

Anti-discrimination laws are a familiar feature of our legal landscape. In the private sector, they operate in certain contexts—for instance, the provision of goods and services, rental housing, and employment—to prohibit us from singling out individuals for less favorable treatment because of certain traits, such as their race, age, gender, or religion.¹ They also sometimes prohibit us from adopting rules or practices that do not explicitly exclude anyone on the basis of a prohibited ground but nevertheless have the effect of disproportionately disadvantaging some groups on this basis, such as aerobic capacity tests that exclude women at a higher rate than men.² And some anti-discrimination laws, such as

I am very grateful for the many helpful comments I have received from the Editors of *Philosophy & Public Affairs*, and from audiences at the Gould School of Law, the McGill Legal Theory Workshop, the NYU Colloquium in Legal, Political and Social Philosophy, the UCLA Legal Theory Workshop, and the University of Toronto Faculty of Law.

1. Some central examples of anti-discrimination legislation in the United States are Title VII of the Civil Rights Act of 1964, 42 U.S.C. §§2000E (2003), the Age Discrimination in Employment Act of 1967, 29 U.S.C. 621–34, and the Fair Housing Act and Fair Housing Amendments Act, 42 U.S.C. §§ 3601–19. In Canada, each province has its own anti-discrimination legislation covering all of the above contexts: see, for a representative example, the Ontario Human Rights Code, R.S.O. 1990, c.H.19. The United Kingdom has recently introduced the Equality Bill (April 24, 2009), a single piece of legislation aimed at harmonizing anti-discrimination laws in the United Kingdom. When passed, it will replace the United Kingdom's current subject-specific anti-discrimination laws, such as the Sex Discrimination Act of 1975 (c.65), the Race Relations Act of 1976 (c.74), and the Disability Discrimination Act of 1995 (c.50).

2. All of the legislation cited above in note 1 also prohibits this second form of discrimination. This form of discrimination does not necessarily involve an exclusionary motive; but the countries cited in note 1 differ somewhat in the ways in which they draw the distinction between it and the kind of discrimination that does. In the United States, the distinction is between “disparate treatment,” which necessarily involves an exclusionary motive, and “disparate impact,” which does not, and can be proven simply by adducing statistical evidence that the rule or practice disproportionately burdens the claimant's

the Americans with Disabilities Act, require us not just to avoid disadvantaging people on the basis of a prohibited ground of discrimination but also to provide the facilities and services they need.³ Such laws are sometimes called *accommodation requirements* because they seem to require us to provide special accommodations to members of certain groups, beyond what is normally required of us by anti-discrimination laws.⁴

Much of the recent theoretical discussion of anti-discrimination laws has focused on the broader political goals that these laws serve. Because they combat the systemic subordination and stigmatization of groups identified by the prohibited grounds of discrimination,

group. In Canada and the United Kingdom, the distinction is between “direct discrimination,” which may or may not involve an exclusionary motive but must involve a rule that explicitly draws a prohibited distinction, and “adverse effect discrimination” or “indirect discrimination,” which involves facially neutral rules or practices that have the effect of disadvantaging some individuals based on a prohibited ground. My argument does not depend on our drawing the distinction in one of these ways rather than another. Indeed, as I shall argue later in the article, my account of discrimination implies that these distinctions are not deep. It is largely irrelevant from the standpoint of justice whether discrimination has been intentional or non-intentional, direct or indirect; although, as I shall explain, the presence of a discriminatory motive may have bearing on whether the claimant is owed additional damages for mental suffering.

3. Americans with Disabilities Act of 1990, as amended by the ADA Amendments Act of 2008, 42 U.S.C. c.126, 47 U.S.C. c.5. See also the Family and Medical Leave Act of 1993, 29 U.S.C. §§ 2601–54 (1994 & Supp. III 1997). In the United Kingdom, the Disability Discrimination Act of 1995 imposed similar accommodation requirements, note 1 above. All of the provincial human rights codes in Canada cited in note 1 impose accommodation requirements with respect to all of the grounds of discrimination. Indeed, the Canadian Supreme Court has defined discrimination in the private sector as the failure to provide reasonable accommodation to those disadvantaged by one’s policies on the basis of a prohibited ground of discrimination: see *British Columbia (Public Service Employee Relations Commission) v. BCGSEU*, [1999] 3 S.C.R. 3 (the “Meiorin” case).

4. Those who claim that accommodation requirements impose more onerous obligations than the rest of anti-discrimination law include Samuel Issacharoff and Justin Nelson, “Discrimination with a Difference: Can Employment Discrimination Law Accommodate the Americans with Disabilities Act?” *North Carolina Law Review* 79 (2001): 307–58; Pamela Karlan and George Rutherglen, “Disabilities, Discrimination and Reasonable Accommodation,” *Duke Law Journal* 46 (1996): 1–41; Mark Kelman, “Market Discrimination and Groups,” *Stanford Law Review* 53 (2001): 833–96; and J. H. Verkerke, “Disaggregating Antidiscrimination and Accommodation,” *William and Mary Law Review* 44 (2003): 1385–1419. Some prominent legal scholars have contested this view: see the articles cited in note 7. The account of discrimination that I shall defend in this article implies that accommodation requirements are not normatively different from the rest of anti-discrimination law. I discuss this point in more detail in Section IV of the article.

anti-discrimination laws have been seen as a tool for effecting distributive justice between these groups and those that are more privileged.⁵ This raises questions both about the scope of the relevant redistributive principles and about whether private parties can fairly be held responsible for implementing them.⁶ Recent discussion has also centered on whether accommodation requirements, which force employers and businesspersons to take costly steps to accommodate the special needs of certain groups, are deeply different from the rest of anti-discrimination law—either empirically, in terms of their actual effects on employers and businesspersons, or normatively, in terms of the redistributive principles that justify them.⁷

My focus in this article is not on the redistributive goals that anti-discrimination laws may serve. Instead, my aim is to elaborate an account of what discrimination involves and why it is unjust that takes seriously a common, but not often discussed, feature of anti-discrimination laws. Anti-discrimination laws are commonly structured

5. For instance, Owen Fiss has argued that anti-discrimination law is based upon a “group-disadvantaging principle,” which prohibits action that aggravates or perpetuates the already subordinate status of certain groups, such as blacks and women: “Groups and the Equal Protection Clause,” *Philosophy & Public Affairs* 5 (1976): 107–77. And Cass Sunstein has defended the view that anti-discrimination law expresses an “anticaste principle,” according to which we must not disadvantage or stigmatize particular groups because of a visible but yet morally irrelevant group-based characteristic: “The Anticaste Principle,” *Michigan Law Review* 92 (1994): 2410–55. See also Robert Post, “Prejudicial Appearances: The Logic of American Anti-Discrimination Law,” *California Law Review* 88 (2000): 1–40.

6. John Gardner has argued that they can: “Liberals and Unlawful Discrimination,” *Oxford Journal of Legal Studies* 9 (1989): 1–22; “Discrimination as Injustice,” *Oxford Journal of Legal Studies* 16 (1996): 353–67. See also Denise Reaume, “Harm and Fault in Discrimination Law: The Transition from Intentional to Adverse Effect Discrimination,” *Theoretical Inquiries in Law* 2 (2001): 349–85.

7. Interestingly, this is a largely U.S. debate: in Canada, as I shall explain later in the article, all of anti-discrimination law is conceptualized in terms of a “duty to accommodate.” In the United States, those who have argued that “accommodation requirements” do not impose obligations of any different or more onerous kind than are already imposed on us by prohibitions on indirect discrimination or disparate impact include Christine Jolls, “Accommodation Mandates,” *Stanford Law Review* 53 (2000): 223–306; and “Anti-Discrimination and Accommodation,” *Harvard Law Review* 115 (2001): 642–99; Samuel R. Bagenstos, “‘Rational Discrimination’, Accommodation and the Politics of (Disability) Civil Rights,” *Virginia Law Review* 89 (2003): 825–923; and Sharon Rabin-Margalioth, “Anti-Discrimination, Accommodation and Universal Mandates—Aren’t They All the Same?” *Berkeley Journal of Employment and Labor Law* 24 (2003): 111–52. For those who have argued the contrary, see the authors cited in note 4.

in such a way as to suggest that discriminators have committed a personal wrong against their victims, akin to a tort.⁸ For they usually employ an individual complaints procedure: that is, they rely on individual claimants to instigate legal proceedings against alleged discriminators, and the process is conducted throughout as an investigation into whether the claimant has been treated in a way that amounts to her personally having been wronged by the alleged discriminator. Carriage of the complainant's case is sometimes given to a legal body that is acting in the wider public interest, such as a human rights commission. But it is still the case that this body's primary role in the proceedings is to bring forward and resolve what is viewed as a personal dispute between the complainant and the alleged discriminator. Moreover, although the claimant's case often has implications for the treatment of a particular group and the claimant must, in cases such as U.S. disparate impact cases, prove her allegations by adducing evidence of how the group that shares her trait was treated relative to other groups, her aim is still to show that she personally has been wronged. The available remedies also suggest that the wrong in question is usually a personal one. For in addition to requiring organizations to change their general rules and practices, most anti-discrimination laws require them to provide specific forms of restitution to the claimant, such as reinstatement, various forms of accommodation, and monetary compensation. Some jurisdictions even offer special damage awards to the claimant, for "injury to dignity, feelings, and self-respect."⁹ The surface structure of anti-discrimination law suggests, then, that the wrong is a personal one, akin to a tort: the discriminator has interfered with some protected interest of the

8. I do not here take up the suggestion that there should actually be a tort of discrimination, because I take it that this question raises complex issues of institutional design that are not relevant to the more basic philosophical question of what discrimination consists in. It may be, for instance, that given that we already have human rights legislation that prohibits private sector discrimination, it would be confusing and even unjust to permit claimants also to bring tort actions for discrimination. See, for such an argument, the Canadian case of *Board of Governors of Seneca College of Applied Arts and Technology v. Bhaduria* (1981) 124 D.L.R. (3d) 193 (S.C.C.). For discussion of whether there should be a tort of discrimination, see Amnon Reichman, "Professional Status and the Freedom to Contract: Toward a Common Law Duty of Non-Discrimination," *Canadian Journal of Law and Jurisprudence* 14 (2001): 79–132; and the Honorable Louis LeBel, "La protection des droits fondamentaux et la responsabilité civile," *McGill Law Journal* 49 (2004): 231–54.

9. For instance, the Ontario Human Rights Code allows for "compensation for injury to dignity, feelings and self-respect," R.S.O. 1990, c.H.19, ss.46(1)1.

victim's and the victim now has a personal claim on him to rectify his wrongdoing.

One difference between discrimination and most legally recognized torts is that in the case of torts, we can easily identify the interests whose injury amounts to a personal wrong toward the plaintiff. For instance, the tort of negligence protects our interests in bodily integrity and property, while the tort of defamation protects our reputation. It is more difficult to locate the interest that is injured by discrimination. Perhaps this is one reason why scholars have been more inclined to see anti-discrimination laws as redistributive policies aimed at bettering the situation of these groups as a whole, rather than as restitution to particular individuals for personal wrongs that they have suffered. I shall argue that the interest that is injured by discrimination is our interest in a set of what I call *deliberative freedoms*: that is, freedoms to have our decisions about how to live insulated from the effects of normatively extraneous features of us, such as our skin color or gender.

The core idea underlying my account is this. In a liberal society, each person is entitled to decide for herself what she values and how she is going to live in light of these values. This means that, in addition to certain freedoms of action, we are each entitled to a set of "deliberative freedoms," freedoms to deliberate about and decide how to live in a way that is insulated from pressures stemming from extraneous traits of ours. Many of us already have these deliberative freedoms. I shall argue that anti-discrimination law attempts to give them to all of us, because each of us has an independent entitlement to them. It attempts to give them to us by preventing our employers, service providers, landlords, and others from acting in ways that deny us opportunities because of these traits, so that when we deliberate about such things as where to work and where to live, we do not have to think about these traits as costs. I shall try to show that, by understanding anti-discrimination law in this way, we can make sense of its surface structure as correct. That is, we can understand why discrimination amounts to a personal wrong against the victim. It is a personal wrong because the discriminator has interfered with the victim's right to a certain set of deliberative freedoms.

I turn now to the task of developing this account of discrimination more fully. I shall do this in four sections. Section I introduces the idea of a deliberative freedom and argues that such freedoms matter to us. Section II explains how anti-discrimination laws can be seen as

protecting each person's entitlement to certain deliberative freedoms, and it discusses both the role of grounds of discrimination on my account and how we are to determine when a given instance of discrimination is "based on" a certain ground. In Section III, I consider the limits of anti-discrimination law. Anti-discrimination law only protects our deliberative freedoms in certain social contexts, such as employment and the provision of goods and services. Even in these contexts, it sometimes allows what would otherwise be discriminatory practices to persist if they are required by 'business necessity' or if providing the necessary accommodations would impose 'undue hardship' on the alleged discriminator. I argue that these limits can be explained on my account as attempts to balance the deliberative freedoms of claimants against the interests of other people, and I give some examples of the other interests that may be at stake. Finally, in Section IV, I consider and respond to a number of objections.

I. DELIBERATIVE FREEDOMS: WHAT THEY ARE AND WHY THEY MATTER TO US

Liberal political theorists agree that each member of a liberal society is entitled to decide for herself what she values and what she wishes to spend her life pursuing. To say this is not to claim that each of us is entitled to unlimited freedom in our deliberations and decisions about how to live. Our decisions are always made within the framework of certain constraints, such as the cost of a certain activity and the resources at our disposal, the needs of our children, and the expectations of those whom we trust. Some of these we think of as constraints whose costs we can rightfully be asked to bear. For instance, we do not think that anyone has a right to be able to decide where to live in a manner that is free from worries about the cost of the apartment and the level of their current income, or that anyone has a right to be able to decide where to work without having to factor in their lack of certain skills. Within the confines of such legitimate constraints, however, we do think that each member of a liberal society has a right to decide for herself how she will live her life. This means that she is entitled to be free not just from certain interferences with her actions but also from pressures on her own deliberations caused by certain kinds of extraneous facts about her—for instance, her sexual preferences and other people's assumptions about how people with these preferences behave, or her back problems, which

would not affect her ability to work at all were it not for her employer's requirement that all cashiers work standing up.

By using the term "deliberative freedoms," I do not mean to suggest that they are just freedoms to think certain things, irrespective of what options are actually available to us or irrespective of what we actually do. On the contrary, they are freedoms to engage in deliberative activities and make decisions in a certain way, a way that is insulated from the pressures or burdens caused by certain extraneous traits. So, insofar as making a decision is doing something, they amount to a type of freedom of action. They also require, as a precondition, that we do actually have available to us the options that we think we have: in order for me to have a particular deliberative freedom with respect to a certain decision, it has to be true not just that I believe I can make that decision without having to worry about pressures from a certain extraneous trait, but that I really am free from those pressures.

In describing deliberative freedoms as freedoms to make decisions about how to live in a way that is "insulated from the pressures or burdens caused by certain extraneous traits of ours," I am using the term "extraneous" in a normative rather than a descriptive sense. I mean that there are certain traits that we believe people should not have to factor into their deliberations, or, more exactly, should not have to factor in *as costs*, even if they are deeply important to the person and relevant to her decisions. For example, to say that a person's race is normatively extraneous to her deliberations is not to ignore the fact that some people's race is extremely important to them and that they care deeply about living near and working with people of the same race. Rather, it is to say that people should not be constrained by the social costs of being of one race rather than another when they deliberate about such questions as what job to take or where to live. So even if someone does wish to consider her race as a factor relevant to where she wants to work, she should not have to think of it as a trait that makes her less able to fulfill the job requirements or less attractive in her employer's or her clients' eyes. Religion is another trait that often seems to us to be normatively extraneous to a person's decisions about where to live and to work. Again, it is not extraneous in the sense that it is factually irrelevant. The religious person may find himself unable and unwilling to make important decisions without considering aspects of his religion. He may, for instance, rule out certain places of work and certain neighborhoods on

the grounds that he would encounter too many people who are dressed immodestly. Yet we still believe his religion to be extraneous to these deliberations in the normative sense that the social costs of practicing it are not costs that we think should have to figure into his deliberations. The private costs of practicing his religion are of course still his to bear: he needs to pay, for instance, for his religious garments and for the cost of his pilgrimage. But beyond the costs to him of doing the things his religion requires him to do, we think he should not have to treat his religion as a burden. In particular, we think he should not have to consider the broader social costs of having that trait, such as the cost to his employer of accommodating his religious practices, or the cost to his landlord of lost business resulting from other tenants' prejudiced attitudes.

I shall say more in Section II about the kinds of traits that seem to be extraneous in this sense, and why they seem to be so. For now, I simply want to note that, as the example of religion suggests, my account does not depend on the assumption that it is only appropriate to relieve someone of the burdens of having a certain trait if he did not choose to have that trait. That is, it is not the case that a trait must be unchosen or beyond our immediate control in order to be "normatively extraneous" in my sense: some normatively extraneous traits, such as race and age, may be unchosen, but others, like religion, are chosen, and still others, like disability, are sometimes the results of our past choices.

I have said that, as members of a liberal society, we value having certain deliberative freedoms; indeed, we believe each of us is entitled to some deliberative freedoms. These freedoms are so much a part of the fabric of most of our everyday lives that we rarely reflect on how fundamental they are. Most of us are, for instance, free to think about which job to accept without having to worry about whether we will be able to practice our religion while working the required shifts. We are free to think about where to work and where to live without having to worry that public transit will be inaccessible to us; free to decide where to eat or shop without having to wonder whether someone might ask us to leave and breastfeed "somewhere more discreet"; and free not to have to think at all about which gender's washroom to use because we fit into one of the conventional gender categories that are used to label public washrooms. We enjoy these freedoms and can afford to take them for granted, however, only because businesses and associations

tend to arrange their affairs and their physical premises in a way that tacitly accommodates the needs and abilities of their usual employees or clients. For instance, fire departments and police forces have historically used admission tests based on the strength and stamina of active men, rather than active women. This went unnoticed during the years in which it was unthinkable for a woman to enter these professions. Similarly, most of our public spaces—parks and monuments, government buildings, public transit, and even public washrooms—have historically been designed in such a way as to accommodate the needs of people who are able-bodied and who fit into conventional gender categories. We have only recently begun to understand the many ways in which these spaces and institutions privilege some people's needs at the expense of others. There has, then, been a constant series of accommodations given to many of us, accommodations that give us deliberative freedoms.¹⁰ Most of the time, these accommodations remain invisible. They come to our attention only when a claimant challenges a rule or practice as discriminatory, as failing to include or accommodate her because of her possession of some extraneous trait.

When a claimant challenges a rule or practice as failing to accommodate her, what she cares about is not just receiving the particular benefit that others have—for instance, a career in the fire department or a trip on public transit—but having the freedom to decide for herself whether to pursue this career or whether and when to use public transit. We care about being able to make these choices ourselves, without pressure from certain extraneous considerations. Since our concern for deliberative freedoms is rooted in a concern for how our own lives are shaped—that is, we care about having deliberative freedoms because we care about being able to shape our lives in our own way—what matters to us is not that we have the same amount of freedom as others happen to have, but that each of us has the amount to which he is entitled. So although it may be true that each of us is equally entitled to the same freedoms, this is not because each of us is

10. As should be clear, this paragraph uses the term “accommodation” simply to refer to the fit between certain rules and practices and the needs or abilities of certain groups, and not to refer only to the subset of anti-discrimination laws that in the United States are known as “accommodation requirements.”

entitled to as much as others happen to have. Rather, it is because each of us has an equal independent entitlement to these freedoms.

In this article I shall not offer a deeper defense of our belief that certain deliberative freedoms are of value. I shall simply assume that there are some deliberative freedoms that we care about and to which we believe each person is entitled. My aim is to argue that we can develop a plausible account of anti-discrimination laws if we see them as a way of guaranteeing everyone a certain set of deliberative freedoms. And I shall try to show that this account of anti-discrimination laws can explain why discrimination is a personal wrong toward the victim.

Before I turn to developing this account, however, I want to clarify two features of deliberative freedoms that might still seem puzzling. First, although I have spoken mostly of deliberations that involve important decisions in our lives (for instance, where to live and where to work), the deliberative freedoms that we value are not limited to those that meet some fixed criterion of importance. On the contrary, because a liberal society leaves it up to each of us to determine which decisions are important and which are not, our entitlement to be free from the pressures of normatively extraneous traits can extend also to decisions that may seem trivial to some of us, such as which route to take to work today, which restaurant to go to tonight, or whether to shop at the candy store tomorrow. Of course, anti-discrimination law does place some limits on the social contexts in which deliberative freedom is protected. For instance, it protects our deliberative freedoms in the context of employment and in the context of the sale of goods and services, but not in the context of familial life or close friendships. I shall discuss these contextual limits in greater detail at a later stage of the article. For the moment, the relevant point is that we do not need to suppose that our entitlements to deliberative freedom are limited to important decisions. On the contrary, part of what these entitlements protect are our opportunities to decide for ourselves what will be important and what will be unimportant in our lives. This is why it is fitting that they hold even in apparently trivial contexts like the sale of goods and services.

The second and final feature of deliberative freedoms that I want to highlight is this. Although I have spoken of the normatively extraneous traits as “traits of ours,” there is debate about whether these traits inhere in the claimant or whether they are in one way or another a social or cultural construct. For instance, many disability rights advocates have

argued that a physical disability is not an inherent limitation in the person: it is just a physical feature of his that differs from those features that other people possess, and it is a limitation only because we have constructed our buildings and transit systems, our workdays, and our expectations in such a way as to favor those with certain bodies rather than others.¹¹ It is beyond the scope of this article to engage deeply with these important and plausible claims. But my arguments are consistent with them. So I shall simply note here that I do not mean to prejudice this issue and that when I speak of the traits in question as though they belong to the claimant, this is only for ease of writing and does not reflect any assumptions about the nature or origins of these traits or the burdens that they often give rise to.

II. HOW ANTI-DISCRIMINATION LAWS PROTECT DELIBERATIVE FREEDOMS

I shall now try to show that we can see anti-discrimination laws as a way of granting to each person the deliberative freedoms to which we believe they are entitled. I shall look first at the way in which anti-discrimination law attempts to carve out certain deliberative freedoms for claimants, and I shall consider some of the commonly recognized grounds of discrimination and the kinds of traits they protect.

Anti-discrimination laws in the United States and the United Kingdom are commonly divided into three types.¹² I shall canvass these only briefly here, so that I can explain how all of them work to give us certain deliberative freedoms. At a later point in the article, I shall explain why, on my account, the differences that are sometimes seen between these three types of anti-discrimination laws are not as deep as is sometimes supposed. And I shall consider in greater detail what sorts

11. See the “social model” of disability first popularized by Michael Oliver in *The Politics of Disablement* (London: Macmillan, 1990). See also Michael Ashley Stein, “Disability Human Rights,” *California Law Review* 95 (2007): 75–121; Anita Silvers, “Formal Justice,” in *Disability, Difference, Discrimination: Perspectives on Justice in Bioethics and Public Policy*, ed. Anita Silvers, David Wasserman, and Mary B. Mahowald (Lanham, Md.: Rowman & Littlefield, 1998); and many of the essays in *Critical Disability Theory*, ed. Diane Pothier and Richard Devlin (Vancouver: University of British Columbia Press, 2006).

12. As mentioned previously, Canada prohibits these types of discrimination, but does not treat them as deeply different.

of revisionist implications the account has for particular countries' anti-discrimination laws. For the moment, I want to focus on some of the similarities and on the way in which all three types of law work to carve out certain areas of deliberative freedom for us. The three commonly recognized types of anti-discrimination laws are (i) prohibitions on what in the United States is known as 'disparate treatment' and in the United Kingdom, as 'direct discrimination'; (ii) prohibitions on what in the United States is known as 'disparate impact' and in the United Kingdom, as 'indirect discrimination'; and (iii) special accommodation requirements for individuals with certain traits, such as disabilities.¹³ Prohibitions on disparate treatment or direct discrimination are aimed at rules that explicitly single out some people for inferior treatment because they possess a certain trait, such as a rule stating that women need not apply for a certain job. Because the exclusion here is explicit and it is difficult for explicit exclusions not to be intentional, such prohibitions are often thought of as prohibitions on intentional discrimination. The second type of anti-discrimination law, which prohibits disparate impact or indirect discrimination, is aimed at rules that, though facially neutral, have a disproportionate impact on the group that possesses this trait and need not involve an intention to exclude. An example of such a rule is a company dress code requiring that all employees wear a particular decorative hat. The rule seems innocuous on its face. But if the hat cannot be worn at the same time as various forms of religious headwear, then the rule will place a greater burden on adherents of these religions than on others. The third group of anti-discrimination laws, the special accommodation requirements, demand that employers or others take special positive steps to provide for people with certain traits: for instance, that they equip the workplace with a braille printer or that they alter individuals' work schedules so that they are able to pray at the required times. All of these laws prohibit only those exclusions (or, in the case of the third group, those failures of accommodation) that are based on a specific list of recognized 'grounds' of discrimination. Race and gender are two examples of prohibited grounds that commonly appear in

13. For examples of legislation prohibiting the first two types of discrimination, see the anti-discrimination laws cited in note 1 above. For examples of the third type of anti-discrimination law, that is, accommodation requirements, see the laws cited in note 3 above.

anti-discrimination laws of the first two types; disability often appears in laws of the third type, as the basis for special accommodations.

How do such laws protect our deliberative freedoms? Anti-discrimination laws of the first two types do so by preventing our employers, goods and service providers, landlords, and others from directly or indirectly denying us opportunities because we possess the normatively extraneous traits that are marked out by the prohibited grounds of discrimination. They thereby free us from having to consider the costs of these traits in our decisions about where to work, what to buy, and how to live. Prohibitions on disparate treatment ensure that these extraneous traits will not explicitly be used to *single us out* for unfavorable treatment. They therefore insulate our deliberations from the costs of these traits: we can decide where to work and where to live without having to think about the low opinion others may have of our race or our gender. Prohibitions on disparate impact ensure that we will not be disadvantaged *even indirectly* because we possess these traits, by rules that are not designed to exclude us but nevertheless tend to exclude people with this characteristic more than others. And accommodation requirements ensure that basic features of institutional design, such as scheduling or physical amenities, will be modified so that we are not disadvantaged by possessing these traits. Hence, they too work to insulate our deliberations and decisions from the costs that would otherwise be associated with possessing these traits.

On this account of discrimination, it is the function of the prohibited grounds of discrimination to tell us which traits we are entitled to be free from considering when we make decisions about how to live (or, more exactly, when we make decisions in the particular social contexts to which anti-discrimination laws apply, such as employment and goods and services). So the grounds of discrimination mark out a list of what I have called *normatively extraneous traits*, traits whose costs we ought not to have to factor into our decisions in these particular contexts. If we think of the grounds that are commonly listed in the anti-discrimination laws of liberal societies—for instance, race, sex, age, and religion—it does seem plausible to suppose that they express a judgment as to which traits it is whose costs we ought not to have to factor into our deliberations in such contexts. Consider religion, for instance. On the account that I am offering, to list religion as a prohibited ground is to say that people should not have to think about their religion as a liability when

deciding whether to accept job offers, or whether to live in a certain apartment, or even whether to participate in a trivial economic transaction; and they should not have to choose between their religion and a given job or apartment or economic transaction. Most of us do think this: we think religion is such an important part of the life of religious persons that they should not have to compromise it in order to have the opportunities that they would have had in these areas, but for their religion. The same is true of sex, but for rather different reasons. Like religion, a person's sex is often a trait that they regard as fundamental to their identity. But it is fundamental in the rather different sense that it determines certain aspects of their appearance and certain of their physical capacities, and it does so regardless of whether they wish it to or not, and regardless of whether they would choose to be of that gender. And it seems important to protect us from the costs of being of one sex rather than another in part because this trait is unchosen.

As these examples suggest, the particular types of traits that we recognize as normatively extraneous to people's deliberations differ markedly. Some, like religion, are chosen commitments. Others, like sex, are an undeniable part of who we are but are unchosen. Still others, such as disability, sometimes express a current commitment of ours (think of the community of those who are "culturally deaf"); sometimes result from past choices (liver disease resulting from alcoholism); and sometimes are completely unchosen and regarded as external to the bearer (such as fibromyalgia). What matters for the purposes of anti-discrimination law, it seems, is not whether the trait is chosen or unchosen, an important part of our life, or something we would rather be rid of. It is that, whatever the nature of the trait, the burdens that result from it are not ones that we believe its possessors should have to factor into their deliberations in these particular social contexts.

If this is a correct analysis of the function of grounds, then it helps to explain why so many human rights tribunals and courts have faltered in attempting to locate a single criterion for something counting as a prohibited ground. On the view I have been suggesting, whether some trait should be recognized as a prohibited ground is a normative question whose answer depends on whether people have a right to make decisions in a manner that is free from the sorts of institutional and attitudinal pressures that are encountered by those with that trait. We must ask: are people entitled to be insulated from these pressures when they

make decisions in these social contexts? There is no reason to suppose that we will answer this normative question by appealing to the same types of reasons in the case of all potential grounds. As I suggested above, if we are asking whether individuals are entitled to deliberate in a manner that is free from the burdens of being excluded from certain jobs or places because of their religion, our answer will likely involve some kind of appeal to how central religious activities are to our conceptions of ourselves and our ability to shape our own lives in our own way. But if we are instead considering the ground of sex, our answer to the normative question will be very different. So on this account of discrimination it seems unlikely that there could be a single principled explanation of what makes each ground of discrimination worthy of the status of a ground. All that we can say is that each ground reflects a judgment that people have a right to make decisions in these social contexts in a manner that is insulated from the burdens imposed by these traits, or by other people's assumptions about them. The normative facts about these traits that form the basis for such a judgment will be diverse.

One might object, at this point, that unless something more general can be said of what ties together the different traits that underlie the prohibited grounds, we will have no reason to think that the victims of discrimination have been *wronged*. For we will have no single principled explanation of that wrong. It is certainly true that this account does not offer a single reductive explanation of the wrong of discrimination—that is, an explanation that traces the wrong of discrimination to some further single kind of normative fact that is operative in all cases. But it does not follow that the account fails to provide us with a principled explanation. Explanations can be principled without appealing only to one further kind of reason. On the account of discrimination that I am proposing, victims of discrimination have been denied an equal set of deliberative freedoms. Which deliberative freedoms we are entitled to depends on which traits should be recognized as normatively extraneous. And the fact that we cannot answer this question by appealing to some single further kind of normative fact in all cases is not a problem. Rather, it reflects the complex nature of the type of injustice that we are trying to explain.

It is worth noting at this point that my account of discrimination does not imply that the lists of grounds currently recognized by anti-discrimination laws in a country such as the United States are correct. I

have used the traits that we commonly recognize as my examples because we do need to start somewhere in elaborating an account of discrimination. And it seems reasonable to suppose that unless an account of discrimination fits *some* of our common practices, it will not be an account of discrimination. But this is not to say that the account must preserve all features of our existing anti-discrimination laws. Once we have arrived at a view of the function of grounds of discrimination that seems to explain some of the grounds on our lists, then we can go back and engage in the kinds of normative inquiry that I have been describing and ask whether each of the recognized grounds should really be recognized and whether there are others we have omitted. We might decide, for instance, that unattractive physical appearance—which is not now on the list of grounds unless it is a recognized illness and so falls into the category of ‘disability’—ought to be a ground of discrimination. Many people have argued that poverty, or at least “social condition,” should similarly be recognized as a ground—and indeed the Canadian province of Quebec’s *Charte des droits et libertés de la personne* does treat it as such.¹⁴ My account is quite open to these possibilities.

I want to turn now to an aspect of my account that is related to the role of the prohibited grounds. So far I have spoken loosely of individuals being excluded or disadvantaged “by” or “because of” or “based on” a trait that is normatively extraneous, or a prohibited ground of discrimination. This might seem troubling. When exactly is someone excluded or disadvantaged “because of” a particular trait? This question is relatively easy to answer in cases of disparate treatment or direct discrimination. Consider the landlord who does not want to lease an apartment to someone because he is of a particular culture, since she would lose the business of other tenants if they had to live near a person of this culture. Here, the “because” involves an appeal to the discriminator’s own reasons or motives. And although there may be difficulties of proof in such cases, it is at least clear what we are trying to prove. But in cases of disparate impact or indirect discrimination, it is much less clear when the “because” requirement has been satisfied. Suppose the operators of a public pool find that the pool is always overcrowded, and so they limit attendance to those who live in the neighborhood. It happens that the

14. R.S.Q. C-12, c. I.1, s.10.

neighborhood is an affluent white neighborhood, and hence that most of the people excluded from the pool are those who live farther away in the poorer black neighborhoods. But it was not the pool operator's intent to exclude blacks when he decided to limit pool attendance. He just wanted to limit the number of people coming to the pool. Have the black swimmers been excluded "because of" their race? How are we to decide this?¹⁵

This question is very similar to a standard question that arises in tort law, the question of whether a particular injury is too remote from the defendant's action to count as his responsibility. Both questions are normative questions, and they both involve judgments about when a particular result is close enough on the causal chain to a particular action to count as something that the agent is responsible for. In my view, we can only address such questions in discrimination law in the same way that we do in tort law—on a case-by-case basis, by asking whether the connection is close enough for the agent to be held responsible for the result. By "close," I do not mean temporally or spatially close, but rather the kind of thing that the agent can justifiably be held responsible for. In the above example, it will depend in part upon whether the character of the neighborhood as affluent and white is obvious to everyone and something the pool owners ought to have known about. (It likely is.) It will also matter exactly why blacks do not live in this affluent white neighborhood. (Is it just because they prefer to live elsewhere? In that case, their race would seem to have nothing to do with their exclusion from the pool. But that seems unlikely.) And how tight is the connection between their exclusion and their race? That is, how many other causal factors intervene, and do we think the operators of the pool can still be held to have wronged these individuals even though the causal chain is so long, and at times, so nebulous? (Probably so, even if the exclusion was caused by poverty, which was in turn caused by lack of education and opportunities, which were in turn only partially the result of other people's assumptions about these individuals' race.) Such difficulties in determining when a disadvantage or an exclusion arises "because of" a particular trait will be faced by any account of discrimination, or at least any account that embraces forms of discrimination other than direct discrimination. They are certainly worth further exploration. But they

15. I am grateful to Seana Shiffrin for pressing this objection.

are not any more problematic for my account than for any other account of discrimination.

III. EXPLAINING THE LIMITS OF ANTI-DISCRIMINATION LAW

I have tried to suggest that the aim of anti-discrimination laws can be seen as giving all individuals certain deliberative freedoms, and giving them these freedoms to an extent roughly equal to those of others. But anti-discrimination law does not protect our deliberative freedoms in all contexts, or at all times. For it applies only to our transactions in certain social contexts, for instance, employment and the provision of goods and services. And even within these contexts, there are defenses available to the alleged discriminator that result, in some cases, in the claimant being denied a deliberative freedom to which she would otherwise have been entitled. In the remainder of this section, I shall try to show that my account of discrimination is consistent with a plausible explanation of these limits. Because of space constraints, I shall not be offering a complete theory of the proper limits of anti-discrimination law. Nor shall I be mentioning all of the values that would have to be considered when discussing the different interests of others that we must balance against the claimant's deliberative freedoms. Rather, I shall simply try to show that my account has the resources to offer such a theory; and I shall explore the kind of reasoning in which we would have to engage, when developing it.

Consider first the social contexts in which anti-discrimination law applies. It standardly applies in the employment setting: in regulating hiring, work scheduling, job promotions, and dismissals. It also applies to the provision of goods and services, rental housing, and education, and to vocational associations such as trade unions and self-governing professions. But it does not apply to our family life. Parents, for instance, are not under a legal duty to give their daughters the same educational opportunities that they give their sons. Although they must give all of their children access to public education where it is available, they are not legally required to give their daughters a private school education or a set of extracurricular activities comparable to that which they give to their sons. Nor does anti-discrimination law apply to our private dealings with friends or, for that matter, with strangers. If we wish to befriend only people of a certain religion, we are free to do so, and if we wish to

give handouts only to homeless people of certain races, we are likewise free to do that. Recreational clubs and special interest groups are interesting marginal cases. Unlike the private interactions of families and friends, they are not always excluded from anti-discrimination law's reach. But like them, they are given greater freedom of association at the expense of certain individuals' deliberative freedom.

In my view, we can see the limitations on the social contexts in which anti-discrimination law operates as attempts to foster certain values other than freedom. One of these may be the value of deep personal relationships. One plausible reason why anti-discrimination law does not regulate personal dealings between family members and friends is that, for caring relationships to flourish, people need to be able to choose whom they wish to be close to and how they are going to treat each other and they need to be able to treat their loved ones in special ways, ways in which they do not treat every other person. By contrast, the purpose of the kinds of public transactions to which anti-discrimination laws apply is not to develop such special relationships. It is to offer some good or service to the public. Another value that seems to be fostered by limiting anti-discrimination law to these public transactions is the value of autonomy. We need to have some spheres of activity in which we can associate with whomever we wish if we are truly to be able to develop and live out our own life plans.

But what about the limited ways in which special interest groups and recreational clubs are regulated? Can my account explain some of these limits? In Canada, most human rights codes have provisions explicitly allowing such groups to restrict membership, employment, and access to their services in ways that would otherwise be discriminatory *if* they are doing so in order to promote the well-being of a needy group that is marked out by a prohibited ground.¹⁶ So, for instance, a hospice for elderly Catholics would be permitted to deny a place to a non-Catholic; and a rape crisis center was recently allowed to deny a job to a male-to-female transsexual on the grounds that her presence would further

16. Many of the Canadian provincial human rights codes have explicit exceptions carved out for special interest organizations that serve the needs of a particular group identified by one of the prohibited grounds: see, for instance, s.18 of the Ontario Human Rights Code, note 1 above, and s. 41 of the British Columbia Human Rights Code, R.S.B.C. 1996, c.210.

traumatize their clients.¹⁷ If it is the goal of anti-discrimination law to protect the deliberative freedom of such groups as non-Catholics and transsexuals, how can we explain such concessions to special interest groups? One explanation is that such groups cannot fulfill their mandate unless they are allowed to exclude some people on the basis of a prohibited ground. It is part of the point of a Catholic hospice to enable seniors to associate with others who share their religious beliefs, just as it is part of the point of a rape crisis center to heal women who have been traumatized by men. So insofar as these are thought of as valuable institutions, the value of their preservation and flourishing provides a reason for limiting the freedoms of the claimants in these circumstances. Moreover, one could argue these are cases in which limiting the freedoms of the claimant increases the opportunities available to those who possess the traits recognized as prohibited grounds of discrimination; and even if this does not increase their deliberative freedom, it will increase their other freedoms. So we can see this type of exception as an attempt to balance the deliberative freedom of the claimant against other values.

An exception involving special interest groups that is more difficult to explain is the one that is sometimes granted to those organizations that do *not* serve a group marked out by a ground of discrimination, such as recreational clubs. In many parts of Canada, such clubs are allowed to restrict membership on the basis of grounds such as age, sex, and family status, but not race.¹⁸ I think we can explain why recreational clubs can legitimately exclude people on the basis of certain grounds by appealing to the values mentioned in the case of family and friends: allowing such clubs to define their membership on their own terms fosters deeper social relations and also promotes the autonomy of their members. But of course the same might be said of clubs that wish to exclude people on the basis of race. If we are willing to allow a singles' social club to restrict membership, why not permit a social club for whites to exclude blacks? Wouldn't it too foster deep social relations and promote autonomy among whites? In my view, the appropriate explanation for this exclusion likewise involves balancing the value of protecting the deliberative freedoms of those excluded from these clubs against such values as the

17. *Vancouver Rape Relief Society v. Nixon*, 262 D.L.R. (4th) 360 (B.C.C.A.).

18. See, for instance, s.20(3) of the Ontario Human Rights Code, note 1 above.

autonomy of the club members and the importance of the relations such clubs foster. Given the history of subordination of various racial minorities in North America and the severity of this subordination, one might argue that their deliberative freedoms deserve greater protection than the autonomy of the members of these clubs.

One might suggest, however, that there is a better explanation of why special interest groups are not permitted to make distinctions on the basis of race. It has to do with a feature of discriminatory actions that I have so far not mentioned. They often send a demeaning message about the victim, a message that she is less worthy than others. One might propose that, given the history of maltreatment of blacks and aboriginal peoples in North America, racial exclusions are demeaning to a much greater extent than other exclusions; and this is why it is unacceptable for special interest groups to make such exclusions. This response is only open to us, however, if we accept that a part of the wrong of discrimination lies in the demeaning message that it sends. As I shall explain in more detail in Section IV, my view implies that the demeaning messages sent by discriminatory actions are a side effect of the wrong rather than a constituent feature of it.

I have argued that we can explain the limits on the social contexts or types of transactions to which anti-discrimination law applies by appealing to the need to balance claimants' deliberative freedoms against other important values. But even in the contexts in which anti-discrimination law does operate, it places constraints on the extent to which the claimant's deliberative freedoms are protected. For it makes certain defenses available to alleged discriminators, which allow them to continue with what might otherwise amount to discriminatory conduct. In the remainder of this section of the article, I shall consider a number of these defenses. I shall argue that in some cases, these defenses allow us to recognize that the claimant's deliberative freedom was not in fact threatened by the allegedly discriminatory practices. In the remaining cases, the defenses work to balance the claimant's deliberative freedom against the freedoms of other people, such as the employer, his clients, or the people whose interests the organization is designed to further. As I shall suggest, the freedoms of others that are balanced against the claimant's deliberative freedoms are like deliberative freedoms in that they are freedoms to make decisions about how to live without having to consider certain factors. So they are not of

a completely different kind from the freedoms of the claimant that anti-discrimination law aims to protect.

Let us first consider a defense available in cases of apparent direct discrimination or disparate treatment, where a rule explicitly excludes the claimant on the basis of a prohibited ground. Most jurisdictions allow the alleged discriminator to try to justify such a rule by showing that it is a “Bona Fide Requirement” (BFR); that is, the rule is “reasonably necessary” for the normal operation of the business or the fulfillment of the organization’s goals.¹⁹ For instance, suppose a theater director hires a male actor to play Hamlet and denies that role to an equally talented female actor. He might argue that the restriction of the role to a male was reasonably necessary to the successful staging of the play. In this type of case, where the exclusionary rule is one of the conventions that we accept as defining the very enterprise that the organization is engaged in, we might wish to conclude that the claimant’s right to an equal set of deliberative freedoms has not really been compromised at all. That is, we might conclude there is no need to balance the claimant’s deliberative freedom against the freedoms of others because *no one* is entitled to this kind of deliberative freedom: in our society, we accept that dramatic characters are normally played by someone of the same gender as the character, and so no one has a right to deliberate about the parts for which they will audition in a manner that is gender-blind.²⁰ But another way to analyze such cases—and this, I think, applies to the majority of cases that fall under the BFR defense, in which the exclusionary rule is simply a more advantageous way for the organization to accomplish a

19. How strict an interpretation of “necessary” is used here depends on the jurisdiction and sometimes on the ground of discrimination at issue. I cannot enter into a detailed discussion of these differences here; but I shall try to show below how my account might explain some of them. (Note that this defense is sometimes referred to as the “Bona Fide Occupational Qualification” defense, or BFOQ; but since “Bona Fide Requirement” is a broader term that encompasses contexts beyond the employment context, I shall use this broader term.) For one example of a provision laying out this defense, see sec. 2000E-2(e) of Title VII of the Civil Rights Act of 1964.

20. I am not claiming here that this is the only reason why a rule could constitute a BFR; I am arguing that it is one such reason, and that in such cases, there is no need to engage in any balancing of the claimant’s interests against others’ interests. I shall go on to consider a different kind of reason for recognizing a rule as a BFR below, one which does require us to engage in this kind of balancing exercise. There too, my claim is not that these are the only kinds of reasons for recognizing a rule, or that these are the only interests relevant in the balancing process.

goal that is not defined by the exclusionary rule—is that the claimant’s deliberative freedom is being balanced against the freedoms of others. So one might argue that the freedom of the female actor in my example *is* in fact at issue, because in spite of our conventions, actors ought not to have their decisions about which parts to play restricted because of their gender. One might then hold that the BFR defense can nevertheless appropriately be invoked in this case because we need to balance the actor’s deliberative freedom against the director’s freedom to stage a play in the manner that he deems appropriate without having to consider the effects of his decisions on the people who are acting in the play. Another example, and one that shows this balancing exercise even more clearly, is mandatory retirement. Mandatory retirement provisions are now explicitly prohibited in the United States; but such age restrictions are permitted as BFRs in some provinces of Canada and also in the United Kingdom. And we can explain why they might be considered BFRs by appealing to the need to balance the deliberative freedoms of older workers against the freedom of choice of others—for instance, the freedom of younger employees to decide to take up certain jobs and try to rise to higher positions, and perhaps also the freedom of the owners or managers of the organization to decide to diversify their organization by being able to hire younger and different workers, without having to expand their workforce.

We can see the same balancing exercise underlying a related defense that is available in the United States in cases of disparate impact or indirect discrimination. Where a practice is facially neutral but has the effect of disproportionately disadvantaging certain groups on the basis of a ground, the alleged discriminator can escape a finding of discrimination by showing that the practice is “consistent with business necessity.”²¹ Now, because my account of discrimination holds that there is no deep difference between cases of disparate treatment and disparate impact (the wrong at issue in both types of case is a denial of someone’s right to equal deliberative freedoms), my account implies that there should be no difference in the stringency of the standard applied in the BFR defense and the standard applied in this “business necessity” defense. This is the position that is taken in Canadian law, which draws no deep distinction between cases of direct and cases of indirect

21. See, for example, s. 2000E-2(k)(1)(A) of Title VII of the Civil Rights Act of 1964.

discrimination, and which allows alleged discriminators to use the same BFR defense in both types of case.²² One might object that it is problematic that my account cannot capture the distinctions between these defenses in U.S. law. I shall consider this objection in Section IV of the article, where I will examine it as part of a broader objection that my account has too many significant revisionist implications for U.S. anti-discrimination law. For now, I want to note that both the existence of some such defense and the kind of reasoning we engage in when we ask about whether it is appropriately invoked can easily be explained on my account, in terms of the same kind of balancing I noted above. Consider, for instance, a basic case of indirect discrimination, such as the adoption by an English tearoom of a policy that waiters must speak only English at all times while on the job, even when on their breaks.²³ This policy would likely disproportionately burden people from certain ethnic groups. Now consider three different scenarios. Suppose, firstly, that eliminating the English-only rule would cause the tearoom to lose some profits, because it would thereby seem somewhat less of an authentically English tearoom to its clients. Suppose, secondly, that the tearoom would lose so many clients for this reason that it would have to close. And suppose, finally, that the clients would go elsewhere not because they felt the tearoom had become less authentic, but rather because they were prejudiced and did not wish to be served by non-English staff. In the first two scenarios, we can understand the relevant question as how to balance the deliberative freedoms of the potential staff against the freedom of choice of the tearoom owners and clients. We might want to claim, in the first scenario, that there is really no balancing necessary: the prospective employees' deliberative freedom is the only freedom at

22. See the "Meiorin" case, note 3 above, and the Ontario Human Rights Code, ss. 11(1) and 11(2).

23. "English-only" workplace policies are quite common, and they have been the source of much litigation and considerable media attention in the United States. U.S. courts are divided on whether and when they amount to disparate impact discrimination, although the U.S. Equal Employment Opportunity Commission has issued a guideline explicitly prohibiting rules that require English spoken at all times. See, for two examples: *EEOC v. Premier Operator Servs., Inc.*, 113 F. Supp. 2d 1066, 1073 (N.D. Tex. 2000), in which the court found that a rule requiring all employees to speak English at all times disproportionately burdened Latino employees; and *Garcia v. Spun Steak Co.*, 998 F.2d 1480, 1487–89 (9th Cir. 1993), in which the court found that bilingual employees were not disproportionately burdened.

issue, because the freedom of the owner to do whatever he thinks will continue to bring in a certain level of profit is not a freedom to which we think any businessperson is entitled, just as the freedom to decide to continue attending a restaurant of a certain character is not one to which any client is entitled. In the second scenario, however, the fact that the abandonment of the English-only rule would result in the tearoom closing down might lead us to think that the restaurant owner's freedom is relevant and has at least some weight. Although no one has a right to a particular level of profit, people may have a right to be able to continue to run their businesses, subject only to the normal pressures of competition in the marketplace. The third scenario might seem more difficult, for it might seem to introduce an element that my account cannot explain. We would surely want to discount the owner's business loss in this case, because it originates from the prejudicial preferences of clients. But if, as I have been suggesting, we must simply balance the competing freedoms when deciding whether a defense applies, it is not clear that my account has the resources to treat this scenario any differently from scenario two. I think, however, that we can explain the difference in the same way that the law often explains its refusal to countenance certain sorts of practices. For anti-discrimination law to treat this loss of business like any other would be implicitly to countenance, and to encourage the persistence of, the prejudicial attitudes that cause this loss. And these are the kinds of attitudes that result in limitations on the freedoms of claimants from these groups. So the law cannot recognize business losses that stem from client prejudice without undermining its goal of promoting and protecting each individual's right to these deliberative freedoms.

As I have mentioned, the third broad group of anti-discrimination laws includes what are called *accommodation requirements*, or requirements that we provide special accommodations to meet the needs of certain people, such as those with disabilities. In the case of accommodation requirements the defense available to alleged discriminators is the related defense of "undue hardship": they are required only to accommodate the claimant up to the point of undue hardship.²⁴ Once again, I want to postpone until Section IV the objection that my account is flawed because it does not treat the wrong or the defense in such cases

24. See the Americans with Disabilities Act, note 3 above, 12112(B)(5)A.

as substantially different from those at issue in disparate treatment or indirect discrimination cases. My aim here is simply to show that I can offer a plausible explanation of how we might reason through an undue hardship defense, in terms of the same sort of balancing involved in applying the other defenses. To see how the defense of undue hardship might work on my account, consider a different kind of restaurant example. Suppose there is a Muslim restaurant that refuses to serve blind clients with guide dogs on the grounds that dogs are impure and they would have to perform a ritual cleansing of the areas the dog had touched after each such client had left.²⁵ Is it undue hardship to require this restaurant to admit the blind clients? This question, on my account, is to be decided by trying to balance the deliberative freedoms of those who are blind against the restaurant owners' freedom to run their business the way they wish and the restaurant owners' freedom to run their business without having to think about the costs that their religion imposes on doing so. If it is a small restaurant with very few staff, we might think that having to cleanse the floor after the odd visit from a guide dog would not pose any difficulty for the operation of the restaurant. So the only freedom that would be compromised, in a relatively minor way, would be the owners' freedom not to have to think about their religion as a cost while running their business. By contrast, the blind client would not be able to attend this restaurant at all. So it seems as though her deliberative freedom would be compromised to a greater extent than the deliberative freedom of the restaurant owner.

I have now argued that discrimination can be understood as a denial of a person's right to have a certain set of deliberative freedoms, and to have these freedoms to an extent roughly equal to those held by others. If we see discriminatory acts in this way, then we can see that they involve personal wrongs toward the victims of discrimination, quite apart from whatever broader harms they may cause toward the group of people who share the trait in question. I have also tried to show that this account of discrimination is capable of explaining certain basic features of anti-discrimination law, such as our appeal to certain grounds of

25. This case is based upon the many recent incidents involving Muslim taxi drivers refusing to carry passengers with guide dogs: see, for instance, "Unclean Guide Dog Banned by Muslim Driver," Report in *The Daily Mail* (U.K.), October 6, 2006; and "Taxi Firm Pays Blind Man after Dog Refused," Report in *The National Post* (Canada), August 16, 2007.

discrimination, and that it is consistent with a plausible explanation of the boundaries of anti-discrimination law. In the next section, I shall turn to several potential objections to my account.

IV. RESPONSES TO OBJECTIONS

One might object to my account on a number of different grounds; and in this section of the article, I shall try to respond to some of these objections. The account might seem, firstly, too focused on deliberation and not focused enough on the actual opportunities that matter to victims of discrimination. Victims care about having a job or being able to access public places, not just about being able to decide to do these things without having to think about their race or their disability. So it may seem that my account fails to capture the reasons why we care about preventing discrimination. Secondly, one might object to the account's focus on freedom: one might argue that in many core cases of discrimination, freedom does not seem to be at issue at all. Thirdly, my account has certain revisionist implications for U.S. anti-discrimination law, and these might seem objectionable. Finally, one might argue that my account overlooks one of the most significant features of a discriminatory act: the demeaning message that it sends about its victim. I shall discuss each of these objections in turn.

On my account, when we discriminate against someone, we wrong her by denying her the freedom to make decisions about how to live in a manner that is insulated from the effects of certain extraneous traits. One might argue that this fails to capture what the victims of discrimination really care about when they seek redress from discriminatory organizations. What they want, one might say, is not just the chance to *decide* to take a certain job or use public transit without having to think about their race or disability, but the chance actually to *do* these things. They care about the opportunities that are made available to them, and they want to have the same opportunities as others. They do not care only about the effects on their mental states of being denied these opportunities because of certain traits.

An initial response to this objection is to point out that it depends on an artificially narrow understanding of what it is to make a decision. For the purposes of this account of discrimination, we do not need to think of decision making as a matter of having a certain mental state. Rather,

we can think of it as making a commitment, taking active steps to bring something about. It is this that anti-discrimination law protects, on my view. It protects our ability to make certain commitments in certain ways, free from the pressures of certain traits. So my account does not imply that claimants do or should care about their mental states. Rather, it implies that anti-discrimination law protects their ability to make commitments in a certain kind of way. And this is something we do care about.

But I think there is also a more powerful response to this objection. This is that it overlooks one of the necessary conditions for our having a particular deliberative freedom. A deliberative freedom is the freedom to make a decision in a certain way. But I am not, in any meaningful sense, 'free' to decide to do something if I do not in fact have the opportunity to do it. I am not, for instance, free to decide to rent an apartment without having to consider my race unless I actually do have the opportunity to rent that apartment. So a necessary condition of my having a particular deliberative freedom is that I really have the opportunity to do the thing I may decide to do. Of course it matters to the victim of discrimination that she actually have the chance to rent the apartment. But this is because, without this chance, she would not really be free to decide whether to rent it or not.

Perhaps, however, this objection about my focus on deliberation could be put in a rather different way. A focus on deliberative freedom, one might argue, renders the wrong of discrimination too subjective, too dependent on the knowledge of the claimant. Surely there can be cases of discrimination where the victim remains oblivious to the wrong that has been done to her. In such cases, there would not be any constraints on her deliberations because she would be unaware of the ways in which an extraneous trait was affecting the opportunities available to her; but we would still want to say that she had faced discrimination. Suppose, for instance, that a company hands out weekly bonuses to all male employees but none to female employees. Because the manager realizes this is discriminatory, he suggests to the male employees that it would be in their interest not to reveal this practice to the female employees. Suppose that the female employees have no idea, for a time, that this is happening. It is probably true that such a policy would eventually be exposed, and so would eventually have an impact on the female employees' deliberative freedom. But we would want to claim that the policy is

discriminatory even before the female employees discover it. Does my account have the resources to explain why? I think it does. For in order for me to have full deliberative freedom with respect to a certain decision, it has to be the case not just that I believe I can make that decision without having to worry about pressures from normatively extraneous traits, but that I really am free from those pressures. If this is so, then the women in the above example do not really have deliberative freedom as long as they are being disadvantaged on the basis of their gender. It is true that they are unaware of their lack of freedom. But as long as having deliberative freedom requires that one actually does have the opportunities that one thinks one has, then the women in this example lack it even though they do not realize that they do.

I have now considered two variants on an objection that focuses on the deliberative aspect of my account of discrimination. But one might object that the problem is not in restricting the account to deliberative freedoms. It is in basing the account on the idea of freedom in the first place. One might argue that there are certain core cases of discrimination that do not seem to involve a denial of freedom at all. Consider the following example. My employer has built a lavish washroom for male employees, complete with expensive plumbing fixtures, granite countertops, large windows, and a gorgeous chandelier. By contrast, the washroom for women is fitted with poor fixtures and cheap sinks, has no natural light at all, and is lit only by a few fluorescent lights in cheap, ugly fixtures. Suppose for the sake of the example that this women's washroom is perfectly functional and does not leave the women any less able to look after their hygiene than the men's washroom leaves the men. It is undeniable that something discriminatory has occurred here. But it seems rather forced to describe the problem as a lack of freedom on the part of women, much less deliberative freedom. After all, they are as free as their male counterparts to use *a* bathroom. And we have set up the example in such a way that their washroom gives them as much freedom to look after their hygiene as the men have. So surely they also have all of the corresponding deliberative freedoms. What makes this act discriminatory must be something else entirely, something different from its effects on their freedom.²⁶

26. I owe this example to Marshall Cohen.

But are they equally free? What exactly is the relevant freedom in this example? Is it really the freedom to use a bathroom and to make decisions about one's hygiene? I think not. It may be true ordinarily that the only freedom that is affected by the washrooms in one's place of employment is the freedom to look after one's hygiene. But when an institution makes a lavish washroom for one sex and not the other, it turns the nature of the washroom into a significant issue in the workplace. Using the washroom then becomes a different kind of activity in the workplace. One might say that it goes from being something that was not an activity in the workplace at all—that is, not a part of your employment at all but just something you did on your breaks—to something that is now a part of what it is to participate fully in the workplace. So in lacking the ability to use the lavish washroom, women lack the ability to participate fully in their workplace. And they lack this because of a normatively extraneous trait, their sex. As a result, they can no longer deliberate about their work (whether to stay with this particular organization, whether to seek a promotion, and so on) in a manner that is free from concerns about their gender. To lack this is to lack an important deliberative freedom.

But what about cases in which it seems as though everyone lacks the same opportunity, so that for this reason we are reluctant to say that anyone's freedom is hindered in a significant way? Suppose that a country authorized restaurants to discriminate against clientele on the basis of religion, provided that within any particular area there were an equal number of restaurants accessible to people of any given religion—for instance, thirty restaurants for Christians, thirty restaurants for Jews, thirty restaurants for Muslims, and so on. To take off the table the situation of people who wish to go out for dinner with friends of different religions, let us suppose there are also thirty restaurants open to everyone. Suppose further that each of these groups of thirty restaurants has the same mix of good restaurants and poorer ones, so that no religious group is left with vastly inferior dining overall. And suppose that it is equally easy for everyone to access the restaurants that are open to members of their religion. This situation still seems discriminatory. Yet everyone seems as free as everyone else: everyone can decide which of thirty restaurants to go to, in the knowledge that they will have the same quality of food to choose from and the same ease in getting there. So how can this be a problem of freedom?

Once again, we need to remember what freedom is at stake here. The relevant freedom in this example is not the freedom to deliberate about which of thirty restaurants to attend. Rather, it is the freedom to make decisions about which restaurant to go to without having to think about one's religion as a barrier to any restaurant. This freedom is entirely lacking in the scenario described above. In fact, what the country in this example has done is to make religion an issue for everyone when deciding where to dine, so now everyone lacks the relevant deliberative freedom. No one can decide whether to go to one restaurant or another without having to consider their religion and whether it will prevent them from entering a particular restaurant. So, in fact, my account is quite capable of explaining why this situation is discriminatory. For my account foregrounds the deliberative freedoms that are absent in this example; and it treats anti-discrimination law as aimed at ensuring not just that everyone has them to an extent roughly equal to others, but that everyone actually does *have* them.

This hypothetical example of religious discrimination by restaurants reveals several important features of the conception of freedom that I am using, and it seems worth pausing to discuss these in greater detail. First, the kind of freedom to which I am appealing is not a matter of having options to choose between, such as a number of restaurants that one could attend; nor is it even a matter of having an adequate range of options. The people in the example lack a deliberative freedom even though each of them has access to a wide array of restaurants. And they would lack it even if we supposed that the number and variety of restaurants increased to the point where everyone could attend some good quality restaurant of any kind they might wish. That is because having a deliberative freedom depends not on the range of options available to a person, but on the absence of certain considerations from his deliberations. Since all of the people in our example would have to treat their religion as a constraint when deliberating about where to eat, they would all equally lack this deliberative freedom even if they all had an equal range of alternative restaurants to attend.

A second feature of my account that this example draws out is that it is never adequate for the alleged discriminator to "level down" and deprive everyone of the relevant deliberative freedom. I explained above that if everyone were to be barred from just as many restaurants as everyone else on the grounds of their religion, this would not eliminate

religious discrimination. On the contrary, everyone would now face religious discrimination, for everyone would lack the relevant deliberative freedom. This point can be generalized to any case of discrimination: it is never adequate, on my account, for a respondent to try to deny everybody the relevant deliberative freedom to an equal extent. This seems to accord with our understanding of the remedies that are appropriate in cases of discrimination. It is never an answer for the respondent to propose that he disadvantage everybody on the basis of the relevant ground of discrimination. And my account of discrimination offers a plausible explanation of why leveling down seems inappropriate in such cases. It is because each of us is entitled to certain deliberative freedoms, and our entitlement is not dependent on whether other people have been granted these freedoms. On this view, then, although it is true that we are all entitled to these deliberative freedoms to an extent roughly equal to others, this is only because each of us has an independent entitlement to the freedoms in question. Our entitlement is not contingent upon whether others also have the relevant deliberative freedoms; though, as I have mentioned, many often already do, and it is often through a comparison with these other people that claimants come to realize that they lack a particular deliberative freedom.

Finally, I want to note that it matters very much on this account that we characterize what is at stake for the victims of discrimination in terms of a freedom rather than in terms of doing some particular thing, such as eating at a particular restaurant, or taking up a certain job, or renting a certain apartment. To see this, consider a different example. Suppose that a company terminates one of its senior employees. She complains to the appropriate human rights adjudication body and it finds that she was discriminated against on the basis of age. Suppose now that, in order to remedy this discrimination, the adjudicator decides to ensure that the complainant keeps her position in the company. So he drafts an order requiring that the company offer her the job back, and he also requires that the claimant accept this job offer and remain at the job for the next five years. The claimant now has her job back. But she clearly does not have what she wanted, or what we normally try to provide through anti-discrimination law. And this suggests that what we care about is not giving a person a particular job, but rather giving her the freedom to decide for herself whether or not to take this job, and for how long. What is at stake in anti-discrimination law, then, is not just certain people's

jobs or places to live or business transactions. It is their freedom to decide whether or not to take a particular job or live in a certain place or enter into a certain business transaction, without having to consider the costs of certain traits of theirs in their deliberations.

The third objection to which I want to respond is that my account has revisionist implications for U.S. anti-discrimination law. In particular, it denies that there is a deep distinction between intentional and non-intentional forms of discrimination, and it denies that there is a deep distinction between disparate impact cases and cases involving “mere” breaches of accommodation requirements. I shall look briefly at each of these distinctions.

My account implies that there is not a deep difference between intentional and non-intentional forms of discrimination: both are wrong because and insofar as they deny equal deliberative freedoms to certain individuals. By contrast, some scholars have argued that these two forms of discrimination are importantly different, at least when intentional discrimination is understood as motivated by a desire to exclude the victim.²⁷ Certainly from a moral standpoint, there seems often to be a difference in how blameworthy one is when one desires to exclude someone on the basis of an extraneous trait (and the exclusion is not a BFR) and when one’s actions simply have the undesired effect of excluding her. And U.S. anti-discrimination law reflects this distinction. Cases of disparate treatment are generally regarded as cases in which the claimant is aiming to prove the discriminator’s objectionable motive, either directly by adducing evidence of the motive itself, or indirectly by pointing to a facially discriminatory rule. And, as I explained in the previous section, the defense available to the alleged discriminator in such cases is more difficult to make out than the laxer standard of “business necessity” applied in cases of disparate treatment, or than the even more relaxed interpretation that is currently given to the U.S. requirement of accommodation to the point of undue hardship, where even certain increases in costs can count as undue hardships. The differences in these defenses seem to reflect a belief that intentional discrimination is worse

27. See Larry Alexander, “What Makes Wrongful Discrimination Wrong?” *University of Pennsylvania Law Review* 149 (1992): 141–219; Matt Cavanagh, *Against Equality of Opportunity* (Oxford: Oxford University Press, 2002); Gardner, “Liberals and Unlawful Discrimination”; and Richard Primus, “Equal Protection and Disparate Impact: Round Three,” *Harvard Law Review* 117 (2003): 493–587.

than non-intentional and is also a wrong of a different kind, so it cannot be negated by showing that the freedoms or interests of the alleged discriminator or his clients would be affected if the discriminatory practice were eliminated.

However, a major problem with this assumption that there is a deep legal distinction between intentional and non-intentional discrimination is that it leaves us without an explanation of why we treat these as two forms of the same injustice: discrimination. On this view, intentional discrimination is readily understandable and amounts to a wrong presumably because of the discriminator's attitudes. But it is a mystery why we would view the non-intentional exclusion of people on the basis of prohibited grounds as similarly objectionable, unless for very different sorts of reasons, such as the group harms of stigmatization and subordination in which they result. So, on this view, the two forms of discrimination really involve two very different kinds of wrongs, and are not actually two forms of the same injustice. Some scholars, such as John Gardner, have openly embraced this apparent disconnect: on Gardner's view, direct discrimination is a personal wrong, and it is a wrong because of the discriminator's motives; but indirect discrimination involves a harm to a group.²⁸ Such a view requires us to give up what I take to be a very fundamental intuition, which is that both forms of discrimination are forms of the same thing, the same kind of injustice.

My account also differs from U.S. anti-discrimination law in implying that there is no deep difference between disparate impact cases, on the one hand, and failures to meet "accommodation requirements," on the other. On my view, the wrong of discrimination is precisely that one has failed to provide to certain people the kinds of accommodations that we often provide to others, and that they are therefore constrained in their deliberations in a way that others are not. So, to return to two of my earlier examples, whether one has adopted an aerobic capacity test that excludes women and therefore needs to formulate a new test and perhaps hire some additional employees, or whether one has failed to accommodate the blind client's guide dog in one's restaurant and (let us suppose) needs to change one's policy and go to the trouble and expense of ritual cleansings, one has committed the same wrong, the wrong of denying someone equal deliberative freedom. Some have argued that

28. See Gardner, "Liberals and Unlawful Discrimination."

there is actually a deep difference between these two kinds of discrimination. But this is a controversial claim. Recently, scholars such as Christine Jolls and Samuel Bagenstos have tried to show that there really is no deep difference between these two types of anti-discrimination law in terms of their effects on employers and businesspersons.²⁹ Both can be equally costly, as my example suggests. And both can impose positive obligations on businesspersons to do certain things, for instance, to redraft the test and hire new employees or to clean the restaurant. Given the difficulty of articulating a principled difference between these two types of anti-discrimination law, it does not seem to be a flaw in my account that it implies they are responses to the same wrong.

There is one final objection that I want to address. One might worry that my account overlooks an important feature of any discriminatory action. This is the message that discriminatory actions send to others about the value of those who have been excluded, the contempt that they express. Actions that exclude or disadvantage someone on the basis of a prohibited ground of discrimination often send the message that such people are not worth considering, and that they have less value than other people. Further, they can send this message even if the discriminator does not himself hold this view or intend to convey it. Following Elizabeth Anderson and Deborah Hellman, we can call this the “expressive dimension” of discriminatory actions.³⁰

My account does not look to the expressive dimension of discriminatory actions as the source of why they are wrong. But it can certainly recognize that discriminatory actions often convey demeaning messages, and it offers us a plausible explanation of why. If it is true, as I have argued, that discriminatory actions prevent some people from having an equal set of deliberative freedoms, then the importance of these freedoms to us is likely to mean that any act or policy that denies some individuals these freedoms will imply that these people are second-class citizens. On my account, then, although discriminatory actions can certainly express demeaning messages, they only do this *because* they

29. See the articles cited in note 7.

30. Deborah Hellman, *When Is Discrimination Wrong?* (Cambridge, Mass.: Harvard University Press, 2008); and “The Expressive Dimension of Equal Protection,” *Minnesota Law Review* 85 (2000): 1–70; and Elizabeth Anderson and Richard H. Pildes, “Expressive Theories of Law: A General Restatement,” *University of Pennsylvania Law Review* 148 (2000): 1503–75.

are restrictions on deliberative freedoms to which a person is entitled. Their demeaning message results from the way in which they wrong individuals. It does not constitute that wrong; rather, it is a side effect of the wrong. Moreover, it is quite consistent with my account to suppose that the demeaning message that some discriminatory actions convey could hurt the victim's sense of self-worth or result in increased stigma attaching to her, and consequently entitle her to greater compensation. For instance, cases of intentional discrimination often express the view that the victim is dispensable, or at least not as important as the discriminator's own objectives. This will often cause additional anguish to the victim, over and above the suffering that results from the exclusion itself; and it may increase the social stigma that attaches to her. The same may occur in some cases of non-intentional discrimination; though here, the emotional suffering and stigma may be somewhat less. My account is open to the possibility that we should make additional damages available in all of these cases, to recognize these further ways in which the claimant has suffered. But these demeaning messages are not, on my view, what makes the discrimination wrongful.

In this article, I have argued that we should understand discrimination in the private sector primarily in terms of wrongful interference with another person's right to a roughly equal set of deliberative freedoms. I have tried to show that my account offers us a principled explanation of the grounds of discrimination: they reflect our judgments about which sorts of deliberative freedoms we can rightfully claim that others give us, and which we cannot. I have argued that we can explain the limits of anti-discrimination law in a manner that is consistent with this account. Furthermore I have urged that, because this account makes sense of the basic structure of most anti-discrimination laws as involving individual complaints of personal wrongs, we have a strong reason to take it seriously.

Of course, there remains work to be done in developing this account of discrimination. One area that requires further exploration is how this view of the personal wrong involved in discriminatory actions might be conjoined with an account of the harms or injustices that are suffered by each of the groups marked out by grounds of discrimination. It may be that some anti-discrimination laws—and in particular, the requirements prohibiting disparate impact—are also designed to rectify injustices suffered by these groups. This is suggested, for instance, by the fact that

enforcement bodies often have the power to order remedies that go beyond what is necessary to redress the wrong suffered by the particular complainant. In addition to giving the claimant the necessary compensation or accommodation, a tribunal can order the organization to alter its policies toward everyone; to change its physical plant; or even to adopt hiring quotas. Any complete account of discrimination law, or of our shared public idea of what discrimination between individuals amounts to, would also have to say something about these group-based injustices or inequalities. It may be that the remedies I have described above do not actually aim at rectifying any injustice toward a group, and that they are instead forward-looking: discrimination law aims primarily at redressing a personal wrong, but it can serve as a convenient vehicle for preventing such injustices from arising in the future by requiring particular organizations to accommodate various groups. Alternatively, it may be that discrimination law aims at redressing two types of past injustice—the type of personal wrong that I have described here and some kind of group-based injustice. If so, it would be worth inquiring whether we can offer a coherent account of discrimination that explains how such group-based injustices are related to the personal wrongs that I have discussed. These are questions that I hope to pursue on another occasion. But I hope this article has provided at least the beginnings of a workable account of discrimination in the private sector and has helped to show why it is a perplexing and philosophically interesting phenomenon.