Hanson & Douglas A. Kysar, *Taking Behavioralism Seriously: The Problem of Market Manipulation*, 74 N.Y.U. L. Rev. 630 (1999).

⁹⁷ See, e.g., Calo, *supra* note 2; Zarsky, *supra* note 2; see also Anthony Nadler & Lee McGuigan, *An Impulse To Exploit: The Behavioral Turn in Data-Driven Marketing*, 35 *Critical Stud. Media Comm.* 151, 151 (2018) (discussing the potential for “consumer surveillance and data-driven advertising” to “exacerbate[] market discrimination and intensify[] marketers’ capacity for social control”); Tal Z. Zarsky, “Mine Your Own Business!”: Making the Case for the Implications of the Data Mining of Personal Information in the *Forum of Public Opinion*, 5 *Yale J.L. & Tech.* 1, 5 (2003) (describing the potential impacts of data mining on privacy, autonomy, and discrimination).

⁹⁸ Frederik J. Zuiderveen Borgesius, Judith Möller, Sanne Kruikemeier, Ronan Ó Fathaigh, Kristina Irion, Tom Dobber, Balazs Bodo & Claes de Vreese, *Online Political Microtargeting: Promises and Threats for Democracy*, 14 *Utrecht L. Rev.* 82, 82 (2018); Ira S. Rubinstein, *Voter Privacy in the Age of Big Data*, 2014 *Wis. L. Rev.* 861, 864; Karen Yeung, “Hypermudge”: Big Data as a Mode of Regulation by Design, 20 *Info. Comm. & Soc’y* 118, 120 (2016).

⁹⁹ See, e.g., *supra* note 6 and accompanying text.

¹⁰⁰ See Daniel Susser, Beate Roessler & Helen Nissenbaum, *Technology, Autonomy, and Manipulation*, 8 *Internet Pol’y Rev.* 1, 3 (2019).

¹⁰¹ *Id.*

¹⁰² *Id.* at 4. Manipulation is not necessarily normatively wrong, although it usually has a negative connotation. As Susser et al. point out, there may be situations where the benefits of influencing behavior outweigh the costs. *Id.* at 5. “Nudging”—using insights about cognitive

markets also emphasize the use of techniques that exploit individuals' decision-making vulnerabilities.¹⁰³ By compromising rational decision making, these forms of manipulation can undermine autonomy and lead to poor choices and worse outcomes for consumers—such as paying too much or purchasing unsafe or unhealthy products.¹⁰⁴

As in consumer and political markets, tech intermediaries in opportunity markets can control information and influence decision making in ways that are not transparent to the participants. Section II.C above explained that platforms' role in the labor market is not simply to transmit information or to make available jobs equally visible to all. Instead, these intermediaries actively intervene to control information flows and to direct opportunities to certain users and not others. By leveraging extensive personal data—some of it voluntarily shared by users, some inferred through analytics—a platform's algorithms determine which opportunities are visible to job seekers and which applicants are most salient to employers. In this way, they influence how participants in the labor market perceive their available choice set. Those perceptions in turn shape behavior, affecting the parties' willingness to enter into bargains (to hire a particular worker or to accept a job offer). As a result, the operation of these predictive algorithms can significantly impact who ultimately gets which opportunity.

Importantly, the choices made by the algorithm are not transparent to the market participants. Although the platforms may not be engaged in any deliberate deception, participants in the market are nevertheless navigating a space they do not fully understand. The typical job seeker will not know what opportunities they are *not* being shown or what positions they have not been recommended for. Similarly, a company might select a target population, but its advertisements will be delivered

biases to encourage choices consistent with good policy—can be seen as a form of manipulation. It is arguably justified when used to achieve socially acceptable policy goals, such as encouraging savings or reducing tobacco use. See Richard H. Thaler & Cass R. Sunstein, *Nudge: Improving Decisions About Health, Wealth, and Happiness* 71–73 (2008).

¹⁰³ E.g., Susser et al., *supra* note 100, at 6–8. They argue that online manipulation “violates its target’s autonomy” because it “alter[s] their decision-making process without their conscious awareness.” *Id.* at 8–9; see also Calo, *supra* note 2, at 1031–32 (explaining how online manipulation exploits consumers’ decision-making vulnerabilities); Cass R. Sunstein, *Fifty Shades of Manipulation*, 1 *J. Marketing Behav.* 213, 217 (2015) (describing manipulation as influence that does not sufficiently engage rational deliberation).

¹⁰⁴ Calo, *supra* note 2, at 1024–25 (describing types of harms that can result from manipulation of consumers); Sunstein, *supra* note 103, at 226–30 (explaining how manipulation can undermine autonomy and reduce welfare).

only to a subset of those in the eligible pool—an audience determined by the algorithm. On a job-matching platform, the employer may see a selected list of candidates without knowing exactly how candidates have been screened. Not only is it difficult for job seekers and employers to learn of any biases affecting ad-targeting or recommendation systems, but they also have limited ability to control what information these systems deliver to or about them.¹⁰⁵ The net effect is that the visibility and salience of different opportunities will vary across individuals in a way that is opaque to participants in the labor market.

Scholars writing about online manipulation have focused on autonomy concerns,¹⁰⁶ overlooking the fact that in opportunity markets manipulation by intermediaries can impact equality as well. More specifically, the risk is that these tech intermediaries will direct opportunities in ways that reflect past patterns of inequality, thereby reproducing or even reinforcing historical disadvantages faced by certain groups. This effect should not be surprising. Predictive algorithms are built to uncover existing patterns in the data, and one of the enduring patterns in American economic life is the unequal distribution of opportunities along the lines of race, gender, and other personal characteristics. As a result, the very feature that makes algorithms so powerful—their ability to recognize patterns and predict behaviors—may cause them to reproduce existing patterns of segregation and exclusion.

The market for labor, for example, has been characterized by segregation along race and gender lines throughout this country's history. Initially, occupational segregation was imposed through law—most obviously, laws that enforced a system of enslaved labor based on race. In addition, beliefs about women's special domestic functions were used to justify legislation limiting their occupational choices.¹⁰⁷ After the abolition of the formal slave labor system, the law continued to protect systems of occupational segregation by tolerating them as the outcome of private choices. Prior to the mid-twentieth century, employers faced few legal impediments if they refused to hire workers based on race or gender

¹⁰⁵ See, e.g., Sharone, *supra* note 26, at 5, 20 (describing how social networking sites such as LinkedIn constrain the ability of job seekers to shape their job searches by, for example, limiting how they can describe their job histories and career goals).

¹⁰⁶ See Calo, *supra* note 2, at 1031–33; Sunstein, *supra* note 103, at 217–20; Susser et al., *supra* note 3, at 34–41.

¹⁰⁷ See, e.g., *Muller v. Oregon*, 208 U.S. 412, 421 (1908); *Bradwell v. Illinois*, 83 U.S. (16 Wall.) 130, 141–42 (1873).

or if they explicitly stratified their work forces along those lines.¹⁰⁸ Only since 1964, when the Civil Rights Act was passed,¹⁰⁹ has the law comprehensively targeted inequality in employment. Title VII of that Act made employer discrimination by private employers illegal on a national basis for the first time. The Civil Rights Act was not merely about constraining private choices; rather, it articulated a national policy aimed at dismantling systems of segregation that were endemic to American economic and political systems.¹¹⁰

After the passage of Title VII, occupational segregation along race and gender lines eased somewhat, but it remains a persistent feature of the labor market. Women's occupational segregation peaked in the 1960s and then declined significantly through the early 1990s.¹¹¹ Progress toward reduced gender segregation stalled in the early part of this century.¹¹² This basic temporal pattern holds true across educational levels, although occupations requiring less education show more gender segregation than those requiring higher levels of education.¹¹³ Using a dissimilarity index, researchers have estimated that about half of female workers would have to change occupations to eliminate the gender disparities.¹¹⁴ Occupational segregation by race and ethnicity—especially for blacks and Hispanics—also remains a notable feature of the labor market. Segregation of black workers declined immediately following the passage of Title VII, but progress toward desegregation appears to have stopped around 1980, and

¹⁰⁸ EEOC, Legislative History of Titles VII and XI of Civil Rights Act of 1964, at 1–7 (1968), <https://babel.hathitrust.org/cgi/pt?id=uc1.32106006452418&view=1up&seq=13> [<https://perma.cc/XW86-9AP2>] (recounting efforts to prohibit discrimination in employment prior to the passage of the Civil Rights Act of 1964).

¹⁰⁹ Civil Rights Act of 1964, 42 U.S.C. § 2000a–2000h-6 (2012).

¹¹⁰ Clay Risen, *The Bill of the Century: The Epic Battle for the Civil Rights Act 1–6*, 240–57 (2014).

¹¹¹ Dawn Michelle Baunach, Trends in Occupational Sex Segregation and Inequality, 1950 to 1990, 31 *Soc. Sci. Res.* 77, 77–78 (2002); Susan M. Carlson, Trends in Race/Sex Occupational Inequality: Conceptual and Measurement Issues, 39 *Soc. Probs.* 268, 268 (1992); Donald Tomaskovic-Devey, Kevin Stainback, Tiffany Taylor, Catherine Zimmer, Corre Robinson & Tricia McTague, Documenting Desegregation: Segregation in American Workplaces by Race, Ethnicity, and Sex, 1966–2003, 71 *Am. Soc. Rev.* 565, 567 (2006).

¹¹² Ariane Hegewisch & Heidi Hartmann, Inst. for Women's Policy Research, Occupational Segregation and the Gender Wage Gap: A Job Half Done 6 (2014).

¹¹³ *Id.* at 7.

¹¹⁴ *Id.* at 6; see also Tomaskovic-Devey et al., *supra* note 111, at 585 (“[W]e can be confident that well more than half of all workers would have to switch jobs to create a sex-neutral employment distribution.”).

some studies suggest that it may have gotten worse in recent years.¹¹⁵ Researchers estimated in the early 2000s that around half of all black workers would need to switch jobs to eliminate occupational segregation.¹¹⁶

Occupational segregation matters because women and racial minorities are not just disproportionately found in certain types of jobs but are also typically overrepresented in occupations that pay less and offer poorer working conditions. Studies suggest that occupational segregation contributes significantly to the gender wage gap,¹¹⁷ and, as a result, pay disparities tend to decrease as occupational segregation declines. Of course, differences in education and training also contribute to occupational segregation, but researchers have found that gender segregation occurs even among jobs requiring comparable levels of education and training.

Thus, it is a fact of economic life that men and women and different racial groups are distributed differently across occupations. Predictive algorithms are likely to reflect those patterns as they disseminate information about available opportunities. Because most predictive systems take observed patterns like these as a given, they will associate certain jobs with certain types of people—for example, pipefitters with males, or nurses with females. These patterns likely reflect past legal controls over women’s labor and ongoing cultures of harassment in certain workplaces, but the reasons they exist are irrelevant to predictive algorithms.

The problem is not solely that the output of these systems will reflect existing patterns of occupational segregation. These predictive technologies may themselves influence worker or employer behavior in ways that reinforce those patterns, stalling or even reversing progress toward greater race and gender integration in the workplace. If women do not receive information about positions as carpenters or electricians, and if

¹¹⁵ Donald Tomaskovic-Devey & Kevin Stainback, *Discrimination and Desegregation: Equal Opportunity Progress in U.S. Private Sector Workplaces Since the Civil Rights Act*, 609 *Annals Am. Acad. Pol. & Soc. Sci.* 49, 51 (2007); Tomaskovic-Devey et al., *supra* note 111, at 583–84.

¹¹⁶ See Tomaskovic-Devey et al., *supra* note 111, at 585.

¹¹⁷ Ariane Hegewisch & Maxwell Matite, *Inst. for Women’s Policy Research, The Gender Wage Gap by Occupation 1* (2013); Ariane Hegewisch, Hannah Liepmann, Jeffrey Hayes & Heidi Hartmann, *Inst. for Women’s Policy Research, Separate and Not Equal? Gender Segregation in the Labor Market and the Gender Wage Gap 1* (2010); Stephen J. Rose & Heidi I. Hartmann, *Inst. for Women’s Policy Research, Still a Man’s Labor Market: The Long-Term Earnings Gap*, at iv (2004).

men do not receive ads for careers in nursing or elementary education, they are less likely to consider these possibilities. Because the information environment is part of the social context that shapes individual preferences and choices, these patterns could over time strengthen patterns of occupational segregation through feedback effects.

To be clear, I am not suggesting that tech intermediaries are deliberately trying to exclude racial and ethnic minorities, women, or other disadvantaged groups from certain opportunities. However, even if discriminatory effects are not intended, they can cause real harm. Workers who are excluded from competing for opportunities will obtain worse outcomes in the labor market, whether their exclusion is the result of a biased human or an algorithm that systematically produces the same result. Anti-discrimination law has long recognized that neutral practices can have discriminatory effects that also warrant a legal response. Thus, the Supreme Court famously held in *Griggs v. Duke Power Co.* that Title VII prohibits “not only overt discrimination but also practices that are fair in form, but discriminatory in operation.”¹¹⁸ The statute, it held, targets “the *consequences* of employment practices, not simply the motivation.”¹¹⁹ Even if tech intermediaries have no invidious motives, by relying on predictive algorithms, they risk reproducing the patterns of segregation and disadvantage that the Civil Rights Act was intended to dismantle. If that occurs, tech platforms should not be absolved of responsibility for the consequences of the systems they have built.

E. Possible Objections

My argument that tech platforms should bear responsibility for discriminatory effects they produce as intermediaries in labor and other opportunity markets is likely to provoke objections. In this Section, I consider and respond to likely objections.

Haven't we seen this before? This objection rests on the observation that the labor market has long been plagued by systemic biases, and argues that there is nothing new or special about what is described here. It is certainly true that some employment agencies and labor unions have referred jobs in discriminatory ways. And employers with homogenous workforces sometimes relied on their current employees to spread the word about job openings, thereby effectively excluding workers with

¹¹⁸ 401 U.S. 424, 431 (1971).

¹¹⁹ *Id.* at 432.

different racial or ethnic backgrounds. Similarly, because some print and broadcast media aimed at niche markets, job postings could implicitly be targeted based on demographic characteristics. Today's tech intermediaries are not replacing some idealized world in which labor markets operated in a completely transparent and unbiased way.

These earlier practices, however, did not escape scrutiny. Where there was bias, the law stepped in. Both unions and employment agencies could be held liable if they deliberately discriminated in allocating job opportunities, as well as when they employed neutral practices that had a disparate impact.¹²⁰ Courts also found firms liable under Title VII for relying on word-of-mouth hiring where the effect was to reproduce a racially homogenous workforce.¹²¹ Similarly, an employer was found in violation of Title VII when it advertised job openings in a local paper in an all-white community, rather than a metropolitan newspaper that reached a readership with a substantial minority population.¹²² Just as these practices led to liability in the past, tech intermediaries should also be responsible for analogous practices today. As the Ninth Circuit has observed, things that are unlawful if they occur offline "don't magically become lawful" because they occur online.¹²³

The fact that labor markets sometimes operated in biased ways in the past does not justify new systems of distributing opportunities that have similar effects. Tech intermediaries are becoming increasingly central to the job search process, and their impact will likely differ in degree, perhaps even in kind, from earlier labor market intermediaries. Their national or global reach and access to vast amounts of personal data give them the potential to scale up pre-existing market biases. Traditional word-of-mouth hiring was inherently limited to the contacts of existing employees. By contrast, tools like recommender systems on job-matching platforms or Facebook's Lookalike Audience feature can be applied to vastly larger populations. In addition, the highly personalized interactions on tech platforms tend to render their effects less visible. Women might not typically read *Esquire* magazine, but they had the ability to pick up a

¹²⁰ 42 U.S.C. § 2000e-2(b)-(c) (2012).

¹²¹ See, e.g., *Thomas v. Wash. Cty. Sch. Bd.*, 915 F.2d 922, 925 (4th Cir. 1990); *Franks v. Bowman Transp. Co.*, 495 F.2d 398, 418-20 (5th Cir. 1974), rev'd on other grounds, 424 U.S. 747, 779-80 (1976).

¹²² *United States v. City of Warren*, 138 F.3d 1083, 1092-93 (6th Cir. 1998).

¹²³ *Fair Hous. Council v. Roommates.com, LLC*, 521 F.3d 1157, 1164 (9th Cir. 2008) (en banc).

copy and see what opportunities were advertised there. By contrast, when information about opportunities is fully customized online, there may be no way to know what opportunities are being shown to others.

Although it may be difficult for individuals to recognize when bias is affecting their access to opportunities, the growing role of tech intermediaries may, ironically, make broad systemic effects more visible. In the past, the interaction of multiple, decentralized actors meant that the causes of systemic disadvantage were often diffuse and difficult to pin down. As the job search process is increasingly concentrated on online platforms, it will expand the possibilities for studying how information and opportunities are distributed as well as which factors influence distributive outcomes. The fact that bias has always existed in labor markets is thus not a reason for accepting biased outcomes that occur online. Instead, to the extent tech intermediaries reproduce past patterns of segregation and inequality, there is an opportunity to better understand and to ameliorate those effects.

That's just the way the world is. Another objection might be that any allegedly discriminatory outcomes simply reflect actual patterns in the world. This argument asserts that predictive algorithms are neutral tools that are capturing real differences among workers. The patterns they uncover might be unfortunate, or even morally troubling, but they are not the responsibility of tech intermediaries.

This argument rests on the mistaken assumption that algorithms are neutral and objective and that the results they produce are inevitable. In fact, as many scholars have pointed out, technical systems are not neutral, and data can be biased.¹²⁴ Predictive algorithms do not just evolve organically. Rather, they are created by humans who design and deploy them for particular purposes. At every step of the way, the humans who build these tools are making decisions—about what problem they are trying to solve, what outcome they are trying to optimize, how the target variable should be measured, what data will be used to build the model,

¹²⁴ Barocas & Selbst, *supra* note 15, at 674; danah boyd & Kate Crawford, Critical Questions for Big Data: Provocations for a Cultural, Technological, and Scholarly Phenomenon, 15 *Info. Comm. & Soc'y* 662, 666–68 (2012); Cynthia Dwork & Deirdre K. Mulligan, *It's Not Privacy, and It's Not Fair*, 66 *Stan. L. Rev. Online* 35, 35 (2013); Joshua A. Kroll, Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson & Harlan Yu, *Accountable Algorithms*, 165 *U. Pa. L. Rev.* 633, 633–34 (2017); Kate Crawford, *Think Again: Big Data*, *Foreign Pol'y* (May 10, 2013, 12:40 AM), <http://foreignpolicy.com/2013/05/10/think-again-big-data/> [<https://perma.cc/9HHU-6UAP>].

and so on.¹²⁵ Any of those decisions may involve value choices, and all will affect the outputs of the model.

YouTube provides a vivid example of how a platform's optimization goals can have secondary consequences. According to former YouTube employees, the company concluded that it could earn more money displaying ads if viewers spent more time on the site. It set out to optimize "engagement"—namely, to hold users' attention as long as possible. Pursuing that goal, it found that users could be induced to spend more time on the site if its recommender algorithm promoted related, but more extreme, content. As a result, the system YouTube developed directs viewers to material with increasingly high shock value, leading them to videos with false information, extremist political views, and violent and sexualized content.¹²⁶ By one metric—the one chosen by YouTube—the algorithm has been wildly successful. Within several years, it grew viewership to achieve the company's goal of one billion hours of viewing a day.¹²⁷

As the YouTube example illustrates, the choices that go into building a model can have significant social impacts. In the case of labor market intermediaries, the risk is that those choices reproduce past patterns of discrimination. Online advertising platforms typically seek to optimize user interaction in order to maximize their revenue. Thus, their ad-targeting algorithms turn on predictions of who will click on an ad, not who is qualified for the advertised position. And if, for example, women are predicted to click frequently on other types of ads, the algorithm may not deliver information to them about career opportunities—not because women would not be interested, but because the platform can earn more money by showing them other kinds of ads.

Similarly, job-matching platforms typically seek to optimize matches, not productivity on the job. To that end, they focus on predicting which job seekers will apply for a job and which types of workers an employer will prefer. These types of predictions are quite different from trying to determine who is capable of doing a job, and because they focus on preferences, they are more likely to incorporate implicitly biased

¹²⁵ Barocas & Selbst, *supra* note 15, at 677–78; Lehr & Ohm, *supra* note 18, at 655.

¹²⁶ Mark Bergen, YouTube Executives Ignored Warnings, Letting Toxic Videos Run Rampant, *Bloomberg* (Apr. 2, 2019, 5:00 AM), <https://www.bloomberg.com/news/features/2019-04-02/youtube-executives-ignored-warnings-letting-toxic-videos-run-rampant> [<https://perma.cc/KCM3-A47K>].

¹²⁷ *Id.*

judgments. To the extent that labor market intermediaries channel information based on something other than who can successfully perform a job, their predictions will not necessarily reflect actual differences in job-relevant skills and abilities.

Of course, job-relevant skills do differ across individuals, and these skills will sometimes vary across demographic groups as well. The mere fact that a recommender or ad-delivery system produces a demographic skew does not mean unfair discrimination has occurred. On the other hand, just because the distribution of information is based on an algorithm does not mean that it is a fair and accurate reflection of actual relevant differences. Tech intermediaries build these systems to optimize their own interests—generally to maximize revenue. Unless they are designed to take equality and fairness concerns into account as well, these systems will produce more clicks or more matches, but not necessarily in a fair and unbiased manner.

A variation on the “*That’s just the way the world is*” objection posits that the patterns uncovered by predictive algorithms are accurate because they reflect what workers actually want. Most men are not interested in being preschool teachers; most women will not apply for jobs as carpenters or pipefitters. Giving them information about opportunities they are uninterested in would be highly inefficient. There are undoubtedly efficiency gains if information can be matched to worker preferences. The systems used by intermediaries, however, do not directly measure preferences, but they instead rely on past behavior to infer a worker’s “true” preferences. Those inferences may be based on a worker’s expressed preferences, but more often they are predicted based on the behavior of other people determined to be similar.

Relying on past behavior to make predictions assumes that workers have fixed preferences that pre-exist their labor market experiences and are stable over time. This thinking is also reflected in the “lack of interest” defense employers have relied on in the past to respond to claims of systemic discrimination. They argued that the low numbers of racial minorities or females in their workforces result from these workers’ lack of interest in the positions, rather than any discriminatory practices on their part. Courts have rejected this argument in a number of race discrimination cases, finding that the reason few black workers applied

was the employer's past discrimination¹²⁸ or its neutral recruitment practices that perpetuated past patterns of discrimination.¹²⁹

On the other hand, courts have been more receptive to the lack of interest defense when they believed low application numbers resulted from societal forces beyond the employer's control. In a number of sex discrimination cases, courts interpreted statistics that few females applied for traditional male jobs as evidence that women simply were not interested in those positions. This rationale led them to accept sex-segregated workplaces as the result not of systemic discrimination, but of women's voluntary choices.¹³⁰

Vicki Schultz has criticized these decisions for failing to recognize that women's work preferences are formed in response to structural and cultural features of the labor market.¹³¹ The lack of interest argument rests on the incorrect assumption that women have stable, pre-existing preferences for certain types of jobs.¹³² She canvasses the sociological literature to show that far from being fixed, women's work aspirations have changed in response to the opportunities available to them, not only across generations, but also within the lifetimes of individual women.¹³³ Given this evidence, she argues that it is a mistake for courts to accept existing patterns of sex segregation as reflecting women's choices and to absolve employers of responsibility.¹³⁴ Schultz's critique applies as well to tech intermediaries in the labor market. By relying on past patterns to infer workers' preferences, platforms overlook the ways in which workers' interests are shaped by their experiences. A woman who would consider non-traditional work in the trades is less likely to express that

¹²⁸ See, e.g., *Jones v. Lee Way Motor Freight, Inc.*, 431 F.2d 245, 247 (10th Cir. 1970); *United States v. Sheet Metal Workers Int'l Ass'n, Local Union No. 36*, 416 F.2d 123, 139 (8th Cir. 1969).

¹²⁹ See, e.g., *NAACP v. N. Hudson Reg'l Fire & Rescue*, 665 F.3d 464, 485 (3d Cir. 2011); *United States v. City of Warren*, 138 F.3d 1083, 1092–93 (6th Cir. 1998); *Thomas v. Wash. Cty. Sch. Bd.*, 915 F.2d 922, 925 (4th Cir. 1990); *Franks v. Bowman Transp. Co.*, 495 F.2d 398, 415 (5th Cir. 1974), rev'd on other grounds, 424 U.S. 747, 779–80 (1976).

¹³⁰ See, e.g., *EEOC v. Sears, Roebuck & Co.*, 839 F.2d 302, 320 (7th Cir. 1988).

¹³¹ Vicki Schultz, *Telling Stories About Women and Work: Judicial Interpretations of Sex Segregation in the Workplace in Title VII Cases Raising the Lack of Interest Argument*, 103 *Harv. L. Rev.* 1749 (1990).

¹³² Some courts have accepted these arguments, often relying on their notion of "common practical knowledge" that certain jobs are just "not attractive to females." *EEOC v. Mead Foods, Inc.*, 466 F. Supp. 1, 3 (W.D. Okla. 1977). This "common practical knowledge" typically reflected stereotyped notions of appropriate work for women and men.

¹³³ See Schultz, *supra* note 131, at 1824–26 & n.288.

¹³⁴ *Id.* at 1831–32.

interest if she never receives information about such opportunities. As tech intermediaries channel the flow of information between employers and job seekers, they can wield significant influence over how these labor market participants perceive their choices and available alternatives. Their influence does not merely reflect the behavior of market participants but likely shapes it as well.

How much of a problem is this, really? A third objection asks whether the risk that tech intermediaries will contribute to discriminatory employment patterns is really all that significant. This is, of course, an empirical question—one on which there has been relatively little study to date. As recounted in Section II.C, there are a handful of studies documenting that discriminatory patterns do occur. And as explored in Section II.D, there are good theoretical reasons to expect that reliance on predictive algorithms can reproduce past patterns of discrimination and exclusion. Nevertheless, I agree that the actual impact of these technologies remains uncertain.

On the one hand, automated systems have the potential to avoid some forms of human bias—for example, where an employment agency relies on human judgments when choosing which workers to refer. And even if platforms sometimes distribute information in biased ways, the growth of online recruiting may create greater competition among intermediaries and more alternatives for workers to learn about job opportunities. If the overall availability of information about crucial opportunities expands, the fact that some systems deliver information in biased ways may not matter as much.

In addition, whether these systems have discriminatory effects depends upon the details of how they are designed and deployed. Data will inevitably reveal patterns of occupational segregation, but designers can make different choices about *which* data to rely on or *how* they are used. Rather than delivering ads and making recommendations that merely reflect observed occupational patterns, system designers might recognize the risks of reproducing bias and impose technical constraints to ensure greater fairness. For example, LinkedIn’s recruiting tool displays results in a manner that reflects the proportion of women in the eligible pool, rather than listing candidates solely in order of predicted “relevance.”¹³⁵ Platforms may develop other ways of removing or countering biases in the data.

¹³⁵ Bogen & Rieke, *supra* note 7, at 24.

In reality, there is not a single labor market but many. Janitors and stockbrokers are not competing in the same market. The role of tech intermediaries, and therefore their effect on the distribution of opportunities, is likely to vary considerably by industry as well as skill level. For example, it is highly likely that recruitment for the most elite jobs will continue to rely on informal networks and offline connections. When jobs require skills that are more common and applicants are numerous, firms may lean more heavily on online intermediaries to screen and sort potential hires. And in other markets, such as for unskilled physical labor, needs may be intermittent and heavily location-based, obviating the usefulness of online intermediaries. In short, the impact of predictive technologies is likely to be quite different for different types of workers.

There are enough moving parts in the overall environment to be uncertain about the exact effects as the recruitment process moves online. The appropriate response to this uncertainty is not to disregard the role of tech intermediaries but to subject them to further empirical study. The platforms themselves are in the best position to undertake these studies. However, because their choices can have significant social impacts, their operations should also be subject to some form of third party or public oversight. In the next Sections, I consider what forms that scrutiny might take, examining first the possibilities under existing anti-discrimination law and then suggesting reasons for turning to regulatory strategies instead.

III. INTERMEDIARIES AND ANTI-DISCRIMINATION LAW

When Congress enacted Title VII of the Civil Rights Act of 1964, it targeted each of the mechanisms that might operate to produce discrimination in an effort to dismantle occupational segregation.¹³⁶ The

¹³⁶ Congress has, of course, passed other laws prohibiting employment discrimination, such as the Age Discrimination in Employment Act of 1967 (ADEA), 29 U.S.C. §§ 621–634 (2012), and the Americans with Disabilities Act of 1990 (ADA), 42 U.S.C. §§ 12101–12213 (2012). This Article focuses on Title VII and the case law interpreting it, necessarily glossing over some differences in other anti-discrimination statutes. For example, the ADEA has a narrower definition of “employment agency.” Compare 29 U.S.C. § 630(c) (“The term ‘employment agency’ means any person regularly undertaking with or without compensation to procure employees for an employer and includes an agent of such a person; but shall not include an agency of the United States.”), with 42 U.S.C. § 2000e(c) (“The term ‘employment agency’ means any person regularly undertaking with or without compensation to procure employees for an employer *or to procure for employees opportunities to work for an employer*”).

statute not only forbids discrimination by employers,¹³⁷ but it also prohibits employment agencies from making discriminatory referrals and unions and apprenticeship committees from discriminating in admitting, referring, or training workers.¹³⁸ In addition, it makes it unlawful for any of these entities to post job advertisements that express a discriminatory preference or limitation.¹³⁹ Congress thus addressed the entities with significant influence over labor market outcomes and the mechanisms known at that time to result in discrimination. Today, the operation of the labor market is radically different from 1964, with online platforms playing a crucial role in shaping how information about available opportunities is distributed and who is considered for which jobs.

The growing influence of tech intermediaries has a number of implications for Title VII doctrine. First, as discussed in Section III.A below, it complicates the question of who is an applicant for purposes of determining employer liability. The answer to that question matters because it determines who can sue and, in class cases, may affect whether plaintiffs can rely on statistical proof of discrimination. It also raises the question of whether tech platforms are the type of intermediary that can be held directly liable for their role in producing discriminatory outcomes. Section III.B examines whether platforms meet the definition of “employment agencies” covered by Title VII. Section III.C considers an alternative theory, asking whether platforms can be held liable for interfering with third party employment relationships.¹⁴⁰

and includes an agent of such a person.” (emphasis added)). Also, unlike Title VII, the ADEA does not permit disparate impact claims by applicants. See *Kleber v. CareFusion Corp.*, 914 F.3d 480, 484–85 (7th Cir. 2019) (en banc).

The ADA also differs from Title VII in important ways. It not only prohibits discrimination on the basis of disability, but it also requires reasonable accommodation of applicants and employees with disabilities. 42 U.S.C. § 12112(b)(5)(A). As applicant recruiting and screening processes move online, the way tech intermediaries structure their platforms may affect the ability of individuals with disabilities to access job opportunities. The impact of tech intermediaries on the application of anti-discrimination laws other than Title VII calls for more sustained analysis than I can provide here.

¹³⁷ 42 U.S.C. § 2000e-2(a)(1).

¹³⁸ *Id.* § 2000e-2(b) (forbidding discriminatory referrals); *id.* § 2000e-2(c) (forbidding labor organizations from discriminatorily admitting and referring members); *id.* § 2000e-2(d) (forbidding discrimination in training programs).

¹³⁹ *Id.* § 2000e-3(b).

¹⁴⁰ This Section does not examine § 704(b) of Title VII, which prohibits employers, agencies, and labor organizations from publishing employment ads that express a discriminatory preference. *Id.* § 2000e-3(b). Whether this prohibition applies to tech intermediaries depends upon whether they are “employment agencies”—the issue discussed

A. Who is an Applicant?

The overwhelming majority of cases under Title VII seek to hold employers liable for discriminatory employment practices. When the alleged discrimination occurred in the hiring process, a crucial question is “who is an applicant?” In both individual and class cases, the answer to this question can determine the outcome. Complainants are required to establish that they have actually applied—usually by showing that they completed an application form or made express inquiries about an open position. Today, however, employer practices like passive recruiting and relying on job-matching intermediaries complicates the question of who should be considered an applicant. Because tech platforms control the flow of information, they can significantly affect whether job seekers apply for or are even aware of particular opportunities.

In individual failure-to-hire cases, a plaintiff must first establish the elements of a prima facie case, including showing “that he *applied* and was qualified for a job for which the employer was seeking applicants.”¹⁴¹ According to the Supreme Court, the prima facie case “serves an important function in the litigation: it eliminates the most common nondiscriminatory reasons for the plaintiff’s rejection.”¹⁴² In other words, if a worker did not apply, the employer’s failure to hire cannot be the result of discrimination but was caused by the worker’s own choices. Courts have dismissed numerous cases on the grounds that the plaintiff never applied and therefore cannot establish a prima facie case.¹⁴³ As one court put it, “we do not believe an employer is obliged to defend its

in Section III.B. Assuming that they are employment agencies, § 704(b) is relevant only when the content of an ad or the context in which it is published expresses a discriminatory intent. Other work has explored how that prohibition applies to online advertising. See Kim & Scott, *supra* note 27. Because the focus of this Article is discriminatory effects that can occur when job opportunities are mediated through predictive algorithms, § 704(b) is only marginally relevant, and I do not discuss it in detail here.

¹⁴¹ *McDonnell Douglas Corp. v. Green*, 411 U.S. 792, 802 (1973) (emphasis added).

¹⁴² *Tex. Dep’t of Cmty. Affairs v. Burdine*, 450 U.S. 248, 253–54 (1981).

¹⁴³ See, e.g., *Velez v. Janssen Ortho, LLC*, 467 F.3d 802, 808–09 (1st Cir. 2006); *Yartzoff v. Thomas*, 809 F.2d 1371, 1375 (9th Cir. 1987). The same principle applies in failure to promote cases. See, e.g., *Williams-Boldware v. Denton Cty.*, 741 F.3d 635, 644 (5th Cir. 2014); *Johnson v. Gen. Bd. of Pension & Health Benefits of the United Methodist Church*, 733 F.3d 722, 728 (7th Cir. 2013); *Petrosino v. Bell Atl.*, 385 F.3d 210, 227 (2d Cir. 2004). In cases in which an employer relies on algorithms to identify promotion prospects rather than allowing employees to apply for open positions, similar concerns arise.

decision not to hire an individual for a position for which she has not specifically applied.”¹⁴⁴

Although the general rule is that a plaintiff must have applied, courts have found an exception where an employer’s discriminatory practices make such an effort futile. In *International Brotherhood of Teamsters v. United States*, the Supreme Court recognized that an employer’s discriminatory policy “can surely deter job applications from those who are aware of it and are unwilling to subject themselves to the humiliation of explicit and certain rejection.”¹⁴⁵ Workers who fail to apply because doing so would be futile are also victims of discrimination and entitled to sue. Non-applicants must show that they were deterred by “a justifiable belief that the employer’s discriminatory practices made application a futile gesture.”¹⁴⁶ Courts differ as to what facts are sufficient to show that a job seeker was deterred from applying, but at the very least, the plaintiff must have known about the opportunity and had specific reasons to believe the effort would be futile.¹⁴⁷

Requiring a plaintiff to apply (or show the futility of applying) in order to establish a prima facie case makes sense when the focus is on employer liability. However, it overlooks the process by which workers learn about employment opportunities and come to apply for them. If discrimination or bias affects the process *prior* to a job seeker submitting an application, it will go unaddressed if the focus is solely on employer behavior. Even the “futile gesture” doctrine will not address this problem, because the exception depends on the *employer’s* behavior, not that of information intermediaries. If job seekers never learn about an opportunity, they will not apply. And because they were wholly unaware of the opportunity, they cannot argue that they were deterred from applying and should still be permitted to sue.

Determining who is an applicant also matters quite a bit for Title VII class cases. Classwide discrimination can be challenged under either a pattern or practice theory, which alleges systematic, purposeful

¹⁴⁴ *Velez*, 467 F.3d at 808.

¹⁴⁵ 431 U.S. 324, 365 (1977).

¹⁴⁶ *EEOC v. Joe’s Stone Crabs, Inc.*, 296 F.3d 1265, 1274 (11th Cir. 2002); see also *McDonald v. Gen. Mills, Inc.*, 387 F. Supp. 24, 37 (E.D. Cal. 1974) (“[I]t is sufficient for plaintiff to allege in her complaint that she was deterred from making application and seeking an interview with defendants because of their employment practices. Where the defendants had indicated an express preference to interview and hire ‘males’ it can reasonably be assumed that plaintiff would believe application for the jobs would be a futile gesture.”).

¹⁴⁷ *Malarkey v. Texaco, Inc.*, 983 F.2d 1204, 1213 (2d Cir. 1993).

discrimination, or a disparate impact theory, which challenges the use of facially neutral employment practices that disproportionately harm workers in protected groups. Both pattern or practice and disparate impact suits rely on statistical evidence that involves comparing the race or gender composition of a company's workforce with the composition of the "relevant labor market." The parties often dispute what the relevant comparison group should be, and the choice can matter quite a bit, sometimes determining whether plaintiffs can establish a prima facie case at all.¹⁴⁸

The Supreme Court has never laid down clear rules for determining the relevant comparison pool in class cases. In *Wards Cove Packing Co. v. Atonio*, it suggested that applicant flow data was most relevant.¹⁴⁹ Elsewhere, however, it has clearly stated that "[t]here is no requirement . . . that a statistical showing of disproportionate impact must always be based on . . . actual applicants."¹⁵⁰ In some cases, it has approved the use of general labor market statistics. For example, in *Teamsters*, it was appropriate to compare the percentage of blacks and Hispanics in the employer's workforce with their percentages in the general population because "the job skill there involved—the ability to drive a truck—is one that many persons possess or can fairly readily acquire."¹⁵¹ In *Hazelwood* as well, the majority agreed that the relevant statistic was the proportion of black teachers in the labor market overall,

¹⁴⁸ For example, *Hazelwood School District v. United States* involved allegations of hiring discrimination based on the low percentage of black teachers—1.4% to 1.8% in the relevant time period—employed by the district. 433 U.S. 299, 305 (1977). Depending upon whether the relevant comparison pool was teachers in the metropolitan area (15.4% black) or teachers only in the surrounding county (5.7% black), *id.* at 310–11, the size of the racial disparity would look quite different. In his concurrence, Justice White suggested that the applicant pool—presumably with an even lower percentage of black candidates—would be a more appropriate comparison. *Id.* at 347 (White, J., concurring). The Supreme Court held that the "proper comparison" was to "the qualified public school teacher population in the relevant labor market," *id.* at 308, but did not decide which of the proffered statistics was the correct one, remanding instead for the district court to decide, *id.* at 313.

¹⁴⁹ 490 U.S. 642, 651–52 (1989). The Court in *Wards Cove* criticized the Court of Appeals' reliance on an improper statistical comparison to find disparate impact, explaining that "[i]f the absence of minorities holding [] skilled positions is due to a dearth of qualified nonwhite applicants," it would be improper to find that the employer's practices caused a disparate impact. *Id.* Absent evidence that there were "barriers or practices" that deterred nonwhites from applying, the relevant comparison was between the racial composition of selected applicants and the composition of the qualified *applicant pool*. *Id.* at 653.

¹⁵⁰ *Dothard v. Rawlinson*, 433 U.S. 321, 330 (1977).

¹⁵¹ *Hazelwood Sch. Dist.*, 433 U.S. at 308 n.13 (explaining the use of labor market statistics in *Teamsters*).

questioning only how the relevant geographic area should be defined.¹⁵² Under Supreme Court precedent, then, there are no hard and fast rules; instead, the appropriate comparison group depends upon the factual circumstances.

Given this fact-based approach, employers responding to class claims often argue that the relevant comparison pool should be the pool of qualified *applicants*.¹⁵³ They point to the low numbers of racial or ethnic minorities or women who have applied and argue that the underrepresentation of these groups in their workforces is the result of the workers' lack of interest in the positions, rather than any discriminatory practices on their part.¹⁵⁴ To the extent that employers rely on applicant data to counter class claims built on statistical evidence, the definition of who is an applicant can have a dramatic effect.

Determining who is an applicant has become increasingly difficult with technological advances. The Office of Federal Contract Compliance Programs (OFCCP) issued regulations in 2006 attempting to define an Internet applicant.¹⁵⁵ That definition focused on an individual's expression of interest and direct interactions with a firm, an approach that is already obsolete. Because of the rise of tech intermediaries in the labor market, applicants and firms may not directly interact until quite late in the sourcing process. Consider, for example, an employer that relies on passive recruiting systems to find candidates. The intermediary may analyze the profiles of hundreds or thousands of persons before suggesting a handful of candidates for the employer to approach. Defining the relevant applicant pool to include anyone whose data was assessed by the system seems far too broad. Yet defining the applicant pool to include only persons who actually submit an application after being recruited will narrow the focus so sharply that any bias occurring in the recruiting process will disappear from view. Similar questions are raised by job-matching platforms. When job seekers post their profiles, should they be considered applicants for all available relevant jobs? Or only those jobs that they actually apply for? What if, due to the matching algorithm,

¹⁵² *Id.* at 310–11; accord *id.* at 315 (Stevens, J., dissenting).

¹⁵³ *Anderson v. Douglas & Lomason Co.*, 26 F.3d 1277, 1286 (5th Cir. 1994).

¹⁵⁴ See, e.g., *EEOC v. Joint Apprenticeship Comm. of the Joint Indus. Bd. of the Elec. Indus.*, 186 F.3d 110, 120 (2d Cir. 1998); *EEOC v. O & G Spring & Wire Forms Specialty Co.*, 38 F.3d 872, 875 (7th Cir. 1994); *EEOC v. Mead Foods, Inc.*, 466 F. Supp. 1, 3 (W.D. Okla. 1977).

¹⁵⁵ 41 C.F.R. § 60-1.3 (2019). The OFCCP developed the regulations to guide federal contractors in meeting their record-keeping and reporting requirements.

information about some jobs for which they are qualified is not shown or is obscured, so that it is practically unavailable to them?

The question of who is an applicant in a world of tech intermediaries is thus a complicated one. For purposes of determining employer liability, it makes sense to focus on processes within employers' control. If an employer's own practices depress minority or female applications, its workforce composition should not escape scrutiny by referencing a narrowly defined applicant pool. If, on the other hand, low numbers of applicants result from a process outside the control or awareness of the employer, then relying on applicant flow statistics may be appropriate when weighing employer responsibility. However, the effect of such a narrowed focus is that the mechanisms contributing to ongoing occupational segregation and unequal access to jobs may occur outside of view. In order to address how sourcing processes may contribute to bias, the law also needs to consider the responsibility of tech intermediaries in shaping labor market outcomes. I turn next to the question of whether intermediaries can be held directly responsible if their choices produce discriminatory effects.

B. Platforms as Employment Agencies

Although most of the litigation under Title VII has focused on employer liability, the statute also covers certain labor market intermediaries—namely, employment agencies, labor unions, and joint labor-management apprenticeship programs.¹⁵⁶ These entities are also liable for discriminatory behavior because they control or influence access to job opportunities and thus can contribute to workplace segregation and inequality. As one court explained, “Congress has determined to prohibit each of these from exerting any power it may have to foreclose, on invidious grounds, access by any individual to employment opportunities”¹⁵⁷ in order to achieve its goal of “provid[ing] equal access to the job market.”¹⁵⁸

Because platforms that operate as labor market intermediaries play a role similar to these traditional intermediaries, they may fall within the statute's definition of an “employment agency.” Title VII defines an “employment agency” as “any person regularly undertaking with or

¹⁵⁶ 42 U.S.C. § 2000e-2(b)–(d) (2012).

¹⁵⁷ *Sibley Mem'l Hosp. v. Wilson*, 488 F.2d 1338, 1341 (D.C. Cir. 1973).

¹⁵⁸ *Id.* (quoting *Diaz v. Pan Am. World Airways, Inc.*, 442 F.2d 385, 386 (5th Cir. 1971)).

without compensation to procure employees for an employer or to procure for employees opportunities to work for an employer.”¹⁵⁹ There are relatively few cases interpreting the boundaries of this definition, perhaps because in the decades immediately following the passage of Title VII, it seemed obvious what an employment agency was.

The first cases interpreting the definition of “employment agency” concerned newspapers. Although Title VII prohibited job advertisements that expressed a discriminatory preference,¹⁶⁰ newspapers continued to publish help wanted ads segregated into male and female columns for years after the statute was passed.¹⁶¹ Women’s advocacy groups brought lawsuits against newspaper publishers, alleging that they were employment agencies subject to Title VII and seeking to enjoin them from printing separate ads for male and female workers. In *Brush v. San Francisco Newspaper Printing Co.*,¹⁶² a district court rejected that argument, concluding that newspapers were not “engaged in the business of procuring employees or employment opportunities” because they merely printed advertising copy provided by employers.¹⁶³

Although *Brush* suggested that newspapers were not covered by Title VII,¹⁶⁴ women’s advocacy groups succeeded a few years later under state and local anti-discrimination laws that made “aiding” and “abetting” employment discrimination an unlawful practice.¹⁶⁵ In 1973, the Supreme

¹⁵⁹ 42 U.S.C. § 2000e(c). Note that the definition of an employment agency under the ADEA is narrower than under Title VII. See *supra* note 136. Depending upon exactly how an intermediary operates, it may be an employment agency for purposes of Title VII but not for the ADEA.

¹⁶⁰ 42 U.S.C. § 2000e-3(b).

¹⁶¹ Part of the problem was that the EEOC initially took the position that sex-segregated help wanted ads did not violate Title VII. After extensive criticism by women’s advocacy groups, it reversed this position. See Nicholas Pedriana & Amanda Abraham, *Now You See Them, Now You Don’t: The Legal Field and Newspaper Desegregation of Sex-Segregated Help Wanted Ads 1965–75*, 31 *Law & Soc. Inquiry* 905, 913–14 (2006).

¹⁶² 315 F. Supp. 577 (N.D. Cal. 1970).

¹⁶³ *Id.* at 580.

¹⁶⁴ Two other cases also addressed whether newspapers that published sex-segregated help wanted ads could be held liable under Title VII, and they reached opposite conclusions. Compare *Greenfield v. Field Enters., Inc.*, No. 71 C 2075, 1972 U.S. Dist. LEXIS 15304, at *17 (N.D. Ill. Feb. 1, 1972) (concluding that they could *not* be held liable), with *Morrow v. Miss. Publishers Corp.*, No. 72J-17(R), 1972 U.S. Dist. LEXIS 10972, at *8–9 (S.D. Miss. Nov. 27, 1972) (concluding that they *could* be held liable and adopting language in *Greenfield*).

¹⁶⁵ See, e.g., Conn. Gen. Stat. § 46a-60(5) (2019); N.J. Stat. Ann. § 10:5-12(e) (West 2019); N.Y. Exec. Law § 296(6) (Consol. 2019); cf. *Passaic Daily News v. Blair*, 308 A.2d 649, 654–57 (N.J. 1973); *Nat’l Org. for Women v. State Div. of Human Rights*, 314 N.E.2d 867, 870

Court in *Pittsburgh Press Co. v. Pittsburgh Commission on Human Relations* upheld the application of such a law to a newspaper. Rejecting the publisher's First Amendment challenge, it held that newspapers could be enjoined from publishing discriminatory employment ads under the local ordinance.¹⁶⁶ Shortly thereafter, newspapers ceased the practice of publishing separate help wanted columns for men and women,¹⁶⁷ and the question of whether they were employment agencies under Title VII became moot.

Since then, a handful of cases have applied Title VII's definition of an employment agency to other labor market intermediaries. Courts have focused on whether the intermediary "regularly undertak[es] . . . to procure employees for an employer" or employment opportunities for an employee.¹⁶⁸ This statutory language indicates that an intermediary is an employment agency if it engages on a regular basis in activities intended to match workers with job opportunities. The term "regularly" suggests a recurring activity rather than a rare or occasional one.¹⁶⁹ It does not imply that such activities are an entity's exclusive function, nor does it suggest that some fixed quantum of such activity is required. Instead, whether an entity is an employment agency should turn on the nature and significance of the intermediary's job-procuring activities and the role it plays in the labor market.

This approach is consistent with how cases have been decided on their facts. For example, the district court in *Kaplowitz v. University of Chicago* found that the University of Chicago Law School was an "employment agency" under Title VII based on its extensive placement activities.¹⁷⁰ Even though the law school was primarily in the business of educating its students, it devoted significant resources to helping them find

(N.Y. 1974); State Div. of Human Rights ex rel Carey v. Binghamton Press Co., 415 N.Y.S.2d 523, 526 (N.Y. App. Div. 1979).

¹⁶⁶ 413 U.S. 376 (1973).

¹⁶⁷ Pedriana & Abraham, supra note 161, at 906 ("By the mid-1970s . . . nearly all [the] nation's largest newspapers had abandoned publishing sex-designated help wanted columns.").

¹⁶⁸ 42 U.S.C. § 2000e(c) (2012).

¹⁶⁹ Wilborn v. S. Union State Cmty. Coll., 720 F. Supp. 2d 1274, 1290–91 (M.D. Ala. 2010) (citing Barbara T. Lindemann, Paul Grossman & C. Geoffrey Weirich, Employment Discrimination Law 1598 (4th ed. 2007)); Scaglione v. Chappaqua Cent. Sch. Dist., 209 F. Supp. 2d 311, 316 (S.D.N.Y. 2002).

¹⁷⁰ 387 F. Supp. 42, 46 (N.D. Ill. 1974). The court distinguished *Brush* by contrasting the "negligible" involvement of a newspaper, which merely printed help wanted ads, with the law school's "significant" involvement in career placement. *Id.*

employment, and the vast majority did so through the school's placement office. Given the law school's "significant" involvement in career placement, the court concluded that it met the statutory definition. Similarly, in *Wilborn v. Southern Union State Community College*, the court held that a truck driver training program "engaged to a significant degree in procuring employment opportunities for program participants."¹⁷¹ Because it actively assisted them in the job search process by helping them complete and submit applications, inviting recruiters, and facilitating interviews, the program was an employment agency.¹⁷²

Conversely, where intermediaries do little more than passively post information about job opportunities, they are not employment agencies. Thus, courts concluded that a voluntary, nonprofit physicians' organization was not an employment agency because it merely received and passively posted job openings.¹⁷³ Absent evidence that it was actively involved in the recruitment process for specific employers or helped to pick candidates, it lacked "a 'significant degree' of engagement in employment-related activities."¹⁷⁴

Under existing case law, tech intermediaries in the labor market may satisfy the definition of an employment agency, depending upon the details of their operation.¹⁷⁵ An online jobs board that merely posts employment opportunities and allows users to review them closely resembles traditional print newspapers. The role of such an intermediary is passive, and it is unlikely to be considered an employment agency.

However, many of the job platforms operating today are far more actively involved in the process of connecting workers and employers. Intermediaries that identify passive candidates are most easily

¹⁷¹ 720 F. Supp. 2d at 1291.

¹⁷² *Id.* at 1291–92.

¹⁷³ See *Stewart v. Am. Ass'n of Physician Specialists, Inc.*, No. 5:13-cv-01670-ODW (DTBx), 2015 U.S. Dist. LEXIS 161149, at *14–16 (C.D. Cal. Nov. 30, 2015); *Radentz v. Am. Ass'n of Physician Specialists, Inc.*, No. CV 13-01486 SJO (OPx), 2014 U.S. Dist. LEXIS 194836, at *12–13 (C.D. Cal. Nov. 10, 2014).

¹⁷⁴ *Radentz*, 2014 U.S. Dist. LEXIS 194836, at *12–13 (quoting *Brush v. San Francisco Newspaper Printing Co.*, 315 F. Supp. 577, 580 (N.D. Cal. 1970)); accord *Stewart*, 2015 U.S. Dist. LEXIS 161149, at *15–16. Importantly, the certification provided by the organization to physicians, while providing a positive signal, was not required to practice medicine.

¹⁷⁵ Case law and EEOC policy statements say that to be an employment agency, the entity must regularly deal with employers that meet the statutory definition under Title VII by employing fifteen or more employees. 42 U.S.C. § 2000e(b)–(c) (2012). This requirement is easily satisfied, as some of the largest employers utilize online advertising and job-matching platforms, so I will not discuss this issue further.

characterized as employment agencies. Their self-described function is to find “top talent” to meet their clients’ staffing needs. By definition, they are “procuring employees” for an employer on a regular basis. Their role in the labor market is closely analogous to that of traditional employment agencies that identified and recruited suitable workers to meet their client-employers’ labor needs.

Job-matching platforms like ZipRecruiter would also appear to satisfy the statutory definition. The recommender systems that power these platforms analyze data about both applicants and employers in order to suggest good matches. They do so by scoring and ranking candidates for employers and recommending that applicants apply to specific employers. These activities are exactly analogous to those of a traditional employment agency. Consider the following description of the Mississippi State Employment Service (MSES), an entity whose status as an employment agency under Title VII was uncontested:

MSES assigns applicants and job orders a nine-digit Dictionary of Occupational Titles Code (Dot Code) . . . [which] identif[ies] a particular occupational group [and] the worker function ratings . . . [U]pon an applicant’s request MSES clerks search job orders for vacancies that match the applicant’s DOT code and, locating an open order with a corresponding code, refer the applicant . . . [T]o fill orders . . . MSES clerks search the active files of applicants with a matching occupational title . . . Ultimately, to match job orders with the most qualified applicant . . . MSES clerks primarily rely on the DOT Codes, the type of job, and the applicant’s experience and job preferences, but as well draw from personal knowledge . . . including subjective judgments (e.g.: an applicant’s attitude, appearance, personality, employment history).¹⁷⁶

Job-matching platforms perform the exact same functions, except that they leverage more data and use sophisticated computer algorithms to aid the matching process. Just like the MSES, tech intermediaries that actively match workers with available opportunities are “regularly undertaking . . . to procure” employees and employment opportunities and should easily meet the definition of an employment agency.

Online advertising platforms like Facebook are a bit more difficult to characterize. When sued for discriminatory advertising practices,

¹⁷⁶ Hill v. Miss. State Emp’t Serv., 918 F.2d 1233, 1234–35 (5th Cir. 1990).

Facebook argued that it was not an employment agency because it merely “[p]rovid[ed] a platform for third parties to publish their ads”¹⁷⁷ like the newspaper in *Brush*. Online advertising platforms, however, are far more actively involved in connecting job seekers and opportunities than traditional print newspapers. While they do not explicitly refer candidates to employers or recommend applicants apply to certain jobs, they do so implicitly by delivering specific job ads to particular users based on their predictions of who is most likely to respond. Because advertisers’ targeting choices do not fully determine their audience, platforms’ advertising algorithms ultimately determine exactly who sees which ad. Thus, although their role is less visible, these intermediaries are functionally matching job opportunities to applicants. Given their control over information flows and their growing importance in the job search process, it is at least an open question whether they “regularly . . . procure employees” for employers.

If a tech intermediary is found to be an employment agency, Title VII makes it unlawful for it “to fail or refuse to refer for employment, or otherwise to discriminate against, any individual because of his race, color, religion, sex, or national origin, or to classify or refer for employment any individual on the basis of his race, color, religion, sex, or national origin.”¹⁷⁸ This language easily covers situations involving intentional discrimination, such as deliberately excluding certain groups from receiving job postings because of race or other protected characteristics. And if a platform’s ad-delivery or matching algorithm produces discriminatory results, the disparate impact theory of liability appears to apply to employment agencies as well.¹⁷⁹ Thus, tech platforms,

¹⁷⁷ Defendant’s Notice of Motion and Motion To Dismiss First Amended Complaint, *supra* note 56, at 19.

¹⁷⁸ 42 U.S.C. § 2000e-2(b).

¹⁷⁹ Uniform Guidelines on Employee Selection Procedures, 29 C.F.R. § 1607.1–1607.13 (2019), which establish standards to evaluate tests and other selection procedures for adverse impact, assume that employment agencies can be liable for disparate impact. 29 C.F.R. § 1607.1(B) (explaining that the Guidelines are “designed to assist employers, labor organizations [and] *employment agencies* . . . to comply” with federal employment discrimination law (emphasis added)). Case law interpreting Title VII is also consistent with this conclusion. In *Eldredge v. Carpenters 46*, the Ninth Circuit held that an apprenticeship committee that employed a facially neutral referral system for new apprentices that had a significantly discriminatory impact on women could be liable under Title VII. *Eldredge v. Carpenters 46* N. Cal. Ctys. Joint Apprenticeship & Training Comm., 833 F.2d 1334, 1337 (9th Cir. 1987). In response to the claim that any discrimination was the result of employer preferences, the court concluded that the committee “cannot avoid liability for the effects of its own admission procedures by pointing to the discriminatory practices [of the employers].”

when acting as employment agencies, can be held liable under Title VII if their facially neutral practices in fact disadvantage protected groups without sufficient justification.¹⁸⁰

C. Interference with Employment Opportunities

Another line of cases is potentially relevant to the question of platform liability under Title VII. In *Sibley Memorial Hospital v. Wilson*, the D.C. Circuit held that entities other than a plaintiff's direct employer may be liable under Title VII if they "control access to . . . employment" and "deny such access by reference to invidious criteria."¹⁸¹ The plaintiff in *Sibley* was a male private-duty nurse who worked directly for his client-patients. He alleged that the defendant hospital barred his access to female patients to whom he had been referred because he was a man.¹⁸² He sued, arguing that even though the hospital was not his employer, it had violated Title VII by interfering with his employment relationships with patient-clients on a discriminatory basis.¹⁸³ The court found that the hospital, like labor organizations and employment agencies, had a "highly visible nexus with the creation and continuance of direct employment relationships" and thus could be liable for discrimination.¹⁸⁴

The *Sibley* court found textual support for its holding in the language of Section 703(a)(1), which states that it is unlawful "for an employer . . . to fail or refuse to hire or to discharge any individual, or otherwise to discriminate against any individual with respect to . . . privileges of employment."¹⁸⁵ The court interpreted the phrase "otherwise to discriminate against any individual" to prohibit a defendant from interfering, on a discriminatory basis, with the employment

Id. Because employment agencies control access to job opportunities in a manner similar to apprenticeship committees, and because Title VII uses similar text to describe liability for the two types of entities, compare 42 U.S.C. § 2000e-2(b), with id. § 2000e-2(d), agencies, like the committee in *Eldredge*, are responsible if their requirements or procedures have discriminatory effects.

¹⁸⁰ When Congress amended Title VII in 1991, Subsection 703(k) was added to the text, stating that a disparate impact is established if "a complaining party demonstrates that a respondent uses a particular employment practice that causes a disparate impact" and the practice is not otherwise justified. 42 U.S.C. § 2000e-2(k)(1)(A)(i) (emphasis added). The statute defines a "respondent" to include employment agencies. Id. § 2000e(n).

¹⁸¹ 488 F.2d 1338, 1342 (D.C. Cir. 1973).

¹⁸² Id. at 1339–40.

¹⁸³ Id. at 1340.

¹⁸⁴ Id. at 1342.

¹⁸⁵ Id. at 1341 (quoting 42 U.S.C. § 2000e-2(a)).

opportunities of anyone, not just the defendant's own employees.¹⁸⁶ Sibley Memorial Hospital clearly satisfied the statutory definition of an employer because it had many employees,¹⁸⁷ and the plaintiff asserted that it had interfered with his employment by a third party (his patient-client) because of his sex. The court permitted his suit to proceed, concluding that he had sufficiently alleged that the hospital had "otherwise discriminated" against him—not by failing to hire or by discharging him—but by preventing him from entering employment relationships with third parties.¹⁸⁸

Although a few courts have rejected the third party interference doctrine,¹⁸⁹ other circuits have followed the reasoning of the *Sibley* court.¹⁹⁰ Thus, courts have held that hospitals may be sued for denying a physician admitting privileges on a discriminatory basis when the effect was to destroy his employment relationships with a third party.¹⁹¹ Similarly, an athletic association that sent referees to work at games hosted by its member schools could be liable for discrimination as well.¹⁹² The member schools employed the referees, but the association decided which referees were assigned to particular games, thereby controlling their access to employment. For courts that follow *Sibley*, "defendant's interference with plaintiff's outside employment relationships may be an independent basis for Title VII liability, where the defendant controls access to employment opportunities with third parties."¹⁹³

This theory of liability for discriminatory interference with employment might also apply to tech intermediaries in the labor market. For the most part, these companies are large enough to satisfy the statutory definition of an employer,¹⁹⁴ and they control information about

¹⁸⁶ *Id.* at 1341–42.

¹⁸⁷ Under the statute, an employer is "a person engaged in an industry affecting commerce who has fifteen or more employees for each working day in each of twenty or more calendar weeks in the current or preceding calendar year." 42 U.S.C. § 2000e(b).

¹⁸⁸ *Sibley*, 488 F.2d at 1342.

¹⁸⁹ See, e.g., *Lopez v. Massachusetts*, 588 F.3d 69, 89 (1st Cir. 2009); *Gulino v. N.Y. State Educ. Dep't*, 460 F.3d 361, 376 (2d Cir. 2006).

¹⁹⁰ See, e.g., *Christopher v. Stouder Mem'l Hosp.*, 936 F.2d 870, 877 (6th Cir. 1991); *Zaklama v. Mt. Sinai Med. Ctr.*, 842 F.2d 291, 294 (11th Cir. 1988); *Gomez v. Alexian Bros. Hosp.*, 698 F.2d 1019, 1021 (9th Cir. 1983).

¹⁹¹ See *LeMasters v. Christ Hosp.*, 777 F. Supp. 1378, 1380 (S.D. Ohio 1991).

¹⁹² *Kemether v. Pa. Interscholastic Athletic Ass'n*, 15 F. Supp. 2d 740, 761–64 (E.D. Pa. 1998).

¹⁹³ *Id.* at 762.

¹⁹⁴ See 42 U.S.C. § 2000e(b) (2012).

available job opportunities. If they exercise that power in a biased manner that interferes with others' employment opportunities, they could be held liable under the *Sibley* line of cases. What remains unclear is whether tech intermediaries exercise sufficient influence over access to jobs to be found liable under this theory.¹⁹⁵ Although Facebook and other online advertising platforms have a great deal of control over information flows about available jobs, they typically do not directly control who actually gets those jobs. On the other hand, job-matching platforms, depending upon how they operate, may exercise a higher level of control, particularly if an employer relies exclusively on a single site to attract and vet applicants.

* * *

Given the prominent role of tech intermediaries in today's labor markets, Title VII's goals of equal access to opportunities and reducing occupational segregation cannot be met solely by targeting employer practices. Achieving truly equal opportunity will also require scrutinizing the role of tech intermediaries in the labor market, as they control the flow of information and intervene between employers and job seekers. As seen in this Section, Title VII's text and doctrine are capacious enough to reach the activities of tech intermediaries, at least in certain circumstances, when they cause discriminatory effects. The next Section considers some of the obstacles to doing so, as well as alternative approaches to holding these platforms responsible for their effects on access to job opportunities.

IV. PLATFORM RESPONSIBILITY

A. Avoiding the CDA Section 230 Defense

Any effort to make tech intermediaries bear responsibility for their discriminatory effects will immediately run into the claim that Section 230 of the Communications Decency Act (CDA) protects them from

¹⁹⁵ For cases addressing the limits of how much control is required, see, e.g., *Anderson v. Pac. Mar. Ass'n*, 336 F.3d 924, 931–32 (9th Cir. 2003); *Wynn v. NBC, Inc.*, 234 F. Supp. 2d 1067, 1110 n.34 (C.D. Cal. 2002) (citing *Gomez v. Alexian Bros. Hosp.*, 698 F.2d 1019 (9th Cir. 1983)).

liability.¹⁹⁶ Section 230 prohibits holding platforms liable “as the publisher or speaker” of information provided by third parties,¹⁹⁷ and online platforms have relied on this defense in a wide variety of circumstances.¹⁹⁸ I argue here that Section 230 does not apply when tech intermediaries distribute information in a discriminatory manner because they are not acting as a “publisher or speaker” when doing so. And when platforms do more than just channel information but also make recommendations about job opportunities or candidates, then they have become content providers themselves and fall outside the protection of Section 230.¹⁹⁹

Congress passed the CDA in 1996 in an effort to restrict minors’ access to sexually explicit material online.²⁰⁰ Section 230, however, provides a safe harbor to online computer services that make some effort to remove offensive materials. Congress wanted to avoid situations in which partial or ineffective measures led to a platform being held liable for material that it had failed to remove.²⁰¹ More specifically, the section was intended

¹⁹⁶ 47 U.S.C. § 230 (2012). It is commonplace to refer to “Section 230 immunity,” but the statute creates a defense, not an immunity from suit, and so I use the former terminology throughout this Article.

¹⁹⁷ *Id.* § 230(c)(1).

¹⁹⁸ Scholars have criticized the broad scope given to § 230 by the courts. See, e.g., Danielle Keats Citron & Benjamin Wittes, *The Internet Will Not Break: Denying Bad Samaritans § 230 Immunity*, 86 *Fordham L. Rev.* 401, 404 (2017) (arguing that § 230’s “overbroad interpretation has left victims of online abuse with no leverage against site operators whose business models facilitate abuse”); Gillespie, *supra* note 28, at 209 (arguing that the safe harbor afforded by § 230 is “increasingly problematic”).

¹⁹⁹ I lay out in the text the argument that § 230 as written does not protect labor market intermediaries from liability for discriminatory effects. Nevertheless, the challenges in parsing the statutory language and applying it in this context suggest that § 230 may simply be inappropriate given the potential for social harms caused by biased markets for labor and other key opportunities. As Gillespie explains, § 230 is premised on a distinction between information conduits on the one hand and media content producers on the other that fails to capture how platforms function today. See Gillespie, *supra* note 28, at 209. Platforms do not fit into either category; they are an unanticipated hybrid. *Id.* at 210. Because they are “neither conduit nor content, then legal arrangements premised on those categories may be insufficient,” *id.* at 211, and platforms “should not get to use Section 230’s protection to avoid laws prohibiting discrimination in employment, housing, or pricing.” *Id.* at 212.

²⁰⁰ See Robert Cannon, *The Legislative History of Senator Exon’s Communications Decency Act: Regulating Barbarians on the Information Superhighway*, 49 *Fed. Comm. L.J.* 51, 52–57 (1996). The core provisions of the statute were ruled unconstitutional in *Reno v. ACLU*, 521 U.S. 844 (1997).

²⁰¹ Jeff Kosseff, *The Gradual Erosion of the Law That Shaped the Internet: Section 230’s Evolution over Two Decades*, 18 *Colum. Sci. & Tech. L. Rev.* 1, 8–9 (2016).

to overrule *Stratton Oakmont, Inc. v. Prodigy Services Co.*,²⁰² a case in which Prodigy, an online service provider, was held to be the “publisher” of defamatory messages posted on its message boards—and therefore to be liable for defamation—because it had voluntarily deleted other messages after judging them to be offensive or in “bad taste.”²⁰³ In enacting Section 230, Congress wanted to encourage websites to develop ways of blocking minors’ access to “objectionable or inappropriate online material” by protecting the platforms from liability when they attempted to do so.²⁰⁴

The core provisions of Section 230 are found in Subsection (c), entitled “Protection for ‘Good Samaritan’ blocking and screening of offensive material.” Subsection (c)(1) states that “[n]o provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.”²⁰⁵ Subsection (c)(2) adds that no computer service provider “shall be held liable” because it voluntarily takes actions in good faith in order to restrict the availability of material it deems offensive.²⁰⁶ Thus, while Subsection (c)(2) provides “Good Samaritan” protection for platforms whose efforts to remove offensive materials are partial or ineffective, Subsection (c)(1) sweeps more broadly, generally protecting them from liability for third party content that they host. This broader protection was apparently motivated by concerns that if online platforms were held responsible for content provided by others, they would face the grim choice of undertaking burdensome reviews of all user-provided content or not permitting it at all.²⁰⁷

Courts have applied the Section 230 defense expansively, permitting Internet platforms to avoid liability for a long list of claims including common law theories like defamation,²⁰⁸ unjust enrichment,²⁰⁹ intentional

²⁰² 1995 WL 323710 (N.Y. Sup. Ct. May 24, 1995).

²⁰³ *Id.* at *4.

²⁰⁴ 47 U.S.C. § 230(b)(4) (2012).

²⁰⁵ *Id.* § 230(c)(1). Although the statute refers to an “interactive computer service,” I use the terms “platform” or “website” for ease of reference.

²⁰⁶ *Id.* § 230(c)(2).

²⁰⁷ See *Doe v. Backpage.com, LLC*, 817 F.3d 12, 18–19 (1st Cir. 2016) (recognizing that “websites that display third-party content may have an infinite number of users” and “holding website operators liable for that content ‘would have an obvious chilling effect’”).

²⁰⁸ See *Bennett v. Google, LLC*, 882 F.3d 1163 (D.C. Cir. 2018).

²⁰⁹ See *Dyroff v Ultimate Software Grp., Inc.*, No. 17-cv-05359-LB, 2017 WL 5665670, at *1 (N.D. Cal. Nov. 26, 2017).

infliction of emotional distress,²¹⁰ and negligence,²¹¹ as well as claims based on the Fair Housing Act,²¹² the Lanham Act,²¹³ and other state and federal statutes.²¹⁴ Crucially, in these cases, the alleged harm suffered by the plaintiff stems from the *content* of the information hosted by the website.

By contrast, the key concern here is the manner in which platforms control the flow of information about critical opportunities. As explained above, platforms' ad-targeting and matching algorithms may distribute information to demographically skewed audiences, even if the targeting criteria are not discriminatory and the advertiser intends that it be distributed on a nondiscriminatory basis. In this situation, the *content* of the information distributed is not harmful in any way; to the contrary, it is an affirmative good. The harm consists of depriving certain groups of equal access to that information. It is precisely because information about employment, housing, and consumer financial services is valuable that demographically skewed distribution of that information raises fairness concerns.

Because the threatened harm stems from the *manner* in which the platform distributes information among users—and not its content—the CDA defense does not apply.²¹⁵ Section 230 protects Internet platforms

²¹⁰ See *Cohen v. Facebook, Inc.*, 252 F. Supp. 3d 140, 147, 158 (E.D.N.Y. 2017).

²¹¹ See *Herrick v. Grindr, LLC*, 306 F. Supp. 3d 579, 593 (S.D.N.Y. 2018).

²¹² See *Chi. Lawyers' Comm. for Civil Rights Under Law, Inc. v. Craigslist, Inc.*, 519 F.3d 666, 671 (7th Cir. 2008).

²¹³ See *Baldino's Lock & Key Serv., Inc. v. Google, LLC*, 285 F. Supp. 3d 276, 278 (D.D.C. 2018).

²¹⁴ See *Doe v. Backpage.com, LLC*, 817 F.3d 12, 29 (1st Cir. 2016) (protecting defendants from liability under the Trafficking Victims Protection Reauthorization Act as well as an analogous Massachusetts state anti-trafficking statute); *Gonzalez v. Google, Inc.*, 335 F. Supp. 3d 1156, 1179 (N.D. Cal. 2018) (same under the Federal Anti-Terrorism Act).

²¹⁵ My argument here differs from the one made by Datta et al., who begin their analysis by *assuming* that imposing liability on platforms for targeted advertising would be “holding them liable as a publisher” and that therefore “Section 230 comes into play.” Datta et al., *supra* note 67, at 11. In contrast, I contend that a claim about a platform's discriminatory targeting algorithm is not holding it liable as a “speaker or publisher” at all. When the alleged illegality does not stem from the content itself, the harm does not turn on what the platform has said or published but on how it treated users in carrying out a separate function. As it distributes information to users, the platform is acting as a conduit, rather than speaking or editing the information that is transmitted. Cf. Tim Wu, *Machine Speech*, 161 U. Pa. L. Rev. 1495, 1497–98, 1521–22 (2013) (explaining the distinction between carriers or conduits that merely transmit speech on the one hand and actors that exercise discretion by selecting and identifying with the content on the other).

from being “treated as the publisher or speaker of any information” provided by others.²¹⁶ If the basis for liability does not turn on the platform’s role as a “publisher or speaker” of harmful content, then Section 230 is irrelevant. The First and Second Circuits require that the claim “is based on information provided by another” and “would treat [the defendant] as the publisher or speaker” for the defense to apply.²¹⁷ The D.C. Circuit similarly determines the applicability of the defense based on “the information for which [the plaintiff] seeks to hold [the defendant] liable . . . as the ‘publisher or speaker.’”²¹⁸

Where liability turns not on the content or information provided but on other actions undertaken by the online platform, the defense should not be available. In *Doe v. Internet Brands, Inc.*,²¹⁹ the Ninth Circuit held that Section 230 was not a defense to a failure to warn claim brought by an aspiring model who was raped after being lured to a fake audition. She alleged that the defendant, which operated a networking website for models, was aware that sexual predators were approaching victims through the site and yet failed to warn of the risks.²²⁰ Because she was not arguing that the defendant had transmitted any harmful information, the court concluded that her claim did not treat the website as a “publisher or speaker,” and therefore the Section 230 defense did not apply.²²¹

Although Datta et al. assume that § 230 is relevant, they assert that online advertising platforms like Google are not protected by that provision for a different reason. As explained *infra*, notes 231–235 and accompanying text, the § 230 defense does not apply where the platform itself “contributes materially to the alleged illegality.” *Fair Hous. Council v. Roommates.com, LLC*, 521 F.3d 1157, 1168 (9th Cir. 2008). Datta et al. argue that when a platform’s targeting algorithms cause discriminatory ad distribution, then the platform is making a material contribution to the alleged illegality and is no longer protected by § 230. Datta et al., *supra* note 67, at 13. I would agree that *if* a platform is considered a “publisher” when deploying a discriminatory ad-targeting algorithm, then it is contributing materially to the illegality, and therefore the protection of § 230 does not apply.

²¹⁶ 47 U.S.C. § 230(c)(1) (2012).

²¹⁷ *Doe*, 817 F.3d at 19 (internal quotation marks omitted); see also *FTC v. LeadClick Media, LLC*, 838 F.3d 158, 173 (2d Cir. 2016) (same).

²¹⁸ *Klayman v. Zuckerberg*, 753 F.3d 1354, 1357 (D.C. Cir. 2014). Other variations similarly suggest that the information provided by a third party is the basis for liability. See, e.g., *FTC v. Accusearch Inc.*, 570 F.3d 1187, 1196 (10th Cir. 2009) (“[I]mmunity can be claimed only with respect to ‘information provided by another information content provider.’”); *Barnes v. Yahoo!, Inc.*, 570 F.3d 1096, 1102 (9th Cir. 2009) (“[W]hat matters is whether the cause of action inherently requires the court to treat the defendant as the ‘publisher or speaker’ of content provided by another.” (emphasis added)).

²¹⁹ 824 F.3d 846 (9th Cir. 2016).

²²⁰ *Id.* at 851.

²²¹ *Id.*

Courts have also found that online platforms cannot rely on Section 230 to avoid complying with generally applicable business regulations. In *HomeAway.com, Inc. v. City of Santa Monica*,²²² the court rejected the defendants' argument that a local ordinance was preempted by Section 230. The ordinance, among other things, limited the types of rentals permitted, required online platforms to collect occupancy taxes, and prohibited them from receiving a fee for booking unlicensed properties.²²³ The court concluded that because the ordinance "does not require the Platforms to monitor third-party content" or to remove it, it does not treat them as publishers and thus falls outside the preemptive scope of Section 230.²²⁴ A key factor was that "the Platforms face no liability for the content of the bookings," only for engaging in unlicensed transactions.²²⁵

Similarly, claims based on platforms' distribution of information about crucial opportunities do not impose liability for harmful content and therefore should not be barred by Section 230. When an online platform like Google or Facebook designs a targeting algorithm to determine which ads are delivered to which users, it is clearly not acting as a speaker of the ad content. Nor is the platform acting as a "publisher" of that content because the targeting algorithms do not involve "traditional editorial functions—such as deciding whether to publish, withdraw, postpone or alter content."²²⁶ The platforms do engage in certain activities similar to traditional editorial functions—for example, restricting the format of ads or determining how content appears in them. However, holding them responsible for discriminatory patterns of information delivery would not impinge on any of those functions.

²²² 918 F.3d 676 (9th Cir. 2019).

²²³ *Id.* at 680.

²²⁴ *Id.* at 682–83.

²²⁵ *Id.* at 684. Section 230 was not intended to "render unlawful conduct 'magically . . . lawful when [conducted] online.'" *Id.* at 683 (quoting *Fair Hous. Council v. Roommates.com, LLC*, 521 F.3d 1157, 1164 (9th Cir. 2008) (en banc)). *Airbnb, Inc. v. City & County of San Francisco*, a case challenging a similar local ordinance, reached the same conclusion. 217 F. Supp. 3d 1066 (N.D. Cal. 2016).

²²⁶ *Zeran v. Am. Online, Inc.*, 129 F.3d 327, 330 (4th Cir. 1997). In *Doe v. Backpage.com, LLC*, the plaintiffs alleged that the platform was liable for harms to minors sex-trafficked on the site. They tried to argue that Backpage was liable—not as a publisher or speaker but because of choices that it made in structuring the website. The court found that the "third-party content is like Banquo's ghost: it appears as an essential component of each and all of the appellants' . . . claims." 817 F.3d 12, 22 (1st Cir. 2016). In contrast, in challenging discriminatory ad targeting, it is not necessary to consider the content of the ads at all. The charge of discriminatory targeting would stand regardless of which jobs were advertised or how the ads were formatted.

Job-matching and recruiting platforms, in contrast to online advertising, often do more than just direct the flow of information between employers and job seekers. Instead, these websites typically curate information, adding context or commentary, such as recommending opportunities to job seekers, ranking applicants, or labeling some a “Good Match” for employers. In those situations, holding the intermediary responsible would more likely entail treating it as a “publisher or speaker.” The Section 230 defense, however, only protects platforms based on “information provided by another information content provider,”²²⁷ not information created by the platform itself.

This limitation on the Section 230 defense was explained by the Ninth Circuit in *Fair Housing Council v. Roommates.com, LLC*.²²⁸ Roommates.com operated a website designed to help connect people who had spare rooms they wanted to rent with people looking for a room. It required users to answer questions about their sex, sexual orientation, and whether they lived with children and then used this information to create a profile page visible to others.²²⁹ The site also elicited users’ preferences regarding prospective roommates—for example, whether they were willing to live with straight or gay males, only straight males, only gay males, or no males at all—and used that information to filter searches and direct emails in order to facilitate matches.²³⁰ The plaintiffs, organizations seeking to enforce fair housing laws, alleged that these activities constituted discrimination on the basis of sex, sexual orientation, and family status in violation of state and federal fair housing laws.

Roommates.com moved to dismiss the case, citing Section 230. The Ninth Circuit noted the limits of the defense: it does not apply where the platform is itself the “information content provider”—i.e., one who is “responsible, in whole or in part, for the creation or development of” the

²²⁷ 47 U.S.C. § 230(c)(1) (2012). An “information content provider” is defined as an entity “that is responsible . . . for the creation or development of information provided through the Internet.” *Id.* § 230(f)(3).

²²⁸ 521 F.3d 1157 (9th Cir. 2008) (en banc). This opinion was issued upon an en banc rehearing of a panel decision in the case. *Fair Hous. Council v. Roommates.com, LLC*, 489 F.3d 921 (9th Cir. 2007). After holding that the suit was not barred by § 230, the en banc court remanded the case to the district court to consider whether the defendant’s actions violated the Fair Housing Act. *Fair Hous. Council*, 521 F.3d at 1175. The case reached the Ninth Circuit again on the Fair Housing Act issue, and the court decided that the Fair Housing Act did not apply given the underlying facts. *Fair Hous. Council v. Roommate.com, LLC*, 666 F.3d 1216, 1222 (9th Cir. 2012).

²²⁹ *Fair Hous. Council*, 521 F.3d at 1161.

²³⁰ *Id.* at 1165.

offending content.”²³¹ As the court wrote: “If [a website] passively displays content that is created entirely by third parties, then it is only a service provider with respect to that content” and is protected by Section 230.²³² However, where the site creates the content or “contributes materially to the alleged illegality,”²³³ then it “is also a *content provider*,”²³⁴ and the defense does not apply. To the extent that Roommates.com required users to provide information about protected characteristics and then used that information to filter listings and direct emails in a discriminatory way, the court concluded that it was “much more than a passive transmitter of information provided by others [but instead was] the developer, at least in part, of that information”²³⁵ and fell within the exception to Section 230’s protections.

To the extent that job-matching and recruiting platforms are affirmatively making judgments about candidates or opportunities, they are not only “speakers” but also creators of that content. When a platform recommends candidates in a biased way, it is “contributing materially” to the illegality, and its actions fall outside Section 230’s protective ambit.²³⁶ Admittedly, numerous plaintiffs have tried to rely on the “information content provider” exception, and most have not succeeded. In those cases, however, the core of the plaintiffs’ claims rested on harmful content provided by others.²³⁷ They are thus distinguishable from claims that would hold a tech intermediary responsible for the discriminatory impacts of *its own* content.

Holding tech intermediaries responsible for discriminatory distribution of information is entirely consistent with the policies articulated in Section 230.²³⁸ Congress recognized that if websites were held liable for

²³¹ Id. at 1162 (citing 47 U.S.C. § 230(f)(3)).

²³² Id. at 1162.

²³³ Id. at 1168.

²³⁴ Id. at 1162 (emphasis added).

²³⁵ Id. at 1166.

²³⁶ As explained in note 215 *supra*, Datta et al. argue that this exception to § 230’s protections applies to online advertising platforms like Google that distribute ads in a discriminatory manner.

²³⁷ See, e.g., *Nemet Chevrolet, Ltd. v. ConsumerAffairs.com, Inc.*, 591 F.3d 250, 260 (4th Cir. 2009); *Batzel v. Smith*, 333 F.3d 1018, 1035 (9th Cir. 2003); *Whitney Info. Network, Inc. v. Xcentric Ventures, LLC*, No. 2:04-cv-47-FtM-34SPC, 2008 WL 450095, at *12 (M.D. Fla. Feb. 15, 2008).

²³⁸ Among the policies it lists are “promot[ing] the continued development of the Internet and other interactive computer services” and “preserv[ing] the vibrant and competitive free market that presently exists.” 47 U.S.C. § 230(b)(1)–(2) (2012). The other two listed policies

content posted by others, it would impose an enormous burden on them. Many online platforms function by allowing users to upload and share information. Given the potentially “infinite number of users generating an enormous amount of potentially harmful content, . . . holding website operators liable for that content ‘would have an obvious chilling effect.’”²³⁹ Imposing liability for content provided by others might threaten the very viability of some website operators. However, holding platforms responsible when they act in another capacity—not as publisher or speaker—does not create the same sort of existential threat. More specifically, scrutinizing how they exercise their power to channel information about opportunities would not require platforms to review all user posts or to make editorial decisions about which posts to permit and which to remove. Instead, it requires them to be attentive to the distributive effects of their choices regarding who sees what information and holds them responsible if their choices produce discriminatory effects.

One of the other stated purposes of Section 230 is to encourage technologies which “maximize user control over what information is received.”²⁴⁰ Algorithms that control the flow of information and determine who sees what are contrary to the vision of “maximiz[ing] user control” articulated in the statute. Protecting platforms from scrutiny over how they distribute information would reduce the transparency of the process, rendering users even less in control of their information environment—an outcome inconsistent with Congress’s stated goals. In short, requiring tech intermediaries to bear responsibility for biased distribution of information about key opportunities is entirely consistent with the policies underlying Section 230.

If the Section 230 defense is avoided, tech intermediaries might respond by claiming that any regulation of their activities in shaping labor markets violates the First Amendment. In an analogous context, scholars have debated whether search engine outputs are constitutionally protected speech.²⁴¹ Although a full discussion is beyond the scope of this Article,

are about developing filtering technologies to restrict children’s access to objectionable online materials and enforcing federal criminal laws on human trafficking. *Id.* § 230(b)(4)–(5).

²³⁹ *Doe v. Backpage.com, LLC*, 817 F.3d 12, 19 (1st Cir. 2016) (quoting *Zeran v. Am. Online, Inc.*, 129 F.3d 327, 331 (4th Cir. 1997)).

²⁴⁰ 47 U.S.C. § 230(b)(3).

²⁴¹ Stuart Minor Benjamin, *Algorithms and Speech*, 161 U. Pa. L. Rev. 1445, 1482, 1486–87, 1494 (2013) (concluding that algorithm-based outputs like search engine results are protected speech when they entail substantive communication between a speaker and listener).

the suggestion that the First Amendment protects tech platforms from any regulation for the discriminatory effects of their practices warrants a brief response. First, it is not at all clear that the algorithms that personalize online advertisements are engaged in constitutionally protected speech. By determining which ads are displayed to which users, the platform (as opposed to the advertiser) conveys no substantive message.²⁴² Rather, the algorithm it deploys is a functional tool designed to distribute the communications of others in a way that maximizes clicks rather than a vehicle for the platform's own expression.²⁴³ In this way, it is similar to the call-routing systems that serve to distribute messages to an intended recipient. The mechanical system that distributes communications created by others does not itself have speech rights.

Other tech intermediaries, such as job-matching platforms or passive recruiting systems, arguably do speak when they endorse certain applicants or curate lists of available candidates or opportunities for firms or job seekers. The First Amendment, however, does not protect this sort of speech from laws of general applicability.²⁴⁴ In the most closely analogous case, *Pittsburgh Press Co. v. Pittsburgh Commission on Human Relations*,²⁴⁵ the Supreme Court rejected a newspaper's First

and do not violate laws of general applicability); Oren Bracha & Frank Pasquale, *Federal Search Commission? Access, Fairness, and Accountability in the Law of Search*, 93 *Cornell L. Rev.* 1149, 1152 (2008) (asserting that the First Amendment does not prohibit regulation of search engines); James Grimmelmann, *Speech Engines*, 98 *Minn. L. Rev.* 868, 912 (2014) (arguing that search engine results are speech but are not categorically protected by the First Amendment); Eugene Volokh & Donald M. Falk, *Google: First Amendment Protection for Search Engine Search Results*, 8 *J.L. Econ. & Pol'y* 883, 884–85 (2012) (arguing in a White Paper commissioned by Google that search engine results are constitutionally protected speech, excluding from analysis paid ad displays that appear next to search results); Wu, *supra* note 215, at 1498 (arguing that the First Amendment protects algorithmic output when it is an expressive “speech product” but not when it operates as a “communication tool”); see also Tim Wu, *Free Speech for Computers?*, *N.Y. Times*, June 19, 2012, at A29 (“[N]onhuman or automated choices should not be granted the full protection of the First Amendment, and often should not be considered ‘speech’ at all.”).

²⁴² Benjamin argues that algorithm-based outputs do not constitute speech unless they are sending a “substantive message.” Benjamin, *supra* note 241, at 1471. At a minimum, this entails a speaker seeking to transmit that message and a listener who can recognize it. *Id.* at 1461.

²⁴³ Wu suggests that the crucial dividing line between protected and unprotected speech is functionality. In his view, First Amendment protections do not apply to carriers or conduits, or to communication tools that “primarily facilitate the communication of another person.” Wu, *supra* note 215, at 1498.

²⁴⁴ See Benjamin, *supra* note 241, at 1487 (pointing out that laws of general applicability like antitrust and tax laws do not implicate the First Amendment).

²⁴⁵ 413 U.S. 376 (1973).

Amendment challenge to a local ordinance forbidding discriminatory job advertising.²⁴⁶ It held that any First Amendment interest was “altogether absent” because the activity regulated—discriminatory advertising—was illegal, and the speech restriction was “incidental to a valid limitation on economic activity.”²⁴⁷ As discussed above, Title VII has long prohibited employment agencies and labor unions from engaging in discriminatory referral practices—practices which inherently involve speaking—and the First Amendment has not shielded them from such regulation. If tech intermediaries discriminate when acting as employment agencies, the First Amendment would afford them no greater protection.

B. Practical Obstacles and Alternative Strategies

Section 230 of the CDA may not be as much of a barrier as first appears, but a number of practical obstacles nevertheless remain. Liability-based regimes like Title VII depend upon aggrieved parties to step forward and challenge unlawful practices. It is, however, exceedingly difficult for individuals to detect when information about opportunities is being distributed in an unequal way. Targeting and recommender algorithms work in the background, and the ways they customize a user’s experience are not readily visible. Some platforms provide a brief explanation, but it is typically quite limited,²⁴⁸ and it does not reveal what

²⁴⁶ Up until the early 1970s, *The Pittsburgh Press*, like many newspapers of that era, published job advertisements in sex-segregated columns (e.g., separate “Male Help Wanted” and “Female Help Wanted” columns). A Pittsburgh anti-discrimination ordinance made it unlawful for employers to publish employment ads that discriminated on the basis of sex, and it separately prohibited any person from aiding in an unlawful employment practice. *Id.* at 378. The ordinance was interpreted to prohibit newspapers from carrying sex-segregated job advertisements, and the trial court ordered *The Pittsburgh Press* to cease this practice. *Id.* at 379–80. The newspaper challenged the order on First Amendment grounds. *Id.*

²⁴⁷ *Id.* at 389. Although this case was decided before the Supreme Court extended First Amendment protection to commercial speech in *Virginia State Board of Pharmacy v. Virginia Citizens Consumer Council, Inc.*, 425 U.S. 748 (1976), the Court’s reasoning in *Pittsburgh Press Co.* did not rest on the commercial nature of the advertisements but on the fact that they were illegal under the local ordinance. 413 U.S. at 388 (“Whatever the merits of [Pittsburgh Press’s] contention that commercial speech should be accorded a higher level of protection] may be in other contexts, it is unpersuasive in this case. Discrimination in employment is not only commercial activity, it is illegal commercial activity under the Ordinance.”). The Court affirmed this understanding of its decision in *Pittsburgh Press Co.* in *Virginia State Board of Pharmacy*. 425 U.S. at 759.

²⁴⁸ For example, Facebook allows users to click on a link to find out “Why am I seeing this ad?”, but it provides quite limited information about why they were targeted. A typical explanation is “[advertiser] wants to reach people ages 18 through 38 who live or were recently

opportunities a user has *not* been shown or why other people were selected to receive that information.

Even if users could clearly demonstrate that information about key opportunities was being distributed in a discriminatory manner, they would likely face standing problems. In order to prove damages, plaintiffs would have to rely on a chain of suppositions—if the targeting algorithm were different, I would have seen the ad; if I had seen the ad, I would have acted on it; and if I had applied, I would have been hired. The difficulty arises because of a mismatch between traditional liability schemes and the harms posed by predictive algorithms. Anti-discrimination law tends to look for specific individual victims, but the discriminatory effects of algorithmically driven processes are *systemic*.²⁴⁹ Predictive algorithms rely on statistical generalizations, and as a result, their effects are manifested across groups of people. While the impact will be felt by actual workers, the systemic nature of the problem makes it difficult to trace the precise effects on a specific person.

The nature of predictive algorithms also poses challenges. These algorithms often use dynamic machine learning techniques to generate predictions, and they are constantly incorporating new data and updating in real time. The “model” used to direct information or recommend matches may not consist of a static decision-making protocol that can be easily described. Efforts to determine liability for past decisions will confront the issue of which iteration of the model or which time period of its operation is relevant to proving discrimination. And determining appropriate injunctive relief will be challenging when dealing with a dynamic model.

All of this is not to argue that liability for tech intermediaries for their discriminatory effects should be abandoned. There may be egregious cases where these obstacles can be overcome and platforms can be held directly liable. And, of course, the risk of litigation can be useful as a spur for intermediaries to examine their own practices and prevent discrimination from occurring. In the context of employer liability, the

in the United States” even though far more data went into the targeting decision. Another part of the explanation is completely uninformative: “One reason you’re seeing this ad is that T-Mobile Careers wants to reach people interested in customer service, based on activity such as liking Pages or clicking ads.” First Amended Class and Collective Action Complaint, *supra* note 54, at 2.

²⁴⁹ See Kim, *Data-Driven Discrimination at Work*, *supra* note 15, at 886–88 (explaining the systemic rather than individual nature of the harm caused by biased algorithms).

Supreme Court has repeatedly emphasized that “the preferred means” of achieving Title VII’s goals is voluntary compliance.²⁵⁰

Nevertheless, because of the limitations of retrospective liability regimes, many observers are now advocating for proactive regulatory approaches to address concerns about fairness and bias in algorithmic systems. The advantage of a regulatory approach is that it can intervene at the design stage to anticipate and prevent problems, rather than after deployment when problematic elements may be more difficult to isolate and extract from a working system. As Lehr and Ohm explain, understanding the harms and benefits of automated decision making requires examining the entire process of development of these predictive tools, beginning with their conception and design, rather than treating them as a “black box” after they are operating in the real world.²⁵¹ They urge scholars and policy-makers to pay more attention to the model development phase because it “provides more opportunities and behavioral levers” for developing potential solutions.²⁵² A backward-looking liability regime misses those opportunities.

Scholars have pointed to different legal frameworks for scrutinizing automated decision systems. Some have looked to harness the Federal Trade Commission’s regulatory power. Dennis Hirsch argues that the Commission should use its authority to prohibit “unfair or deceptive” trade practices in order to curb the use of socially harmful algorithms.²⁵³ Danielle Citron and Frank Pasquale similarly call for the Commission to police the unfair use of consumer data in credit decisions under the “unfairness” prong of its regulatory authority.²⁵⁴

²⁵⁰ See, e.g., *Local No. 93, Int’l Ass’n of Firefighters v. City of Cleveland*, 478 U.S. 501, 515 (1986).

²⁵¹ Lehr & Ohm, *supra* note 18, at 657; accord Barocas & Selbst, *supra* note 15, at 731.

²⁵² Lehr & Ohm, *supra* note 18, at 657; accord Barocas & Selbst, *supra* note 15, at 676–77; Kroll et al., *supra* note 124, at 643–44, 701.

²⁵³ Dennis D. Hirsch, *From Individual Control to Social Protection: New Paradigms for Privacy Law in the Age of Predictive Analytics*, *Md. L. Rev.* (forthcoming 2020) (manuscript at 41–42) (on file with author).

²⁵⁴ Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 *Wash. L. Rev.* 1, 23–24 (2014); accord Mark MacCarthy, *New Directions in Privacy: Disclosure, Unfairness and Externalities*, 6 *I/S: J.L. & Pol’y for Info. Soc’y* 1, 1 (2011).

Others have suggested that the European Union's General Data Protection Regulation (GDPR)²⁵⁵ may be a model.²⁵⁶ Although the GDPR is commonly understood as protecting data privacy, it also has a number of provisions addressing the use of automated decision making. The Regulation requires companies to inform data subjects about the existence of such systems as well as to provide "meaningful information about the logic involved."²⁵⁷ An entity relying on automated processing or profiling must also implement "suitable measures to safeguard the data subject's rights and freedoms and legitimate interests,"²⁵⁸ which include freedom from discrimination. In addition, the Regulation requires the use of Data Protection Impact Assessments in order to prospectively identify the risks of harm, including discrimination, to data subjects.²⁵⁹

The idea of using impact assessments to evaluate and mitigate harms caused by predictive algorithms is gaining traction.²⁶⁰ The literature has identified numerous benefits to requiring them. They can force entities to "think hard" about why they are adopting an automated system and what "the collateral effects" of doing so will be.²⁶¹ If assessments occur before an automated decision system is deployed, they can permit interventions in the design and model-building stages, thereby avoiding sources of unfairness or bias before they are baked in. And a rigorous impact assessment will produce information, making clear the assumptions underlying a model and the data used to build it as well as anticipating its effects. That information, if shared, will allow for greater accountability and enable public engagement and input.

Policy-makers have begun to take note. In early 2019, a bill was introduced in Congress that would direct the Federal Trade Commission to develop regulations regarding the use of automated decision

²⁵⁵ Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016, 2016 O.J. (L 119) 1 [hereinafter GDPR].

²⁵⁶ See, e.g., Bryan Casey, Ashkon Farhangi & Roland Vogl, Rethinking Explainable Machines: The GDPR's "Right to Explanation" Debate and the Rise of Algorithmic Audits in Enterprise, 34 *Berkeley Tech. L.J.* 143, 187 (2019) ("[T]he path of least resistance for many companies will likely entail treating the GDPR as the new 'gold standard.'"); Kaminski, *supra* note 18, at 1530.

²⁵⁷ GDPR, *supra* note 255, art. 13.

²⁵⁸ *Id.* art. 22.

²⁵⁹ *Id.* art. 35.

²⁶⁰ See, e.g., Casey et al., *supra* note 256, at 171–72; Kaminski, *supra* note 18, at 1551, 1557–58; Andrew D. Selbst, Disparate Impact in Big Data Policing, 52 *Ga. L. Rev.* 109, 118–19 (2017).

²⁶¹ Selbst, *supra* note 260, at 171.

systems.²⁶² The Algorithmic Accountability Act would empower the Commission to require companies that use algorithms to conduct impact assessments.²⁶³ These assessments would mandate study of how an algorithm was developed and the benefits and costs of its use, including an assessment of the risks posed to individuals, such as the risk of discrimination.²⁶⁴ Beyond these process requirements, however, the proposed legislation offers little in the way of substantive standards, leaving it to the Commission to work out the details through the rule-making process.

While this bill appears unlikely to pass in this Congress, it represents a first step toward developing processes for governing algorithmic decision systems. Of course, the effectiveness of such a system of regulation will depend a great deal upon the details ultimately adopted. Scholars have already identified a number of critical procedural elements. Impact assessments should generate information that can be used to meaningfully assess an algorithm's impact when in operation. This means building into systems ways to audit their performance.²⁶⁵ In addition, some level of transparency and public accountability is necessary. Kroll et al. caution that insisting on transparency may jeopardize other important interests, such as protecting trade secrets, protecting the privacy of individual data subjects, and preventing strategic gaming of automated systems.²⁶⁶ However, as many others have noted, complete transparency is not necessary for meaningful forms of oversight. Methods can be developed to provide enough information to allow for public scrutiny while accommodating competing interests. Lastly, because machine learning algorithms are dynamic systems, assessments of impact must be ongoing and feed back iteratively into design choices.

Apart from procedural requirements, any regulatory system will also have to develop substantive standards. When considering the role of tech intermediaries, there will be tradeoffs between the level of personalization and discriminatory impacts. Prohibiting any form of personalization in order to ensure equality is far too drastic and likely unworkable. Some reasons for narrowing an audience may be legitimate, even if they have a

²⁶² Algorithmic Accountability Act of 2019, S. 1108, 116th Cong. (2019).

²⁶³ Id. § 3(b)(1).

²⁶⁴ Id. § 2(2).

²⁶⁵ See Pauline T. Kim, Auditing Algorithms for Discrimination, 166 U. Pa. L. Rev. Online 189, 191 (2017) (explaining the importance of auditing algorithms to detect and counter bias).

²⁶⁶ Kroll et al., *supra* note 124, at 638–39.

disparate impact—such as limiting the pool to individuals possessing a required license or level of education. Others may be nothing more than proxies for protected characteristics with no connection to job-related skills. Determining the appropriate substantive standards necessary to ensure equal opportunity will entail difficult and likely contested judgments.

The point of creating a regulatory regime is to provide a framework for working through hard questions like these, so it would be premature to lay out specific solutions here. Nevertheless, potential tools for mitigating bias do exist. Computer scientists are developing tools that impose fairness constraints on automated decision systems. Data limitations could bar the use of variables that act as proxies for protected class characteristics but have no causal relationship to job performance. Interfaces can be designed to reduce the risks that users will make discriminatory choices.²⁶⁷ Recommendations could be delivered in a way that minimizes bias—for example, by banding groups of similarly qualified candidates rather than rank ordering them.²⁶⁸ The goal here is not to justify the adoption of any specific proposals but rather to argue for a framework for developing these and other possible solutions that encompass technical expertise, design choices, and empirical evidence while also ensuring public accountability.

V. CONCLUSION

Tech intermediaries are playing an increasingly central role in the markets for crucial resources like employment, housing, and credit, but their impact on the distribution of these opportunities is not well understood. The ability to channel information and match participants offers the possibility of vastly improving the efficiency of these markets and opening opportunities to all on a fair and equal basis. However, there are significant risks of unfairness as well, as illustrated by this study of labor market intermediaries. In particular, we should be wary of the possibility that predictive algorithms will simply reproduce a past history of segregation and exclusion. Although many of today's most important intermediaries are private entities, they play a significant role in allocating

²⁶⁷ See, e.g., Levy & Barocas, *supra* note 19, at 1220–34 (surveying the ways that platforms' design choices can enable or mitigate the discriminatory bias of users).

²⁶⁸ See, e.g., Bogen & Rieke, *supra* note 7, at 24.

access to opportunities. As a result, their activities can have significant social impacts for which they should be accountable.

Law can play a role in ensuring that these technologies work toward a more equitable society rather than reinforcing disadvantage. Existing anti-discrimination laws may apply to these intermediaries, incentivizing them to take care to avoid discriminatory effects. In many ways, however, a liability model is not well suited to addressing the challenges of algorithmically mediated markets. Instead, policy-makers should look to build a framework for examining the impacts of these intermediaries and regulating their design to mitigate social harms.