

ANTIDISCRIMINATORY PRIVACY

*Ignacio N. Cofone**

ABSTRACT

Law often blocks sensitive personal information to prevent discrimination. It does so, however, without a theory or framework to determine when doing so is warranted. As a result, these measures produce mixed results. This article offers a framework for determining, with a view of preventing discrimination, when personal information should flow and when it should not. It examines the relationship between precluded personal information, such as race, and the proxies for precluded information, such as names and zip codes. It proposes that the success of these measures depends on what types of proxies exist for the information blocked and it explores in which situations those proxies should also be blocked. This framework predicts the effectiveness of antidiscriminatory privacy rules and offers the potential of a wider protection to minorities.

TABLE OF CONTENTS

I. INTRODUCTION	140
II. CASE STUDIES.....	142
A. TWO FIELD EXPERIMENTS ON INFORMATION BLOCKING.....	142
B. SCOPE OF AN INFORMATION APPROACH.....	144
C. PREVENTIVE INFORMATION RULES	146
III. PRIVACY VERSUS ANTIDISCRIMINATION	147
A. CURRENT VIEWS	147
B. IDENTIFYING WHEN INFORMATION FOSTERS DISCRIMINATION	149
IV. DISCRIMINATION AS AN INFORMATION PROBLEM	151
A. THE CRUCIAL ROLE OF PROXIES.....	151

* Assistant Professor, McGill University, Faculty of Law. I'm thankful to Abi Adams, BJ Ard, Ian Ayres, Jack Balkin, Delfina Braguerie, Adelle Blackett, Kiel Brennan-Marquez, Goran Dominioni, John Gardner, Eric Goldman, Claudia Haupt, Richard Hynes, Sebastien Jodoin, Angus Johnston, Robert Leckey, Kirsten Martin, Helen Nissenbaum, Paul Ohm, Julia Powles, Jeremias Prassl, Jason Schultz, Katherine Strandburg, Ari Waldman, Felix Wu, and Angela Zorro Medina for their helpful comments. The paper also benefited from comments received at the Privacy Law Scholars Conference (Berkeley), the Canadian Law & Economics Association Conference (Toronto), the Northeast Privacy Scholars Workshop (New York), and internal presentations at the NYU School of Law Privacy Research Group, McGill University, the Technical University of Munich, and the Yale Law School Information Society Project. I also thank Fabian Bargout for his extraordinary research assistance and the editors of the *SMU Law Review*.

B.	CHOOSING WHICH PROXIES TO BLOCK	153
C.	A TYPOLOGY: TRANSFER, REDUCING, AND EXPANDING PROXIES	154
V.	EXPANDING EQUAL PROTECTION	158
A.	ANTICLASSIFICATION AND ANTISUBORDINATION	158
B.	DISPARATE IMPACT-LIKE PROTECTION UNDER DISPARATE TREATMENT	160
C.	FACIALLY NEUTRAL RULES	162
D.	LIMITATIONS: AFFIRMATIVE ACTION	164
VI.	BUILDING A PRECAUTIONARY APPROACH	166
A.	THE DISTINCTION BETWEEN PREJUDICE-BASED AND STATISTICAL DISCRIMINATION	166
B.	A BIAS-DRIVEN AND RATIONAL-CHOICE DISCRIMINATION DISTINCTION	168
C.	IMPORTANCE OF PREVENTING HARM	171
VII.	CONCLUSION	172
VIII.	APPENDIX: IDENTIFYING PROXIES THROUGH CROSS ELASTICITY	174

I. INTRODUCTION

IN a job interview, there are several things that potential employers are prohibited from asking job applicants. Think of rules precluding employers from asking about an applicant's age, disabilities, marital status, or intention to have a child. The logic behind these rules is that decision-makers will be unable to discriminate if they lack the sensitive information needed to do so. These rules operate in several areas of the law, such as those that govern landlord-tenant relationships, healthcare, school admissions, loans, and, perhaps most prominently, employment. If properly understood, they could operate in many more.

This dynamic shows an often overlooked approach to discrimination. The traditional approach to discrimination acts after a discriminatory act produces harm. It is reactive in the sense that it responds to how a decision-maker *used* information about a protected class. But sometimes the law takes a proactive approach to discrimination. It prevents the decision-maker from acquiring the information about an individual's protected class in the first place, which prevents her from taking an action that antidiscrimination law would deem unlawful.

Discrimination is better avoided than compensated. Legal scholars widely recognize that, while litigation is key to providing redress to those who have been discriminated against and to discourage discrimination *ex ante*, preventive measures can address the aim of eradicating discrimination more directly.¹ Antidiscrimination law regulates behavior by curbing

1. MARIE MERCAT-BRUNS, *DISCRIMINATION AT WORK: COMPARING EUROPEAN, FRENCH, AND AMERICAN LAW* 108 (Elaine Holt trans., 2016) (stating that “[a]lthough litigation is important for bringing to light purportedly objective requirements perpetuating

how people *use* information about others to establish discriminatory practices. Regulating the *acquisition* of such information next to these traditional, reactive efforts can yield enormous potential benefits.²

The law does this, however, without a coherent theory or framework to determine when this approach is warranted—and why. Shooting in the dark in such a way has led to mixed results, sometimes solving discrimination in a particular scenario and sometimes backfiring in unforeseen ways, even escalating to scandals like the one surrounding “Ban the Box.”³ In this paper, I address the method of regulating personal information to prevent discrimination. To do so, I explore the information dynamics of discrimination to identify the conditions for designing effective antidiscriminatory information rules, or antidiscriminatory privacy rules. I offer a framework to determine when information should flow, and when it should not flow, to effectively reduce discrimination against minorities.

To identify such conditions, I analyze two canonical case studies often used to illustrate the tension between information blocking (privacy) and discrimination. The first is a well-known study conducted on auditions by female musicians to examine a method for reducing gender discrimination in orchestra hires. Here, blocking information through privacy rules protected against discrimination. The second is the attempt to protect convicts by banning a “box” in job application forms that asks whether the candidate has a criminal record. Here, privacy rules gave rise to further discrimination.

Until now, we do not know why one of the policies worked so well and the other worked so poorly. In other words, we have no theory to explain

workforce segregation, prevention is key to eliminating systemic discrimination” and arguing that there is an interest in institutional changes that focus “on mechanisms of inclusion over causes of exclusion: exploring other measures inciting people to take preventive action against the causes of discrimination or to establish institution-wide safeguards”).

2. For example, this approach has been explicitly taken by the Canadian provinces of Quebec and Manitoba. See Act Respecting the Protection of Personal Information in the Private Sector, C.Q.L.R., c P-39.1, s. 20 (Can. Que.) (“In the carrying on of an enterprise, authorized employees, mandataries or agents or any party to a contract for work or services may have access to personal information without the consent of the person concerned only if the information is needed for the performance of their duties or the carrying out of their mandates or contracts.”); Personal Health Information Act, C.C.S.M., c. P33.5, s. 13(1) (Can. Man.) (“A trustee shall not collect personal health information about an individual unless (a) the information is collected for a lawful purpose connected with a function or activity of the trustee; and (b) the collection of the information is necessary for that purpose.”); see also Patrik Florencio & Erik Ramanathan, *Secret Code: The Need for Enhanced Privacy Protections in the United States and Canada to Prevent Employment Discrimination Based on Genetic and Health Information*, 39 OSGOODE HALL L.J. 77, 111–12 (2001) (stating that “[a]t least two Canadian provinces, Quebec and Manitoba, have nevertheless enacted legislation that is designed to govern the collection of personal information in the private sector, and similar legislation has been proposed in other provinces,” criticizing the Manitoba statute because it only applies to health institutions thereby excluding collection of health information by employers and insurers, and criticizing ambiguities in the Quebec statute that favor employers and insurers).

3. See *infra* Section II.A.

when more information will help antidiscrimination efforts and when it will harm them.

The argument proceeds as follows. Part II discusses two case studies to frame the analysis of the debate and examines the commonalities between the cases to determine when privacy rules could prevent discrimination. Part III focuses on the paper's theoretical implications. It suggests that the privacy versus antidiscrimination debate is premised on the false dichotomy that information either always or never helps fight discrimination. It then addresses the economic literature on information and discrimination to introduce three scenarios in which a privacy-based antidiscrimination approach will be more effective than its *ex ante* alternatives: skewed samples, processing errors, and intentional discrimination.

Part IV suggests that discrimination can also be viewed as an information problem. It shows how proxies determine when an antidiscriminatory privacy rule will be effective. It then evaluates the scope of this method by presenting a typology with three types of proxies (transfer proxies, reducing proxies, and expanding proxies) and discusses which proxies should be blocked.

Part V applies this analysis to the Equal Protection Doctrine. It shows that information rules can be used to provide protections similar to disparate impact when there is no applicable statute that recognizes disparate impact. It addresses how information rules can be used to address facially neutral discriminatory decisions and introduces their limitations to address substantive inequality, particularly regarding affirmative action. Part VI reconsiders the predominant classification between prejudice-based and statistical discrimination in light of these policy considerations and proposes rather to distinguish between bias-driven and "rational" discrimination for information purposes. It then evaluates the general advantages of building a precautionary approach to antidiscrimination: harm is better avoided than compensated.

II. CASE STUDIES

A. TWO FIELD EXPERIMENTS ON INFORMATION BLOCKING

The two case studies analyzed are the cases of orchestra auditions and "Ban the Box." The first is a field experiment conducted on symphony orchestra auditions to examine a method of reducing gender discrimination.⁴ The second is an experiment that studies the attempt to help people who have been in prison reintegrate into society by banning the "box" in employer forms that asks whether the candidate has a criminal record.⁵ As I will show, these cases illustrate when rules that regulate the flow of

4. This study has been used to show privacy's alleged compatibility with antidiscrimination. *See infra* Part III.

5. This study has been used to show privacy's alleged tension with antidiscrimination. *See infra* Part III.

personal information (privacy rules)⁶ are compatible with antidiscrimination efforts and when they are not.

In 2000, a group of economists conducted a field experiment that provided a novel way to test for gender-biased hiring in symphony orchestra auditions.⁷ Most of the prior economic literature on discrimination had focused on disparities in earnings,⁸ but few were able to address actual hiring practices.⁹ Even though orchestra directors are the kind of well-trained professionals that one expects not to have gender biases, and even though symphony orchestras had a fairly transparent hiring procedure, a gap prevailed between the proportion of female elite music school graduates and female elite orchestra hires.¹⁰

After auditions were held behind a physical screen (often a curtain), preventing those hosting auditions from knowing the auditionee's gender, female hires increased by one third.¹¹ The screens increased the *ex-ante* probability of each woman to pass the initial round by 50%.¹² The study generated a research context where discrimination in the workplace seemed pervasive.

Another field experiment that speaks to this issue, although concluding in an opposite direction, studies the Ban the Box policy. The Ban the Box policy prohibits employers from asking applicants before interviews whether they have a criminal record—a question often found in application forms accompanied by yes-or-no checkboxes. This policy is an attempt to protect people who have been in prison and help them reintegrate into society, given that having a criminal record is a substantial barrier for obtaining employment.¹³ In many states, the initiative also explicitly aimed to protect black men to the extent that they are disproportionately likely to have criminal records, so they should disproportionately benefit when criminal records are not considered in employment applications; the policy would thereby level employment possibilities for people who hold those records.¹⁴

By sending fictitious job applications to entry-level positions in states that implemented the policy, Amanda Agan and Sonja Starr showed that

6. See generally HELEN NISSENBAUM, *PRIVACY IN CONTEXT: POLICY, TECHNOLOGY, AND THE INTEGRITY OF SOCIAL LIFE* 231 (2009) (defining privacy in terms of information flows).

7. Claudia Goldin & Cecilia Rouse, *Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians*, 90 *AM. ECON. REV.* 715, 715 (2000).

8. See, e.g., GARY S. BECKER, *THE ECONOMICS OF DISCRIMINATION* 9–10, 14 (1957).

9. Goldin & Rouse, *supra* note 7, at 715.

10. See *id.* at 715–16.

11. *Id.* at 716 (stating that “the screen increases the probability a woman will be advanced out of a preliminary round when there is no semifinal round”).

12. *Id.* at 738.

13. See Harry J. Holzer, Steven Raphael & Michael A. Stoll, *Perceived Criminality, Criminal Background Checks, and the Racial Hiring Practices of Employers*, 49 *J. L. & ECON.* 451, 452–53 (2006); Devah Pager, *The Mark of a Criminal Record*, 108 *AM. J. SOC.* 937, 938–39 (2003).

14. Amanda Agan & Sonja Starr, *Ban the Box, Criminal Records, and Statistical Discrimination: A Field Experiment*, 113 *QUARTERLY J. ECON.* 191, 193–95 (2018).

Ban the Box backfired by fostering discrimination based on race.¹⁵ This adverse effect had been anticipated before,¹⁶ and it had been argued that “empirical estimates indicate that employers who perform criminal background checks are more likely to hire black applicants than employers that do not.”¹⁷ However, until 2016, that effect remained to be proved experimentally.

Agan and Starr showed that, in the absence of the policy, white applicants received 7% more callbacks than similarly qualified black applicants; however, after the policy went into effect, the gap increased to 45%.¹⁸ This more than six-fold increase in the gap indicates that the policy was detrimental for black applicants without a criminal record and beneficial for white applicants with one.¹⁹ “[W]hen employers lack individualized information about criminal history, they tend to statistically generalize that black applicants are likely to have records and white applicants are likely not to have them,”²⁰ although the impact of the policy in hiring practices exaggerates actual racial differences in criminal records.²¹

B. SCOPE OF AN INFORMATION APPROACH

These case studies are useful for exploring how privacy rules can overcome their limitations for antidiscrimination aims. As I mentioned above, statutes and regulations often block information flows to prevent discrim-

15. *Id.*

16. See Holzer, Raphael, & Stoll, *supra* note 13, at 452 (“If accessibility to criminal history information is limited (because of cost or legal prohibitions), employers may infer the likelihood of past criminal activity from such traits as gender, race, or age.”); see also Shawn D. Bushway, *Labor Market Effects of Permitting Employer Access to Criminal History Records*, 20 J. CONTEMP. CRIM. JUST. 276, 278, 284–85 (2004) (using data from the 1997 cohort of the National Longitudinal Survey of Youth to argue that, in states without availability of criminal records, African Americans have slightly lower wages due to statistical discrimination); Keith Finlay, *Effect of Employer Access to Criminal History Data on the Labor Market Outcomes of Ex-Offenders and Non-Offenders*, in STUDIES OF LABOR MARKET INTERMEDIATION 89–90 (David H. Autor ed., 2009) (using the same database to argue that the availability of criminal records worsen ex-offenders position in the labor market but do not improve the position of non-offenders from highly offending groups); Lior Jacob Strahilevitz, *Privacy Versus Antidiscrimination*, 75 U. CHI. L. REV. 363, 367–71 (2008) (explaining the differences between both studies).

17. Holzer, Raphael, & Stoll, *supra* note 13, at 474.

18. Agan & Starr, *supra* note 14, at 195; see also Mike Vuolo, Sarah Langeson & Christopher Uggen, *Criminal Record Questions in the Era of “Ban the Box”*, 16 CRIMINOLOGY & PUB. POL’Y 139, 159 (2017) (finding that race gaps in callbacks is lower when job applicants can signal not having a criminal record); Jeffrey Selbin, Justin McCrary & Joshua Epstein, *Unmarked? Criminal Record Clearing and Employment Outcomes*, 108 J. CRIM. L. & CRIMINOLOGY 1, 48–50 (2018) (finding that record clearing positively affects short-term employment rate and earning capacity); Jennifer L. Doleac & Benjamin Hansen, *Does “Ban the Box” Help or Hurt Low-Skilled Workers? Statistical Discrimination and Employment Outcomes When Criminal Histories are Hidden* (Nat’l Bureau of Econ. Research, Working Paper No. 22469, 2016).

19. See Agan & Starr, *supra* note 14, at 195.

20. *Id.*

21. *Id.*

inatory outcomes but, so far, a framework is lacking—and is needed—to determine when this approach is warranted.

To determine the scope of application of privacy rules with antidiscriminatory purposes, the first question to ask regarding the two case studies is what they have in common. Both of them involved a decision-maker with a certain degree of discretion, a one-time decision (being hired from the audition, or being called for an in-person interview from the resume), and a minority status that the law did not want considered. The approach explored here, accordingly, can be used in any instance in which there is a decision-maker who has a certain degree of discretion, and the law wants to prevent the use of one of the data points available to make a decision—such as minority status.

Given the clarity of these employment examples, I use them to center the analysis on the dynamic between discrimination and disclosure of information. Employment discrimination is particularly relevant because employment is a clear way to distribute wealth and opportunities in our society. It also encompasses a normatively distinct set of decisions since it is covered by Title VII of the 1964 Civil Rights Act,²² which prohibits employment discrimination and includes the Disparate Impact Doctrine.²³ However, employment is not the only socially relevant and institutionalized context in which discrimination takes place, and the antidiscriminatory information approach studied here can be applied to a wide range of discriminatory practices.

People's lives can be significantly affected by discrimination in these decisions, such as in the housing or insurance market.²⁴ Other domains where this framework applies, for example, are healthcare, law enforcement, or university admissions, among many others.²⁵ Casual instances of discrimination that are individually difficult to detect and address, but can be measured in aggregation, can be as ubiquitous as they can be im-

22. See Title VII of the Civil Rights Act of 1964, Pub. L. 88-352, 78 Stat. 241 (codified as amended in 42 U.S.C. § 2000e (2012)).

23. In this sense, employment discrimination in the U.S. is more similar to that of other jurisdictions than is antidiscrimination law in other areas. See Brit. Columbia (Pub. Serv. Emp. Relations Comm'n) v. B.C. Gov't Serv. Emps.' Union, [1999] 3 S.C.R. 3, 3-5 (Can.) [hereinafter Meiorin] (establishing a test to determine the extent of bona fide occupational requirement justification, equivalent to the U.S. business necessity concept).

24. See Lior Jacob Strahilevitz, *Reputation Nation: Law in an Era of Ubiquitous Personal Information*, 102 Nw. U. L. REV. 1667, 1167-78 n.31 (2008) (explaining that the New York housing market had traditionally discriminated against musicians, because they are loud, and lawyers, because they know their rights too well); see also Donald McNeil Jr., *He Took a Drug to Prevent AIDS. Then He Couldn't Get Disability Insurance*, N.Y. TIMES (Feb. 12, 2018), <https://www.nytimes.com/2018/02/12/health/truvada-hiv-insurance.html> [<https://perma.cc/GKT7-H4XC>] (telling the story of a urology resident who was denied disability insurance for taking a pre-exposure prophylaxis drug for HIV).

25. See, e.g., Mary Crossley, *Discrimination Against the Unhealthy in Health Insurance*, 54 U. KAN. L. REV. 73, 73-74 (2005); R. Richard Banks, *Racial Profiling and Antiterrorism Efforts*, 89 CORNELL L. REV. 1201, 1203 (2004); *The Model Minority Is Losing Patience*, THE ECONOMIST (Oct. 3, 2015), <https://www.economist.com/briefing/2015/10/03/the-model-minority-is-losing-patience> [<https://perma.cc/AFC8-5Q48>].

pactful.²⁶ For example, research indicates that doctors are more likely to prescribe painkillers to white patients than they are to black patients—racial disparities in pain assessment lead to different treatment recommendations.²⁷ While it is difficult to regulate this behavior through antidiscrimination law because it is a casual instance of discrimination that is difficult to see at an individual level, one could devise different ways to block information to prevent it, such as by hiding patients' ethnicity in initial stages.

C. PREVENTIVE INFORMATION RULES

As the case studies illustrate, blocking an information flow is successful in combating discrimination in some cases, but not in all of them. One could be tempted to see the interactions in these case studies as not involving a privacy rule because they do not refer to the larger societal values that privacy often serves, such as autonomy or personhood. The reason why one can and should refer to it a privacy rule is that the information dynamic is the same as that of privacy: a channel of personal information flow deemed undesirable is blocked.²⁸ Although the social goals here are different from in most cases in which privacy is established, the interaction between two people is intervened in a similar vein.²⁹ Both the screen in the symphony orchestra study and the Ban the Box policy produced a rule that blocked an information flow. By doing so, the rules forced decision-makers to make their choice without using a data point that the policy-maker did not want considered.³⁰

Whether we choose to call this mechanism “privacy” is not central to adopting the framework that I propose here. One could read this proposal as one that studies “antidiscriminatory information rules” instead of “antidiscriminatory privacy rules” while leaving the proposed framework unchanged. But becoming conscious of this informational parallel allows one to see a new dynamic between privacy and antidiscrimination. Privacy has long been known to safeguard societal interests broader than

26. See Colleen Sheppard, *Institutional Inequality and the Dynamics of Courage*, 31 WINDSOR YEARBOOK OF ACCESS TO JUSTICE 103 (2013) (discussing institutionalized inequality and systemic discrimination).

27. See Kelly M. Hoffman et al., *Racial Bias in Pain Assessment and Treatment Recommendations, and False Beliefs About Biological Differences Between Blacks and Whites*, 113 PROCEEDINGS NAT'L ACAD. SCI. 4296, 4296 (2016).

28. See Jessica L. Roberts, *Protecting Privacy to Prevent Discrimination*, 56 WM. & MARY L. REV. 2097, 2146 (2014) (“Prohibiting requests for information related to race, sex, ethnicity, national origin, religion, age, or disability could bypass discrimination in at least some instances.”); see also Robert Post, *Prejudicial Appearances: The Logic of American Antidiscrimination Law*, 88 CALIF. L. REV. 1, 14–16 (2000).

29. Lisa Austin, *Privacy and the Question of Technology*, 22 LAW & PHIL. 119, 144 (2003) (“instead of waiting for such discrimination to occur and then providing a remedy, we may want to prevent the discrimination in the first place. Granting a privacy right over one’s HIV status can in this way function as a kind of anticipatory remedy”).

30. See generally HELEN NISSENBAUM, *PRIVACY IN CONTEXT* 4, 236–37 (2010); Helen Nissenbaum, *Privacy as Contextual Integrity*, 19 WASH. L. REV. 119, 143–49 (2004) (conceptualizing privacy as information flows).

privacy interests.³¹ For example, Robert Post has said that privacy attempts to protect rules of civility and the identities of individuals in a community,³² and I have argued elsewhere that privacy law also protects reputational interests.³³

Similarly, privacy rules can be used to further society's antidiscrimination interests.³⁴ When used for this purpose, privacy rules have a preventive nature.³⁵ This preventive nature of these rules allows us to avoid the harms created by discriminatory conduct, rather than repairing them.³⁶

III. PRIVACY VERSUS ANTIDISCRIMINATION

A. CURRENT VIEWS

The framework I propose challenges the universal applicability of two opposing beliefs: that privacy and antidiscrimination are always in opposition to one another, sometimes found in economics, and the belief that they are always complementary. I argue that neither scholarly perspective presents a full account of the dynamic between privacy and discrimination: privacy can protect against discrimination as well as enable a discriminatory dynamic. This Part identifies the conditions that give rise to each of these effects.

Often in economics, discrimination is described as a problem of not having enough information about others.³⁷ According to this account, having insufficient information leads people to resort to heuristics to judge others, which can easily result in false opinions. These false opinions, in turn, are attributed to everyone who falls under the heuristic, resulting in beliefs that could be racist, sexist, or homophobic.³⁸

Making information about oneself more available would therefore avoid the need for such heuristics and reduce discrimination, while having more privacy would worsen it.³⁹ Representing the standard economic account, Lior Strahilevitz has argued that "by increasing the availability of information about individuals, we can reduce decision-makers' reliance on information about groups"⁴⁰ and that, therefore, "there is often an

31. Austin, *supra* note 29, at 144 ("By granting a privacy right over the information, we protect an individual against wrongs that are not themselves best described as invasions of privacy").

32. Robert Post, *The Social Foundations of Privacy: Community and Self in the Common Law Tort*, 77 CALIF. L. REV. 957, 1008 (1989).

33. Ignacio N. Cofone & Adriana Z. Robertson, *Privacy Harms*, 69 HASTINGS L.J. 1039, 1058–61 (2018).

34. See Scott Skinner-Thompson, *Privacy's Double Standards*, 93 WASH. L. REV. 2051, 2052–54 (2018) (arguing that principles of constitutional equality should be used to rejuvenate privacy law, which may otherwise have disparate outcomes along the lines of privilege in its application).

35. Austin, *supra* note 29, at 144 ("privacy functions more like an anticipatory remedy than a description of the wrong").

36. See *infra* Section VI.C.

37. See, e.g., Strahilevitz, *supra* note 16, at 364, 380.

38. See *id.*

39. See *id.*; see also Bushway, *supra* note 16, at 285.

40. Strahilevitz, *supra* note 16, at 364.

essential conflict between information privacy protections and antidiscrimination principles, such that reducing privacy protections will reduce the prevalence of distasteful statistical discrimination.”⁴¹

The underlying idea of this account is that, in many contexts, society must tolerate statistical discrimination because the only way to dispense with it is to provide decision-makers with more information, but there is a normative reason for which they should not have access to it.⁴² Increasing privacy, therefore, could only make matters worse.

This idea permeates interactions outside of the academic discourse as well. For example, the LGBTQ movement implicitly used this rhetoric to advocate coming out of the closet as a political strategy. The idea was that increased visibility would lead the general population to see that they interact with LGBTQ people every day and, consequently, lead to a decline in prejudices.⁴³ Their understanding views privacy and efforts to halt discrimination as existing in tension.

Providing more information may be effective as an individual strategy when information is not blocked for the decision-maker; individuals have incentives to disclose information that would help them avoid discrimination. In states that did not subscribe to Ban the Box, for example, job applicants with a criminal record could present further information such as lack of recidivism or other post-conviction jobs in the hopes of a more individualized consideration. In some jurisdictions, women who already have children mention them in their resume to signal that they are less likely to have more.⁴⁴ However, while effective for some individuals, these measures cannot eliminate statistical discrimination—they can only make it narrower.⁴⁵

Some legal scholars, however, have argued that, in specific institutional contexts, limiting information can aid in antidiscriminatory efforts, particularly for landlord-tenant relationships,⁴⁶ genetic information,⁴⁷ and disa-

41. *Id.* at 364; *see also* Strahilevitz, *supra* note 24, at 1682–88.

42. *See* Strahilevitz, *supra* note 24, at 1669–70.

43. *See, e.g.*, MICHELANGELO SIGNORILE, QUEER IN AMERICA: SEX, THE MEDIA AND THE CLOSETS OF POWER 82–83 (1993); RANDY SHILTS, THE MAYOR OF CASTRO STREET: THE LIFE AND TIMES OF HARVEY MILK 224–25 (1988) (describing Harvey Milk’s famous Hope Speech and reporting that Milk stated “I ask my gay sisters and brothers to make the commitment to fight. For themselves, for their freedom, for their country . . . [W]e will not win our rights by staying quietly in our closets . . . We are coming out to fight the lies, the myths, the distortions. We are coming out to tell the truths about gays, for I am tired of the conspiracy of silence, so I’m going to talk about it. And I want you to talk about it. You must come out.”).

44. *See* Joni Hersch & Jennifer Bennett Shinall, *Something to Talk About: Information Exchange Under Employment Law*, 165 U. PA. L. REV. 49, 54–56, 83 (2016) (finding that concealing family information lowers female applicants’ hiring prospects).

45. *See infra* Part VI.

46. *See* Rudy Kleysteuber, *Tenant Screening Thirty Years Later: A Statutory Proposal to Protect Public Records*, 116 YALE L.J. 1344, 1369–72 (2007).

47. *See* Florencio & Ramanathan, *supra* note 2, at 111–12 (2001); Roberts, *supra* note 28, at 2127–40.

bility law.⁴⁸ In particular, Jessica Roberts has made the more general claim that privacy and antidiscrimination are complementary.⁴⁹

Taking the Genetic Information Nondiscrimination Act (GINA) as a case study, Roberts argues that discriminators need information to discriminate, and the law can strip them from it.⁵⁰ Unlike previous statutes, GINA's privacy provision is explicit in its antidiscriminatory efforts (even including them under the employment discrimination title), but it illustrates a wider mechanism that Roberts calls the privacy/antidiscrimination symbiosis.⁵¹

The framework I propose here questions the universal applicability of both beliefs: that privacy and antidiscrimination are opposing values and that they are complementary ones. Privacy rules, I suggest here, have an appropriate scope in aiding antidiscriminatory efforts. To identify such scope, the next Section asks how and when privacy rules can be used as a method to prevent discriminatory decisions. After that, the following Part argues that the effectiveness of privacy rules crucially depends on the type of proxies that they force decision-makers into using.⁵²

B. IDENTIFYING WHEN INFORMATION FOSTERS DISCRIMINATION

The question that emanates from the case studies and the discussion outlined above is under which circumstances the use of privacy rules is an effective method to combat discrimination. After seeing that, in some cases, privacy has aided in antidiscriminatory efforts and, in others, it worsened discrimination, we will see how the existence and type of proxies determine under which conditions blocking an information flow will be successful in preventing discrimination. But before doing so, the next question to address is whether and when the economic approach described might be mistaken in believing that more information would solve the discrimination problem. Thus, blocking information might be effective but unnecessary. This leads to a typology. Privacy rules reduce discrimination in a way that more information does not under three conditions.

First, blocking future information is useful when information samples are expected to be skewed and only a non-infinite amount of information can be gathered. For example, if the New York Police Department decided to engage in predictive policing based only on prior arrest numbers, and if it had more of its police force in the Bronx than in other boroughs, then it would be likely to make more arrests there than anywhere else.

48. See Nicholas Caivano, *Inaccessible Inclusion: Privacy, Disclosure and Accommodation of Mental Illness in the Workplace*, 5 CAN. J. HUM. RTS. 97, 117–27 (2016); Roberts, *supra* note 28, at 2157–58 (focusing on disability law but also extrapolating those conclusions to make a general argument).

49. *Id.* at 2113 (clarifying that this is despite the different normative values held by privacy law and antidiscrimination law).

50. *Id.* at 2126–27.

51. *Id.* at 2127–40.

52. See *infra* Part IV.

This would lead to more comparative police presence in the Bronx, leading to more comparative arrests, and so on and so forth, independent of the actual crime rates. More information would be unhelpful for the New York Police Department in that case, and they would be better off if that information flow were blocked.

Second, blocking certain data points can be useful when people do not update prior beliefs as cleanly as neoclassical economic models would suggest—perhaps due to representativeness heuristics and confirmation biases.⁵³ People have limited time and attention to receive information, even when the information is unlimited.⁵⁴ Moreover, once people form a belief, most do not update their prior beliefs as cleanly as an ideal rational actor would when new information arrives.⁵⁵ The importance of this effect will be different depending on context and on the informational demands of the decision. Simple decisions that require little information are likely to be better updated with new information than complex decisions either rooted in deep beliefs or subject to an information overload because they are informed by an already prodigious amount of information.⁵⁶

Moreover, people often make assumptions and generalizations subconsciously, and these assumptions and generalizations about groups permeate their decisions about individuals both consciously and subconsciously.⁵⁷ In the orchestra auditions example, orchestra directors were probably not discriminating against women willingly, but their gender biases led them to underestimate women's performance. This effect is most relevant when people make decisions based on facts that are neither entirely correct nor mistaken, but rather hold an exaggerated "kernel of truth."⁵⁸ For example, employers in the Ban the Box case were correct about the general effect of race as a predictor of having a criminal record, but on average they exaggerated its magnitude; the actual correlation was neither nonexistent nor as high as employers estimated—employers overreacted to an existing gap.⁵⁹

53. See David Grether, *Bayes Rule as a Descriptive Model: The Representativeness Heuristic*, 95 Q. J. ECON. 537, 537–38 (1980) (describing representativeness heuristic); Raymond Nickerson, *Confirmation Bias: A Ubiquitous Phenomenon in Many Guises*, 2 REV. G. PSYCHOL. 175, 198 (1998) (describing confirmation bias).

54. See JEFFREY ROSEN, *THE UNWANTED GAZE: THE DESTRUCTION OF PRIVACY IN AMERICA* 9–10 (2000); Lawrence Lessig, *Privacy and Attention Span*, 89 GEO. L.J. 2063, 2063–64 (2001).

55. See Jennifer L. Eberhardt et al., *Seeing Black: Race, Crime, and Visual Processing*, 87 J. PERS. SOC. PSYCHOL. 876, 876–93 (2004); Brian A. Nosek et al., *Pervasiveness and Correlates of Implicit Attitudes and Stereotypes*, 18 EUR. REV. SOC. PSYCHOL. 36, 36–53 (2007).

56. See Ignacio N. Cofone & Adriana Z. Robertson, *Consumer Privacy in a Behavioral World*, 69 HASTINGS L.J. 1471, 1489–90 (2018).

57. See Eberhardt et al., *supra* note 55, at 876 (“Merely thinking about Blacks can lead people to evaluate ambiguous behavior as aggressive, to miscategorize harmless objects as weapons, or to shoot quickly, and, at times, inappropriately.”); Nosek et al., *supra* note 55, at 7.

58. Pedro Bordalo et al., *Stereotypes*, 131 Q. J. ECON. 1753, 1791 (2016).

59. Agan & Starr, *supra* note 14, at 195.

A third scenario in which blocking information points can be a useful preventive mechanism is when there are no problems in collecting or processing information but, due to larger societal values, the law does not want an information point considered. This is the case of intentional statistical discrimination, where a decision-maker purposely employs a heuristic for information cost saving, but that heuristic discriminates against a protected class.⁶⁰

These scenarios help determine when a privacy-based antidiscrimination rule is warranted. For all three scenarios, the key element in establishing an effective privacy-based antidiscriminatory rule is noting that, when a data point is blocked with the intention of preventing a discriminatory decision, often there are proxies available for such blocked information, and the rule should address those proxies accordingly.⁶¹

IV. DISCRIMINATION AS AN INFORMATION PROBLEM

A. THE CRUCIAL ROLE OF PROXIES

Proxies are a tool that the law has been using and fearing for decades and a framework to evaluate them properly can be useful to update the conversation about them. The considerations above allow one to examine the set of conditions under which privacy rules will worsen discrimination and the ones under which they will prevent it. Why did a privacy rule work to stop discrimination against female musicians but backfired in the Ban the Box case?

For an answer, we can build on an insight from the literature on algorithmic discrimination. In evaluating how algorithms discriminate, Solon Barocas and Andrew Selbst show that algorithms can pick up data points that act as proxies for proscribed categories, thereby leading to discriminatory results.⁶² This conclusion can be extended to human beings, who can also identify, consciously or unconsciously, proxies for information points that are proscribed by the law. Identifying when decision-makers are able to do this is crucial when evaluating privacy-based antidiscriminatory rules.

60. See *infra* Section VI.A (developing this idea). Viewing Phelps's statistical discrimination model, explained below, in light of this statement, the law might want to prevent people with purple hair from being discriminated against.

61. See *infra* Part IV.

62. See Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CALIF. L. REV. 671, 675 (2016) ("Even in situations where data miners are extremely careful, they can still effect discriminatory results with models that, quite unintentionally, pick out proxy variables for protected classes."); *id.* at 691–92 ("Cases of decision making that do not artificially introduce discriminatory effects into the data mining process may nevertheless result in systematically less favorable determinations for members of protected classes. This is possible when the criteria that are genuinely relevant in making rational and well-informed decisions also happen to serve as reliable proxies for class membership."); see also James Grimmelman & Daniel Westreich, *Incomprehensible Discrimination*, 7 CALIF. REV. ONLINE 164, 171–72 (2017); Jane Bambauer & Tal Zarsky, *The Algorithm Game*, 94 NOTRE DAME L. REV. 1, 2–4 (2018) (explaining how algorithmic decision-making works with proxies in a dynamic setting).

In terms of information dynamics, the difference between one case and the other is the existence of proxies that allowed the decision-makers to gauge the data that the rule attempted to block. Orchestra directors could not know through any other information point whether, for example, a violin player was male or female (they might have guessed had it been an opera audition). Employers in states that implemented the Ban the Box policy, on the other hand, knew or believed that black males had a higher probability of having a criminal record than other applicants.

In the Ban the Box case, privacy rules were less effective not only because they failed to protect most of the group that they intended to protect, but also because they shifted the existing discrimination from a disfavored group to proxies that affected another group disproportionately. This was especially damaging because the second group (black males) was socially disadvantaged and historically considered worthy of protection as well.⁶³

Shifting discrimination in such a way is only possible when proxies for the blocked piece of information are available. Where the proxies are believed to be accurate and are easily observable, the risk of shifting discrimination only increases. The decision-maker can use those proxies with equally negative or even more devastating results than would have resulted from using the blocked information in the original discriminatory decision. Therefore, it is crucial for any effective antidiscrimination-motivated privacy rule to identify and block those information flows as well. Checking for the existence of these proxies is necessary to predict the effectiveness of any antidiscriminatory privacy rules and is of crucial importance to design better ones.

This analysis would have helped predict that the Ban the Box policy would not achieve its objectives. Moreover, it could have predicted that it would backfire. This prediction, in turn, would have saved not only the policy's administration costs but also the social costs of the discrimination it unintentionally generated.

It also illustrates that blocking those proxies as well could have reversed the regulation's unintended effect. For example, the regulation could have mandated the use of initials (and not full names) in resumes, making it more difficult to identify race before the interview.⁶⁴ This would have effectively impeded gender- and race-based discrimination until the interview stage by eliminating available proxies, thereby achieving the policy's aim.

63. See Strahilevitz, *supra* note 23, at 1669–70, 1685–86 (arguing that the case shows that some types of statistical discrimination must be tolerated).

64. See Marianne Bertrand & Sendhil Mullainathan, *Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination*, 94 AM. ECON. REV. 991, 991–92 (2004) (showing that names work as accurate predictors of race, and that people with black-sounding names receive fewer callbacks than people with white-sounding names).

B. CHOOSING WHICH PROXIES TO BLOCK

A crucial task, then, is to identify which pieces of information are proxies for others. A first approximation could lead one to believe that one should measure these objectively, and to consider that proxies are about correlation between two data points. Under this conception, if Abby does not know X about Ben, but she knows Y , and Y happens to correlate with X , then she could see Y as a proxy for X . However, it is not the objective relationship between X and Y , but their relationship in Abby's subjective perception what will determine the extent to which she is likely to use X instead of Y to judge Ben. Abby's subjective perception might or might not trace an objective relationship. This idea is further developed in the Appendix.

One policy outcome that could stem from these considerations is the following: the law should block information to prevent discrimination when there are no available proxies for those information points, or when it can acceptably block those proxies as well.⁶⁵ The problem with such a statement is that, because there will always be some subjective relationship between different information points, one might have to block information *ad infinitum*. One would never cease to find more information points that, to some degree, people perceive to be predictive of each other and would need to be blocked. This would become unfeasible not only because it would be administratively difficult, but also because some of those information points that are perceived to predict protected class will also be genuinely useful to measure job performance—so blocking them might ultimately be undesirable. For example, education can predict race in some institutional settings, but it is also relevant for job performance.

One approach to such a problem would be to implement a discretionary cutoff rule. Such a rule would block information points only above a certain level of perceived predictiveness.⁶⁶

This approach could be complemented with a balancing test between the information point's perceived predictiveness of the protected category and its predictiveness of successful job performance. This is analogous to the business necessity rule, which provides for an employer's duty of accommodation to be sidestepped if it creates undue hardship.⁶⁷ Discriminating against a deaf individual, for example, is proscribed, but it would be reasonable to allow a call center not to take such applications.

65. This relates to the wider argument that both disability law and antidiscrimination law place burdens on decision-makers. Because blocking any proxy might be inconvenient for decision-makers, doing so poses burdens for them. See Christine Jolls, *Antidiscrimination and Accommodation*, 115 HARV. L. REV. 642, 684–86 (2001) (showing that antidiscrimination law, like disability law, contains duties of accommodation).

66. In the terms set in the Appendix, this would mean setting the cutoff point at some level of cross elasticity. The cutoff point could be $E=0.5$, since this is the point in which the information points are closer to being substitutes than to being independent, but higher cutoff points should be considered depending on the number of proxies identified.

67. An even closer analogy exists in Canadian law under the second stage of the analysis established in Meiorin. See *supra* note 23, at 5.

Oftentimes it is difficult to determine what successful job performance entails, but this is a general limitation and such obstacle is not new to employment discrimination law.

There are cases in which one should not give any weight to whether the information that serves as a proxy is also useful to determine job performance.⁶⁸ There are often situations in which policymakers can reasonably believe that there is intentional discrimination that falls under disparate treatment but, because discriminatory intent is difficult to prove, disparate impact is used as an evidentiary approach.⁶⁹ In those cases, one might simply want to block the proxy regardless of its usefulness to measure job performance. For example, imagine a group of decision-makers that cannot perceive race given an antidiscriminatory privacy rule, but they are considering race when making decisions, either consciously or unconsciously, so they use any proxy that they have available. In that case, if as a matter of policy the legislature does not want them to consider race at all, it will also want to prevent them from gauging race through any proxies at all.⁷⁰

Nonetheless, legislatures might want to balance these concerns with perceived predictiveness to avoid unduly burdening decision-makers when lawmakers believe that there is no intentional discrimination (no disparate treatment) but, rather, discrimination with an actual disparate impact.

A different approach, rather than simply blocking all proxies or relying on balancing, stems from analyzing the effects that different proxies have on decisions.

C. A TYPOLOGY: TRANSFER, REDUCING, AND EXPANDING PROXIES

Given what has been discussed so far, I propose to classify the proxies that decision-makers can be forced into when an information point is

68. *See infra* Part VI (describing statistical discrimination).

69. *See infra* Section V.A (describing disparate treatment and disparate impact); *see also* Mark S. Brodin, *Costs, Profits, and Equal Employment Opportunity*, 62 NOTRE DAME L. REV. 318, 357–58 (1987) (arguing that the Supreme Court established disparate impact theory partly due to the difficulty of proving discriminatory intent); George Rutherglen, *Disparate Impact Under Title VII: An Objective Theory of Discrimination*, 73 VA. L. REV. 1297, 1331 (1987). *Contra* Michael Selmi, *Was the Disparate Impact Theory a Mistake?*, 53 UCLA L. REV. 701, 740–41, 768–73 (2006) (arguing that courts do not use disparate impact theory in a manner to sidestep the intentional requirement of disparate treatment); *but see* EEOC v. Joe's Stone Crab, Inc., 220 F.3d 1263, 1281–82 (11th Cir. 2000) (concluding that there was disparate treatment and no disparate impact); Joseph Seiner, *Disentangling Disparate Impact and Disparate Treatment: Adapting the Canadian Approach*, 25 YALE L. & POL'Y REV. 95, 107–09 (2006) (presenting the argument that disparate impact has become an antiquated doctrine).

70. Another scenario in which balancing would be undesirable is proxies that can predict job performance but are redundant to measure it because there are other information points that can predict it as well—and which are not a proxy for protected class. If there are information points that are redundantly informative of job performance, legislatures could mandate that decision-makers use those that do not function as a proxy for protected categories.

blocked into three categories: transfer proxies, reducing proxies, and expanding proxies.

I will call transfer proxies those that shift discrimination from one group to another, non-identical, group. This was the case, for example, under Ban the Box, under which discrimination shifted from people with a criminal record to black men. The groups “people with a criminal record” and “black men” overlapped, but not completely.

Transfer proxies make one group better-off and another group worse-off. Individuals in the information point group that do not overlap with the proxy group see their situation improved, while members of the proxy group that do not overlap with the information group see it worsened. In other words, transfer proxies benefit one group to the detriment of another group that has some overlap with the first. In the Ban the Box case, the shift made “people with a criminal record that are not black men” better-off and “black men without a criminal record” worse-off. The members of group with an overlap (black men with a criminal record) would not have their situation substantively modified.




Figure 1: Illustrates Transfer Proxies

Reducing proxies are those in which the proxy group is a subset of the information group. Instead of shifting discrimination like transfer proxies do, reducing proxies benefit the members of the information group that are not also members of the proxy group.

Imagine, for example, an industry that has shown an aversion towards hiring Latinos, where being Latino is unobservable at a callback stage but can be ascertained through proxies such as having a Spanish first name or last name. A system with initials for first names in resumes would force decision-makers into only using the last name as a predictor of being Latino. In such a way, it would prevent discrimination at a call back stage for individuals with a Spanish first name but without a Spanish last name. It would improve the situation of those in the “Spanish first name or Spanish last name” information group that are not also members of the new “Spanish last name” proxy group. In other words, it would not solve the problem, but at least some individuals would be protected, creating an improvement over the prior situation.




Figure 2: Illustrates Reducing Proxies

Expanding proxies are the opposite of reducing proxies. With expanding proxies, the information group is a subset of the proxy group. Instead of focusing discrimination in a subgroup of the group that was initially discriminated against, expanding proxies diffuse instances of discrimination among the members of a larger group.

The difference in the mechanics between expanding and reducing proxies is that, when a proxy expands an information group, this move is not categorical but probabilistic.

Imagine an employer that wants to avoid hiring someone who might soon take maternity leave. Most jurisdictions prohibit asking this question directly.⁷¹ Many have rightfully argued that, in that case, the employer might disadvantage all women instead—or, more likely, young women who do not already have children.⁷² However, the employer will not disadvantage this larger group of candidates as much as it would have disadvantaged candidates whom it knows will take a maternity leave if it could ask them so, because any of these candidates have only a probability of doing so.⁷³

Because the process is probabilistic, the groups' relative sizes matter. It might not be feasible for an employer to avoid hiring all women, so the

71. *King v. Trans World Airlines, Inc.*, 738 F.2d 255, 258–59 (8th Cir. 1984) (holding that unlawful bias may be found where a job interview included questions about pregnancy, childbearing, and childcare, which were not regularly asked of either male or female applicants); *Barbano v. Madison Cty.*, 922 F.2d 139, 141–43 (2d Cir. 1990) (holding that questioning a female applicant about whether she “would get pregnant and quit” is discriminatory under Title VII); *Bruno v. City of Crown Point, Ind.*, 950 F.2d 355, 358, 362 (7th Cir. 1991) (holding that questioning a female applicant about whether “it was time to have more children” is discriminatory under Title VII); U.S. EQUAL EMP. OPPORTUNITY COMM’N, PRE-EMPLOYMENT INQUIRIES AND MARITAL STATUS OR NUMBER OF CHILDREN, https://www.eeoc.gov/laws/practices/inquiries_marital_status.cfm [<https://perma.cc/MKV2-DL2F>] (advising employers against asking the applicant about her future child bearing plans). For protections against this question at the state level, *see, e.g.*, 804 MASS. CODE REGS. 3.02 (2015) (prohibiting “[i]nquiries into whether applicant has children, plans to have children, or has child care arrangements”).

Similar protections exist in Canada. *See, e.g.*, Charter of Human Rights and Freedoms, C.Q.L.R., c C-12, s. 18.1 (Can. Que.); *Vaid v. Freeman Formal Wear*, 2009 HRTO 2273, ¶ 24–31 (Can. Ont.) (holding that a question relating to a female applicant’s “plans to have a family” contravenes the Human Rights Code, R.S.O. 1990, c. H.19, s. 5(1)).

72. One strategic way for women who already have children to partially counteract this effect would be to add this information in their resume to signal that, since they already have children, they are less likely to have more. *See, e.g.*, Hersch, *supra* note 44, at 56, 83 (finding that concealing family information lowers female applicants’ hiring prospects).

73. *See infra* Section VI.A (describing statistical discrimination).

move from the information point to the proxy will not lead to female applicants being categorically excluded. However, it might be feasible for an employer to avoid hiring people from a minority religion. For example, imagine an employer wanted to avoid hiring women who wear a veil and a policymaker wanted to protect these candidates from such a decision. The policymaker should evaluate how many Muslim women choose to wear a veil: if the majority do, then blocking the information might lead employers to shift to discriminate all Muslim women, against the policy's aim. The question of whether to make the probabilistic shift from women who want children to all women may seem like the question of whether to make the shift from Muslim women who wear a veil to all Muslim women, but the groups' relative sizes should lead to opposite policy conclusions.




Figure 3: Illustrates Expanding Proxies

When an information point is blocked for antidiscriminatory aims, and there are proxies available for such information point, blocking these proxies will be most effective. However, blocking those proxies is often not possible or, while possible, is unfeasible—for example, because they are also legitimate predictors of job performance. But that should not be the end of the story. In those cases, the law should evaluate what is the information dynamic that the proxies introduce.

Shifting discrimination from one group to another (transfer proxies) will generally be considered undesirable. However, privacy-based antidiscrimination rules can sometimes force decision-makers into using proxies that are narrower than the information they were previously relying on (reducing proxies), preventing at least a section of that group from being targeted. These rules can also force decision-makers into using proxies that are diffused enough so that they cannot target any group directly, but must address all members of a larger group probabilistically (expanding proxies).

Whether it is desirable to force a decision-maker into reducing proxies or expanding proxies will depend on the relationship of the proxy with groups that have been historically protected by legislatures or are considered particularly vulnerable. It will also depend on how much one be-

lieves that traditional antidiscrimination law will be effective at addressing the remaining discrimination.

One example of this mechanism is the effort of the several jurisdictions that have historically engaged, even absent this framework, in the third type of proxy to diffuse discrimination against women who might want to have children and are applying for jobs.⁷⁴

Another example are the recent laws in Massachusetts, New York, and Philadelphia that prohibit employers from asking job applicants about previous salaries (and, for Massachusetts, requiring employers to state compensation figures upfront) in order to prevent gender discrimination.⁷⁵ With this prohibition, employers may assume that all female (or minority female) applicants had lower salaries than their male counterparts did, especially by using information from gender pay gap statistics in the applicants' industry or territory. This would force employers into an expanding proxy that diffuses the targeted discrimination. In line with this prediction, evidence from field experiments suggests that employers who cannot observe wage history consider a wider variety of candidates.⁷⁶

V. EXPANDING EQUAL PROTECTION

A. ANTICLASSIFICATION AND ANTISUBORDINATION

In the constitutional context, courts and scholars have struggled to define the appropriate scope of antidiscrimination in a way that mirrors the tension between statutes that include disparate impact and statutes that limit discrimination to disparate treatment.⁷⁷ The information approach analyzed here avoids some of these hurdles that antidiscrimination law faces.

The struggle over the understanding of equal protection since *Brown v. Board of Education*⁷⁸ has been articulated as falling under two competing conceptions of equal protection: anticlassification and antisubordination.⁷⁹ The anticlassification principle holds that one should not classify

74. See Hersch & Shinall *supra*, note 44, at 56, 83.

75. MASS. GEN. LAWS ch. 149, § 105A(c)(2) (2016); PHILA., PA., CODE § 9-1131(2) (2017); see also *Rizo v. Yovino*, 887 F.3d 453, 456 (9th Cir. 2018) (ruling that salary history cannot be used to justify a wage difference between male and female employees). Cf. Chamber of Com. for Greater Phila. v. City of Phila., 319 F. Supp. 3d 773, 812 (E.D. Pa. 2018) (forbidding the use of such information to justify wage gaps but allowing asking the question on free speech grounds).

76. Moshe A. Barach & John J. Horton, *How Do Employers Use Compensation History?: Evidence From a Field Experiment* 3 (CESifo, Working Paper No. 6559, 2017), <http://john-joseph-horton.com/papers/WageHistory.pdf> [<https://perma.cc/7GWK-M4KA>].

77. See, e.g., Patient Protection and Affordable Care Act § 1557, 42 U.S.C. § 18116 (2012); Age Discrimination in Employment Act § 4, 29 U.S.C. § 623 (2012); Equal Credit Opportunity Act, 15 U.S.C. § 1691 (2012).

78. 347 U.S. 483, 488–92 (1954).

79. See Reva B. Siegel, *Equality Talk: Antisubordination and Anticlassification Values in Constitutional Struggles over Brown*, 117 HARV. L. REV. 1470, 1474–75 (2004); Jack M. Balkin & Reva B. Siegel, *American Civil Rights Tradition: Anticlassification or Antisubordination?*, 58 U. MIAMI L. REV. 9, 12–14 (2003); see also Catherine A. MacKinnon,

on the basis of protected categories, such as race and gender. The antistatutory principle holds that one should not disadvantage or aggravate historically oppressed groups, such as people of color and women.⁸⁰

The logistics of antistatutory and antistatutory overlap with the distinction between disparate treatment and disparate impact in antidiscrimination law. The conventional view is that courts predominantly take an antistatutory position, limiting the scope of equal protection to treating all persons with equal civility and respect, but refraining from using equal protection as a tool to redistribute, accommodate, or object to disparate impact.⁸¹ Reva Siegel's historical account argues that American law has shifted ambivalently from one principle to the other.⁸²

A question that arises in this context is whether privacy-based antidiscrimination rules operate under the logic of antistatutory or antistatutory. The answer is that it operates under both. One does not need a notion of antistatutory to block a flow of information. This is an advantage because it makes this proposal viable under the current, antistatutory-dominated, Equal Protection Doctrine, even outside of the scope of statutes that recognize antistatutory, such as Title VII.⁸³ At the same time, because this method can operate alongside standard antidiscrimination law mechanisms—and is not a competing way to address discrimination—it is also fully compatible with the logic of antistatutory.

The logics of antistatutory and antistatutory fall, to some extent, onto the division between disparate treatment and disparate impact in antidiscrimination law. Disparate treatment forbids decision-makers

Unthinking ERA Thinking, 54 U. CHI. L. REV. 759, 765 (1987) (discussing ERA and proposing that the liberal interpretation mistakenly reduced “the problem of the subordination of women to men to a problem of gender classification by law”).

80. Owen M. Fiss, *Groups and the Equal Protection Clause*, 5 PHIL. & PUB. AFF. 107, 108, 157 (1976); see also Cass R. Sunstein, *The Anticaste Principle*, 92 MICH. L. REV. 2410 (1994). Note that this analysis can also be extrapolated to other jurisdictions. For example, equal protection in the U.S. is the functional equivalent to equal rights in Canada and the European Union, and its principles as well can be translated into Canadian law's recognition of direct and adverse effect discrimination, as well as the European Union's recognition of direct and indirect discrimination.

81. See Bradley A. Areheart, *The Antistatutory Turn in Employment Discrimination Law*, 63 ALA. L. REV. 955, 960–67 (2012) (noting, however, that employment discrimination statutes have an antistatutory orientation); see also *Washington v. Davis*, 426 U.S. 229, 246–48 (1976) (declining to extend the disparate impact doctrine from Title VII to the equal protection context, thus taking an antistatutory approach, and holding that equal protection requires discriminatory purpose and disparate impact can be evidence of such purpose).

82. See Siegel, *supra* note 79, at 1544–46. Siegel also traces a third understanding, in between these two, under which equal protection strives not to achieve color blindness or protection from subordinating practices, but protection from the threat of society's balkanization. Reva B. Siegel, *From Colorblindness to Antibalkanization: An Emerging Ground of Decision in Race Equality Cases*, 120 YALE L.J. 1278, 1300–03 (2011) (arguing that courts concerned with antibalkanization focus on diversity more than on equality and that some antistatutory-based strategies might generate further divisions in society).

83. See Title VII of the Civil Rights Act of 1964, Pub. L. 88–352, 78 Stat. 241 241 (codified as amended in 42 U.S.C. § 2000e (2012)).

from making distinctions based on protected categories, such as choosing not to hire women or black men. Disparate impact forbids decisions that impact protected categories disproportionately, such as choosing to hire only people above a certain height or only people who can shave.

While in other jurisdictions, such as Canada and the European Union, both disparate treatment and disparate impact are covered by antidiscrimination law,⁸⁴ in the U.S. disparate impact has a limited scope.⁸⁵ It is applied only when it is explicitly recognized by a statute,⁸⁶ such as Title VII or the Federal Housing Act.⁸⁷

A benefit of the *ex-ante* regulatory approach to discrimination presented here, therefore, is that it allows disparate impact concerns to be addressed with a disparate treatment logic, thus targeting discrimination in areas of the law where disparate impact is not recognized. The next Sections explore this idea further.

B. DISPARATE IMPACT-LIKE PROTECTION UNDER DISPARATE TREATMENT

Privacy rules will be most useful in cases in which the Disparate Impact Doctrine cannot be applied, leaving groups under-protected by traditional antidiscrimination law.⁸⁸ That is, when there is no statute explicitly incorporating disparate impact, when groups are excluded from such protection, or when probatory difficulties exist.

Few statutes prohibiting discrimination contain disparate impact provisions as Title VII and the Federal Housing Act do. Many daily decisions in life, while arguably not as crucial as employment or housing, significantly affect our quality of life and depend on decision-makers who might

84. See Katerina Linos, *Path Dependence in Discrimination Law: Employment Cases in the United States and the European Union*, 35 YALE J. INT'L L. 115, 117 (2010); Joseph A. Seiner, *Disentangling Disparate Impact and Disparate Treatment: Adapting the Canadian Approach*, 25 YALE L. & POL'Y REV. 95, 98–99 (2006). These exist under the categories of direct and adverse effect discrimination (Canada) and direct and indirect discrimination (European Union).

85. *Griggs v. Duke Power Co.*, 401 U.S. 424, 436 (1971) (holding that Title VII invalidates facially-neutral requirements with a disparate impact on a protected category even without discriminatory intent unless there is a proven relationship between job requirements and performance).

86. See, e.g., George Rutherglen, *Disparate Impact, Discrimination, and the Essentially Contested Concept of Equality*, 74 FORDHAM L. REV. 2313 (2006); Richard A. Primus, *Equal Protection and Disparate Impact: Round Three*, 117 HARV. L. REV. 493 (2004); Samuel R. Bagenstos, *The Structural Turn and the Limits of Antidiscrimination Law*, 94 CALIF. L. REV. 1, 4 (2006); Selmi, *supra* note 69, at 732–45.

87. See National Housing Act of 1934, H.R. 9620, Pub. L. 73–479, 48 Stat. 1246 (codified as amended at 12 U.S.C. §§ 1701–1749 (2006)); see also *Tex. Dep't of Hous. & Cmty. Affairs v. Inclusive Cmty. Project*, 135 S. Ct. 2507, 2523–24 (2015) (holding that the Fair Housing Act includes disparate impact claims).

88. Cf. *Roberts*, *supra* note 28, at 2124–26 (stating that privacy belongs to the realm of anticlassification and antisubordination requires providing more information). I would disagree with Roberts on this point. Once we start thinking of antidiscriminatory privacy rules in terms of proxies, we can see that whether blocking information can help situations of disparate impact depends on whether the rule blocks only protected categories but also their proxies. There are, however, additional limitations explored in Section D.

unintentionally discriminate against protected categories.⁸⁹ For example, some argue that men tend to mentor more men than women—a phenomenon that should raise concerns especially because more men than women occupy positions of power.⁹⁰ There is also evidence that doctors tend to provide pain medication to white patients more than to black patients.⁹¹ Most notably, there is no disparate impact protection for decisions on loan applications, which significantly affect people's ability to obtain housing.⁹²

GINA provides an example of this. Despite the importance of genetic discrimination, GINA explicitly prohibits disparate impact analysis.⁹³ This generated criticism arguing that disparate impact assessments should be incorporated to the law.⁹⁴ One plausible rationale for this rule is that blocking information to prevent discrimination, like GINA does, already achieves efforts of antistatutory discrimination similar to disparate impact antidiscrimination.

Privacy rules will also be useful when disparate impact statutes apply but do not afford protection to certain minority groups. These rules are particularly relevant, for example, in efforts to protect LGBTQ individuals from employment discrimination, given that courts have held that Title VII protects discrimination based on race but not sexual orientation.⁹⁵ In the same vein, privacy rules are useful for efforts to protect any other minority that is not considered a protected category under Title VII, such as resident legal aliens.⁹⁶

Moreover, this method is useful in cases in which disparate impact is

89. See Sheppard, *supra* note 26 (arguing that retroactive legal remedies are ineffective at addressing institutionalized inequality and systemic discrimination).

90. See, e.g., KIM ELSSESSER, *SEX AND THE OFFICE: WOMEN, MEN, AND THE SEX PARTITION THAT'S DIVIDING THE WORKPLACE* 159–61 (2015).

91. Hoffman et al., *supra* note 27, at 4296; Sophie Trawalter, Kelly M. Hoffman & Adam Waytz, *Racial Bias in Perceptions of Others' Pain*, 7(11) PLOS ONE 1, 1 (2012).

92. While the Supreme Court has not recognized disparate impact claims under the Equal Credit Opportunity Act (ECOA), the Consumer Financial Protection Bureau has considered that disparate impact does apply to ECOA. See U.S. CONSUMER FED. PROT. BUREAU, CFPB BULL. NO. 2012-04, *LENDING DISCRIMINATION* (2012).

93. Genetic Information Nondiscrimination Act § 208, Pub. L. No. 110-223, 122 Stat. 881 (2008) (“[disparate impact] does not establish a cause of action under this act”).

94. Ifeoma Ajunwa, *Genetic Data and Civil Rights*, 51 HARV. C.R.-C.L. L. REV. 75, 100–12 (2016).

95. See *Boy Scouts of Am. v. Dale*, 530 U.S. 640, 660 (2000) (holding that forcing the Boy Scouts to include a homosexual man violates their First Amendment freedom to express that homosexuality is inappropriate); *DeSantis v. Pacific Tel. & Tel. Co.*, 608 F.2d 327, 329 (9th Cir. 1979) (holding that sexual identity is not covered by Title VII and emphasizing that Congress chose not to pass amendments to extend Title VII to cover sexual preferences). *But see Rene v. MGM Grand Hotel, Inc.*, 305 F.3d 1061, 1068 (9th Cir. 2002) (extending protection because the claim included unwanted physical conduct, considered by the court to be always of a sexual nature and therefore sex discrimination). Canada took a different approach. See *Vriend v. Alberta*, [1998] 1 S.C.R. 493, 496–97 (Can.) (ruling that, even if Alberta's human rights legislation omitted sexual orientation as a ground for discrimination, firing a college professor for being gay violated s. 15 of the Charter and therefore homosexuality should be read as grounds for discrimination).

96. See *Espinoza v. Farah Mfg. Co.*, 414 U.S. 86, 89–96 (1973) (holding that Title VII protection on national origin does not extend to alienage or citizenship).

recognized but there are probatory difficulties.⁹⁷ Discriminatory intent, in particular, is sometimes required to establish discrimination but it is often difficult to prove.⁹⁸ By focusing on the acquisition of suspected information rather than on its use, and by depriving decision-makers of these information points, antidiscriminatory privacy avoids the need for proving what is in the mind of the employer.⁹⁹ This is a significant advantage since this proof is difficult to obtain as employers have incentives to hide discriminatory intent and can do so especially when they have mixed motives.¹⁰⁰

C. FACIALLY NEUTRAL RULES

Antidiscriminatory privacy would be effective at dealing with facially neutral screening rules that serve ulterior motives. As we saw above, the key to using privacy to fight discrimination is the identification and elimination of easily observable proxies for protected categories.¹⁰¹ Therefore, a facially neutral screening rule that covertly aims to discriminate can be addressed by identifying whether the seemingly facially neutral information used as the basis for a decision works as a proxy for a protected category. If it does, then the discriminatory decision rule can be prevented by blocking the information flow, therefore achieving the same result as striking it down for its disparate impact, but all within a disparate treatment logic.

For example, in *New York City Transit Authority v. Beazer*, the Supreme Court decided that it was constitutional not to hire candidates under methadone treatment even if it created a disparate impact against black applicants.¹⁰² The decision hinged on whether the claim was a Title VII one—and the dissent argued that it was—as Title VII would have recognized the decision as discriminatory due to such disparate impact, but the Equal Protection Clause would not require so. Similarly, in *NAACP v. North Hudson Regional Fire & Rescue*,¹⁰³ a fire department's residency requirement was questioned for having a disparate impact against black candidates. In this case, the decision and its residency requirement were considered to fall under Title VII, and the Third Circuit held that it failed to have established business necessity.¹⁰⁴ With very sim-

97. See Roberts, *supra* note 28, at 2149–55.

98. See *Washington v. Davis*, 426 U.S. 229, 247–48 (1976) (holding that equal protection requires discriminatory purpose).

99. See *Watson v. Fort Worth Bank & Tr.*, 487 U.S. 977, 986–87 (1988).

100. *Fuller v. Phipps*, 67 F.3d 1137, 1141–42 (4th Cir. 1995); *Tyler v. Bethlehem Steel Corp.*, 958 F.2d 1176, 1181 (2d Cir. 1992); *Price Waterhouse v. Hopkins*, 490 U.S. 228, 246–47 (1989); see also *Desert Palace, Inc. v. Costa*, 539 U.S. 90, 98–100 (2003) (holding that overt discrimination is not always required in mixed-motive cases).

101. See *supra* Part IV.

102. *New York City Transit Auth. v. Beazer*, 440 U.S. 568, 584–94 (1979).

103. 665 F.3d 464 (3d Cir. 2011).

104. When disparate impact is shown, Title VII allows a defense if the employer can demonstrate that the challenged practice is “job related for the position in question and consistent with business necessity.” See *Lanning v. Se. Penn. Transp. Auth.*, 181 F.3d 478, 489 (3d Cir. 1999) (constructing the language in Title VII to mean minimum qualifications

ilar fact patterns, the cases reached opposite results, although there might be policy reasons to treat them the same.

An antidiscriminatory privacy rule blocking the information point about a candidate's use of methadone or a candidate's residency in the area would have avoided the efforts necessary to frame the case under Title VII, which were successful in the former case and unsuccessful in the latter. In such a way, this method allows us to address facially neutral forms of discrimination against minorities not protected by Title VII and discriminatory decisions outside of an employment relationship (therefore not covered by its disparate impact branch). When State action is involved, these are also currently unprotected by the Equal Protection Clause when interpreted under a disparate treatment logic.¹⁰⁵

This feature becomes especially relevant in overcoming the distinction between mutable and immutable traits. In our current legal context, courts starkly distinguish between personal traits based on whether they are mutable or immutable.¹⁰⁶ This, according to Kenji Yoshino, traces to an artificial distinction between being a member of a protected category and behaving like one.¹⁰⁷ Courts protect individuals from distinctions based on immutable traits such as skin color and chromosomes, but rarely from those based on the behaviors typically associated with those groups.¹⁰⁸ According to Yoshino, the underlying logic of this separation is that minority members can suppress their behavior to blend into the mainstream, so less protection is warranted for them.¹⁰⁹

In such context, more subtle kinds of discrimination exist towards groups that do not assimilate—either because they refuse to or because they cannot.¹¹⁰ Assimilation is the way to avoid this subtler form of dis-

that are necessary to perform the job successfully); *see also* *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425 (1975) (setting a high standard for job relatedness, arguing that to show job relatedness for business necessity in a standardized test the employer must show empirical evidence on correlation between test scores and job performance); Elizabeth Bartholet, *Application of Title VII to Jobs in High Places*, 95 HARV. L. REV. 945, 978 (1982) (arguing that courts have been strict on low-level employment regarding business necessity requirements and lenient with high-level employment for the same requirements).

105. *Village of Arlington Heights v. Metro. Hous. Dev. Corp.*, 429 U.S. 252, 265 (1977) (“Proof of racially discriminatory intent or purpose is required to show a violation of the Equal Protection Clause.”). Moreover, “Official action will not be held unconstitutional solely because it results in a racially disproportionate impact.” *Id.* at 264–65.

106. *See, e.g., Andrews v. Law Society of B.C.*, [1989] 1 S.C.R. 143, 195 (Can.) (illustrating the logic of immutable traits by ruling that, when a ground is not covered by provincial legislation, it should be read in through s. 15 of the Canadian Charter of Rights and Freedoms if it is “not within the control of the individual and, in this sense, immutable,” meaning that it qualifies as “a characteristic of personhood not alterable by conscious action”).

107. *See* Kenji Yoshino, *Covering*, 111 YALE L.J. 769, 772, 774–75 (2002).

108. *Id.* at 876–77. Also note that courts use different types of protections for individuals that have different immutable characteristics. Strict scrutiny is applied to race, but intermediate scrutiny is applied to gender, while disability and sexual orientation are not considered protected categories.

109. *Id.*; KENJI YOSHINO, *COVERING* 22–23, 172 (2006).

110. Yoshino, *supra* note 107, at 839, 913 (referring to “queers,” by essence those who refuse to cover and pregnant women, by nature who cannot cover). YOSHINO, *supra* note 109, at 19 (“As the gay rights movement gained strength, the demand to convert gradually

crimination and, at the same time, it is precisely its effect.¹¹¹ “[A]fter all, the logic goes, if a bigot cannot discriminate between two individuals, he cannot discriminate against one of them.”¹¹²

Privacy rules are an unexpected way out of the dilemma. What Yoshino could have said is that if a bigot cannot *distinguish* between two individuals, he cannot discriminate against one of them. When applied to concrete, potentially discriminatory decisions, privacy rules can provide coverage without forcing people to cover; they provide assimilation without requiring people to actually assimilate. This also serves long-term goals of reducing biases by furthering the inclusion of minority members: if more minority members are hired, decision-makers will progressively see the idea of hiring members of such class as less unusual.¹¹³

D. LIMITATIONS: AFFIRMATIVE ACTION

While being operable under an anticlassification paradigm is an advantage for privacy-based antidiscrimination rules in the sense that it makes them compatible with the mainstream of antidiscrimination doctrine and case law, this compatibility also implies their most important limitation. While antidiscriminatory privacy rules work well to avoid disparate treatment, their effectiveness is limited when dealing with some aspects of disparate impact. In particular, antidiscriminatory privacy rules are not useful to address cases in which affirmative action is considered desirable.¹¹⁴

This is because the method can fix a decision-making process that involves discrimination, but cannot fix a selection set that is in itself unequal. In other words, as it can be seen by comparing both case studies, since the method only addresses the decision-making process, it is unlikely to solve significant structural inequalities.

The problem of inadequacy in the decision-making process and the problem of a candidate pool with embedded inequality are different, and blocking information from the decision-maker can only address the former. The main limitation is when inequality is substantive (socio-economic conditions, generations of unequal educational opportunities) so

ceded to the demand to pass. This shift can be seen in the military’s adoption in 1993 of the ‘Don’t ask, don’t tell’ policy, under which gays are permitted to serve so long as we agree to pass. Finally, at millennium’s turn, the demand to pass is giving way to the demand to cover—gays are increasingly permitted to be gay and out so long as we do not ‘flaunt’ our identities.”).

111. Yoshino, *supra* note 107, at 877 (“Put differently, the descriptive claim that the group *can* assimilate because of the mutability or invisibility of its defining trait transmutes into the prescriptive claim that the group *should* assimilate with very little intervening investigation by a court.”).

112. Kenji Yoshino, *The Pressure to Cover*, N.Y. TIMES (Jan. 15, 2006), <http://www.nytimes.com/2006/01/15/magazine/the-pressure-to-cover.html> [<https://perma.cc/AP23-C9VF>].

113. *See generally* Grutter v. Bollinger, 539 U.S. 306, 308 (2003) (arguing that the inclusion of different races in an educational program helps break down stereotypes).

114. Roberts, *supra* note 28, at 2124 (“Specifically, to assess whether a policy disproportionately impacts a particular group, or to allow affirmative action, the covered entity must have some knowledge of the relevant information.”).

that neutral decision parameters still lead to discriminatory outcomes. If a decision-maker must select from a pool of candidates with very few minority candidates, the chosen candidates are less likely to be minority members even if the decision-maker is impartial. Addressing embedded inequalities requires different methods, such as affirmative action. But affirmative action cannot be implemented if the decision-maker does not have access to the relevant data point.

When there are embedded inequalities, and decision-makers might also be biased against protected categories, the question of whether blocking information will be useful for minority members or an affirmative action program (which requires allowing for the information) will be better for them is a question of how much substantive inequality they face. For example, to choose a Supreme Court justice, will ignoring gender or considering gender lead to equal opportunities for women? If we ignore gender and use only information such as employment history, this might lead to unequal outcomes depending on whether positions such as circuit judge or appellate judge were inaccessible to women, for example due to implicit sexism or loss of opportunity due to prior choices to take maternity leaves. Mechanisms that build on antistatutory principles (such as affirmative action) would provide a more robust protection in this situation. In the orchestra auditions example, blocking gender worked because women had similar access to a variety of musical instruments and similar access to elite music schools. But if women had historically been confined to playing the piano and being educated through private lessons at home, having auditions behind a curtain would not have resolved such substantive inequality problem.

There are some alternatives for making this method compatible with affirmative action and other tools that address diversity concerns under an antistatutory logic. One possibility would be to condition the information flow instead of banning it directly. When dealing with explicitly diversity-concerned decision-makers, information could be released under the condition of a specific use: if active diversity measures are to be established. For example, if regulation requires companies to have a gender quota for the board of directors, but not one for interns, and if gender discrimination was suspected for lower hierarches as well, regulation could also require the company to receive gender-blind resumes.¹¹⁵ Then, if one of these companies wanted to go further in their diversity efforts and establish an affirmative action program for women in the workplace for which they would need information on gender, regulation could allow that company to access applicants' gender information conditionally on

115. Note that, historically, quotas were not considered affirmative action. Information on race or gender can be included to increase diversity as a soft variable in a holistic analysis. See *Ricci v. DeStefano*, 557 U.S. 557, 589–593 (2009); *Fisher v. Univ. of Tex.*, 570 U.S. 297, 304 (2013); *Fisher v. Univ. of Tex.*, 136 S. Ct. 2198, 2207 (2016).

the existence of such program.¹¹⁶

Another possibility would be choice bracketing. A system could be designed to block information at stages where decision-bias could be more prevalent and allow information at stages where antisubordination principles are likely to play a larger role. For example, gender could be blocked at a first resume submission stage, and then considered when choosing the candidates that would move from the first to the second round of interviews.

VI. BUILDING A PRECAUTIONARY APPROACH

A. THE DISTINCTION BETWEEN PREJUDICE-BASED AND STATISTICAL DISCRIMINATION

Broadly speaking, economists distinguish two reasons why a decision-maker might engage in discrimination. The first occurs when the decision-maker has a bias or animus against a particular group of people.¹¹⁷ I will call this “prejudice-based discrimination.”¹¹⁸ The second occurs when the decision-maker uses someone’s gender, race, sexual orientation, or any other characteristic to make a statistical inference. Although it does not assume animus or bias against the protected category, it still involves the unequal treatment of equals.¹¹⁹ This is commonly referred to as statistical discrimination.¹²⁰

The standard economic narrative is that a large amount, if not most, of discrimination is of the statistical type.¹²¹ This idea is sometimes found in legal scholarship as well.¹²² It is an established idea in economics that, when a decision-maker (in our prior example, Abby) wants to learn information about another individual (Ben) and such information is impossible or costly to acquire, the decision-maker will rely on statistical

116. Of course, some degree of monitoring would be needed to ensure that the program was legitimate and not a covert way to find the information to install a discriminatory practice.

117. See BECKER, *supra* note 8, at 8–9; Cass R. Sunstein, *Why Markets Won't Stop Discrimination*, 8 SOC. PHIL. & POL'Y 22, 22 (1991).

118. This is commonly referred to as taste-based discrimination in the economics literature. See Jonathan Guryan & Kerwin Kofi Charles, *Taste-Based or Statistical Discrimination: The Economics of Discrimination Returns to Its Roots*, 123 ECON. J. 417 (2013) (tracing the evolution of empirical work in economics concerning discrimination).

119. See Edmund S. Phelps, *The Statistical Theory of Racism and Sexism*, 62 AM. ECON. REV. 659, 659–61 (1972); Kenneth Arrow, *The Theory of Discrimination*, in DISCRIMINATION IN LABOR MARKETS 3–4 (Albert Rees & Orley Ashenfelter eds., 1973).

120. Hanming Fang & Andrea Moro, *Theories of Statistical Discrimination and Affirmative Action: A Survey*, in SOCIAL ECONOMICS 140–42 (Jess Benhabib, Matthew Jackson & Alberto Bisin eds., 2011).

121. See Stewart Schwab, *Is Statistical Discrimination Efficient?*, 76 AM. ECON. REV. 228, 228–29 (1986); David A. Strauss, *The Law and Economics of Racial Discrimination in Employment: The Case for Numerical Standards*, 79 GEO. L.J. 1619, 1622 (1991); Holzer, Raphael, & Stoll, *supra* note 13, at 452.

122. Rachel F. Moran, *Whatever Happened to Racism?*, 79 ST. JOHN'S L. REV. 899, 908 (2005).

generalizations of the groups to which that individual belongs.¹²³

In Edmund Phelps' seminal statistical discrimination model,¹²⁴ an employer cannot observe workers' level of skill, q , drawn from a normal skill distribution,¹²⁵ but it can observe their group identity and a noisy productivity signal.¹²⁶ We can define such identity in any way, such as P for purple hair and G for green hair.¹²⁷ Under the model, the question of statistical discrimination is the question of why two workers with the same productivity signal, but from different groups, are treated differently.¹²⁸

This can take place under two scenarios: stereotyping and differential observability. Stereotyping takes place when the groups' signals are equally informative of individuals' productivity but the employer believes that one group, P , has lower average human capital investments leading to lower average skill. It does not matter whether the identity is probative of skill level, but only whether the employer believes so.¹²⁹ In that case, employers will consider employees from P to have a lower expected productivity even when they have the same signal as G . Therefore, P workers will receive a lower salary under the same signal.

Differential observability takes place when the skill distributions are identical but the signals for P workers' skills are less informative than those of G workers. Because this will mean that the expected productivity of a P worker with any signal will be closer to that of the population average, this will lead highly qualified P workers to receive a lower salary than their G equivalents, and low qualified P workers to receive a higher salary than their G equivalents.¹³⁰

Twenty years ago, we began to see empirical evidence of this process. In a field experiment involving fictitious automobile buyers, Ian Ayres and Peter Siegelman showed that car dealers quoted significantly lower prices to white men than to female and black buyers.¹³¹ Under Phelps'

123. See Phelps, *supra* note 119, at 659; Dennis J. Aigner & Glen G. Cain, *Statistical Theories of Discrimination in Labor Markets*, 30 *IND. LABOR RELAT. REV.* 175, 176–77 (1977); see also Cofone & Robertson, *supra* note 33, at 1049–50 (explaining the relationship between privacy and statistical generalizations).

124. Phelps, *supra* note 119, at 659.

125. $N(\mu_j, \sigma_j^2)$

126. θ , where $\theta = q + \varepsilon$, q being skill and ε being a normally distributed zero-mean error.

127. $j \in \{P, G\}$

128. See Shelly Lundberg & Richard Startz, *On the Persistence of Racial Inequality*, 16 *J. LABOR ECON.* 292, 293, 306 (1998) (introducing a model showing that statistical discrimination in competitive markets and without differences in average human capital introduces inefficiencies in the system. Basically, due to statistical discrimination minorities face lower incentives to invest in human capital, community social capital is lowered, and they develop lower levels of productivity.).

129. See *infra* Appendix.

130. Fang & Moro, *supra* note 120, at 137–40.

131. Ian Ayres & Peter Siegelman, *Race and Gender Discrimination in Bargaining for a New Car*, 85 *AM. ECON. REV.* 304, 304 (1995) (clarifying that the empirical evidence could be motivated by both theories of discrimination); see also Ian Ayres & Joel Waldfogel, *A Market Test for Race Discrimination in Bail Setting*, 46 *STAN. L. REV.* 987, 991–92 (1994)

model, the car dealers probably made an inference of the first type and assumed that white men are more informed about the value of cars or are better equipped to negotiate bad offers than other groups. In a groundbreaking field experiment ten years later focusing on race, Marianne Bertrand and Sendhil Mullainathan showed that resumes with white-sounding names have a fifty percent greater chance of receiving a callback than identical resumes with black-sounding names.¹³² This would be, once again, the first type of statistical discrimination under Phelps' model. As resumes are equally informative on average, employers seemed to have thought that black applicants were on average less qualified.

Title VII of the Civil Rights Act bans the use of these types of statistical generalizations to exclude some historically disadvantaged groups along a set of dimensions, such as race and gender. This is also the case when the generalization is true.¹³³

B. A BIAS-DRIVEN AND RATIONAL-CHOICE DISCRIMINATION DISTINCTION

The information dynamic identified above leads to the conclusion that, even though prejudice-based discrimination and statistical discrimination are often considered to work very differently,¹³⁴ and in many contexts they do, for the purposes of privacy rules they operate under the same information dynamic.

There are relevant legal implications to determining whether a discriminatory decision was made under a prejudice-based classification pattern or a statistical inference that has a disparate impact.¹³⁵ However, both types of discrimination have the same information dynamic. While a distinction between prejudice-based and statistical discrimination can be useful to determine whether one should apply disparate treatment or disparate impact doctrine,¹³⁶ it does not determine (*ex ante*) which informa-

(using data on bail amounts to show race discrimination in bail setting in New Haven, by showing lower bond rates (bond fee divided by bail amount) for minority defendants).

132. Bertrand & Mullainathan, *supra* note 64, at 992 (showing that, for employers in Boston and Chicago, resumes with white-sounding names have fifty percent more chances of receiving a callback).

133. *See* City of Los Angeles Dep't of Water & Power v. Manhart, 435 U.S. 702, 708 (1978).

134. *See, e.g.,* Strahilevitz, *supra* note 16, at 373–75 (arguing as a pragmatic concern that disclosing previously private information will help antidiscriminatory efforts only as long as statistical discriminators significantly outnumber prejudice-based discriminators).

135. Disparate treatment applies to both prejudice-based and statistical discrimination, as long as the groups are treated differently. One could have a classification pattern based on statistical discrimination. For example: "I will not hire people with purple hair because they have lower average human capital than people with green hair." However, the decision-maker in that example is concerned with human capital, not hair color. So, if hair color as a proxy for human capital in that decision is not made explicit (i.e., "I will hire people that I perceive to have high human capital") it will be treated as a disparate impact problem, not a classification one.

136. *See supra* Sections V.A and V.B.

tion flow should be blocked. The same privacy-based preventive solution should be applied to both.

This leads to the possibility of making a similar distinction instead: one between bias-driven discrimination and rational-choice discrimination—both of which can work under disparate treatment or disparate impact doctrine.¹³⁷ Rational-choice discrimination would concern “correct” statistical inferences, while bias-based discrimination would concern choices made based on animus, prejudice, or otherwise incorrect inferences.¹³⁸ From the two scenarios described for statistical discrimination, the first (belief of lower average human capital) would be bias-driven, while the second (different informativeness of signals) would be rational-choice driven.

Initially, one would think that this distinction is also unimportant in discrimination-preventive privacy rules. In the Ban the Box example, for instance, whether employers were correct in believing that black men are more likely to have a criminal record (making it rational-choice discrimination) or they were mistaken (making it bias-driven discrimination) is irrelevant for the purposes of eliminating the information that enables a discriminatory act. However, to prevent discrimination by blocking information flows one must understand decision-makers’ perceptions.¹³⁹ People do not always update their prior beliefs when they obtain new information.¹⁴⁰ The way that a decision-maker processes information will be different if she does so through an irrational heuristic or one that is only an information cost saving mechanism.

Moreover, privacy rules affect bias-driven discrimination with an added wrinkle that does not exist in the way they affect rational-choice discrimination. For bias-driven discrimination, privacy rules can often effectively shelter the disadvantaged groups that they attempt to protect. However, this comes at the cost of not reducing those biases against them, which generate the need for protective rules in the first place. That is, privacy rules can prevent discrimination due to bias, but cannot fix the bias itself.

137. See Linda Hamilton Krieger, *The Content of Our Categories: A Cognitive Bias Approach to Discrimination and Equal Employment Opportunity*, 47 *STAN. L. REV.* 1161, 1169–71 (1995) (arguing that discriminatory employment decisions often derive not from discriminatory intent but from judgment errors that are unintentional, and suggesting to follow a cognitive process approach to discrimination instead).

138. See Bordalo et al., *supra* note 58, at 1753–54 (“We present a model of stereotypes based on Kahneman and Tversky’s representativeness heuristic. A decision-maker assesses a target group by overweighting its representative types, defined as the types that occur more frequently in that group than in a baseline reference group.”); see also David Arnold, Will Dobbie, & Crystal Yang, *Racial Bias in Bail Decisions*, 133 *QUARTERLY J. ECON.* 1885, 1886–89 (2018) (using a similar distinction that separates correct and incorrect inferences of risk consistent with the theory of stereotyping).

139. See *supra* Section VI.B; see also Christine Jolls & Cass R. Sunstein, *The Law of Implicit Bias*, 94 *CALIF. L. REV.* 969, 982 (2006) (describing the role of biases in decision-making); Roberts, *supra* note 28, at 2170–72.

140. See *supra* Section VI.C; see also Eberhardt et al., *supra* note 55, at 876–92; Nosek et al., *supra* note 55, at 50–52.

This produces a similar tradeoff to that of “covering,” where minorities’ civil rights face the tension of their short-term interest to blend in and their long-term interests to promote identity politics.¹⁴¹ The effectiveness of a privacy rule that blocks information will depend on the timing of the protection. For a minority with the ability to cover, such as the LGBTQ community or religious minorities,¹⁴² an effective privacy rule that hides their status and protects them from a discriminatory decision (being hired, or having a cake baked for their wedding) will be easier to implement than for a protected group without such ability, such as racial minorities or women.

However, a rule that blocks information in a long-term relationship will be more difficult for the individual and potentially more detrimental to the long-term interests of the group. The rule would be difficult and damaging for the individual because she would have to exert a constant effort to cover. It would be detrimental to the long-term interests of the group because members of society would not get used to seeing members of that group integrate—it reduces those biases that generate the need to cover in the hope that one day that need will disappear.¹⁴³ In the LGBTQ example, this was sometimes referred to as the normalizing effect of coming out as a group.

This dynamic illustrates that privacy rules will be useful to address discrimination for one-time decisions, rather than ongoing relationships. It also illustrates that these information rules will work to combat unconscious biases differently than do traditional antidiscrimination rules.

In sum, a traditional distinction in economics separates prejudice-based and statistical discrimination, where the first is driven by an intent to discriminate against the protected class and the second is not.¹⁴⁴ I propose that, while prejudice-based and statistical discrimination work differently in many ways, they have the same underlying information dynamic. Moreover, I propose that there is another relevant distinction between bias-driven and rational-choice discrimination. When bias-driven discrimination is involved, those who design and implement antidiscriminatory

141. See Yoshino, *supra* note 107, at 837–38; see also MARTHA MINOW, MAKING ALL THE DIFFERENCE: INCLUSION, EXCLUSION, AND AMERICAN LAW 42, 48 (1991) (introducing the dilemma of difference, under which people must choose whether to hide their differences with a ruling majority to blend in or emphasize their differences in order to be compensated for the adverse consequences of such differences, but then also exacerbate those adverse consequences by making the difference more evident); Adam Cureton, *Hiding a Disability and Passing as Non-Disabled*, in DISABILITY IN PRACTICE: ATTITUDES, POLICIES, AND RELATIONSHIPS 15 (Adam Cureton & Thomas Hill Jr. eds., 2018) (discussing when people with non-visible disabilities choose to pass as non-disabled or identify themselves as having a disability).

142. See Yoshino *supra* note 107, at 774–80 (explaining how different minorities cover differently).

143. This of course assumes that, given the option, at least some minority members would choose to integrate, were they free to do so.

144. See, e.g., Strahilevitz, *supra* note 16, at 365 (“Statistical discrimination is based, not on irrational animus, but on the use of heuristics by decisionmakers who believe—correctly or not—that observable hallmarks of membership in a group correlate with some undesirable characteristic.”).

privacy rules should be aware of the potential long-term detrimental effects of these rules, usually but not always offset by their short-term gains. There is a tradeoff between assimilation and avoiding the need for assimilation.¹⁴⁵ For rational-choice discrimination, this potential long-term problem is not present.

C. IMPORTANCE OF PREVENTING HARM

The empirical studies on discrimination mentioned above¹⁴⁶ raised critical questions about the effectiveness of Title VII in preventing discrimination in the workplace.¹⁴⁷ Antidiscrimination law, some argue, is ineffective at dealing with unconscious biases—discrimination remains nonetheless.¹⁴⁸

Independently of where one stands in this assessment, the claim that antidiscrimination law has limited effectiveness in dealing with unconscious bias leads to the question of what alternative and complementary methods exist to combat discrimination. By focusing on reproachable conduct in the use of information, antidiscrimination law necessarily has an *ex post* orientation. Thus, it can prevent discrimination only insofar as decision-makers factor in the expected costs of sanctions *ex ante* and avoid the undesirable conduct in order to avoid those sanctions—in the same way that tort law prevents accidents. However, relying on people to avoid a conduct because they anticipate sanctions is less effective when the reproachable conduct is deciding based on biases that are unconscious. In some situations, people might not know that they are engaging in a discriminatory conduct, or might tell themselves that they are not doing so. This inefficacy highlights the importance, particularly for antidiscrimination, of a complementary *ex ante* approach that focuses on prevention like the one this paper discusses.

Having a precautionary approach to antidiscrimination that complements traditional antidiscriminatory efforts presents two central advantages. First, its ability to deal with unconscious biases. Second, its ability to address not only classification but also subordination problems without generating a worry of breaking social cohesion—insofar as the subordination problem is in the decision-making process and not in the embedded inequality of the selection pool.¹⁴⁹

145. Cultural assimilation can be described as the process according to which a minority group, or some of its members, starts converging in cultural resemblance towards the majoritarian group that is dominant in their society.

146. See Ayres & Siegelman, *supra* note 131, at 305; Ayres & Waldfoegel, *supra* note 131, at 993–94; Bertrand & Mullainathan, *supra* note 64, at 991–93.

147. Angela Onwuachi-Willig & Mario L. Barnes, *By Any Other Name?: On Being “Regarded as” Black, and Why Title VII Should Apply Even If Lakisha and Jamal Are White*, 2005 WIS. L. REV. 1283, 1284 (2005).

148. Lu-in Wang, *Race as Proxy: Situational Racism and Self-Fulfilling Stereotypes*, 53 DEPAUL L. REV. 1013, 1017 (2004).

149. See Siegel, *supra* note 79, at 1300–03 (introducing the antibalkanization perspective as one in which courts are concerned with social cohesion and worry that some antisubordination-based strategies might generate further divisions in society).

Because antidiscrimination law focuses on sanctioning reproachable conduct in the *use* of information about protected categories, it can only operate *ex post facto*, after a discriminatory act has taken place. Therefore, it can prevent discrimination only to the extent that decision-makers factor the expected costs of sanctions. This incentive-setting is less effective when the reproachable conduct is based on unconscious biases.

Because discrimination driven by unconscious biases can be rooted in a decision-maker's "common sense" through heuristics built on generalizations that could even have some statistical value, the decision-maker might not always know that she is employing them. Then, sanctioning the behavior *ex post* will have little effect in her *ex ante* incentives. Regulating information in this scenario is particularly important as a preventive antidiscriminatory effort.

Finally, given the *ex-ante* nature of privacy-based antidiscrimination rules, they also avoid *ex post* harm better than traditional antidiscrimination efforts.¹⁵⁰ First, they avoid the harm of discriminatory acts that, even if addressed *ex post facto* by antidiscrimination law, are difficult to fully compensate.¹⁵¹ Second, because they do not rely on deterrence, they can prevent discrimination produced by unconscious biases. Third, due to their preventive nature, they avoid the worries of antibalkanization to which disparate impact measures are sometimes subjected.¹⁵²

VII. CONCLUSION

There are different ways to fight discrimination alongside claims under legislative frameworks that are operative *ex post*, such as Title VII claims. There is, for example, affirmative action in education and quotas for board membership. By exploring information flows that are common in privacy norms and antidiscrimination norms, I develop another way to fight discrimination: by using privacy rules.

Privacy rules can offer short-term protection from discrimination by blocking information that the law deems harmful in a decision-making process. When decision-makers' samples are skewed, when they have processing errors such as behavioral biases, or when discrimination is intentional, blocking information might be more effective at preventing discrimination than allowing it. In this way, privacy rules can aid antidiscriminatory efforts.

The crucial task for preventing discrimination through privacy rules is identifying and blocking data points that are believed to be proxies for categories that the law wants to protect. That is, for antidiscriminatory

150. See Roberts, *supra* note 28, at 2102 ("Antidiscrimination law prohibits discriminatory actions by outlawing certain types of conduct, but privacy law renders the offensive conduct practically impossible by impeding access to the information necessary for the unfavorable differentiation."); Austin, *supra* note 29.

151. See *id.* at 2155-56.

152. See Siegel, *supra* note 79, at 1299, 1302-03 (arguing that courts concerned with antibalkanization focus on diversity more than on equality and that some antisubordination-based strategies might generate further divisions in society).

privacy rules to be effective, they must block information that can be used as proxies that shift discrimination to groups other than those protected by the regulation. They must also identify proxies that reduce or expand the group of people that are discriminated against and gauge those proxies' usefulness in helping protected categories.

This preventive method is useful to expand protection from discrimination. First, it widens the scope of antidiscrimination protection for those cases in which disparate impact is not recognized. Second, it is useful in those cases where disparate impact is recognized but there are probatory difficulties: (a) intent is required to establish discrimination but it is difficult to prove it, or (b) disparate impact is admitted as an evidentiary standard but it is also difficult to prove (such as housing). Third, this preventive method can attend to worries about antibalkanization.

At a policy level, by virtue of operating *ex ante*, it avoids the social harms created by discriminatory conduct—a problem that *ex post* compensation only partially solves. Moreover, it can address unconscious biases better than *ex post* antidiscriminatory efforts.

In sum, the framework introduced here helps to determine the effectiveness of antidiscrimination measures based on information restrictions, and it explains how to design privacy rules that prevent discrimination effectively.

VIII. APPENDIX: IDENTIFYING PROXIES THROUGH CROSS ELASTICITY

I mentioned before that the objective relationship between information points X and Y does not determine the extent to which someone is likely to use X instead of Y to judge someone else.¹⁵³ Their relationship in that person's mind determines it, and her perception might or might not trace an objective relationship. Therefore, the question of whether a data point qualifies as a proxy for another data point for our purposes is a question of cross elasticity.

Cross elasticity evaluates the relationship between two goods as independent, complementary, or substitute depending on how the purchased quantity of one of them changes upon a price change in the other. For example, imagine Abby and Ben have a daily budget to buy a combination of apples and bananas. Upon an increase in the price of bananas, Abby moves to buy only apples, and Ben continues to buy the same proportion of both in smaller quantities. While Abby sees these fruits as substitutes, Ben views them as independent goods.

Let A be Abby, B be a company called Boss Inc., X the unavailable information that Boss would like, and Y the information that Boss could use instead. Let c^X be the cost for Boss of acquiring information X about Abby, and c^Y the cost for it to acquire information Y about her. Let q be quantity. Let c_1 and q_1 refer to the cost and quantity of information at time 1, and c_2 and q_2 refer to the cost and quantity at time 2. When the cost of acquiring X increases, the proxy relationship between information X and information Y is given by the change in the quantity of Y that Boss attempts to acquire divided by the change in cost of the now costlier piece of information X .¹⁵⁴

$$E = \frac{c_1^X + c_2^X}{q_1^Y + q_2^Y} * \frac{\Delta q^Y}{\Delta c^X}$$

If the elasticity between the data points is negative, this denotes that X and Y are complementary pieces of information, while if it is positive, it denotes that they are substitutes.¹⁵⁵ A piece of information is a proxy for another for Boss if it treats them as substitutes.

An example might clarify this. Imagine Abby applies for jobs with four companies, called Boss, Chief, Director, and Employer. Imagine these companies only care about productivity, and they want to know whether Abby is likely to ask for a maternity leave during the duration of the employment (X). Imagine state laws protect Abby from having them ask her this directly.

153. See *supra* Section IV.B.

154. As with any other measure of elasticity, elasticity cannot be measured when there is no change in cost.

155. Formally, $E \in (-1,1)$

Before the law's implementation, the companies could just ask Abby this directly, which would take them about two minutes, assuming one minute for their question and one minute for her answer. Assume each minute to cost them one imaginary utility point. After the law is implemented, they need to ask five indirect questions instead of one to get that same information, taking them five times as long during their brief interview and, therefore, taking them ten minutes instead. So, state law increased their cost of information acquisition X by 500%.¹⁵⁶ They can now choose between incurring this increased cost and spending the amount of time for one question (the original cost of two minutes) by asking a different question: whether Abby is engaged or recently married (Y). Imagine Boss still asks Abby about her potential maternity leave through five indirect questions, and that Chief, Director, and Employer only inquire about marital status instead.

The price of X changed from two to ten, and the quantity of X demanded changed from four to one. Therefore, the cross elasticity of X and Y is

$$E = \frac{2 + 10}{4 + 1} * \frac{3}{8} = 0.9$$

In this example, the companies treated marital status (on average) very much as a proxy for the applicant's likelihood of eventually asking for a maternity leave. An elasticity of one would have meant that they treated them as perfect substitutes, and an elasticity of zero would have meant that they treated them as completely independent. Given the elasticity of 0.9, if state law really wants to protect Abby, it should prohibit companies from asking her about her marital status as well.

Alternatively, imagine that Boss hires a human resources company to interview job candidates. The company offers two services: since it is forbidden to ask the candidates whether they want to have children, a basic service asks them whether they recently married, while a more complete service estimates their likelihood of having a child directly through an algorithm with a series of estimators with almost complete accuracy. How much is Boss willing to pay for the second service compared to the first? If it is not willing to pay much for the difference, we can infer that it considers the data points to be substitutes.

156. The cost of X will be determined by its Shannon entropy $H(p)$, where p is the probability distribution of X and $H(p) \in (0,1)$. A variable's entropy is a measure of its uncertainty. That is, of how much information, on average, one needs to know its value. X has a higher entropy than Y when we need more yes/no questions to know its value. See THOMAS M. COVER & JOY A. THOMAS, *ELEMENTS OF INFORMATION THEORY* 12–15 (2012); see generally Claude Shannon, *A Mathematical Theory of Communication*, 27 *BELL SYST. TECH. J.* 379, 10–12 (1948). Since, for these purposes, we care about elasticity, the relative entropy of X and Y , which is the distance between their probability distributions, is relevant as it will determine cost differences. The relative entropy of X and Y , $D = (p \parallel z)$, will describe the inefficiency of describing X (with a probability mass function p) through Y (with a probability mass function z). See COVER & THOMAS, *supra*, at 18–21.

The relationship between these data points can be identified in a similar way when the law makes X not more difficult but impossible to acquire. In this situation, costs and prices are absent. After the state law is enacted, Boss, Chief, Director, and Employer are unable to incur a higher cost to acquire X as Boss did in the previous example by asking Abby about a potential maternity leave through a larger set of questions. In this new scenario, imagine that, when questions about future pregnancy were allowed, 75% of employers did so (Boss, Chief, and Director), and 25% did not (Employer). When the law forbids it, 50% ask about marital status (Boss and Chief), and 50% do not (Director and Employer). The degree of substitution of the data points is given by:

$$E = \frac{\Delta Q^Y}{\Delta Q^X}$$

In the example,

$$E = \left| \frac{0 - 50}{75 - 0} \right| = 0.66$$

As it can be seen, determining whether a data point is a proxy for another in terms of antidiscrimination is an empirical question; it is not one of objective correlation, but one of subjective perception, measurable through cross elasticity. A sound empirical strategy to interrogate whether a concrete data point functions as a proxy, therefore, would not use quantitative analysis to measure predictiveness but a field experiment or survey to measure perceptions, and evaluate cross elasticities based on those.

Of course, we are not always going to be able to test the cross elasticity of information. But even when we cannot, this framework is useful because it points to what questions we should ask.

We should not worry about objective correlations, but about decision-maker's subjective perceptions of them; therefore, it is valuable to have empirical research that explores these perceptions. Moreover, on occasions we might be able to build heuristic approaches. Even if we ignore the degree to which a decision-maker sees two information points as replaceable, we might know from anecdotal evidence whether the average decision-maker in the field views them as completely replaceable, fairly replaceable, or not replaceable at all.