

Faces of Inequality:
A Theory of Wrongful Discrimination

Sophia Moreau

For Chris, Kathy, and Rebecca

Acknowledgments

One of the best things about being an academic is that it gives you a chance to have conversations that stretch over many years, with people who live in many different places, about the things that you think matter most. The ideas in this book have grown out of such conversations –conversations with philosophers, legal academics, lawyers, judges, and also with people who have faced discrimination. I am grateful to everyone who has engaged with me in these ongoing conversations about equality rights and discrimination, and particularly to Rebecca Cook, John Gardner, Debbie Hellman, Tarun Khaitan, Niko Kolodny, Kasper Lippert-Rasmussen, Arthur Ripstein, Larry Sager, and Seana Shiffrin. I also want to thank Tim Scanlon, my Ph.D. supervisor at Harvard: although I did not write my dissertation on anything to do with equality, his teaching fundamentally shaped the way I approach philosophical problems like the ones discussed in this book.

I have been immensely fortunate to have a wonderful group of colleagues at the University of Toronto, who have discussed these ideas with me over the years. I am grateful to all of them. I also want to thank Dean Iacobucci for the Faculty of Law’s generous support of my research, and the Social Sciences and Humanities Research Council of Canada for research grants that helped cover the costs of research assistance.

Many other people have read draft chapters of this book and generously gave me extensive comments on them. I am very grateful to all of them, including: Cheshire Calhoun, Colm O’Cinneide, Joshua Cohen, Hugh Collins, Chris Essert, David Estlund, Barbara Herman, Greg Keating, Brian Langille, Alan Patten, Jonathan Quong, Larry Sager, Anthony Sanguiliano, Samuel Scheffler, Adam Swift, Victor Tadros, Ilias Trispiotis, Daniel Viehoff, Wil Waluchow, Gary Watson, and Jacob Weinrib. I am also very grateful to Timothy Endicott, John Gardner and Les Green, the editors of the Oxford Legal Philosophy series, for their helpful suggestions; to Jamie Berezin for all of his editorial help; and to two anonymous referees from Oxford University Press, for their comments. Finally, I want to thank my excellent research assistants Maggie O’Brien, Matteen Victory, and Andy Yu, who provided first rate research support.

Draft chapters of the book were presented at a variety of workshops and conferences, and I want to thank these audiences for their helpful suggestions: the Workshop in Law, Philosophy and Social Theory at the University of California, Berkeley; the Department of Philosophy at the University of Copenhagen; the N.Y.U. Workshop on Political Equality in 2017; the 5th Annual Workshop for Oxford Studies in Political Philosophy; the Law and Philosophy Workshop at the U.S.C. Gould School of Law; the UCL Faculty of Laws; the Harvard Law and Philosophy Workshop; the Princeton University Political Philosophy Colloquium; the Stanford Political Theory Workshop; the McMaster Philosophy Series; the Oxford/Toronto/UCLA Colloquium in Moral, Political and Legal Philosophy; the CELPA seminar at the University of Warwick; the UCLA Legal Theory Workshop; the Queen’s University Colloquium in Legal and Political Philosophy, and the Fordham Legal Theory Workshop.

Several chapters of the book draw on previously published material. An earlier version of Chapter Two was published as Chapter Five in the *Oxford Studies in Political Philosophy*, Volume 5, 2019; and some of my critical comments on Khaitan’s work in Chapter Three were previously published as part of a review of his book, entitled “Discrimination and the Freedom to Live a Good Life,” *Law and Philosophy* 35(5) (2016), 511-527.

I am also very grateful to Robert Davidson, *Guud Sans Glans*, a Haida artist from Canada’s Northwest coast, for allowing me to use his beautiful serigraph “I am You and You Are Me” on the cover of this book. For a discussion of the symbols in this serigraph and the ways in which I think they relate to the arguments of the book, see the book’s Conclusion.

This book is dedicated to three women who have, in very different ways, inspired my work. Chris Korsgaard was a role model for me when I was a graduate student; and when, some years later, I made a concerted effort to develop my own philosophical voice, it was Chris’ writing, confident and unabashedly

directed towards a “you” without any rigid formalities, that helped to guide me. Chris was also the first person who suggested to me that philosophical analysis at its best is really the narrating of a story. I hope the stories in this book are worthy of her. Madam Justice Kathy Feldman from the Ontario Court of Appeal befriended me when I was just an undergraduate and has been a source of unconditional support ever since, meeting with me regularly to talk and to encourage me, regardless of what I was working on. And Rebecca Cook, my colleague here at the University of Toronto, has done more than almost anyone I know to create a true community of equals among scholars working on women’s equality around the world, in which everyone’s contribution is recognized and no one is deemed less important because of the country or institution at which they work. Chris, Kathy and Rebecca: I will always be grateful to you for your wisdom and our friendship.

Lastly, I owe a lifetime of thanks to my husband Stephen for his boundless love, patience, and faith in us and our family, and to our children, Gabrielle, Sebastian, Nicholas and Emma, who teach me more every day than I could ever write about.

Table of Contents

Chapter One - A Question of Inequality	
1.1 Discrimination and Inequality	
1.2 Wronging Someone and Acting Wrongly: An Important Distinction	
1.3 Two Forms of Discrimination	
1.4 Monism and Potential Problems with Monist Theories	
1.5 The Relevance of Victims' Experiences and the Relevance of Law.....	
1.6 Challenges Facing a Pluralist Theory of Discrimination.....	
Chapter Two - Unfair Subordination	
2.1 Unfair Subordination: A Plausible Starting Point.....	
2.2 Discrimination that Subordinates: Restaurant Dress Codes	
2.3 Hellman's Expressivist View: Subordinating by Demeaning.....	
2.4 Towards an Account of Social Subordination	
2.5 How Direct and Indirect Discrimination Subordinate	
Chapter Three - The Relevance of Deliberative Freedom	
3.1 Why the Wrongness of Discrimination Extends Beyond Subordination	
3.2 What is Deliberative Freedom, and When Do We Have a Right to It?	
3.3 The Perfectionist Challenge	
3.4 Why Indirect Discrimination, Too, Can Interfere with Deliberative Freedom	
3.5 Deliberative Freedom and the Role of the Prohibited Grounds.....	
Appendix to Chapter Three	
Reply to Lippert-Rasmussen.....	
Reply to Campbell and Smith.....	
Chapter Four - Access to Basic Goods	
4.1 A Third Form of Wrongful Discrimination	
4.2 Basic Goods: Further Clarification	
4.2.a To identify a good as "basic" is not to claim it is objectively good	
4.2.b Some basic goods are privately appropriable; others are public	
4.2.c "Basic" in relation to particular people in a particular society	
4.2.d Importance of the discriminatee's perspective	

4.2.e	Something can be a basic good for some people even if no others need it.....
4.2.f	What counts as “denying” someone a basic good?.....
4.3	Why This is a Problem of Inequality, and a Distinctive Form of Wrongful Discrimination.....
4.4	Basic Goods, Prohibited Grounds, and Responsibility
Chapter Five - A Pluralist Answer to the Question of Inequality.....	
5.1	Why the Theory is both Pluralist and Unified.....
5.2	Resolving Worries about Arbitrariness.....
5.3	Nothing Distinctively Wrong With Discrimination?
5.4	Each Wrong Sufficient for Wrongful Discrimination
5.5	Personal Wrongs and Group Wrongs
5.6	Cases of Conflict and the Relative Weight of Different Reasons
5.7	Advantages of this Pluralist Theory
5.7.a	It makes possible a nuanced analysis of cases
5.7.b	It resolves “the comparative puzzle”
5.7.c	It resolves “the puzzle about individuals and groups”
Chapter Six - Indirect Discrimination	
6.1	What My Theory Implies about Indirect Discrimination
6.2	Is Indirect Discrimination Easier to Justify?
6.3	Responsibility for Cost
6.4	Culpability in Direct and Indirect Discrimination
6.5	The Negligent Discriminator
Chapter Seven - The Duty to Treat Others as Equals: Who Stands Under It?	
7.1	Situating the Question
7.2	A Seemingly Plausible Answer
7.3	The State’s Duty to Treat Those It Governs As Equals
7.4	The Individual’s Duty to Treat Others As Equals.....
7.5	Reasons for Not Legally Prohibiting Discrimination in Private Contexts
7.6	The Value of Freedom and the Duty to Treat Others as Equals
Conclusion	

Chapter One

A Question of Inequality

1.1 Discrimination and Inequality

Many of us care passionately about eliminating wrongful discrimination. And we share a sense of disappointment and indignation at its recent instances: harassment of and violence against Muslims; medical staff that refuse to treat transgendered people or the infant children of gay couples; the persistent gender wage gap; the lack of safe drinking water on indigenous reserves in countries such as Canada, when other communities in these same countries have easy access to it.

These cases are troubling for a number of reasons, including some that can be explained without reference to the concept of discrimination. Harassment and unprovoked violence are morally problematic whomever they are directed against, and on whatever basis. We do not need the idea of discrimination in order to condemn them as wrongful. We can appeal to each person's right to bodily integrity or security of the person. Similarly, one explanation of why it is wrong for medical staff to refuse to treat trans people or the children of gay couples is that such patients have an independent right to certain forms of medical treatment. If they do, then it is wrong to deny them such treatment –and we can conclude this without needing to make any comparisons between these people and other people whom the staff have treated or would treat. To the extent that the gender wage gap results from failing to pay women for the full value of the work that they have done and failing to give them the kind of fair chance at promotion that each person is independently entitled to, it too can be understood as a violation of certain prior rights.¹ And there are several ways to explain what is troubling about the water crisis on indigenous reserves in countries such as Canada which likewise do not mention discrimination. One might suggest, for instance, that everyone has a moral right that their government provide basic necessities such as clean drinking water; or one might appeal to the fact that, having forcibly resettled these indigenous groups on the least arable and most inhospitable tracts of land, governments owe them a special duty to provide basic infrastructure such as piped and purified water. In all of these cases, it seems that we can appeal to a prior moral right, a legal right, or a special duty. There is no need to mention discrimination, no need to compare these people to actual or hypothetical others who have or would have received

¹ I do not think that the gender wage gap can be fully explained in this way, however. For my own analysis of it, see Section 3 of this Chapter.

different treatment, no need to consider the traits on the basis of which these people were disadvantaged.

One might take this as evidence that there is no independent wrong of discrimination. Perhaps all apparent cases of wrongful discrimination are really just wrong for some other reason, concerning an infringement of some prior and independent right. This view was defended by scholars such as Peter Westen and Joseph Raz, when Anglo-American philosophers were just beginning to think about equality and discrimination.² Their view might, in fact, be correct. But whether it is correct can't be settled simply by noting that there are other ways of conceptualizing the reason why such acts are wrongful. We need, instead, to think deeply about the further moral concern that we seem to have about apparent cases of wrongful discrimination, and to see whether we can give a coherent and systematic account of this moral concern. For the cases of wrongful discrimination that I mentioned above appear, at least, to be wrong for an additional reason, a reason that leads us to think of them as cases of wrongful *discrimination*.

What is this reason? What is the further concern that leads us to think of these as cases of wrongful discrimination? There are two importantly different ways of conceptualizing this concern, a broader way and a narrower way, and they lead us to two different ways of formulating the central question that a theory of wrongful discrimination must try to answer.

The broader way of thinking about our concern in such cases is as a concern with treating some people differently because they have certain traits. If you refuse to allow people from predominantly Muslim countries to enter your country because they are Muslim, or if you pay women less because they are women, you have treated the members of these groups differently because of their race or their gender. Of course, saying only this much does not explain *why* it is wrong to treat people differently on the basis of such traits as race and gender, nor does it tell us which sorts of traits it is wrong to use as the basis for treating people differently. These are the sorts of questions that different theories of wrongful discrimination will answer in different ways. Nevertheless, on this way of understanding our broad concern underlying wrongful discrimination, all such theories are attempts to answer the following question, which we can call:

The wrongful differentiation question: When and why do we wrong people by treating them differently on the basis of certain traits?

It may seem obvious that our moral concern with wrongful discrimination, *qua* discrimination, must just be a concern with wrongful differentiation.³ After all, doesn't

² Joseph Raz, *The Morality of Freedom* (Oxford: Clarendon Press, 1986) at pp. 217–44; and Peter Westen, "The Empty Idea of Equality," *Harvard Law Review* 95(3) (1982), pp. 537–597.

³ As John Gardner suggests in "Discrimination: The Good, the Bad and the Wrongful," *Proceedings of the Aristotelian Society* 118(1) (2018), pp. 55–81.

“discrimination” just mean differentiation? Surely, if an act is to wrong someone by virtue of being an act of discrimination, this must be because of the way in which the agent has treated this person differently, on the basis of certain traits?

There is, however, an important respect in which the wrongful differentiation question is too broad. To see this, notice that although it is possible to answer the wrongful differentiation question by appealing to the importance of equality, it is also possible to answer this question in ways that do not invoke the value of equality. For instance, some legal scholars have argued that much of the American jurisprudence surrounding the 14th Amendment is best understood as embodying the view that discrimination is wrongful because it classifies people on the basis of arbitrary or irrelevant traits.⁴ Denying a student admission to a University because he is black is, on this view, wrong because a person’s race or perceived skin colour is irrelevant to whatever the appropriate criteria are for university admissions. If this is correct, then it follows that it is also wrong, and wrong for the same reason, to deny a white student admission on the basis of his race or skin colour. As long as the white student has been classified on the basis of an irrelevant trait, he too has suffered at least a *pro tanto* wrong—even if his social status is not thereby lowered, nor his well-being, decreased, and even if it is only by denying him and other white students admission that we are able to raise the position of certain disadvantaged minorities. So this particular answer to the wrongful differentiation question presupposes no necessary connection between wrongful discrimination and inequality.

Why might this then suggest that the wrongful differentiation question is too broad? It might, for several reasons. First, because most anti-discrimination laws—both at the international and at the national levels—explicitly use the language of equality. They use it both in their preambles and in the wording of their prohibitions. They present these prohibitions on discrimination as a way of ensuring that governments, and others who owe similar duties of non-discrimination, treat everyone whom they affect as equals. For instance, the anti-discrimination protections in the Universal Declaration of Human Rights

⁴ See e.g. Paul Brest, “Foreword: In Defense of the Antidiscrimination Principle,” *Harvard Law Review* 90 (1976), pp. 1–55; and Owen Fiss’s description of the anti-classificationist approach, which he referred to as “the anti-discrimination principle” (though note that he went on to reject it): “Groups and the Equal Protection Clause,” *Philosophy & Public Affairs* 5(2) (1976), pp. 107–177. See also cases such as *Adarand Constructors, Inc. v. Peña*, 515 U.S. 200 (1995), in which the U.S. Supreme Court held that a standard of strict scrutiny applied to the use of race in awarding government contracts, even where the purpose of considering race was to try to eliminate racial subordination; and Justice Roberts’ judgment in *Parents Involved in Cmty. Sch. v. Seattle Sch. Dist. No. 1*, 551 U.S. 701, 748 (2007), in which he wrote that “[t]he way to stop discrimination on the basis of race is to stop discriminating on the basis of race” (i.e. not to use racial classifications). However, scholars such as Riva Siegel, Jack Balkin, Ruth Colker and Randall Kennedy have suggested that many American cases that seem on the surface to apply anti-classificationist principles are in fact aiming at least in part to eliminate unequal *status* relations. See, for instance, Jack Balkin and Riva Siegel, “The American Civil Rights Tradition: Anticlassification or Antisubordination?,” *University of Miami Law Review* 58 (2003), pp. 9–34; Ruth Colker, “The Section Five Quagmire,” *UCLA Law Review* 47(3) (2000), pp. 653–702 at p. 688; and Randall Kennedy, “Persuasion and Distrust: A Comment on the Affirmative Action Debate,” *Harvard Law Review* 99 (1986), pp. 1327–1346.

are presented as equality rights, aimed at guaranteeing every person equal status. Article 7 states that “All are equal before the law and are entitled without any discrimination to equal protection of the law.”⁵ Canada’s *Charter of Rights and Freedoms* contains a similarly worded provision in section 15, which presents anti-discrimination laws as part of a guarantee of equal treatment.⁶ The U.K.’s anti-discrimination laws are laid down in legislation that is called the *Equality Act*, and one of the aims stated in its preamble is “to increase equality of opportunity.”⁷ And this is typical of anti-discrimination laws around the world: they are normally presented as ways of implementing some kind of equal treatment.

Of course, the interpretation of all of these laws is complex and contested; and there is a limit to what the wording or presentation of a law can tell us about the ways in which that law ought to be interpreted. But there are other considerations, too, that suggest that what I have called “the wrongful differentiation question” is too broad a place to start, when we are looking to explain what makes discrimination wrongful, and that we need in some way to invoke the value of equality. For instance, a focus on inequality seems to make better sense of the moral criticisms that we make of those who engage in apparently wrongful discrimination. When we find it troubling that some pediatricians refuse to treat the infant children of lesbian couples, what concerns us is not the mere fact that these children have been classified on the basis of their parents’ sexual orientation, but rather the fact that they have been treated as second-class citizens on the basis of beliefs about the immorality of their parents’ relationship. When we object to medical staff refusing to treat members of the trans community for complications resulting from gender-reassignment surgeries, our objection cannot be that the medical staff have made their decision on the basis of an irrelevant trait: both sides in this debate agree that the patients’ gender identity is certainly relevant to these decisions. Our objection is rather that the medical staff have treated these patients as though they were not the equals of all other patients, perhaps on the basis of the belief that their bodies are unnatural. Similarly, those who criticized the American government for denying entry visas to citizens of predominantly Muslim countries were not only concerned with the use of race and religion in determining who is allowed to enter a certain country: they were, at a more basic level, concerned that this policy relegated citizens of these countries to an inferior status. In all of these cases, our worries seem best conceptualized, not just as concerns with inappropriate differentiation, but as concerns with differentiation that is inappropriate because it fails to treat certain people as equals. The problem is not just that distinctions are drawn in certain ways rather than others. It is that they are drawn in ways that fail to treat certain individuals and groups as equals.

⁵ U.N. General Assembly, *Universal Declaration of Human Rights*, 10 December 1948, 217 A (III), Article 7.

⁶ Canadian Charter of Rights and Freedoms, s. 15, Part I of the Constitution Act, 1982, being Schedule B to the Canada Act 1982 (U.K.), 1982, c. 11.

⁷ Equality Act 2010 (U.K.).

Of course, the authorities mentioned in these examples would certainly dispute my characterization of their acts. Doctors who refuse to treat trans patients argue that they are not casting aspersions on trans people but are simply acting in accordance with their own religious obligations. Similarly, American officials have stated that they are not implying that people from these countries are inferiors; they are just protecting their country from the risk of terrorism. But these claims support rather than call into question the point that I am making here. For they assume, rather than denying, that the relevant question, in determining whether discrimination is actually wrongful, is whether the agent has treated people differently *in ways that fail to treat them as equals*. They accept, that is, that whatever is wrong with wrongful discrimination concerns a certain sort of inequality. They just deny that their particular acts or policies actually fail to treat people as equals in the relevant sense.

A final consideration that nudges us toward a narrower, more equality-focused question than the wrongful differentiation question is the fact that the wrongful differentiation question does not itself point us in the direction of any moral concerns. As Dworkin might say, it does not present us with a recognizable moral ideal, some principle or value the departure from which might explain *why* discrimination is wrongful. By contrast, the idea of failing to treat someone as an equal does present us with a recognizable moral ideal. This is not to say that the idea of equality can, in isolation from other values, yield principles telling us how to act.⁸ On the contrary, as I shall be arguing in this book, there are different interpretations of what it is to treat someone as an equal, and these interpretations appeal to other values, values such as respect, recognition, deference, freedom, and social participation. So we can think that the wrongful differentiation question is too broad, and needs to be narrowed to include reference to the value of equality, without supposing that equality on its own must do all of the work in explaining why discrimination is wrongful.

If I am right that the wrongful differentiation question is too broad, and that one of our main concerns in cases of wrongful discrimination is a concern with inequality, then we need to think of a different question to structure our inquiry. My question will be:

The question of inequality: When we treat some people differently on the basis of certain traits, when and why do we wrong them by failing to treat them as equals?

I have called this “the question of inequality.” But you may be thinking, on the basis of the examples discussed above, that “inequality” is not the best word to use to describe what troubles us about cases of wrongful discrimination. The term “inequality” is often used in philosophical work to describe a difference in the amounts of something that different people have, whether it is well-being or opportunities or status. On this usage of

⁸ As Sam Scheffler reminds us in “The Practice of Equality,” in *Social Equality*, ed. Carina Fourie, Fabian Schuppert et al., (Oxford: Oxford University Press, 2015). Speaking of the relational conception of equality, he notes that equality is “an ideal that itself draws on a variety of other values.”

the term, if it is Pi Day and I give one pie to each of my colleagues except for two people, one of whom I give two pies and the other of whom I give no pie, I have treated both of these two people unequally, relative to all of my other colleagues. I have obviously treated the person who gets no pie unequally. But I have also treated the person who gets two pies unequally, because I have given her *more* pies than any other person. Why does this suggest that “inequality” might be the wrong word for us to use in cases of wrongful discrimination? Because no one worries about wronging those who are given more opportunities or a higher status than others, even though these people have technically been treated unequally. We do not worry if a doctor lavishes extra attention on a particular patient, provided this extra time is taken from the doctor’s own private time rather than from some other patient’s appointment; and if the extra time is taken from another patient’s appointment, it is *that* patient whom we feel has been unfairly treated, not the patient who received the extra attention. It was concerns such as this that led Harry Frankfurt and Derek Parfit to argue that although we may think we value equality, we really do not.⁹ They proposed that what we care about is not inequality *per se*; it is only certain kinds of inequalities, such as those that leave some people below a threshold of “sufficiency,” or those that make the “worst off” in society even worse off. The same might be said of our concern in cases of discrimination. Whatever is wrong with wrongful discrimination, it cannot consist simply in creating an inequality. We are not troubled by all inequalities, or all differences. We are only troubled by some. In cases of wrongful discrimination, what troubles us is not just any inequality. It is that some people are treated *as inferiors*.

I think the substance of this objection is correct. But I hope I can persuade you that it is not in fact an objection to my way of formulating what I have called “the question of inequality.” It is true that our concern in cases of wrongful discrimination is not with every inequality, but only with some—those that constitute treating people as an inferior. But it does not follow that we are mistaken when we appeal to the value of equality in explaining why discrimination is wrongful. That is because the legal meaning of the term “equality” is importantly different from the specific philosophical meaning that, according to philosophers such as Parfit and Frankfurt, does not capture our concern in such cases. When anti-discrimination laws invoke the value of equality, they do not do so in order to insist that everyone should be treated in exactly the same way, given the same amount of whatever good or opportunity that is at issue. Rather, they require that everyone be respected, with no one treated as though they had a status below that of others. So we can say, I think, that the legal ideal of equality combines two ideas: first, that everyone should be treated as though they were just as deserving of respect as others, and second, that

⁹ See Derek Parfit, “Equality and Priority,” *Ratio* 10(3) (1997), pp. 202–221; and Harry Frankfurt, “Equality as a Moral Ideal,” *Ethics* 98(1) (1987), pp. 1–13.

everyone should be treated as though they were deserving of respect, in absolute terms.¹⁰ It is this combination of ideas that I think Ronald Dworkin meant to invoke when he urged us to think about the state's duty, in relation to equality, as a duty of treating those whom it governs "as equals."¹¹ For Dworkin, treating people "as equals" meant treating them as well as others, in a context in which one was *already* treating others with *sufficient* concern and respect. This is why I have used Dworkin's phrase in my definition of the question of inequality. Recall that my question of inequality asks: "When we treat some people differently on the basis of certain traits, when and why do we wrong them by failing to treat them as equals?" This question deliberately speaks, in Dworkinian terms, of failing to treat others "as equals." It does this in order to highlight that what is troubling about acts of wrongful discrimination is not that certain people have been treated differently, but that they have been treated differently in a way that leaves them inferior to others. They ought to have been treated as well as others, in a context in which others were already being treated well; but instead they were treated as inferiors.

I argued earlier that the question of inequality does a better job than the broader wrongful differentiation question at capturing both the purpose of anti-discrimination laws and the concerns that underlie our objections to some of the most troubling cases of wrongful discrimination. I want now to suggest that, in addition, the question of inequality seems to me to focus our gaze in the right place, as we search for the features that make certain acts of discrimination wrongful. Anti-classificationist theories, which try to answer the wrongful differentiation question by giving us an explanation of when it is wrongful to differentiate between people on certain bases, often focus on the discriminator's process of reasoning. And, perhaps because of the initial influence of such theories, many philosophical theories of discrimination, too, treat the problem of wrongful discrimination as being primarily a problem on the side of the discriminator, a problem with the reasons he or she has acted upon.¹² But if what makes acts of wrongful discrimination wrongful is that they fail to treat certain people as equals, then the problem lies more in the impact of the discriminatory act on the *discriminatee*. We can certainly appeal to facts about the discriminator's reasons, in understanding what has happened to the discriminatee and in understanding how her relationship with the discriminator has been affected. But an explicit invocation of the value of equality has the advantage of bringing the discriminatee

¹⁰ See Jeremy Waldron's similar arguments about dignity in *Dignity, Rank, and Rights* (Oxford: Oxford University Press, 2012), Ch. 1 at p. 34.

¹¹ See Ronald Dworkin, *Taking Rights Seriously* (Cambridge, MA: Harvard University Press, 1978) at pp. 226–227; and Ronald Dworkin, *Sovereign Virtue* (Cambridge, MA: Harvard University Press, 2002).

¹² See e.g. Gardner, "Discrimination: The Good, the Bad and the Wrongful," *supra* note 3; Richard Arneson, "What is Wrongful Discrimination?" *San Diego Law Review* 43(4) (2006), pp. 775–806; Larry Alexander, "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies," *University of Pennsylvania Law Review* 141 (1992), pp. 149–219; and Matthew Cavanagh, *Against Equality of Opportunity* (Oxford: Clarendon Press, 2002).

into the centre of our gaze, reminding us that what matters is primarily what happens to her, and what happens to her relationship with the discriminator.

For all these reasons, the question I shall attempt to answer in this book is the one that I have called “the question of inequality,” rather than the broader “wrongful differentiation question.” I hope, in the process of answering the question of inequality, to build up a coherent and systematic explanation of when and why discrimination is wrong.

You may conclude, when you finish the book, that the answers I try to offer to the question of inequality are not sufficiently coherent or systematic, or that they have no independent explanatory power and seem simply to be restating the question. You may therefore decide that there is no satisfactory answer to the question of inequality—that is, no answer good enough to suggest that our intuitive responses to cases of apparently wrongful discrimination represent an inchoate grasp of some moral truth, rather than a mistake. If that is right, and if I am right that attempts to answer the wrongful differentiation question without appealing to inequality do not make good sense either of our laws or of our moral intuitions about discrimination, then we may have to admit that Raz and Westen were correct after all. Perhaps acts that seem to be wrongful because they wrongfully discriminate are really just wrongful for other reasons, or perhaps not wrongful at all.

But I do not think we shall have to admit this. I shall try to persuade you that there is a systematic and coherent answer to the question of inequality, and that it is an answer with genuine explanatory power. Or rather, I shall suggest that there are *several* answers to the question of inequality, several reasons why in disadvantaging a certain person or group on the basis of certain traits, we can fail to treat them as equals. For the theory that I am going to defend in this book is a pluralist theory of discrimination. I argue that there are at least three different ways in which a practice can treat some people differently on the basis of certain traits and thereby fail to treat them as an equal. It can subordinate some people to others; it can deny some people deliberative freedoms in circumstances where they have a right to these freedoms; and it can leave some people without access to certain “basic” goods, goods that one needs to have access to, in a particular society, if one is to participate as an equal in the life of that society. I argue that each of these is, on its own, sufficient to explain why discrimination wrongs people in certain cases; though I shall suggest that many cases of discrimination wrong people for more than one of these reasons. I shall not be claiming in this book that these are the only reasons why discrimination wrongs people. But I shall try to show that they are some of the main reasons—that, together, they can help us to understand many of the complaints of those who allege discrimination, and can help us to make sense of our own reactions to cases of apparent wrongful discrimination.

2.1 Wronging Someone and Acting Wrongly: An Important Distinction

I have been speaking of “wrongful discrimination.” And when I set out the question that this book aims to answer—the question of inequality—I spoke of how we “wrong” people by discriminating against them. I did not speak of our “doing the wrong thing” or “acting wrongly.” This was deliberate. My main concern in this book is with the ways in which we *wrong other people*, when we discriminate against them. In many, if not most of these cases, we thereby act wrongly. That is, wronging people by discriminating against them is most often wrong, all things considered. But I want to leave room for the possibility that in some cases, even though we wrong other people by discriminating against them, there are other pressing or urgent needs that justify us in continuing the discriminatory practice that wrongs these people. I want to leave room, that is, for the possibility that an act can wrong people by failing to treat them as equals and yet nevertheless may not be wrong all things considered, if certain justificatory considerations are present.

The moral distinction between “wronging someone” and “doing what is wrong, all things considered” has a legal parallel. Some constitutions, such as the Canadian constitution, allow that even when a particular constitutional right such as an equality right has been violated, this rights violation can be justified if certain special tests are met—for instance, if it can be shown to be a proportional means of achieving a legitimate and pressing objective.¹³ Under these constitutions, if a rights violation is justified, it does not follow that there has been *no* rights violation, no legal wrong. On the contrary, we still recognize that certain people have suffered a legal wrong: their equality rights, for example, have been violated. But this legal wrong is deemed justified, all things considered. I am appealing here to a similar distinction, in the moral realm. Discrimination, I have suggested, wrongs people when it fails to treat them as equals. But it does not follow from this that when it does so, it is always wrong, all things considered. There may be special circumstances in which it is not; which circumstances count as special justifications may vary, depending on whether the agent of discrimination is the state or a private individual. I shall say more about the relevant justifying factors, and about the difference that the type of agent makes to our assessment of which factors are relevant, in Chapters Six and Seven.

1.3 Two Forms of Discrimination

Before I turn to the task of developing my answers to the question of inequality, I need to back up a little. I have been speaking so far as though there were a single phenomenon that we collectively understood as “discrimination.” You might dispute this, for one of two reasons.

¹³ See s. 1 of Canada’s *Charter of Rights and Freedoms*, cited above in note 7.

First, you might contend that different countries have very different views of what discrimination is and why it wrongs people, as are evidenced by the different legal frameworks they use for identifying and rectifying problematic sorts of discrimination. And consequently, you might hold that there is no point in speaking of “our” concept of discrimination.

This view seems to me mistaken. There is a great deal of legal writing and philosophical writing on discrimination that presupposes that there is at least enough of a common core to the anti-discrimination laws of different countries that we can think of them all as a response to the same phenomenon: discrimination. Moreover, recently, scholars such as Tarunabh Khaitan have provided systematic accounts of the common features of anti-discrimination laws in countries such as the United States, the U.K., Canada, Germany, and India.¹⁴ So I think we can assume on the basis of such work that we can coherently speak of at least some shared features of anti-discrimination laws. Moreover, for the purpose of building my own account, I shall not be taking as basic any unusual or highly controversial features of particular countries’ laws. Mostly, I shall be appealing to several widely shared features of different countries’ anti-discrimination laws, such as that they typically recognize wrongful discrimination only in cases where it has occurred on the basis of a certain kind of personal trait, which I shall call a “prohibited ground”; that they typically recognize two forms of discrimination, direct and indirect; that they are often structured in such a way as to suggest that the discriminator has committed a personal wrong against the discriminatee; and that they do not require proof of whether a practice is morally valuable in order to protect us from exclusion from it. And even in the case of these widely shared features of anti-discrimination law, I shall not assume that they always reflect the truth. On the contrary, although I shall use these features as some of the preliminary data in developing my theory of when and why discrimination is wrongful, I shall go back and scrutinize them after I have developed the theory. I shall argue that the moral truth about discrimination is somewhat more complicated than these legal features might suggest.

There is, however, a second and better reason for doubting that we have a single idea of wrongful discrimination. This is that most countries that have laws prohibiting wrongful discrimination recognize *two* forms of wrongful discrimination. In the U.K., Canada, and Europe, they are referred to as “direct discrimination” and “indirect discrimination”; in the U.S., “disparate treatment” and “disparate impact.”¹⁵ Most countries’ laws do not specify precisely what moral significance this distinction is supposed to have; and it is not clear, either in the law or in our ordinary moral lives, what the precise boundaries of each concept are. So when we theorize about wrongful discrimination, we

¹⁴ Tarunabh Khaitan, *A Theory of Discrimination Law* (Oxford: Oxford University Press, 2015).

¹⁵ Canada is one of the few countries that does not treat the difference between these forms of discrimination as being of moral or legal significance, and I shall argue in Chapter 6 that this approach is largely correct. See, for example, *British Columbia (Public Service Employee Relations Commission) v. BCGSEU*, [1999] 3 SCR 3.

can take either of two approaches. We can begin, as some scholars do, by trying to develop our own more precise definition of each type of discrimination, a definition that we think is morally robust but that may differ from the legal definitions of these concepts.¹⁶ Alternatively, we can begin, as other scholars do, with the ideas that the law gives to us. We can ask what theory of wrongful discrimination, if any, might make sense of these ideas.¹⁷ I will be taking the latter approach. This is because, as I shall argue in more detail later in this chapter, our ideas of discrimination seem to me to owe so much to our legal frameworks that any theory of discrimination that does not start from the rough contours of some of the basic ideas about discrimination given to us by the law risks not being a theory of *discrimination* as we understand it, and so risks irrelevance to our public discourse and our moral lives.

So I shall, in this section of the chapter, introduce two rather rough definitions of direct and indirect discrimination gleaned from the law, and a rather rough idea of when they are wrongful. In later chapters, I will discuss both forms of discrimination in much more detail, as I develop my own account of the reasons why each sometimes wrongs us. But for now, my aim is simply to introduce the distinction, in broad brushstrokes, as we know it from anti-discrimination law.

Consider first what we call “direct discrimination.” It consists of an act or a practice that explicitly singles out a person or group that possesses a certain trait and treats them less favourably because of that trait. Many of the most commonly recognized and most seemingly outrageous acts of discrimination are instances of direct discrimination. Think back, for instance, to the first two examples of wrongful discrimination with which I began this book: the singling out of Muslims for harassment and violence and the denial of medical treatment to transgendered persons. In both cases, there is an act, or what I am calling a “practice”—a set of acts or a combination of acts and omissions directed at a certain end, that might be written out as a formal policy or rule or might just be generally understood as “what we do around here”—that treats a particular group less favourably, on the basis of a particular trait, than it would treat those who lack this trait. This is what I shall understand as:

Direct Discrimination: A practice directly discriminates against a person, P, if the practice treats P less favourably on the basis of some trait, *t*, than it would treat those who lack *t*, either by explicitly singling out people with *t* or by singling out those who have a different trait, *u*, that is in some way very closely connected to *t*

¹⁶ Scholars who take this approach include Kasper Lippert-Rasmussen, *Born Free and Equal: A Philosophical Inquiry into the Nature of Discrimination* (Oxford: Oxford University Press, 2014); Benjamin Eidelson, *Discrimination and Disrespect* (Oxford: Oxford University Press, 2015); and Deborah Hellman, *When is Discrimination Wrong?* (Cambridge, MA: Harvard University Press, 2013).

¹⁷ Scholars who take this approach include Khaitan, *A Theory of Discrimination Law*, *supra* note 14; and Gardner, “Discrimination: The Good, the Bad and the Wrongful,” *supra* note 3.

(for instance, only those who have t can have u , or many who have t have u and many who do not have t do not have u).

Anti-discrimination laws generally treat direct discrimination as wrongful when all of the following conditions obtain: (i) trait t is, and ought to be, a prohibited ground of discrimination; (ii) the agent is a government or government agency, or has taken on what we might call a public role, by being an employer, or a provider of goods and services to the public; and (iii) there are no relevant justifying factors. Some jurisdictions recognize no relevant justifying factors at all in cases of direct discrimination; whereas others allow for some.¹⁸ I shall be exploring the rationale for, and the justifiability of, each of these conditions, in subsequent chapters.

What about indirect discrimination? Two examples of it are provided by the remaining two cases of wrongful discrimination with which I began the book: the gender pay gap and the practices that result in indigenous communities lacking clean water when other communities in the same country have it. These are not instances of direct discrimination, where a certain person or group is explicitly singled out and treated differently on the basis of a trait that constitutes a prohibited ground of discrimination. The gender pay gap, for instance, is not for the most part caused by policies that assign women lower salaries directly on the basis of their gender. There are, of course, exceptions: in some industries, such as the tech industry, women are commonly offered lower starting salaries specifically because they are women. But for the most part, the gender wage gap is caused by promotions practices that deny women a fair chance of promotion into senior and more lucrative positions, and by educational practices that discourage women from entering more lucrative professions. And interestingly, these practices do not usually explicitly mention gender, nor do they disadvantage women directly because of their gender. The causal chain is longer, mediated by other things, and so the connection between these practices and gender is more difficult to spot. In fact, many practices disadvantage women for reasons that present themselves as specific to the individuals in question rather than as related to their gender. Keiko is perceived as “not aggressive enough” in negotiations; Medveh is seen as “too emotional”; Alice’s teachers think that she “isn’t intellectually suited for” a career in a STEM subject. Each of these, taken on its own, is the sort of assessment that one might also make about a man—indeed, that we do make about many men. It is only if we move from the individual case to consider the situation of women as a group that we are able to see such assessments as part of a pattern, in which stereotypes about women may be colouring people’s perceptions of what capacities particular women have and what roles they are capable of taking on.

¹⁸ See, for instance, the U.K. Equality Act, *supra* note 7, and Article 2 of the Racial Equality Directive binding the Court of Justice of the European Union, neither of which allow for the justification of direct discrimination; contrast with Canada’s various human rights statutes, which allow that the same factors can justify direct and indirect discrimination.

Similarly, the indigenous populations in Canada that lack clean drinking water on their reserves do not lack it because they have been explicitly singled out and denied piped or purified water because they are indigenous. Rather, although the Canadian government has the responsibility for providing 80% of the funding for water infrastructure on reserves, it has provided unpredictable and insufficient funding, and it has not investigated whether particular indigenous groups are actually able to make up the additional 20% required.¹⁹ Meanwhile, the federal government has regulated water management extensively in off-reserve contexts, and has collaborated with the provinces to ensure that high water quality is maintained off-reserves.²⁰

So both the practices surrounding women's promotion and education and the practices surrounding government funding of water on indigenous reserves amount to what we call:

Indirect Discrimination: A practice indirectly discriminates against a person, P, on the basis of trait *t*, if P has *t*, P is disadvantaged by the practice, and although the practice does not explicitly single P out because of *t* or some related trait, *u*, it nevertheless disproportionately disadvantages those who have *t* relative to those who do not.²¹

¹⁹ Human Rights Watch Report, "Make It Safe: Canada's Obligation to End the First Nations Water Crisis," 2016, available at: <https://www.hrw.org/report/2016/06/07/make-it-safe/canadas-obligation-end-first-nations-water-crisis>.

²⁰ Note a complexity here: the actual *provision* of drinking water for the rest of the Canadian population actually falls under provincial jurisdiction. But there is still much that the federal government is doing to facilitate clean water off reserves that it is not doing on reserves.

²¹ One might argue that, according to these definitions, any case of indirect discrimination can be re-described as a case of direct discrimination on the basis of some other trait. So, for instance, the kind of indirect discrimination on the basis of gender that I discuss in this paragraph might also be described as direct discrimination on the basis of "not being aggressive enough" or "being too emotional" or "not being suited to a STEM career." John Gardner claims this in "Discrimination: The Good, the Bad, and the Wrongful," *supra* note 3. However, I shall argue later in the book that this is not true of all cases of indirect discrimination. Some do not involve distinguishing between people on the basis of any trait; so they cannot be re-described as cases of direct discrimination. But more importantly, even for those that can be so re-described, the re-description is not morally helpful to us. For when cases of indirect discrimination are re-described as cases of direct discrimination based on other traits, we lose the reference to the particular trait that ought to be a prohibited ground of discrimination and that makes these into cases of *wrongful* discrimination; and relatedly, we lose the reference to what all of these cases *have in common*. So, for instance, if we look at the discrimination experienced by Keiko as "direct discrimination based upon not being aggressive enough" and the discrimination faced by Medveh as "direct discrimination based upon being too emotional" and the discrimination faced by Alice as "direct discrimination based upon not being suited to a STEM career," we risk losing our grasp on the fact these are all cases of gender-based discrimination, involving gender-based stereotypes that are closely related to each other. I shall argue in Chapter 6 that we can only explain why such cases are wrongful, and can only explain what they share as wrongs, if we look at them under the description of wrongful *indirect* discrimination.

Anti-discrimination laws, generally treat indirect discrimination as wrongful when all of the following conditions obtain: (i) trait *t* is, and ought to be, a prohibited ground of discrimination; (ii) the agent is a government or government agency, or has taken on what we might call a public role, by being an employer, or a provider of goods and services to the public; and (iii) there are no relevant justifying factors.²²

The main differences between our legal conceptions of wrongful direct and indirect discrimination, as I have laid them out here, are twofold. First, wrongful direct discrimination explicitly singles out a certain group or person using a trait that is and ought to be recognized as a prohibited ground of discrimination (or some trait that is closely connected to such a ground), whereas practices that discriminate indirectly do not. The latter are apparently neutral, seemingly applying the same apparently innocent criterion or criteria to everyone. But they nevertheless have a disproportionately disadvantageous effect on a group that shares a trait that ought to be recognized as a prohibited ground of discrimination. Second, many jurisdictions assume that it is harder, or even impossible, to justify direct discrimination, and that indirect discrimination is easier to justify. What moral difference there really is between these two forms of discrimination, and whether indirect discrimination ought indeed to be regarded as easier to justify, are questions I shall explore in Chapter Six.

Notice that the definitions of direct and indirect discrimination that I have given here leave open a crucial question. It is a question that our laws, too, leave unanswered. Why, in order for both direct and indirect discrimination to be wrongful, must the trait on the basis of which a person or group is treated less favourably or disadvantaged be the kind that is and ought to be recognized as a *prohibited ground* of discrimination? What is it that all of the traits that ought to be recognized as prohibited grounds of discrimination have in common, which makes it wrong to treat people less favourably or disadvantage them on the basis of such traits? Our answer to this question will vary depending on the particular reasons why discrimination is wrong in different cases, so I shall not attempt a single answer to it here. Rather, as the book progresses, we will discover three related answers to it, three explanations of the moral relevance of prohibited grounds of discrimination, which correspond to three reasons why discrimination wrongs people by failing to treat them as equals.

I have now laid out the rough ideas of direct and indirect discrimination that are present in many anti-discrimination laws, and the circumstances in which the law holds these two forms of discrimination to be wrongful. We have seen that these forms of discrimination involve treating someone with certain traits less favourably than others, or disadvantaging them, on the basis of a particular trait, where this trait is the kind of trait that ought to be a prohibited ground of discrimination. I have explained already that the

²² See, for instance, the *Equality Act*, *supra* note 7.

legal prohibitions on these two forms of discrimination are presented, within the relevant legislation, as ways of achieving equality, or equal status for all members of society. So if our concern is to figure out when and why it is wrongful to discriminate in either of these ways, we can helpfully do this by asking the question of inequality: When and why, in treating people differently on the basis of certain traits, does a practice fail to treat them as equals? Moreover, when we have an answer, or several answers, to this question, these answers should tell us which traits really ought to count as prohibited grounds of discrimination and which ought not.²³ So by asking the question of inequality, we can help fill in the gaps in our loose, legally derived definitions of direct and indirect discrimination.

When we treat the question of inequality as the question that a theory of discrimination must answer, we are also able to eliminate an apparent puzzle about direct and indirect discrimination. Our laws treat wrongful direct and indirect discrimination as though they were two forms of the same phenomenon, discrimination. Now, if the wrongful differentiation question is the right question for us to ask, then it can seem puzzling how direct and indirect discrimination could really be forms of the same thing, and wrong for the same kinds of reasons. Recall that the wrongful differentiation question asks why it is wrong to treat people differently “because of” certain traits. Many scholars who have asked this question have assumed that this “because” refers to a causal chain that extends from the practice or policy back *through the discriminator’s mind* to the particular mental processes that led him or her to factor in the trait that ought to constitute a prohibited ground of discrimination—the woman’s gender, the indigenous group’s race, and so on. It might work to interpret wrongful direct discrimination in this way. We have seen that wrongful direct discrimination “treats P less favourably on the basis of” a trait that ought to be a prohibited ground of discrimination. The phrase “on the basis of” could refer to a person’s mental states. (Though importantly, it need not: it could alternatively refer to the causal chain that runs from the practice to the discriminatee via a particular trait of hers, where the policy would not have disadvantaged her had she not possessed that trait.) Wrongful indirect discrimination, however, need not involve a causal chain that extends back to any objectionable mental states or processes in the discriminator. The relevant causal chain can extend simply from the practice to the individual or the group that has the trait in question, and it is often inferred from proof that the discriminatee is disadvantaged by the practice, that she has the trait, and that the practice disproportionately disadvantages the group who have that trait, relative to others who do not have it. So if the question we are asking is the wrongful differentiation question, and if, as some scholars assume, the best way to interpret “because of” in this question is in relation to the discriminator’s mental processes, then wrongful indirect discrimination will

²³ The process is of course a bit more complicated. We will need to start with at least a preliminary sense of the kinds of considerations that we commonly treat as prohibited grounds of discrimination. But the list can be treated as revisable: once we develop a particular explanation of why discrimination is wrong, we can then look back and see whether certain traits ought to be added to this list, and whether others, that are currently on it, really don’t belong.

not seem to be, as it were, the right kind of wrong. It will not seem to be the same kind of wrong that is captured by our question.

This conclusion is explicitly endorsed by some theorists of anti-discrimination law.²⁴ They start from the premise that only direct discrimination is properly thought of as “wrongful discrimination.” In their view, indirect discrimination is misdescribed: it is not really a form of wrongful discrimination at all. It is either wrong for very different reasons as direct discrimination, or it is not wrong at all. But while we may end up forced back upon this conclusion if our attempts to answer the question of inequality are unsuccessful, it seems an odd place to start. Why start from the position that, when the law identifies these two phenomena as forms of the same thing, it must be mistaken? By contrast, if we take the question of inequality as our starting point, it is more natural to think of the two forms of discrimination as continuous with each other. When we ask when and why, in disadvantaging people on the basis of certain traits, a practice fails to treat them as equals, we can see easily that we do not need to interpret “on the basis of” as referring to a causal chain that goes through a particular person’s mind. The causal chain might go through someone’s mind. But it might not. It might extend simply from the practice to the discriminatee and the group that shares the relevant trait with her. This alternative way of thinking of the phrase “on the basis of” gives us a unified interpretation of direct and indirect discrimination. It allows us to recognize that the two forms of discrimination are certainly in some respects different: the former explicitly singles out a person or group by means of a trait that ought to be a prohibited ground of discrimination, whereas the latter does not; the latter results in an at least initial disadvantage to the individual and group in question, whereas the former may not; and the causal chains in the case of direct discrimination will usually go through the mind of the discriminator, whereas in the case of indirect discrimination they may not. But when either form of discrimination is wrongful, it is wrongful for the same kinds of reasons: namely, because this particular way of treating people less favourably, or of disadvantaging them, is one that fails to treat these people as the equals of others.

1.4 Monism and Potential Problems with Monist Theories

I have now defined “direct discrimination” and “indirect discrimination,” and I have argued that pursuing the question of inequality will help us understand how both forms of discrimination could be wrongful, and wrongful for the same kinds of reasons. The theory of wrongful discrimination that I shall go on to develop in this book is not, of course, the only way of answering the question of inequality. Indeed, one reason why I think it is helpful to appeal to the question of inequality to orient us is that doing so enables us to see

²⁴ See Eidelson, *Discrimination and Disrespect*, *supra* note 16; and Hellman, *When is Discrimination Wrong?*, *supra* note 16.

different theories of wrongful discrimination as answers to the same question. Recently, a number of philosophers and legal scholars have developed theories of what makes discrimination wrong. Some of these theories have been presented by their authors as moral theories, theories of why discrimination is morally wrong.²⁵ Others have been presented as legal theories, theories that try to explain when and why discrimination is a “legal wrong,” not just in the positivist sense that there are laws that prohibit it, but in the Dworkinian sense that it is justifiably prohibited by law, at least in certain contexts.²⁶ Most of these theories can be understood as attempts to answer the question of inequality, to explain why, in disadvantaging a person or group on the basis of certain traits, we fail to treat them as equals. Some theories, which we might call *subordination theories*, appeal to the fact that discriminatory acts subordinate the discriminatee, either because they send a demeaning message about her, or because they lower her social status, or because the discriminator fails to give proper weight to the discriminatee’s needs and interests.²⁷ Other theories, which their own proponents have labelled *desert-prioritarian theories*, focus on the undeserved harms that discrimination causes, particularly to those social groups who are already underprivileged.²⁸ Still other theories, *freedom-based theories*, have foregrounded the fact that discriminatory policies deny some people equal freedom, perhaps by failing to guarantee them access to the conditions necessary for autonomy.²⁹ And *equality of opportunity theories* foreground the denial of equal opportunities to victims of discrimination.³⁰

One important feature shared by these recently developed theories is that they are *monist*. That is, each of them traces the wrongness of discrimination to some single further disvalue, which is supposed to explain why discrimination on the basis of any of the recognized prohibited grounds fails to treat people as equals. So, for instance, for Hellman, discriminatory acts are always wrong because they demean people, in the special sense that they both send the message that the discriminatee is inferior and also work to lower the discriminatee’s social status. For Khaitan, anti-discrimination law as a whole is justified by the general aim of protecting everyone’s equal freedom, of ensuring that everyone has

²⁵ See, for instance, Alexander, “What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies,” *supra* note 12; Arneson, “What is Wrongful Discrimination?,” *supra* note 12; and David Edmonds, *Caste Wars: A Philosophy of Discrimination* (New York: Routledge, 2006).

²⁶ See Elizabeth Anderson and Richard Pildes, “Expressive Theories of Law: A General Restatement,” *University of Pennsylvania Law Review* 148(5) (2000), pp. 1503–1575; Fiss, “Groups and the Equal Protection Clause,” *supra* note 4; Khaitan, *A Theory of Discrimination Law*, *supra* note 14; and Sophia Moreau, “What is Discrimination?,” *Philosophy & Public Affairs* 38(2) (2010), pp. 143–179.

²⁷ See Anderson and Pildes, “Expressive Theories of Law: A General Restatement,” *ibid.*; Eidelson, *Discrimination and Disrespect*, *supra* note 16; and Hellman, *When is Discrimination Wrong?*, *supra* note 16.

²⁸ See Arneson, “What is Wrongful Discrimination?,” *supra* note 12; and Lippert-Rasmussen, *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*, *supra* note 16.

²⁹ See Khaitan, *A Theory of Discrimination Law*, *supra* note 14.

³⁰ See Shlomi Segall, *Equality and Opportunity* (Oxford: Oxford University Press, 2013); and Joey Fishkin, *Bottlenecks: A New Theory of Equality of Opportunity* (Oxford: Oxford University Press, 2014).

relatively equal access to the goods necessary for an autonomous life. Particular acts of discrimination amount to moral and legal wrongs, on Khaitan's view, when they violate such laws.

As both of these examples show, to say that a theory is "monist" is not to suggest that it ignores social context, or that it does not appeal to different facts about discriminatory acts, in explaining when the relevant value is engaged. Hellman, for instance, has a very nuanced view of the many social conventions and the many kinds of social relationships which are relevant in determining whether an act demeans someone. Likewise, Khaitan has a list of a number of quite different conditions that, on his view, must be satisfied if a particular social group is to have equal access to freedom or autonomy—including negative freedom, self-respect, and a range of valuable options from which to choose. That many different kinds of facts that must be considered in determining whether an act of discrimination is wrong on either of these theories does not make the theories any less "monist" in my sense of the term. For it is still true that each of these theories traces the wrongness of discrimination to some single further disvalue—in Hellman's case, demeaning others, and in Khaitan's case, failing to provide the necessary conditions for autonomy. For both theories, this one disvalue explains, in all cases, why discrimination is wrong.

By contrast, a *pluralist* theory such as the one that I shall be defending does not trace the wrongness of discrimination to some single disvalue. Rather, it allows that discrimination can be wrong for different reasons, reasons that we can trace to fundamentally different kinds of problems with particular acts or practices. It may be that one and the same practice can be wrong for several of these different reasons. But the reasons are not reducible to some single sort of disvalue.

I want to pause here for a moment, to clear up a potential confusion. It might seem that any answer to what I have called "the question of inequality" will have to be a monist answer. After all, the question assumes that discrimination, when it is wrongful, always fails to treat a certain person or group as an equal. Isn't "failing to treat someone as an equal" a single kind of disvalue; and so won't any answer to the question of inequality be a monist answer?

No, it isn't; and no, it won't. Of course, any theory of wrongful discrimination that aims to answer the question of inequality will, at the highest or most abstract level, trace the wrongness of discrimination back to some kind of failure to treat others as equals. But, as the many philosophical debates about the value of equality in recent years attest, the bare idea of "failing to treat others as equals" is open to *many interpretations*. So any theory of discrimination will have to give this idea some moral content, and that content will need to be provided by some value beyond the bare idea of inequality. What makes a theory of wrongful discrimination "monist" or "pluralist" in my sense is whether it gives content to the idea of failing to treat others as equals by appealing to some single type of

inappropriate treatment (such as demeaning someone, on Hellman's view, or denying them the conditions for positive freedom, on Khaitan's) or whether it instead gives content to the idea of failing to treat others as equals by appealing to very different kinds of inappropriate treatment, as my pluralist theory does. So yes, it is true, and unsurprising, that any answer to the question of inequality will appeal to the value of equality. But this is not what determines whether the theory is monist or pluralist, because it is not the value of equality that explains *why* certain forms of treatment constitute failing to treat others as equals. What determines whether the theory is monist or pluralist is the kind of treatment that the theory invokes, to interpret the idea of equality—is it a single kind of treatment across all cases, or does the theory appeal to multiple sources of inappropriate treatment, which are irreducible to each other?

The idea that a helpful theory of wrongful discrimination must be monist and not pluralist is never explicitly defended by the proponents of recently developed monist theories. But it does seem to be presupposed by most recent theories. And I think it is a presupposition that we should question. Why should we assume that a successful theory of discrimination must be monist? Discrimination is, after all, a large and unwieldy moral concept. We have already seen some of its breadth, in laying out the definitions of direct and indirect discrimination. But it becomes even broader, and even messier, when we start to think about the variety of traits that constitute and ought to constitute prohibited grounds of discrimination: for instance, race, gender, gender identity, religion, creed, disability, and, more controversially, social condition or poverty, and physical appearance. These traits range from those that are deeply important to many of us and in some sense under our control, such as our religion, to traits that we cannot change at will and often view as impediments to our achievements, such as many disabilities. Some are features inherent in a person, such as her age. Others are in large part defined socially, such as race. Others, such as gender, have a biological and a social component.

Probably unsurprisingly, all of the recently developed monist theories seem capable of explaining wrongful discrimination only by disregarding or re-describing some of the complex features of discrimination. Each requires us to bracket some of the lived experiences of victims of discrimination and some of the goals of grassroots organizations fighting to eliminate discrimination, and each requires us to re-interpret or ignore certain key features of our anti-discrimination laws.

For instance, as we have seen, our legal concepts of wrongful direct and indirect discrimination extend only to disadvantage that occurs on the basis of certain traits, but not to all traits. And in virtually no legal system does this list of traits include poverty. This would seem rather inefficient, if the point of anti-discrimination law were to ensure that we bring about the most valuable outcomes consistent with respecting every one's desert, (where in assessing what is most valuable, we give priority to the interests of groups that are worst off). Nor do we tend to think that people's entitlements to non-discrimination

depend in any way on desert; in fact, moralizing judgments about whether someone deserves the harm that has befallen her seem to many scholars to be out of place within anti-discrimination law, one of whose aims is generally agreed to be that of allowing misunderstood, undervalued groups to find their own voice and portray themselves to us in their own light, without judgment from us.³¹

Freedom-based theories of discrimination, by contrast, seem to overlook a different aspect of our ordinary moral thought about discrimination. This is that our moral outrage in cases of discrimination seems to be at least in part an outrage at social subordination. Even if it is true that all victims of discrimination have had some freedom, or the necessary conditions for some freedom, denied to them, part of what seems wrong about discrimination is not just that it denies people these freedoms or their conditions, but the fact that it places some wrongly above others.

And although equal status theories have tried to capture this fact about discrimination, they too seem incomplete. They tend to offer more individualistic analyses of subordination, which focus too narrowly on what discriminatory acts express about the victim—without taking account of the broader social groups to which the discriminator and the discriminatee belong, the relationships between these groups, and the many very real ways in which the subordinated group is affected by a discriminatory act, beyond the effects of the message that this act sends. Moreover, it seems inaccurate to suggest that what victims of discrimination always care most about is eliminating subordination, and that the freedoms they fight so hard for matter only as a way of achieving equal social status. Those same-sex couples who want access to the institution of marriage don't just want an end to this particular kind of social subordination; though of course they do want this. They want to be free to define themselves and their union in their own way, without having to navigate around our assumptions about what people of their sexual orientation can and can't do. And they want access to the institution of marriage because they believe that it is only if they are publicly recognized as eligible to marry that they will be, and be seen as, true equals in their societies.

1.5 The Relevance of Victims' Experiences and the Relevance of the Law

One might at this point object that it's hardly obvious that a failure to accord with the experiences of victims or explain the basic features of anti-discrimination law is a deep flaw in a theory of discrimination. Our laws may presuppose a misguided picture of what

³¹ See Sophia Moreau, "Equality and Discrimination," *The Cambridge Companion to Philosophy of Law* (Cambridge: Cambridge University Press, forthcoming); Elizabeth Anderson, "What is the Point of Equality?," *Ethics* 109(2) (1999), pp. 287–337; and Elizabeth Anderson, "The Fundamental Disagreement between Luck Egalitarians and Relational Egalitarians," *Canadian Journal of Philosophy: Supplementary Volume on Justice and Equality* 36 (2010), pp. 1–23.

discrimination is or of why it is wrongful. Moreover, there may be good pragmatic reasons for shaping legal incentives and disincentives in ways that do not accurately reflect the structure of the moral wrong that they are trying to address.³² And why should we look to the lived experience of victims of discrimination for an indication of why it is wrongful? Allegations of wrongful discrimination arouse great resentment and bitterness in our society. This might give us all the more reason to pause and question whether the views of victims or the concerns of grassroots organizations are really the best places to look for dispassionate guidance on the nature of this wrong. So does the fact that these monist theories fail fully to accommodate the lived experiences of victims and fail to accord with certain features of our laws really count against them?

This is an important question for me to ask because its answer explains much of the methodology of this book. In my view, the fact that a particular theory of discrimination fails to accord with certain basic, shared features of anti-discrimination law does count against it. And this is because, whatever kind of injustice is involved in wrongful discrimination, it seems true that our moral understanding of it has been deeply shaped by our legal regimes for regulating it. In this respect, discrimination is arguably different from certain other moral wrongs, such as failing to keep promises or murdering. We could imagine developing a detailed and accurate conception of what a promise is even without consulting contract law, or a deep understanding of what murder is and why it is morally wrong without looking at the structure of any criminal prohibitions on murder. But it is arguable that our shared public views of what discrimination is and why it is wrong have, in large part, been shaped by domestic and international human rights law over the past fifty years. I am not sure how a moral account of wrongful discrimination could expect to be accurate if it did not explain the most basic of the shared features of anti-discrimination laws across different jurisdictions—such as the fact that most jurisdictions treat discrimination as wrong only if it occurs on the basis of a limited list of prohibited grounds, or the fact that most jurisdictions recognize some kind of distinction between direct and indirect discrimination. Of course, it is still open to proponents of a given theory to re-describe the function of a particular legal rule or legal distinction, arguing that in fact it serves some other more pragmatic purpose and does not in fact reflect any deep moral fact about what makes discrimination wrong. But if a theory fails to give us any good explanation of why certain widely shared features of these laws exist, or if a theory requires such a radical revision of our laws that the phenomenon it is describing bears very little resemblance to what we think of as discrimination, then this seems to me good

³² See John Gardner, “On the Ground of Her Sex(uality),” *Oxford Journal of Legal Studies* 18 (1998), pp. 167–187; John Gardner, “Liberals and Unlawful Discrimination,” *Oxford Journal of Legal Studies* 9 (1989), pp. 1–22; John Gardner, “Discrimination as Injustice,” *Oxford Journal of Legal Studies* 16(3) (1996), pp. 353–363; and Sophia Moreau, “Discrimination as Negligence,” *Canadian Journal of Philosophy: Supplementary Volume on Justice and Equality* 36 (2010), pp. 123–149.

grounds for doubting that it really is an account of the wrongness of this phenomenon that we collectively call “discrimination.”

This same interdependence of the legal and the moral in questions of discrimination seems to me also to necessitate giving somewhat greater credence than we might otherwise do to people’s lived experiences of discrimination, in assessing the adequacy of our theories. Our experience of discrimination is shaped by our society’s understanding of it, which is in turn deeply influenced by our legal prohibitions on it. And this suggests that victims of wrongful discrimination are unlikely to be wholly mistaken about the nature of their complaints. They might be mistaken about whether the facts that make discrimination wrongful really obtain in their case, or about the weight that their complaint carries, relative to other people’s interests, and hence, about whether the act that wrongs them is nevertheless justifiable, all things considered. But it is difficult for me to see how they could be wholly mistaken about the nature of their complaints –about what makes discrimination wrongful, in the first place. When victims decry discrimination for inappropriately subordinates certain people to others, or when they object that it generates obstacles to their freedom that they should not have to face, I do not think we can simply dismiss these claims as mistaken. We need, at the very least, to investigate whether there is a way of interpreting them according to which they can form part of a coherent and plausible theory of why discrimination is wrong. So, given the kind of moral phenomenon that discrimination is, it seems to me that the perceptions of victims of discrimination are appropriately treated as one kind of check on theories of discrimination.

For related reasons, this book does not, like many of its philosophical counterparts, use many hypothetical examples in an effort to provide more precise tests of the particular moral principles it puts forward. Rather, as far as possible, it uses real cases of discrimination—among them, cases that have been litigated and discussed by courts or tribunals, cases that have settled before reaching the courtrooms but that have been discussed by the media, cases that have been taken up by grassroots political organizations, who are fighting to have them recognized as genuine cases of discrimination. These cases are messy and sometimes difficult to think about. Yet I think it is very important to use such real cases when theorizing about discrimination. Part of the moral philosopher’s method in invoking these cases is to test proposed theories against our moral intuitions: would we really consider this instance of discrimination wrong for these kinds of reasons or those ones? When we try to test a theory of discrimination by appealing to happenings in fictitious societies, such as the story of the flat-earthians,³³ or to scenarios of discrimination that could not possibly occur, in which the same policy, for the same reason, discriminates directly against one gender and indirectly against another gender,³⁴ we

³³ See Niko Kolodny, “Rule Over None II: Social Equality and the Justification of Democracy,” *Philosophy & Public Affairs* 42(4) (2014), pp. 287–336.

³⁴ See Lippert-Rasmussen, *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*, *supra* note 16 at p. 187.

bracket the complex social contexts in which real acts of discrimination occur. And these social contexts are, I shall argue, the key to understanding discrimination. Most real cases of discrimination impose many different harms on their victims, and most result from the interaction of a variety of explicit and implicit policies, assumptions, and structures that together work tacitly to accommodate certain groups' needs while disadvantaging others, and that also work to stereotype certain groups in certain ways. If we are to see and analyze these processes, we need to do so by looking at real cases of discrimination, not by invoking hypothetical cases in which there is no extended social context to analyze.

An analogy might help here. One of the greatest innovations in European studies in biology occurred in the 18th century, when naturalists such as Gilbert White began observing live animals in their natural habitats. This may seem unremarkable to us now, but it was startlingly different from the prevailing scientific methods of the time, which involved dissecting specimens in the lab and making minute observations of their skeletal structure, their muscles, and so on. The problem with the dissection-based approach, of course, was that it offered no insights into how animals interacted with their environment. It helped to give scientists one view of what a bird was. But it could not show them how birds interacted with each other, how one bird's song differed from another, what the purpose of their songs was, or—as was crucial to later biologists such as Darwin—how they changed over time. It was only when naturalists began observing animals actually living in their environments that they were able to think of birds as complex, evolving animals whose behaviour and actual physical structure depended on their environment. I am suggesting that it is only if we take as our data the real cases of discrimination and the real responses of discriminatees in their full, rich social contexts, that we will be able, similarly, to have as full and accurate a picture of wrongful discrimination.

Moreover, there is a serious risk to trying to analyze discrimination in abstraction from its social context and without paying particularly close attention to the views of those who have experienced it. We risk frustrating an important aim of anti-discrimination law. This is to help us, as a society, give underprivileged groups a chance to have *their* voices heard—a chance to be considered for who they are rather than for who we think they are, and a chance to become the people who do the considering, who are in positions of power and who determine the agendas for our workplaces, the policies for our educational institutions, the values for our communities. Anti-discrimination laws do not aim only to improve the situation for such groups: they also aim to give these groups a voice in determining what an improvement for them might look like. If, in our academic discussions of discrimination, we set aside the real dilemmas faced by these groups and substitute our own more carefully crafted hypotheticals and our own more useful descriptions of people whom we think are like them, then we risk perpetuating both their silence and our own habits of not hearing them when they do speak.

1.6 Challenges Facing a Pluralist Theory of Discrimination

If a sound theory of wrongful discrimination must accord with the basic features of our laws and capture the experience of victims, and if monist theories fail to do this, then why have so many theorists felt the pressure to offer monist theories? As I noted earlier, this is a choice, but it is never one that is defended at any length. Why should we think that a sound theory of wrongful discrimination cannot be pluralist? There are some real challenges facing pluralist theories of concepts such as discrimination; but there are also some worries that seem to me spurious. And so it seems worth weeding out those challenges that are spurious, so that we can focus in later chapters on answering the real challenges.

First, one might think that only a monist theory of wrongful discrimination could be *coherent*. But we don't suppose that theories of other important political concepts must be monist if they are to be coherent. A coherent theory of justice can consist, as many liberal theories of justice do, in the conjunction of different but complementary principles, such as principles appealing to basic liberties and principles requiring some form of equal treatment or equal recognition, principles which most scholars do not try to trace back to some single further value, beyond suggesting that they are all "principles of justice."³⁵ Nor do we suppose that coherent accounts of particular "thick" moral virtues and vices must always be monist. No one would think it necessary, for instance, to give an account of cruelty that was monist. Like wrongful discrimination, some acts of cruelty are deliberately hurtful and some are perpetrated by agents who are best described as negligent. Like wrongful discrimination, cruel acts seem to be cruel both because of the kinds of harms they inflict on their victims and because of the kind of relationship that the cruel agent thereby sets up between himself and his victim. But there is no one further value that we feel obliged to invoke in all cases, in order to explain what makes an act cruel. Why then should we suppose that the wrongness of discrimination must be reducible to a single further value or single type of explanation, in order to be theoretically coherent—beyond the more abstract idea that such acts all treat some people differently on the basis of certain traits, in a way that fails to treat them as equals?

Perhaps the worry is that unless a theory of wrongful discrimination is monist, the kinds of disvalue that it invokes to explain the wrongness of discrimination, and to flesh out the idea of "failing to treat others as equals," will appear *arbitrary*.³⁶ But what is meant by "arbitrary" here? It might mean that we have no greater reason to appeal to *these* ones than to *others*. If so, then my earlier reflections on the need to take seriously the structure of

³⁵ Such theories include that of John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971). However, for a notable exception, see Ronald Dworkin, *Sovereign Virtue: Equality in Theory and Practice* (Cambridge, MA: Harvard University Press, 2000).

³⁶ See Lawrence Blum, "Racial and other Asymmetries: A Problem for the Protected Categories Framework for Anti-discrimination Thought," in Deborah Hellman and Sophia Moreau (eds.), *Philosophical Foundations of Discrimination Law* (Oxford: Oxford University Press, 2013), pp. 182–202.

anti-discrimination law and the complaints of real victims of discrimination seem to provide an answer to this worry. Surely a pluralist theory that adequately explains the law and accommodates our lived experiences of discrimination is not “arbitrary” in this sense. Both the law and the experiences of victims can provide us with good reasons to appeal to certain features of discriminatory acts and practices, and certain resultant forms of disvalue, as the relevant ones. So if I can show in later chapters that the kinds of disvalue to which I appeal help us to make sense of certain basic features of anti-discrimination law and of certain facts about the complaints of victims, I will have answered this version of the objection.

On the other hand, perhaps “arbitrary” means *unconnected* and *unexplained*. Perhaps the deeper concern motivating monists is that if we appeal to a number of different sorts of disvalue to explain why discrimination is wrongful and what it is, exactly, to fail to treat others as equals, then even if these different sorts of disvalue do capture complainants’ experiences and help us make sense of the basic features of our anti-discrimination laws, they will nevertheless still seem unconnected and unexplained unless we can tie them together by appealing to some other, more foundational value.³⁷ They will be a mere *list* of harmful effects of discrimination, rather than a theory of discrimination.

I think this is a real worry. And it is two-pronged. There is a worry about a lack of connection here, and a worry about a lack of explanatory power. I hope to show in later chapters that these worries can be satisfactorily addressed. I shall argue that the abstract idea of “failing to treat others as equals” is all that we need, in order to connect the different reasons why discrimination is wrong into a coherent whole. This is something they all share: they are all ways in which we can fail to treat others as equals, by disadvantaging them or treating them less favourably on the basis of certain traits. And, I shall argue, this is all we need, by way of a coherent connection between them. The worry about explanatory power is perhaps more difficult. I have already suggested that the idea of failing to treat others as equals is too general and abstract to do much explanatory work on its own. Most of the important explanatory work in my theory will be done by the various different sorts of disvalue that I invoke, as ways of fleshing out this more abstract idea. Will these explanations still seem too different, too disjointed, to be coherent? We can only answer this question once we see the different explanations, the different components of this pluralist theory.

1.7 *Structure of the Book*

³⁷ See, for instance, Patrick Shin, “Is There a Unitary Concept of Discrimination?”, in Deborah Hellman and Sophia Moreau (eds.), *Philosophical Foundations of Discrimination Law* (Oxford: Oxford University Press, 2013), pp. 163–181.

In the next three chapters of the book, I shall lay out the three components of my pluralist theory, the three main reasons why, on my view, discrimination is wrong. In subsequent chapters, I shall explain how these components fit together into a pluralist but coherent theory of wrongful discrimination; I shall discuss the relevance of the distinction between direct and indirect discrimination; and I shall return to the question of whether and why wrongful discrimination is sometimes nevertheless justified all things considered. I shall also consider the different obligations that are had by the state and by individuals to treat people as equals.

Chapter Two, “Unfair Subordination,” explores the idea that acts of discrimination are wrong because they subordinate some people to others. I argue here that if we are to understand why and how discrimination subordinates, we need to think of subordination as *social subordination*—that is, as something that happens to a person by virtue of her membership in a certain social group. And I develop a detailed account of what social subordination involves, and of how discrimination contributes to it. My account differs considerably from the accounts recently developed by proponents of equal status theories, who think of subordination in a much more individualized way.³⁸ I argue that these accounts of subordination focus too much on the isolated act of the discriminator and on the power of the discriminator over the discriminatee, without looking at the broader power differentials between the social groups to which each of them belongs. What we really need is an account of subordination that focuses on both parties and on the social groups to which they belong. Building on the work on subordination done by feminists and disability theorists, I try to offer such an account. I argue that it enables us to make sense of both direct and indirect discrimination as wrong because they subordinate; and I try to show that this conception of subordination helps us make sense of the role of prohibited grounds of discrimination, providing us with a compelling idea of what it is to differentiate between some people and others, in ways that fail to treat them as equals.

Chapter Three, “The Relevance of Deliberative Freedom,” begins by considering a number of recent legal cases of discrimination in which we cannot understand the concerns of the claimants unless we think of the wrongness of discrimination as extending beyond social subordination. I argue that a large part of the wrongness in these cases involves an infringement of a right to what I call the “deliberative freedom” of the claimant. I explore the nature of deliberative freedom and distinguish it from the perfectionist conception of freedom employed by Tarunabh Khaitan in his different freedom-based theory of discrimination.³⁹ I explain that both direct and indirect discrimination can deprive people of deliberative freedom in circumstances where they have a right to it. And I argue that, given the importance in our society of treating others as beings capable of autonomy,

³⁸ See Eidelson, *Discrimination and Disrespect*, *supra* note 16; and Hellman, *When is Discrimination Wrong?*, *supra* note 16.

³⁹ See Khaitan, *A Theory of Discrimination Law*, *supra* note 14.

infringing someone's right to deliberative freedom is a way of failing to treat them as an equal.

Chapter Four, "Access to Basic Goods," turns to a third way in which discriminatory practices can be wrong. Discriminatory practices do not just subordinate people and deny them deliberative freedom in circumstances where they have a right to it. They also leave some people without access to certain resources or to certain social institutions—goods that are "basic" in the sense that access to them is necessary if one is to participate fully and equally in a particular society.⁴⁰ Many of the most prominent recent political battles over discrimination have been battles over just such goods, precisely for the reason that access to them is basic in this sense. When same-sex couples fight for the legal right to marry, for instance, or women fight for the right to breastfeed in public, they are not just fighting for these particular opportunities. Rather, they feel that without these opportunities, they are not able to participate fully and equally in society. It might seem that the wrongness of discrimination in these cases must really lie in social subordination, or that deliberative freedom is really reducible to a kind of basic good. But I argue in Chapter Four that there is a distinctive kind of wrongness involved when discrimination leaves someone without access to a basic good, different from the wrongs explored in previous chapters.

In Chapter Five, "A Pluralist Answer to the Question of Inequality," I explain how the three different components of the theory fit together, and I attempt both to answer the challenges that a pluralist theory faces and to highlight the advantages of the theory. I draw on the arguments of the first four chapters to explain how each of the three features of discrimination that I have discussed is a way of failing to treat others as an equal. So the three components of the theory can be understood as parts of a single coherent account. But what does much of the moral work, in explaining why discrimination is wrong, are the three different ways of failing to treat others as an equal—namely, by contributing to their social subordination, by infringing their right to deliberative freedom, and by failing to give them access to a basic good. I argue that the theory is therefore genuinely pluralist, but nonetheless coherent. Chapter Five also addresses some difficult questions about the relationship between these different wrong-making features of discrimination. One of these questions is whether a discriminatory act could be wrong even if it lacked one or two of these three features. I argue that it could. However, I try to show that most wrongly discriminatory acts possess at least the first two features. That is, most subordinate the discriminatee, though some make a greater contribution to social subordination than others; and most infringe the discriminatee's right to deliberative freedom. Many also deny

⁴⁰ Note that to identify a good as "basic" in this sense is not to claim that it is *objectively good*; nor is it to claim that access to it is necessary in *all* societies, as a precondition for full social and political participation. Rather, goods that are "basic" in my sense of the term are simply the goods access to which is necessary for full social and political participation in a particular society. Access to a certain good might be necessary for full participation in one society but not in another; and it might be necessary for full participation in a particular society even if it is not plausibly thought of as objectively good.

an already underprivileged group access to certain important goods; though not all wrongly discriminatory acts do so.

Another question that I explore in Chapter Five concerns the different kinds of reasons that each of these features gives us to avoid or rectify discrimination. I argue that denials of a right to deliberative freedom and denials of a basic good can constitute personal wrongs toward the victims of discrimination—that is, wrongs that give that victim a personal claim for rectification against the discriminator, and that place the discriminator under a corresponding duty to rectify the wrong himself. By contrast, it is not clear that all contributions to social subordination amount to personal wrongs toward the victim. Some may be wrong, but may only generate a duty to take steps to give members of the subordinated group greater opportunities in the future. If I am right about this, then it has the significant implication that it matters very much not just that we ascertain *whether* a particular act of discrimination is wrong or wrongful, but that we ascertain *why*. For it is only when we know the particular source or sources of the wrongness in a given case that we will be able to determine which kinds of duties the discriminator stands under. Finally, Chapter Five presents a number of advantages of my pluralist theory, arguing that it both explains and helps us resolve certain persistent controversies about wrongful discrimination.

In Chapter Six, “Indirect Discrimination,” I make explicit and draw together the implications that my theory has for our understanding of indirect discrimination. For interestingly, this pluralist theory suggests that in certain ways, the distinction between direct and indirect discrimination is helpful, but that it does not carry the moral weight that we often think. The two forms of discrimination work in different ways to subordinate people, so whether a particular practice directly or indirectly discriminates can be relevant, if the wrong in question involves contributing to social subordination. But if the wrong involves a denial of deliberative freedom or leaving someone without access to a basic good, then it may not make a difference whether a practice discriminates directly or indirectly. I then discuss certain revisionist implications of the theory. For certain legal regimes, such as the U.K.’s Equality Act, suggest that direct discrimination can never be justified, whereas indirect discrimination can. By contrast, my theory implies that both direct and indirect discrimination may, in some cases, wrong people but nevertheless be all things considered justifiable. Direct and indirect discrimination are on a par in this respect: both can wrong people by failing to treat them as equals, and both can sometimes be justified even though they wrong people. I also argue that it is a mistake to think that there is a difference in the moral responsibility of agents of wrongful indirect discrimination and agents of wrongful direct discrimination. I distinguish between responsibility in the sense of “responsibility for costs” and responsibility in the sense of “culpability.” I argue that in neither sense of “responsible” is it true that agents of wrongful direct discrimination are somehow more responsible than are agents of wrongful indirect discrimination. And I suggest that in many cases of wrongful direct discrimination and

many cases of wrongful indirect discrimination, we can see the agents as culpable for a form of *negligence*.

Chapter Seven, “The Duty to Treat Others as Equals: Who Stands Under It?” turns to a series of questions that are not often addressed in the philosophical literature on discrimination, concerning the different obligations of the state, on the one hand, and individuals, on the other, to treat others as equals. It is quite plausible to suppose that the state always stands under a duty to treat those whom it governs as equals—but there are different ways of understanding what grounds this duty, and I explore some of them. I then turn to questions about the obligations of individuals. Some have argued that, as individuals, we do not normally stand under a duty to treat others as equals, and that we acquire such a duty only when we step into certain institutional roles –the role of employer, for instance, or provider of goods or services or accommodation to the public. I try to show that this view is problematic, and I defend a different view: namely, that we *always* have an obligation to treat others as equals, but that there are often good reasons, in the context of personal decision-making about our friends and our families, for the state not recognize a legal duty corresponding to this obligation, and not to impose sanctions on its violation. I try to show that my view enables us nevertheless to recognize each person’s interests in freedom of association and contract; and I argue that we need to think further about the ways in which the state can, *without* using the machinery of anti-discrimination law, nevertheless assist individuals in complying with our duty to treat others as equals and, ultimately, in creating a society of equals.

Chapter Two

Unfair Subordination

2.1 Unfair Subordination: A Plausible Starting Point

An obvious place to start, in answering what I have called “the question of inequality,” is with the concept of unfair subordination. When we treat some people differently on the basis of certain traits, either through practices that explicitly single them out or through practices that indirectly disadvantage them on the basis of these traits, we often seem to be failing to treat them as equals in the sense that we are *unfairly subordinating* them, putting them beneath others, marking them as inferior to others. Indeed, the kinds of discrimination that usually give rise to the greatest moral indignation involve the creation or perpetuation of different classes of people, with some having a superior status and others an inferior one, in circumstances where we think that everyone ought to have an equal status. Consider, for instance, the Jim Crow laws, which turned African-Americans into second-class citizens; or dress codes for waitresses or female retail employees that mark them out as sexual objects, lacking the full and independent agency that we ascribe to men.

The idea that unfair subordination might help to explain the wrongness of certain forms of discrimination does not only have a hold on our moral imaginations. It is also deeply rooted in the law. Both the United States’ Fourteenth Amendment and the constitutional equality rights in section 15 of the *Canadian Charter of Rights and Freedoms* have been understood—by courts, and also by academics—as prohibiting government policies that unfairly subordinate people based on certain traits.¹ And of course, when private sector anti-discrimination law was first developed in these two countries in the 1960s and 1970s, it was treated as a form of quasi-criminal law that aimed to eliminate acts of prejudicial subordination, acts that deliberately denied certain privileges or benefits to members of certain social groups on the grounds that these groups were less worthy

¹ For seminal discussions of unjust subordination in the context of the United States’ Fourteenth Amendment jurisprudence, see Owen Fiss, “Groups and the Equal Protection Clause,” *Philosophy & Public Affairs* 5(2) (1976), pp. 107–177; Ruth Colker, “Anti-Subordination Above All: Sex, Race, and Equal Protection,” *New York University Law Review* 61 (1986), pp. 1003–1044; Mari J. Matsuda, “Voices of America: Accent, Antidiscrimination Law, and a Jurisprudence for the Last Reconstruction,” *Yale Law Journal* 100(5) (1991), pp. 1329–1407; J. M. Balkin, “The Constitution of Status,” *Yale Law Journal* 106(8) (1997), pp. 2313–2374; and Jack M. Balkin & Reva B. Siegel, “The American Civil Rights Tradition: Anti-classification or Anti-subordination?,” *University of Miami Law Review* 58 (2003), pp. 9–33. Canadian courts adjudicating equality rights cases tend to speak of “stereotyping” rather than of subordination, but their discussions of stereotyping are often discussions about subordination. See, for instance, Colleen Sheppard, *Inclusive Equality* (Montreal: McGill-Queen’s University Press, 2010).

than others.² But what exactly does unfair subordination involve? How do some discriminatory acts and policies work to subordinate certain social groups unfairly, and to sustain this subordination? In this chapter, I shall try to answer these questions.

Rather than referring throughout to “unfair subordination,” I shall often refer, simply, to “subordination.” This is not because I am assuming that subordination in every form is always unfair. All of what I shall go on to say is quite consistent with recognizing that the subordination of certain individuals in certain isolated contexts is permissible, and even necessary. It surely must be true that, at least in a certain sense of “subordinate,” a soldier is rightfully subordinate to his commanding officer, just as a company’s employees are rightfully subordinate to its managers. But these particular cases of subordination seem rightful only because the differences in status hold only within a particular organization, rather than across a number of different social contexts; and they also seem fair only to the extent that the differences in power, authority, and deference between the superior and the subordinate are in fact justified by the demands of these individuals’ roles, relative to each other. A commanding officer could not organize his unit unless he had the power to make his soldiers obey him. But he cannot rightfully make female soldiers walk ten paces behind their male counterparts.³ Similarly, a store manager cannot manage her employees unless she has the power to give directions, assess performance, and impose sanctions on work that is poorly done, and her employees do not have these powers. But she does not rightfully have the power to dictate how her employees behave in their spare time, outside of the office. So certain forms of subordination may be justified within a particular organization, when they are necessary for the operation of that organization. But subordination is often unfair when it is based, not upon the demands of a particular social role within a particular organization, but upon perceptions about certain personal (or allegedly personal) traits, and when the group of people that are alleged to have that trait are subordinated, not just within a particular organization, but across a variety of different social contexts.

Most of the legal scholars who analyze discrimination in terms of its contribution to subordination invoke a relatively under-specified, intuitive idea of subordination. For instance, when Owen Fiss first urged that the U.S. Equal Protection Clause was best interpreted not as preventing arbitrary classifications but as eliminating unfair subordination, he suggested that subordination was a “status harm” that involved perpetuating the lower social position of persistently disadvantaged social groups.⁴ But it was not his aim to develop a general account of what that status harm involved. More recently, Reva Siegel and Joel Balkin have examined the ways in which courts, in cases of discrimination, are motivated by concerns with “social stratification” and “the secondary

² See Denise Réaume, “Harm and Fault in Discrimination Law,” *Theoretical Inquiries in Law* 2 (2001), pp. 349–385.

³ An example from the U.S. military cited in Valerie Vojdik, “The Invisibility of Gender in War,” *Duke Journal Of Gender Law & Policy* Volume 9:261 2002, p. 262.

⁴ Fiss, “Groups and the Equal Protection Clause,” *supra* note 1 at p. 157.

social status of historically oppressed groups.”⁵ But they do not explain in detail what “social stratification” involves. Their interest is largely in charting how concerns about social stratification motivate judges to adopt particular legal doctrines and to decide certain cases in certain ways. My aim in this chapter is different. I want here to develop a more detailed account of subordination, one that can help us understand some of the ways in which different forms of discrimination subordinate people and some of the reasons why they might be wrong, in virtue of contributing to such subordination. I shall not try to specify a set of individually necessary and jointly sufficient conditions for social subordination. But I shall, in Section 4 of this chapter, lay out in some detail four conditions that seem to be satisfied in most cases where the subordination of a certain group persists across different social contexts for some extended period of time and seems unjust. I shall then turn in Section 5 to considering how discrimination contributes to subordination, and I shall argue that it is plausible to think that this is one of the reasons why discrimination is wrong, when it is.

The scholars of anti-discrimination law who come closest to offering a philosophical analysis of subordination and of how discrimination subordinates are expressivists Deborah Hellman, Elizabeth Anderson, and Richard Pildes.⁶ They have argued that an act is wrongfully discriminatory when it subordinates a person to others in the sense that the act demeans her, or sends the message that she is of less value than others. I shall consider Hellman’s account in Section 3 of this chapter. While this account is extremely helpful in drawing our attention to the ways in which discriminatory acts send messages about the inferior status of certain groups, I shall suggest that, at least in its current form, the account is too individualistic. It focuses too much on the individual power dynamic between the discriminator and the discriminatee, when in fact we need to look at the relative amounts of power possessed by the different social groups to which these people belong. I shall also argue that expressivism offers us too narrow an understanding of subordination. Subordination is not only a function of the social messages sent by particular acts or policies. Rather, it is kept in place by a variety of effects that discriminatory acts have on different social groups, such as perpetuating differences in power and authority between them and rendering certain social groups or their needs invisible in certain contexts.

In addition to focusing on social groups, my own account of social subordination places special emphasis on one feature of it that is not often foregrounded. This is the fact that social subordination often depends, for its persistence, on what I shall call “structural accommodations.” These are policies, practices, and physical structures that tacitly accommodate a more privileged group’s needs at the expense of the subordinate group or groups. Normally, within anti-discrimination law, we use the term “accommodation” to

⁵ Balkin and Siegel, “The American Civil Rights Tradition: Anti-classification or Anti-subordination?”, *supra* note 1 at p. 9.

⁶ Elizabeth Anderson and Richard Pildes, “Expressive Theories of Law: A General Restatement,” *University of Pennsylvania Law Review* 148(5) (2000), pp. 1503–1576, especially at pp. 1533–1544; and Deborah Hellman, *When Is Discrimination Wrong?* (Cambridge, MA: Harvard University Press, 2008).

refer to a special measure that must be adopted in order to give the subordinated group an opportunity equal to that of the more privileged group. And we assume that the subordinate group requires an accommodation because that group has certain special needs. So, for instance, when a Muslim employee requests an altered work schedule so that he can pray at the times that his religion requires, we treat the altered schedule as an “accommodation” to which he is entitled, because of the special demands of his religion. As we have learned from feminists, critical race theorists, and disability theorists, however, at least part of the reason why these groups require an accommodation in the first place is that our social environment has been constructed in such a way as tacitly to accommodate the needs of more privileged groups.⁷ I shall be proposing that we need to think of *these* prior policies, practices, and structures as “accommodations”—accommodations to the more privileged social groups, which make their interests and needs seem normal and the interests of other groups seem exceptional. And I shall argue that we cannot understand the subordination of one group by another, or the real contribution of discrimination to subordination, unless we consider these “structural accommodations.” For they serve indirectly to rationalize the greater power and *de facto* authority that are held by these groups and the greater deference we pay to them. We can, I shall argue, see many of the policies that constitute wrongful indirect discrimination as “structural accommodations” that contribute to unjust subordination. So my account of subordination will give us a way of explaining in detail not just how direct discrimination subordinates, but how indirect discrimination, too, subordinates.

2.2 Discrimination that Subordinates: Restaurant Dress Codes

Before I turn to accounts of unfair subordination, I want to lay out several examples of discrimination that seem to be wrong at least in part because they involve unfair subordination. I shall focus on discrimination involving restaurant dress codes for employees. I shall lay out a number of examples in this section and shall return to them at various points in the chapter.⁸

Many restaurants and bars impose a gendered dress code on their employees. Rather than providing a gender-neutral set of options and allowing each employee to choose what suits them, these restaurants require female employees to wear tighter fitting clothing designed to show the shape of their bodies, low cut tops, shorter skirts, and high

⁷ See the works cited in note 16, *infra*.

⁸ The examples discussed in this section are taken from a recent report by the Ontario Human Rights Commission on common practices surrounding restaurant dress codes and their impact on a variety of underprivileged groups. See *Not on the Menu: Ontario Human Rights Commission Inquiry Report on Sexualized and Gender-based Dress Codes in Restaurants* (March 2017), available at <http://www.ohrc.on.ca/en/not-menu-ohrc-inquiry-report-sexualized-and-gender-based-dress-codes-restaurants>.

heels. Men, by contrast, are usually permitted to wear more comfortable, looser fitting clothing that is not revealing. If you have thought about such dress codes, you have probably already reflected on the messages they send about the appropriate social roles of men and women; the ways in which they mark women out as inferior; and the additional physical and health burdens they place on women, both through the tight clothes that restrict their movement and through the high heels that cause foot and back pain. They are standard examples of direct discrimination, since they explicitly single out women for a different and less comfortable uniform, and thereby on the basis of their gender.

The dress codes for employees adopted by many restaurants also indirectly disadvantage transgendered people, pregnant women, and members of certain religions, though none of these groups is explicitly singled out by the policies. Trans employees may have to pigeon-hole themselves into the uniform for one gender or the other, even if this does not express their current gender identity. Pregnant women are often disproportionately burdened because at some point their pregnancy makes it impossible for them to fit into a tight-fitting uniform, and they then face the difficulty of having to tell their employer sooner than they would have liked. Similarly, members of religious groups that require particular modes of attire may find it difficult or even impossible to adopt the dress codes while practising their religion, and so might be unable to work in this industry. So, in addition to constituting direct discrimination against women as a group, the dress codes constitute indirect discrimination.

There are also two related practices worth examining, in connection with restaurant dress codes; for these practices constitute an even more subtle form of indirect discrimination than those I have already examined. First, many restaurants do not stock any uniforms of a kind that might be easily put on and worn by people with muscular disabilities (disabilities that make it difficult for them to put on or wear tight clothing, for instance, or to do up all of the buttons on button-up shirts). And this means that such people, when hired, are placed in the difficult position of being unable to put on their uniform without assistance, or of having to step forward and ask for a different uniform and so present themselves as “abnormal.” Second, even restaurants that have officially adopted gender-neutral, disability-friendly dress codes often hand new employees a training manual that has pictures only of young, svelte, conventionally attractive women dressed in the most feminine uniform options. Insofar as such manuals reinforce the image of people as useful and employable only insofar as their physical appearance is “normal,” and insofar as they suggest that it is part of a woman’s role as a waitress to use her body to gratify men, they disproportionately burden women, pregnant women, people with disabilities, and trans people.

Let us now consider what expressivist theories of discrimination would say about such practices, looking in particular at Deborah Hellman’s theory of discrimination as wrongful because it demeans.

2.3 Hellman's Expressivist View: Subordinating by Demeaning

Hellman has argued that discrimination is wrong when and because it puts someone down, treating them as though they are “not fully human” or “not as worthy as others.”⁹ She uses the term “demean” to refer to the kind of subordination that she has in mind. One might think that to demean someone is simply to act in a way that sends the social message that another person is less worthy of respect. And this is certainly how traditional expressivists have understood it. But Hellman uses “demean” in a special way, to refer to a sub-group of those acts that send the message that someone is less worthy of respect than others. For she notes, quite rightly, that not all acts that send such messages actually do have the capacity to affect the social status of others. For instance, she suggests, when an employee spits at her boss or a child taunts her classmate, their acts send an inferiorizing social message but do not normally have the capacity to lower the status of the person insulted.¹⁰ In Hellman's view, in order for a discriminatory act actually to constitute objectionable subordination, two things must be true of it. First, it must send an inferiorizing message about someone. Second, it must have the power to put that person down.¹¹ Whether it has this power depends in large part, she suggests, on the discriminator's status relative to the discriminatee: as she says, “in most instances, demeaning also requires that the speaker hold a higher status than the person demeaned.”¹² She does not state exactly what kind of “status” she has in mind here. But in her analysis of her examples, she speaks mostly of status that is given to us by our roles in particular social institutions –for instance, the higher status of an employer relative to an employee, the higher status of a parent relative to a child. So the employee cannot put her boss down in Hellman's example because she lacks the power to do so, given her lower status in that company; the child cannot put her classmate down, given that they have an equal status in the classroom, as classmates.

It seems right that many acts of discrimination that subordinate do so in part because of the social messages they send—messages about the inferiority of some groups or the superiority of others. And it also seems right that, in order to subordinate someone, one must be capable of affecting their status in the world, rather than simply expressing something about them. But there are at least two respects in which the expressivist account seems to me incomplete.

First, I do not think Hellman places enough emphasis on the importance of looking at the social groups to which people belong and the relationships between these groups, when assessing whether someone has the capacity to put another person down. In her

⁹ Hellman, *When Is Discrimination Wrong?*, *supra* note 6 at p. 35.

¹⁰ Hellman, *When Is Discrimination Wrong?*, *ibid.* at pp. 35–36.

¹¹ Hellman, *When Is Discrimination Wrong?*, *ibid.* at p. 38.

¹² Hellman, *When Is Discrimination Wrong?*, *ibid.* at p. 36.

examples, she focuses a great deal on the institutional role that the discriminator has, relative to the discriminate, and on the status and power that this institutional role gives the discriminator over the discriminate: such as the teacher's power over the student, and the employer's power over the employee. But whether one person's expression of disrespect for another has the power to put that other person down seems often to depend, more broadly, on the power and authority that the *social groups* to which they belong have, relative to each other. Consider again Hellman's examples of spitting at one's boss or insulting a classmate. These acts could sometimes put the person down. If I am a white employee and you are my African-American boss and I spit at you in full view of all of my other fellow white employees, I can indeed lower your social status. I can't, of course, affect your employment status—that is, your status as my boss. But I can certainly lower your social status by spitting at you, even if I am just your employee. I can do that precisely because, in our version of the example, I am white and you are black, and whites and blacks in our society have a certain history, relative to each other, which my act invokes. Similarly, if the child insulted by a classmate in Hellman's example is indigenous and he is called a "drunk Indian" by a white child, the act can lower his status in the classroom, or perpetuate his already low status. That is because the group to which the white child belongs has historically possessed, and continues to possess, a great deal of power and authority over the ways in which indigenous children are portrayed in our society. So, in order to assess whether someone's act puts down another, we need to look, not just, and not primarily, at the institutional roles of these two people relative to each other and the powers that these roles confer on each of them, but also to the social groups to which each belongs and the relationship of these groups to each other.

I think Hellman is aware of the relevance of the power relations between different social groups; but I do not think she places nearly enough emphasis on it. It is not a central feature of her discussions. She does qualify her examples by noting that her analyses would be different if there were "unequal status or hierarchy" in the classroom, or if there were "unusual circumstances" at work.¹³ But by leaving these as brief qualifications applicable to exceptional situations, she suggests that situations in which power relations between groups are relevant as unusual or special. As I hope the rest of this chapter will show, such power dynamics characterize most of the social situations in which we find ourselves. They are not unusual or special. On the contrary, they are what define a person's "status," in the sense relevant to subordination. They should be considered in all cases, not just in some.

But there is a second problem I see in Hellman's view. It is a more serious problem for the view, since it concerns, not a query about emphasis, but a problem with the expressivist core of the view. As I shall now argue, it is unclear that it is only the social message sent by a particular act or policy that determines whether it unjustly subordinates.

¹³ Hellman, *When Is Discrimination Wrong?*, *ibid.* at pp. 35–36.

To see this, let us look at some of the common policies concerning restaurant dress codes that I mentioned in Section 2.

First, consider the explicitly gendered dress codes adopted by many bars and restaurants, that require female employees to dress in sexualized and revealing ways, while permitting men to wear more comfortable, non-revealing clothing. It is true that such dress codes send a message about women needing to appear in a sexualized way so as to please male clients, and that, given the overall context, this message really does have the effect of confirming women's already inferior status. But such dress codes also seem to do many other things to women –and importantly, they do not seem to do these things because of the social message they send. Requiring women to dress in tight clothing and heels hampers their ability to move, thereby giving them less power in the workplace than their male colleagues and male clients. So such dress codes lessen women's power, quite independently of the social message they send about women. Such dress codes also give women less authority over their own self-presentation than men have --and again, they do this not because they send a message about women's inferiority, but because they deny women a choice that they give to their male colleagues. These dress codes also present a certain conception of how women ought to dress as preferable to others conceptions, thereby elevating the women who dress this way above those women who do not, and so creating a hierarchy *within* the class of women --a hierarchy that in turn helps to perpetuate many women's subservience to many men's attitudes about them. And these dress codes render invisible those women who have figures that do not conform to this image. Although we do not yet, at this stage in the paper, have a detailed conception of subordination to help confirm the relevance of these facts to women's subordination, it seems implausible to suppose that all of these facts are irrelevant to unfair subordination. But the expressivist account gives us no way of recognizing the relevance of such facts, except insofar as they can be reinterpreted as aspects of, or effects of, the social message of inferiority sent by the discriminatory policy.

Now consider several ways in which such gendered dress codes also seem to discriminate indirectly, both against women and against other groups such as trans people, and members of certain religious minorities. Trans employees may have to pigeon-hole themselves into the uniform for one gender or the other, even if this does not express their current gender identity. Pregnant women are often disproportionately burdened because at some point their pregnancy makes it impossible for them to fit into a tight-fitting uniform, and they then face the difficulty of having to tell their employer sooner than they would have liked. Similarly, members of religious groups that require particular modes of attire may find it difficult or even impossible to adopt the dress codes while practising their religion, and so might be unable to work in this industry. But it is unclear that these dress codes send a message about the inferiority or lack of worth of these groups: they seem rather simply to overlook their situations. So it is unclear that an expressivist could recognize this as discrimination of a sort that wrongfully subordinates.

Hellman would likely reply that in her view, direct and indirect discrimination are two different kinds of wrongs. Direct discrimination demeans, and so wrongfully subordinates. By contrast, indirect discrimination does not usually demean but is wrong for some other reason—for instance, as Hellman has recently suggested, because it compounds past injustices.¹⁴ Hellman might argue that compounding a past injustice is a way of contributing to subordination; it is just different from demeaning someone. However, if both of these forms of discrimination seem to confirm or perpetuate the lower status of certain social groups, then, rather than drawing a bright line between those that demean and those that do not, perhaps we should see whether there is a single account of subordination that can allow us to explain, in rich and detailed ways, how different forms of discrimination work to create or confirm a person's lower status. And, as I argued earlier, we should try to do so in a way that clearly acknowledges the role of the social groups to which that person and the discriminator belong. In the next section of the paper I shall try to develop such an account.¹⁵

2.4 *Towards an Account of Social Subordination*

What we need is an account of subordination that considers the broader relationship between the social group (or groups) to which the discriminatees belong and the social group to which the discriminator belongs. We need an account of what I shall call *social subordination*—that is, the state of affairs in which one social group has a standing in society as a whole that is lower than that of another social group.

As a preliminary to developing such an account, I should clarify what I mean by a “social group.” As I am using the term, a “social group” is an entity that has an existence apart from any particular member: one can speak about the group without reference to those who happen to be its current members. A social group shares, or is presumed to share, a certain trait. But it is not just any group of people who happen to share a certain

¹⁴ Deborah Hellman, “Indirect Discrimination and the Duty to Avoid Compounding Injustice,” in Hugh Collins and Tarunabh Khaitan (eds.), *Foundations of Indirect Discrimination Law* (Oxford: Hart Publishing, 2018), Ch. 5 at pp. 105–121.

¹⁵ My account owes much to the work of Niko Kolodny in “Rule Over None II: Social Equality and the Justification of Democracy,” *Philosophy & Public Affairs* 42(4) (2014), pp. 287–336. I also draw on the insights of philosophers who write on subordination within certain specific contexts. For instance, Rae Langton and Catharine MacKinnon have developed theories of subordination within discussions of pornography: Rae Langton, *Sexual Solipsism* (Oxford: Oxford University Press, 2009) and Catharine MacKinnon, *Only Words* (Cambridge, MA: Harvard University Press, 1993). And relational egalitarians such as Elizabeth Anderson and Samuel Scheffler have given analyses of the kinds of non-subordinating relationships that are required in order for us to live in a true society of equals—that is, a society in which everyone stands, to every other person, in a relationship of equality. See Elizabeth Anderson, “What is the Point of Equality?,” *Ethics* 109(2) (1999), pp. 287–337 and “The Fundamental Disagreement between Luck Egalitarians and Relational Egalitarians,” *Canadian Journal of Philosophy: Supplementary Volume on Justice and Equality* 36 (2010), pp. 1–23; and Samuel Scheffler, “The Practice of Equality,” in Carina Fourie, Fabian Schuppert, and Ivo Wallimann-Helmer (eds.), *Social Equality: On What It Means to Be Equals* (Oxford: Oxford University Press, 2015), pp. 21–44.

trait, such as “people with bushy eyebrows,” that count as a social group in the sense that I am concerned with. Rather, the kinds of social groups that usually have a lower status across a number of different social contexts either possess, or are presumed by others to possess, a particular type of trait, a trait that is *socially salient*, in the sense that others in society take that trait to have implications for the character and behaviour of members of the group, and for the social roles that they are capable of occupying. And it is often by this socially salient trait, or a combination of it and other traits associated with a sub-group of those who possess this trait, that we identify members of that social group. So, for instance, “wearers of blue, crimson, or scarlet velvet” does not, in our current society, mark out a group of people through their shared possession of any socially salient trait. But in Tudor England, where Sumptuary Laws regulated the materials and colours that people from each social stratum could wear, it marked out the group of noblemen who stood at or above the level of a Knight of the Garter. This was an important social marker in those times; and it was a moral marker of sorts too, marking out people believed to be of superior moral fibre. So this group of people would, at that time and in that place, have counted as a “social group” in my sense.

On this understanding of a social group, all of the groups that are marked out by the sorts of traits that our laws commonly treat as prohibited grounds of discrimination constitute “social groups”—for instance, women, Jews, Haidas, and people with hearing-impairments. But notice both that we need not assume that members of a social group, so defined, identify closely with each other or that their well-being is in some way bound up with their group identity. Nor should we assume that social groups are homogenous, either in the aspirations of their members, or in their needs or abilities. It may be that one sub-group within a particular social group is affected quite differently from another by a given policy. For instance, to return to our gendered dress code example, these codes impose one set of burdens on non-religious women, and an additional set of burdens on those women whose religions require them to dress in ways incompatible with the dress codes. So if we are to understand how such codes subordinate women, we may need to look at a variety of different sub-groups within this broader social group.

We can now go on to consider what it is for a social group, so understood, to be subordinated to others. What does this involve? First, in most situations of social subordination, members of the subordinated social group have less power than members of other groups, across a variety of social contexts. Not just less political power, but less social power as well, and not just less power in the sense of a diminished capacity to do certain things on their own, but also less power in the sense of a diminished capacity to compel others to do what they want them to do. There are, of course, difficult questions here about how we are to conceive of power—whether it is relational or can be conceptualized as a kind of resource that could be distributed; whether it makes sense to analyze how much power particular individuals have or whether it must be analyzed

structurally and systemically. But I do not think that my argument requires me to take a stance on these questions, so I shall leave them open.

Subordinated social groups also generally have less *de facto* authority than others, across a variety of social contexts. Having *de facto* authority over others is different from having power over them.¹⁶ In order to get you into the place that I want you to be in, it is enough for me to have the power to move you there: all I need is a large and strong enough army of helpers and a means of confining you. But I can only get you to do what I want of your own volition if I have *de facto* authority over you. So *de facto* authority includes the power to get you, of your own volition, to obey me.

It is common within political philosophy to think of *de facto* authority largely in terms of the power to secure others' obedience. But within the context of social subordination, I think it is important for us to think of *de facto* authority as involving a broader set of powers, including the power *to be listened to*, and to be taken seriously when one brings a complaint against another. One fascinating effect of the many recent successful complaints of sexual harassment against prominent film producers such as Harvey Weinstein and prominent actors such as Bill Cosby is that it has made us collectively aware of one kind of authority that women and many racial minorities have lacked. For before the recent successful complaints of harassment, some actresses did try, unsuccessfully, to bring complaints against these same people. And they were not believed. As women, they lacked the authority to speak and to be assumed to be telling the truth, and were too often assumed to have been overreacting in an emotional way or misinterpreting the meaning of men's actions. In fact, women in the film and theatre industries are often in something uncomfortably analogous to the position of women in cultures in which the testimony of two to three women is legally required in order to equal the weight of the testimony of one man. Their voices simply do not carry the same credibility as a man's. So it is important not to have too narrow a conception of what *de facto* authority involves, when we think of the kind of authority that subordinated social groups lack. They do not only lack the power to have other people to obey them when they issue orders. Before one is even in a position to get other people to obey, one needs to have authority in the sense that other people are ready to listen, ready to assume that you are telling the truth rather than overreacting emotionally or misreading other people's actions. And that prior authority is the kind of authority that members of subordinated groups quite often lack.

Niko Kolodny has helpfully described a further feature of social subordination using the term "consideration."¹⁷ In situations of social subordination, Kolodny argues, the group that possesses more power and *de facto* authority may be ascribed certain attributes or personal traits that, within that particular society, attract positive responses of deference

¹⁶ Kolodny emphasizes the importance of power and authority in understanding subordination, in "Rule Over None II: Social Equality and the Justification of Democracy," *ibid.*

¹⁷ Kolodny, "Rule Over None II: Social Equality and the Justification of Democracy," *ibid.* at pp. 296–97.

and respect. Importantly, these responses are directed, not just at these attributes or traits, but at the people who possess them: people with certain features are more likely to be shown greater deference and respect, and their interests are likely to be given greater priority, even in situations where they ought to be weighed equally with those of others. Moreover, the traits themselves are not just regarded as pleasing or as important (as, for instance, is athleticism in some circles, or intellectual acuity, in others), but as traits that mark people out as in some sense belonging to a higher or better class of people. So when one shows deference to someone on the basis of such a trait, or excludes another person because she lacks it, one is contributing to a pattern of responses that mark some people out as higher, or lower, than others.

It seems right, and deeply insightful, that in many cases of subordination, certain traits attract greater deference of just this sort. Though I think it's important to add that subordinated groups aren't just perceived to *lack* such traits. Often, the subordinated group is defined in terms of a corresponding trait that comes to be regarded as worthy of *censure*, because it has been identified with patterns of action or dispositions of behaviour that are perceived as worthless, or worse, as vices. For instance, Muslims living in the United States at the moment don't just suffer from a lack of deference or consideration, based on perceptions of their religion. Rather, this trait—their religion—is in certain social and political circles regarded as a sign that they are likely either to be terrorists or to be connected with terrorists or, at the very least, to be unpredictable religious extremists. So the trait “being Muslim” functions in certain social circles to mark people out as deserving of condemnation and ostracism. When we think of subordination, then, we should think not just of the absence of consideration towards the disempowered group, but of the use of corresponding traits to condemn, publicly humiliate, or ostracize this group. I shall use the term “censure” to refer to these negative public attitudes.

Kolodny does not investigate how some of the traits ascribed to certain social groups come to attract greater consideration than the corresponding traits of other groups, because he does not need to investigate this for the purposes of his own argument. But if we are trying to develop a picture of social subordination that will help us understand how discrimination subordinates, I think it matters very much that we examine how certain traits come to attract greater consideration. For of course, certain races and religions don't randomly or arbitrarily attract greater consideration, while the corresponding traits of others happen to attract censure. Rather, particular traits come to be associated with dispositions to behave in certain ways, with certain talents or lack of talents, and with certain social roles. And it is through this association that the traits come to acquire greater consideration or greater censure. So, for instance, as I suggested above, in certain social and political circles in the United States, the Muslim religion has come to be associated with religious extremism and with a propensity to engage in terrorist activity. These associations of certain traits with particular dispositions, patterns of behaviour, and roles are what we commonly call “stereotypes.”

Stereotypes, as I understand them here, are generalizations about particular social groups that ascribe most of their members certain desires, dispositions of behaviour, or obligations, simply because they possess whatever trait defines that group, as a group: Muslims are thought more likely to be religious extremists, simply by virtue of being Muslim; women are held to be under an obligation to beautify themselves, because that's what women are for.¹⁸ Some of these generalizations may be false; others may be true, and may still contribute to subordination. What is important about them, for the purposes of subordination, is that they serve to rationalize the differences in the power and *de facto* authority given to the groups marked out by these traits, and the differences in the consideration and censure they attract. By "rationalize" I do not mean that they actually justify them, but rather that they constitute the kind of proposed justification that is plausible enough that many people in fact accept them. And in some cases, stereotypes seem to work by making us think that *there is no need to justify* certain ways of treating others: stereotypes make the connection between a certain trait and a certain social role or a certain kind of treatment seem so obvious that we feel we do not need to give any reasons for placing someone with that trait in that social role or for treating her in a certain way. So, whether by rationalizing certain acts or by apparently obviating the need to justify them in the first place, stereotypes play an important role in the persistence of disparities of consideration and censure, and in the perpetuation of unequal power and *de facto* authority.¹⁹

If we were just to stop here—thinking of subordination in terms of disparities in the power and *de facto* authority held by certain social groups, and in the degree of consideration or censure they attract, based on certain traits—we would omit an important fact about social subordination. Differences in power and *de facto* authority are not only held in place by habits of conscious or explicit consideration or censure or by the stereotypes that support such consideration and censure. Perhaps even more importantly—because more silently, and more insidiously—they are kept in place by apparently neutral policies, practices, and physical structures that privilege the interests of the dominant group, while overlooking those of the subordinate group. Particular such

¹⁸ I have learned much about stereotypes from Rebecca Cook and Simone Cusack's pioneering discussions of stereotypes and stereotyping in *Gender Stereotyping: Transnational Legal Perspectives* (Philadelphia: University of Pennsylvania Press, 2010). For my own analysis of the role of stereotypes in cases of discrimination, see Sophia Moreau, "Equality Rights and Stereotypes," in David Dyzenhaus and Malcolm Thorburn (eds.), *Philosophical Foundations of Constitutional Law* (Oxford: Oxford University Press, 2016), pp. 283–304.

¹⁹ Must stereotypes be false, to contribute to subordination? I think not. Some of the gender stereotypes that contribute to women's subordination, for instance, may be statistically true; but they are presented or used as though they are true as a matter of biology, when in fact they reflect the ways in which women in certain communities are socialized and they limit the opportunities open to women and the careers that they are able easily to enter and to conceive of as possible for themselves. Similarly, "statistical discrimination"—that is, differential treatment of different groups that results from using group averages to determine which policies are rational—is often based on statistically sound generalizations. What is problematic in these cases isn't the truth value of the generalizations. It's the fact that using the generalizations in these ways perpetuates differences in the power and authority enjoyed by these different groups, results in undue deference being given to some and undue censure to others, and renders certain groups invisible in certain contexts.

structures have been examined by legal scholars working on indirect discrimination, by feminists such as Rae Langton and Catharine MacKinnon working on pornography, and by critical race theorists and disability theorists trying to expose the ways in which apparently neutral policies and political concepts work to perpetuate the privileged status of certain groups and the disadvantaged status of others.²⁰ But no one has, to my knowledge, developed a general philosophical theory of subordination across these different contexts that gives a place to these structures; and there is, quite strikingly, no general scholarly term for them. I shall call them “structural accommodations.” This term is intended to highlight two important facts about them. First, unlike consideration and censure, they are not attitudes or dispositions of behaviour of the discriminator or the public at large: they are real *structures* in our social and physical environment. In some cases, as we shall see, they are literally physical structures. In other cases, they are structures in the sense that they are policies and practices that structure our workplaces, our homes, and our shared social environment. And second, they work by tacitly *accommodating* the needs or interests of one group, and overlooking those of others—with the result that the needs of the dominant group come to seem normal and natural, whereas the different needs of the subordinate group come to seem exceptional and even odd.

Consider first a very literal example of structural accommodation: certain standard features of the buildings in which we live and work. Most houses have a short flight of steps leading up to the front door, and most storefronts facing onto commercial streets standardly have a single step leading up to the door. This easily accommodates those of us arriving on foot, but poses obvious difficulties for people in wheelchairs or for those with certain muscular difficulties. Light switches are standardly placed four feet above the ground, and bathroom mirrors at a similar height—again, perfectly within reach of many adults, but out of reach, and out of sight, for people in wheelchairs. Tobin Siebers, a disability rights theorist, has written quite movingly about the ways in which such structures not only exclude people with certain disabilities from these spaces, but implicitly send a message about the normal human body who is expected to reside there and the normal guest or client who is welcome there.²¹ Our houses and our stores presuppose a certain kind of human body and tacitly invite inside those who share such a body, while not issuing invitations to those who do not share it. This ideal or normal body remains invisible until someone with a non-standard body appears. When that happens, we might add, it can look as though it is the person with the disability who requires some

²⁰ See, for instance, Kimberlé Crenshaw, “Race, Reform and Retrenchment: Transformation and Legitimation in Antidiscrimination Law,” *Harvard Law Review* 101(7) (1988), pp. 1331–1387. See also Langton, *Sexual Solipsism*, *supra* note 15; MacKinnon, *Only Words*, *supra* note 15; Catharine MacKinnon, *Towards a Feminist Theory of the State* (Cambridge, MA: Harvard University Press, 1991); and Catharine MacKinnon, “Difference and Dominance: On Sex Discrimination,” in Ann Cudd and Robin Andreasen (eds.), *Feminist Theory: A Philosophical Anthology* (Oxford: Blackwell, 2005), pp. 392–402.

²¹ Tobin Siebers, “Disability Studies and the Future of Identity Politics,” in Linda Alcoff, Michael Hames-Garcia, Satya Mohanty, and Paula Moya (eds.), *Identity Politics Reconsidered* (London: Palgrave Macmillan, 2006), pp. 10–30. See also Tobin Siebers, *Disability Theory* (Ann Arbor: University of Michigan Press, 2008).

“special” accommodation. But this is only because houses and storefronts have already been built in such a way as to accommodate the needs of the rest of us.

Siebers’ point is not that the construction of houses and stores is an act of deliberate exclusion or deliberate deference to certain body types and censure of others. He recognizes, quite rightly, that houses and storefronts are constructed this way because this answers to the needs of the majority of the adult population. For this reason, I don’t think this particular structural accommodation can be accurately re-described as just another form of “consideration” in Kolodny’s sense. It is not an instance of conscious deference or respect for people without disabilities. It reflects a quite neutral, pragmatic effort to build in a way that is efficient and in demand. But it contributes significantly to the subordination of those with disabilities—by making it physically impossible for them to enter into certain buildings, by thereby making it more difficult for them to enter into certain social and commercial relationships, and by making their bodies seem invisible and unnatural.²²

Consider next the example of public washrooms that are segregated by gender, with washrooms for men bearing a large sign on the door that represents a man in trousers, and washrooms for women bearing a large sign on the door that represents a woman in a skirt. We are now aware of the ways in which such signs and practices marginalize transgendered persons and place them at greater risk of being taunted or bullied. But at the time when many such washrooms were built, most of the people commissioning them thought that it was normal and natural to segregate people in this way. This was not intended as an expression of respect or deference for people whose body fits their gender identity, or of censure for those who are not in this position; yet it has had the effect of normalizing the divide, and of rendering invisible those who do not fit on one side of it or the other. This is another example of what I am calling a “structural accommodation”—in this case, a feature of our built environment that accommodates the needs of the majority and constitutes them as normal, while overlooking the needs of a less privileged social group.

²² One might object here that the exclusion of people with disabilities from traditional buildings is not a true case of *discrimination*, so this is not a helpful example. In American law, the “failure to accommodate” certain disabilities or religions is treated as something distinct from wrongful discrimination—both are prohibited, but there are separate bodies of law that govern them. This distinction between failures to accommodate and wrongful discrimination has, however, been criticized. Legal scholars such as Samuel Bagenstos, Christine Jolls, and Sharon Rabin-Margalioth have argued that the obligations imposed on us by accommodation requirements are no different in kind or degree of onerousness from those imposed on us by anti-discrimination law. See Samuel Bagenstos, “Rational Discrimination, Accommodation and the Politics of (Disability) Civil Rights,” *Virginia Law Review* 89 (2003), pp. 825–923; Christine Jolls, “Antidiscrimination and Accommodation,” *Harvard Law Review* 115 (2001), pp. 642–699; 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), pp. 111–152.

I have given two quite literal and physical examples of structural accommodations. But “structural accommodations” in my sense need not actually be physical structures, and they need not function to exclude the subordinated group in quite such a literal way. Think of the many policies in your own academic faculty or department that accommodate the needs of the average male junior faculty member, while posing some obstacles for young female junior faculty members. A tenure clock that runs out four or five years after one’s first appointment is perhaps a good idea for someone whose wife can bear their children; but if you have to bear them yourself, and you have to do so within these particular five years because you are getting older, it is more difficult. In some departments, faculty meetings run from 4–6pm, which means that a woman who has children, and who is responsible for picking them up from day-care, has no choice but to exit the meeting early, in full view of her colleagues, who know exactly where she is going and who sometimes view it as a sign that she isn’t able to be fully attentive to her work. When we invite guest speakers to give a talk, we often take them for drinks (and if you were Ronald Dworkin, you would take them for raw oysters). There is a tacit expectation that each of us will have a drink or consume a few raw oysters, partly out of collegiality and partly to demonstrate our sophistication. This poses a dilemma for those women who are, or are trying to become, pregnant. They may not want to partake, and may not want to disclose why; yet if they don’t partake and don’t offer some explanation, they appear at best less than collegial, and at worst provincial.

None of these practices is designed to disadvantage women. They do not seem well described as expressions of greater consideration or deference towards men or censure of women. They just happen to accommodate the needs of men who either have no children or have a partner who can bear and take care of them, because this particular social group formed the majority of faculty members at the time that these practices were developed. So they are, in my sense, “structural accommodations”—features of our environment that tacitly accommodate the needs of certain groups, while also normalizing them and rendering the more marginalized group invisible or seemingly exceptional. I should also add that, as the day-care pick-up example shows, the needs that are accommodated do not need to be natural or biological needs: they can be needs that arise because of the social burdens that are placed on one group or another, the way women tend to bear more of the burdens of taking children to and from childcare.

I hope I have given enough examples to explain why, in my view, states of social subordination need to be thought of not just as involving differences in power and *de facto* authority and lesser consideration or censure, across a variety of different social contexts, but also as involving a variety of structural accommodations that both deny certain opportunities and resources to the subordinated group and serve to render their different needs invisible or abnormal. Because they serve this normalizing function, structural accommodations seem to me to stand in a special, supportive relationship to the other features of subordination. Because they help to constitute the needs of the superior group

as normal and natural, they serve indirectly to rationalize the differences in power and *de facto* authority between these groups and those that are subordinate to them, and also indirectly to provide further support for the various expressions of deference and consideration that are given to these groups in other contexts. If, as the gender-segregation and labelling of public toilets implies, it is normal and natural to be born one gender or the other and to have the gender identity that corresponds to the body you were born with, then those who don't have this are unnatural—and perhaps they don't deserve the kind of consideration given to the rest of us. If it is normal and natural for a smart, high-powered academic to produce a book within their first few years, then it looks as though women who can't manage this aren't capable enough to hold power and don't deserve as much deference.

There is another reason why structural accommodations help to rationalize differences in power, *de facto* authority, and consideration or censure. This is that they, just like the patterns of consideration and censure we examined earlier, are bound up with stereotypes about the subordinated group. One such stereotype is that when a woman has young children, she becomes unable to focus on anything except her children; whereas when a man has young children, he is able properly to compartmentalize them and remain a serious scholar. Because of this stereotype, the structural accommodation of holding meetings from 4–6pm has particularly serious effects on women—because it colours our interpretation of what members of the subordinated group are trying to do, when they try to work around this particular accommodation. When a woman walks out of a departmental meeting at 5:50pm, she is not just a scholar leaving the meeting early, as her male colleague might be seen to do. She is much more likely to be seen as a mother abandoning her work for her children; and this in turn is often taken as evidence that she must not really have been fully focused on her work, even during the time when she was at the meeting. So the structural accommodation and the stereotype work together to paint her action in a particular light, to reinforce the stereotype, and to rationalize the differences in power and *de facto* authority that put the subordinated group in this position to begin with.

I have suggested that structural accommodations are bound up with stereotypes, with differences in power and *de facto* authority, and with practices of censuring certain groups, or giving them less consideration, on the basis of certain traits. But it is worth noting that it is quite possible for a structural accommodation to be innocuous, if it is unconnected with these other features of subordination. So there is nothing inherently objectionable in structural accommodations per se. They become implicated in unjust subordination only because, and only to the extent that, they are bound up with stereotypes, differences in power and *de facto* authority, and practices of assigning censure and lesser consideration to certain social groups.

To see this, consider one structural accommodation: the fact that most stores are open during daylight hours and close at night, rather than being open all night and closed during the day. This is convenient for the majority of us, who are awake during the day and who sleep for some part of the night. But it adds hardship to the lives of those employees who work night shifts: if they need to make purchases at stores other than all night convenience stores, or if they wish to shop together with friends, they have to disrupt their normal sleep time, which is during the day. And this disruption likely affects their bodies more than it would ours, given that their natural sleep rhythms are already disrupted. So the disadvantage they suffer as a result of this policy is more than trivial. It seems also worth noting that this is a group that is already disadvantaged, since night shift work increases one's risk of suffering from a host of health problems, such as high blood pressure and metabolic syndrome. So we have here a structural accommodation that imposes more than a trivial disadvantage on an already disadvantaged group. Nevertheless, in certain societies, this structural accommodation would not seem problematic. Suppose that the only people who worked night shifts in a particular society were people in relatively prestigious professions: emergency physicians and nurses at hospitals, lawyers who burned the candle at both ends, judges who were on call all night. The fact that most stores were only open during the day would not then perpetuate practices of censure towards, or lesser deference towards, these night-shift workers; nor would it support stereotypes about them being less able to handle regular work or less well educated, nor would it perpetuate differences in social or political power or *de facto* authority between this group and day workers. In fact, it might even have positive effects on how others viewed the members of this group by adding to the mystique and aura surrounding them: these professionals somehow still manage to get their groceries purchased even though most stores aren't open while they are at work! By contrast, in a society such as our own, in which many night jobs involve menial labour, require little education, and have much less prestige attached to them (jobs such as janitorial work, cleaning, garbage collecting, security enforcement) and tend overwhelmingly to be held by immigrants who are already mistrusted and misunderstood, the shared practice of only opening stores during daylight hours starts to look more problematic. In order to be implicated in social subordination, then, structural accommodations need to be supported by, and in turn perpetuate, stereotypes, habits of censure and consideration, and differences in power and *de facto* authority between different social groups.

Because structural accommodations, like the differences in power and authority possessed by different social groups, can be innocuous or justified, there is an important difference between these features of subordination and the expressions of consideration or censure. Consideration and censure involve taking the praise or criticism that is due to a certain *trait* and transferring it to the *person* in a variety of other contexts. So they are always unjustified. Structural accommodations, by contrast, and differences in power and *de facto* authority, and even stereotypes, may sometimes be innocuous. They become

problematic only when they work together to consign certain social groups to a lesser status in society.

Thus far, I have laid out a number of common and morally relevant features of social subordination. I have argued that one social group is unjustly subordinated to another when:

- (i) The members of that group have, across a number of social contexts, less relative social and political power and less relative *de facto* authority than the other group, and
- (ii) The members of that group have, or are ascribed, traits that attract less consideration or greater censure than the corresponding traits of the empowered group, and
- (iii) These traits are the subject of stereotypes, which help to rationalize the differences in power and *de facto* authority, the habits of consideration and censure, and the structural accommodations, and
- (iv) There are structural accommodations in place in society that tacitly accommodate the needs of a dominant group while overlooking the needs of at least some members of the subordinate group; and these accommodations work together with stereotypes to rationalize the differences in power and *de facto* authority, and the differences in consideration or censure.

I offer these four conditions as a set of common and morally salient features of situations involving the unjust subordination of one social group by another, features that, as we will see in the next section, are relevant in understanding when and why discrimination subordinates. All four conditions will often be satisfied when one social group is subordinated for some substantial period of time. But I do not think we need to suppose that they are individually necessary conditions, nor that they are jointly sufficient conditions. A complete philosophical account of unjust subordination might require additional stipulations. And it seems quite possible that, at certain early stages in the subordination of a particular social group, some of these conditions could be satisfied but not others, even though over time all four will likely be satisfied.²³ But it does not follow

²³ For instance, it may be that a certain structural accommodation, initially innocuous, comes gradually to support stereotypes about a certain group that rationalize excluding them from certain prestigious professions, and that over time these stereotypes, combined with the persistence of the structural accommodation, lead in turn to expressions of censure of this group. Or it might be the case that expressions of censure, without any accompanying structural accommodations, lead certain groups to become regarded as so inferior that they are effectively invisible in certain social circles or certain areas of life; and that this in turn nurtures stereotypes about them, and leads to structural accommodations that privilege the needs of others and fail to consider the needs and capacities of this group. In both of these examples, although all four conditions do come to be satisfied over time, there is an interim period in which one or two are not satisfied. During some of that time, we might want to say that there was no unjust subordination. But we might well

that these features of social subordination are unimportant. They are present in most cases of ongoing subordination. And as I shall go on to argue in the next section, they give us a good basis for understanding how discrimination can sometimes subordinate particular individuals and groups, and why it is wrongful when it does. They also help us understand the differences between the ways in which direct discrimination subordinates and the ways in which indirect discrimination subordinates.

2.5 How Direct and Indirect Discrimination Subordinate

Now that we have an account of what social subordination involves, we can go on to consider how direct and indirect discrimination might contribute to such subordination.

Let us look first at direct discrimination. Recall that, according to our earlier definition, a policy directly discriminates against a person, P, if the policy treats P less favourably on the basis of some trait, *t*, than it would treat those who lacked *t*. And recall that, as we noted in Chapter One, policies that directly discriminate in a wrongful way either explicitly single out people with a certain trait that is a prohibited ground of discrimination, or single them out on the basis of some trait that is very closely connected to such a trait. In order to see how such policies can unfairly subordinate certain social groups, it helps to note an important fact about the prohibited grounds of discrimination. We hold, both as a matter of law and in our own moral thought, that not just any trait can constitute a prohibited ground of discrimination. Rather, those traits that are justifiably treated as prohibited grounds—race, gender, sexual orientation, and religion, for instance—are traits on the basis of which at least one, and often quite a number of social groups have been denied equal power and *de facto* authority over others; have been subjected to greater censure or lesser consideration, in the sense that they have been condemned or thought of as less worthy of respect than others; have been stereotyped; and have had their needs overlooked by certain structural accommodations that cater to the needs and circumstances of more powerful social groups. To say this is not to claim that in any particular case of direct discrimination, the use of such a trait or its proxy will necessarily perpetuate all of conditions (i) through (iv). But it is highly likely to perpetuate a number of them, given the past history of these traits and the social uses to which they have been put.

Consider, as an example, the Jim Crow laws briefly mentioned at the start of this chapter, which left blacks in the United States with separate and inferior schools, hospitals, prisons, washrooms, seating areas in public transit, and even water fountains. These laws

conclude that there was unjust subordination of the group for some of the time, even in the absence of any censure of them, or even in the absence of structural accommodations excluding them.

used the trait “black” in order to accord blacks less power and *de facto* authority, and they used it in such a way as to ascribe to blacks a variety of undesirable traits, because of their alleged blackness—laziness, stupidity, incivility, uncleanness, and so on. So they helped to perpetuate the disparities of power and *de facto* authority between blacks and whites and the stereotypes that held such disparities in place, and they thereby helped to rationalize the many structural accommodations that privileged the needs of whites over blacks.

Our account of social subordination, then, helps us to understand that policies that are directly discriminatory against groups marked out by a prohibited ground of discrimination can play an important causal role in sustaining the four conditions of unjust subordination.

But there is a second way in which direct discrimination can subordinate. It can also *constitute* an expression of censure, of the kind mentioned in condition (ii), a statement that a particular group is inferior and can justifiably be treated as inferior. During the Jim Crow era, even water fountains were segregated. The signs above white fountains read “Drinking fountain: Whites only.” The sign above fountains for blacks read: “Drinking fountain: Colored.” These signs did not just function to tell people where to drink, nor did the water fountains just provide water. Perhaps more importantly, they marked out “Colored” as the inferior group. They did so partly because the term “only” was attached only to the sign for “Whites,” implying that no one would want to drink from the fountain for “Coloreds” if they were eligible to drink from the “White” fountain. But they also did so through their association with stereotypes such as “Coloured people are unclean” and through their association with the many other separate and inferior public facilities which this group was assigned.

Our account of subordination, then, allows us to conceptualize two rather different ways in which direct discrimination can subordinate a certain social group. It can (a) play an important causal role in sustaining some or all of the four conditions of subordination. But in addition, it can (b) constitute an expression of censure of the subordinated group, or an expression of lack of deference towards them, a way that marks out this group as inferior.

Consider, as another example, the gender-specific dress codes I discussed earlier. Recall that these codes explicitly prevent women from wearing certain allegedly “male” uniform options, and require them instead to wear tight, body-fitting, and revealing clothing. This perpetuates the stereotype that women are sexual objects without independent agency, and that part of their function, not just as waitresses but as women, is to be beautiful in the eyes of men. Such dress codes thereby mark women out as inferior. They imply that men have independent agency and need to dress as such, but that women need to dress in such a way as to please men. Unlike the segregated water fountain example, the gendered dress codes seems problematic less because they involve censure

and more because they give lesser authority and lesser consideration to women. But this is still a case of one social group being branded or stigmatized as inferior to another.

What about cases of wrongful indirect discrimination? As I indicated in Chapter One, such practices do not specifically single out a person or group because of some trait that amounts to a prohibited ground of discrimination; but they do disproportionately disadvantage those who have a trait that amounts to a prohibited ground of discrimination, relative to those who do not have this trait. Indirect discrimination can seem puzzling, and its moral status unclear, partly because it is less easily interpreted as the kind of expression of censure or denial of equal consideration that is involved in direct discrimination. But my account of subordination has the resources to explain why indirect discrimination, too, can subordinate people. For my account of subordination focuses not just on expressions of censure or lesser consideration, but also on the “structural accommodations” in condition (iv) that work tacitly to disadvantage groups marked out by certain traits, and on the stereotypes that rationalize these accommodations and seem to rationalize our not looking for viable alternatives. As I shall now explain, many instances of wrongful indirect discrimination can be seen as structural accommodations—and moreover, as the kinds of structural accommodations that are problematically bound up with stereotypes, differences in power and *de facto* authority, and practices of censure and lesser consideration of subordinated groups.

As an initial example of wrongful indirect discrimination, consider the cases involving tests for promotion within a certain occupation, such as tests for firefighters or police—tests that do not draw any explicit distinctions along racial lines, but are failed in far greater proportions by blacks and Latinos than by whites.²⁴ In some cases of this type, the differential results are due to prejudicial grading or “buddy systems” and networks of nepotism within the profession that give whites an edge. These cases look rather more like direct discrimination. So let us consider those cases in which only the tests themselves are responsible for the difference: the test questions use situations and analogies and bits of information that, in a particular community, whites are more likely to have encountered already. This is still, I take it, an example of the kind of indirect discrimination or disparate impact that many would find wrongful. My account of subordination allows us to explain why. These tests are an instance of a “structural accommodation” inadvertently given to white employees. They privilege the interests and knowledge of whites over those of blacks, and even though they do so completely unintentionally and without malice, they nevertheless serve to perpetuate differences in power and *de facto* authority, and they work together with stereotypes about blacks (they are so lazy that these results must be accurate; they couldn’t be competent enough to do well on these tests anyway) to rationalize the persistence of these structural accommodations.

²⁴ See, for instance, *Ricci v. DeStefano*, 557 U.S. 557 (2009).

This account of subordination, then, gives us at least one plausible way of understanding how indirect discrimination causally contributes to social subordination. Policies that discriminate indirectly can constitute the kind of structural accommodation that privileges other groups over a given group, reinforces stereotypes about that group, and indirectly rationalizes habits of censure and lesser consideration of them. So indirect discrimination, like direct discrimination, can play an important causal role in sustaining conditions (i) through (iv).

It might seem, however, as though there is no analogue in the case of indirect discrimination to the capacity of direct discrimination to contribute to subordination in a further way, by marking out certain groups as inferior. After all, didn't we see earlier that structural accommodations, unlike expressions of censure and lack of deference, are not inherently problematic? Don't they only become problematic through their association with certain stereotypes, differences in power and *de facto* authority, and practices of ascribing censure and consideration? Perhaps, on this account, all that we can say about indirect discrimination is that it plays some causal role in sustaining subordination—but, unlike direct discrimination, it does not literally mark out certain groups as inferior.

Some scholars might be quite content to claim this. Indeed, many believe that indirect discrimination *is* significantly different from direct discrimination. Indirect discrimination is, on their view, either an injustice of a different and less serious kind, or it is not an injustice at all, but simply an unfortunate state of affairs for those who are disadvantaged.²⁵ And if you take this view, you might think it is actually quite plausible to suggest that indirect discrimination doesn't, in fact, mark out certain groups as inferior, but only indirectly contributes to states of affairs in which one social group is socially subordinated to another.

But I want to resist this view. I think that indirect discrimination, too, can fail to treat people as equals by subordinating them to each other. And I think that my account of subordination gives us the resources to explain why indirect discrimination does not merely play an indirect causal role in sustaining subordination, but can actually mark out a group as inferior. It is of course true that indirect discrimination does not explicitly *classify* subordinated groups using the traits that are the basis for lesser consideration or censure of them. But I shall try to argue in what follows that the structural accommodations that are at issue in many unjust cases of indirect discrimination serve in an important way to render subordinated groups invisible, and thereby to mark them out as inferior.

To see this, let us turn back to my example of restaurant dress codes and the practices associated with them. As I mentioned earlier, it is a common practice for

²⁵ See, for instance, Hellman, "Indirect Discrimination and the Duty to Avoid Compounding Injustice," *supra* note 14; Benjamin Eidelson, *Discrimination and Disrespect* (Oxford: Oxford University Press, 2015) at p. 39; John Gardner, "Liberals and Unlawful Discrimination," *Oxford Journal of Legal Studies* 9 (1989), pp. 1–22; and John Gardner, "Discrimination as Injustice," *Oxford Journal of Legal Studies* 16(3) (1996), pp. 353–367.

restaurants not to stock uniforms for pregnant women, and not to stock any uniforms for those who have muscular disabilities, such as difficulties doing up buttons. This is generally not done out of prejudice towards people with disabilities or pregnant women. Rather, adapted uniforms are not easily and conveniently available, and most people who apply for waitressing jobs are not pregnant; so it isn't economically efficient for restaurants to keep a stock of such uniforms on hand. But this structural accommodation works together with certain stereotypes about pregnant women and people with disabilities (stereotypes such as "pregnant women aren't able to work very efficiently or to focus on their work" and "people with disabilities are not beautiful, so who would enjoy being served by them?") to mark them out as inferior. We can say the same about the practice, common even in restaurants that do offer female employees a choice of more and less revealing uniforms, of filling their training manuals with pictures only of svelte women wearing feminine uniform options. The absence of pictures of people who look different and have made different choices functions to render these groups invisible, and to deny their claim to equal status in no less real and forceful a way than would a sign that read "Keep out!" In fact, in an interesting way, the absence of uniforms for those who are pregnant or disabled, and the absence of pictures of people wearing non-revealing uniforms in the manuals, mark these groups as inferior even more effectively than a sign would—and even more effectively than the signs on the water fountains do, in my earlier example of direct discrimination. For a sign at least names the subordinated group and so calls attention to their existence. By contrast, the absence of the uniforms, and the absence of pictures of pregnant women or women wearing non-revealing uniforms, quite literally serves to render them invisible as potential candidates for the job of waitress. They simply do not exist in this particular part of our social world—and so neither do their needs. So indirect discrimination, too, can mark out a social group as inferior. It does so by working together with associated stereotypes and habits of censure or lesser consideration to render a group invisible.

One might at this point object that there is something paradoxical, and therefore problematic, about my claim that indirect discrimination both renders a group invisible and marks them as inferior. How could a policy really do both of these things?²⁶ In order to mark a group out as inferior, doesn't a particular policy have to call attention to them in some way, the way that direct discrimination does? Or, otherwise put, if a structural accommodation really did render a group invisible to us, wouldn't we simply stop seeing them, rather than see them *as inferior*?

But the paradox here is only apparent. In societies ordered by social castes, the lowest caste, such as the Dalit caste in India, is both invisible and inferior. Indeed, the full extent of their inferiority is demonstrated by their invisibility. Although others "see" them in the sense that they see human bodies occupying a certain space, they do not "see" them in the sense of recognizing them as full citizens, capable of participating fully in society and

²⁶ I am grateful to Cheshire Calhoun for this objection.

deserving of all of the rights that others are given in that society. Nor do we need examples as extreme as caste systems in order to see that people can be at once made invisible in a particular context and yet at the same time marked out as inferior through that very invisibility. We are just beginning to have public conversations about the ways in which many commonplace business practices still treat the average male body as the norm, and thereby render women's bodies invisible, and, in the process, devalue their work. For instance, NASA had to cancel a planned all-female spacewalk early in 2019 because it realized that it did not have any spacesuits in a size small, and had only two in a size medium: it had stopped ordering smaller suits in the 1990's in order to cut costs, because most men wouldn't fit into those sizes. Similarly, safety equipment for science labs, diving equipment for marine biologists, and protective armour for the military are all designed to fit male bodies rather than female ones. This treats female scientists and soldiers as though they do not exist, and thereby implies that their work is not as important as the work of their male colleagues.²⁷

For an even more vivid example of how the same act can render a group invisible and inferior, consider the artistic program launched by the transit commission in my home city of Toronto. Local artists have been commissioned to sketch members of the public riding the city's buses, streetcars, and underground trains, and copies of the sketches are put up inside these vehicles, under the heading "Sketching the Line." However, although the majority of people who ride public transportation in Toronto are from visible minorities, and a huge number are women, the sketches that have been put up so far consist largely of identifiable white men and of a few non-descript, unidentifiable people with their backs turned to the viewer, who might be members of visible minorities but might also be white. At the time of my writing this, there is only one drawing of a woman, and she is obviously white. When I telephoned the Director of the program to discuss the absence of visible minorities and women in these sketches, he replied, slowly and with a tone of barely controlled exasperation: "Ma'am, the artists are just drawing . . . what . . . they . . . see." In one sense, this is false: the drawings are quite stereotyped caricatures even of the white men who are pictured, with older men falling asleep and teenage boys with acne-covered skin eating potato chips. So one doubts that the artists are actually "drawing what they see." But in another sense, the Director's claim is disconcertingly accurate. Perhaps the artists who made these sketches did ride trains packed with Chinese, Haitian, Korean, and Pakistani women, and yet these artists "saw" only the one white male teenager sitting in the corner. For this is often what happens to members of "visible" minorities in Toronto. Even though the city celebrates its multiculturalism, it is nevertheless true that racial minorities, and especially women from racial minorities, are still not regarded as normal or typical, even when they outnumber white men in a particular place. They are, in this sense, invisible; and the "Sketching the Line" program unfortunately perpetuates their invisibility. And in doing so, it marks them out as inferior. Even though they are the

²⁷For these and other examples, see Caroline Perez, *Invisible Women: Data Bias in a World Designed for Men* (New York: Abrams Press, 2019).

overwhelming majority of riders, they are not “normal” enough to be sketched as representative riders.

I have now used my account of social subordination to suggest a number of ways in which direct and indirect discrimination can work to subordinate social groups. Direct discrimination can sometimes constitute an expression of censure towards, or lesser consideration for, a subordinated group, as is mentioned in condition (ii). And it can sometimes causally perpetuate the conditions described in (i), (iii), and (iv), sustaining differences in power and authority between the subordinated group and more privileged groups, supporting stereotypes that in turn rationalize inferior treatment of the subordinated group, and keeping in place problematic structural accommodations. Indirect discrimination can sometimes contribute to subordination, similarly, by playing an important causal role in sustaining conditions (i) through (iii). And although it does not normally constitute an expression of censure, it can sometimes serve to render certain social groups invisible in certain contexts, thereby marking them out as inferior.

One might at this point object that talk of “marking out people as inferior” and “rendering people invisible” sounds very much like talk of the expressive meaning of an act or policy. What is it, really, to “mark out as inferior” or to “render invisible” if it is not to send the message that a particular social group is inferior or invisible? So it might seem as though, in spite of my earlier criticisms of the expressivist view, my account has an important expressivist dimension to it, even though it goes beyond this and looks also at a broader range of effects of discriminatory acts.

But although part of what it is to “mark out someone as inferior” is to send a message about this person, this idea cannot be entirely parsed in terms of the social message that a policy sends. Rather, marking someone out as inferior or invisible (or both) involves doing things in the world to that person. It involves altering their situation in certain ways, imposing additional costs on certain opportunities, creating certain disincentives that the rest of us do not have to worry about. When all storefronts have a step leading up to them and this renders Jean and his disability invisible, the step doesn’t just send a social message about Jean. It literally prevents him from accessing the store by himself: it is a physical barrier. It also reinforces our shared assumption that the normal shopper is someone who is not in a wheelchair. This in turn creates disincentives for Jean to come forward and ask for the same opportunities as others. For if he were to come forward, he would face costs that the rest of us do not. He would have to present himself as different from the rest of us, different in ways that, given our society’s stereotypes about disability, have pejorative connotations attached to them. He would have to present himself as less capable than the rest of us, and as dependent on our help. Of course it is true that these costs depend in complex ways upon the social messages that the steps, and other practices in our society, send about people with disabilities. But it does not follow

that the costs for John are reducible to these social messages. They are very real costs in the world—as real as the physical barrier that the steps pose.

Similarly, when female soldiers have to wear body armour that is designed for men, we cannot understand their invisibility solely in terms of the social message that is sent about them. It certainly does include this message. But it also includes very real health risks that women must bear, which men do not have to bear. And it includes a complex array of other costs—as is shown by the case of Rebecca Lipe, who was an American Air Force Judge Advocate General working in Iraq in 2011.²⁸ Lipe explained in a recent military hearing that the American military's standard issue "ballistic vest" is not designed to fit a woman's body, and that female soldiers have to jerry-rig it by removing its side panels. This leave them without protection if they are shot in the side. Lipe found that in order to make the vest fit, she needed to hike it up and overtighten it around her waist to reduce the slack. But 5 ½ months of wearing the vest this way caused her a pelvic herniation –which was misdiagnosed by the military doctors, who assumed that because she was a woman, the most likely explanation for her pain was menstrual cramps or an STD. They suggested she was exaggerating her pain; and then began testing her for STD's and questioning her on whether she had been having extra-marital affairs. Interestingly, the medical officers in this case believed that they were paying proper attention to Lipe as a woman. But their misdiagnoses were based upon false, stereotyped assumptions about women's injuries and habits of behaviour –which helped to make Lipe's actual injuries, and her actual pain, invisible to the doctors. This case shows particularly vividly how stereotypes about particular groups work together with practices that render the group invisible, so that even when members of this group come forward to report their problems, others are often unable to see these problems clearly. Instead, others see what they had expected to see in a person of that type. And I hope this case also shows that what it is to be rendered invisible can't be analyzed simply in terms of the social message that is sent by a particular practice. We need to think of a much broader, and more serious, set of costs that are placed upon those who are rendered invisible.

I considered the examples of Lipe and Jean in order to highlight the ways in which my account of subordination differs from the expressivist account. As I hope these two examples have shown, my account takes the social meaning of acts and policies to be but one component of a full analysis of the ways in which those acts and policies mark out some people as inferior or render them invisible. My account also looks at a broader range of effects of discriminatory acts than do expressivist accounts. And my account considers

²⁸ Subcommittee on Military Personnel Hearing: "Feres Doctrine – A Policy in Need of Reform?" Tuesday, April 30, 2019. Transcript of the hearing available at: <https://armedservices.house.gov/2019/4/feres-doctrine-a-policy-in-need-of-reform>.

subordination to be something that happens to a person as a member of a group, because of a socially salient trait that they share with others.

My account of subordination also enables us to see more clearly the ways in which different sub-groups within a particular social group can be subordinated in relation to each other, and subordinated in different respects and to different extents. On an expressivist account, all acts of discrimination that wrong the members of a particular group wrong them in the same way—by demeaning them as individuals. But my account of social subordination gives us a richer way of describing what is going on, and encourages us to think about the differences between what happens to one sub-group and what happens to another. Think back to my example of the training manual. Because the manual depicts only women with svelte figures wearing conventionally feminine uniform options, it implies that women waitresses are objects of beauty to be enjoyed by men, and so marks out all such women as inferior to men. But, as we saw earlier, it also creates a hierarchy within the class of women, between those who measure up to conventional standards of beauty and those who do not—and it subordinates the latter in a further way. My account of subordination offers us a rich set of concepts with which to analyze this nested form of subordination, and also with which to analyze the ambiguous position of the women who meet these standards and the precariousness of their status relative to the more subordinated women—which does not seem adequately described simply by saying that they are not demeaned relative to these other women but are demeaned relative to men. The account encourages us to explore the particular ways in which directly and indirectly discriminatory policies reinforce different patterns of consideration or censure towards different sub-groups, support different stereotypes about each sub-group, and rationalize differences in power and authority between these sub-groups as well as between the group as a whole and other, more privileged groups.

My account of subordination also offers us a plausible explanation of the function of prohibited grounds of discrimination—though, as I shall later explain, this is not, on my view, the only explanation of their function. All of the traits that are commonly recognized in our laws as prohibited grounds—for instance, race, gender, sexual orientation, religion, and disability—mark out traits on the basis of which at least one, and often quite a number of social groups have been denied equal power and *de facto* authority over others; have been subjected to greater censure or lesser consideration, in the sense that they have been condemned or thought of as less worthy of respect than others; and have had their needs overlooked by certain structural accommodations that cater to the needs and circumstances of other, more powerful social groups. In other words, they are traits that not only mark out particular subordinated groups but also frequently help to explain why they have faced subordination in certain contexts. So, in cases where discrimination is wrongful because it subordinates, the requirement that discrimination must have occurred on the basis of a prohibited ground helps to ensure that the discriminatees really are part of a group that is socially subordinated, and may point us toward an explanation of why

and how they have been subordinated in this case—though it does not, of course, guarantee that they are subordinated in that particular context by that particular discriminatory practice—so a further contextual analysis of whether and how the practice contributes to subordination is always going to be necessary. If this is the rationale for our list of prohibited grounds, or rather, the rationale for it in those particular cases where the wrongfulness of discrimination stems from unfair subordination, then it is arguable that our lists of prohibited grounds should also include other traits—traits that are not commonly added to such lists, such as physical appearance, for instance, and social condition or poverty. Moreover, my account implies that what traits should or should not be on the list will vary from one society to another, and may vary over long periods of time, depending on what forms of social subordination exist and persist in each particular society.

In this chapter, I have tried to show that one of the ways in which discrimination can wrong people by failing to treat them as equals is by subordinating them to others. I have explored two senses in which direct and indirect discrimination subordinate. First, both types of discrimination play an important causal role in sustaining unfair social subordination. And second, both types of discrimination can mark out a particular social group as inferior to others—direct discrimination, by constituting an act of censure or lesser consideration of a certain group, and indirect discrimination, by rendering a certain social group invisible and thereby creating real barriers to their participation in certain social institutions.

But this is not, in my view, the only reason why a particular case of discrimination can be wrongful. Discriminatory policies, and the stereotypes that underlie them, have also been challenged as wrongful for very different reasons. One of these reasons is that they sometimes deny people the freedom to shape their lives according to their own values, without constantly having to factor in other people's assumptions about certain traits of theirs, and without having to treat these traits as costs.²⁹ Although we do not always have a right to this kind of freedom, I shall be arguing in the next chapter that we sometimes do. When discriminatory acts and policies deny us such freedom in circumstances where we have a right to it, they fail to treat us as beings capable of autonomy. And given the value that our societies place upon autonomy, to treat us in this way is to fail to treat us as equals. So there is a second way in which discriminatory practices can wrong us, a second way in which they can fail to treat us as equals, over and above whatever contribution they might make to the kind of social subordination that I have been discussing in the current chapter. I shall turn in the next chapter to the task of exploring this second way.

²⁹ See, for instance, Tarunabh Khaitan, *A Theory of Discrimination Law* (Oxford: Oxford University Press, 2015).

Chapter Three

The Relevance of Deliberative Freedom

I argued in the previous chapter that one way in which certain discriminatory practices fail to treat people as equals is by subordinating them to others, whether by marking them out as inferior to others, or by contributing to the social subordination of a group to which they belong. But the wrongness of many discriminatory acts and policies does not seem to be exhausted by their contribution to subordination. I want to turn now to two cases in which a discriminatory policy *also* does something else—something that, when we ask ourselves how these people have been wronged by this instance of discrimination, calls out for inclusion in our explanation. I shall suggest that this other factor has to do with the impact of discriminatory acts and policies on a certain kind of freedom, to which people sometimes have a right. It will be the task of the rest of this chapter to explain exactly what this kind of freedom is, why we have a right to it in certain circumstances and not others, and why it is so important to think of the wrongness of certain kinds of discrimination as stemming in part from infringements of a right to this freedom.

3.1 Why the Wrongness of Discrimination Extends Beyond Subordination

Why might we think that the wrongness of discriminatory acts and policies extends beyond their contribution to subordination? Consider first the sprinter Dutee Chand's 2015 challenge to the "Hyperandrogenism Regulations"¹ that had been laid down by the International Association of Athletics Federations (I.A.A.F.).² The Regulations stipulated that female athletes whose natural testosterone tested higher than certain parameters would have to take measures to lower their levels of testosterone, or would be banned

¹ Hereinafter "The Regulations." For extremely helpful discussions of the Regulations, I am very grateful to Bruce Kidd, who has been involved in the C.A.S. litigation on behalf of Dutee Chand.

² These regulations have since been replaced by the I.A.A.F.'s new "Eligibility Regulations for the Female Classification (Athletes with Differences of Sex Development)," often referred to as the "DSD Regulations," which were published on 23 April 2018, suspended during a challenge by Caster Semenya, and enforced when the Court of Arbitration for Sport ruled against Semenya on May 1 2019. I have chosen to discuss the older Hyperandrogenism Regulations because, although the issues posed by the two sets of regulations are in many respects the same, the new regulations target only athletes competing in events between 400m and a mile, and only athletes who are intersex. So although they still affect Semenya, they do not apply to Dutee Chand. And it is not just Semenya's comments but also Chand's understanding of her complaint that is important for my argument. Focussing on the older, broader Regulations allows me to consider both Semenya's and Chand's complaints.

from competing as a woman in international athletic events. In 2015, Chand successfully challenged these Regulations, and they were suspended for two years by the Court of Arbitration for Sport (C.A.S.).³ Chand's natural levels of testosterone are higher than the parameters set for women in the Regulations, and she had been suspended by the I.A.A.F. because she had refused to take the required measures. She claimed that the Regulations were unfairly discriminatory, as they were essentially an attempt to regulate "womanhood" in the context of athletics and served no legitimate athletic purpose. She argued that testosterone was just one factor in an athlete's performance, and that other factors—other genetic advantages, national and personal income, and access to coaching and facilities, none of which are regulated—play a far greater role. And she noted that there are no parallel regulations for natural hormone levels in men: men are always permitted to compete as men, regardless of the levels of hormones in their bodies. Neither are there regulations governing any of the other biological or genetic variations that might be advantageous for an elite athlete, even though more than 200 genetic variations have been identified that might provide advantages, such as genes affecting blood flow to muscles, muscle structure, oxygen transport, and lactate turnover. Moreover, the only women to whom the Regulations have been applied are women from the global south—such as Chand, from India, and African sprinters such as Caster Semenya. And the medical procedures that these women have been urged to undergo include not just oral contraceptives that would lower their levels of testosterone, but intrusive procedures that have an obvious connection with gender identity—for instance, feminizing vaginoplasty, estrogen replacement therapy, and clitoris reduction.

These Regulations and the mode of their application seem to contribute to the subordination of all female athletes from the global south, as a class. They also mark out as inferior the sub-class of these athletes that are deemed hyperandrogenist, implying that these athletes are not real women.⁴ So at least part of what is morally troubling about the Regulations does seem to stem from their contribution to social subordination on the basis of race, gender, and gender identity.⁵ Indeed, this seems to be an example of the phenomenon that I noted at the end of Chapter One, in which a sub-class within an already underprivileged group is further subordinated by a particular policy, and the subordination

³ *Dutee Chand v. Athletics Federation of India (AFI) & The International Association of Athletics Federations (IAAF)* (2015), CAS 2014/A/3759 (Arbitrator: Court of Arbitration for Sport).

⁴ Indeed, in the extensive media coverage surrounding the legal challenge of the Regulations and the athletes who have been affected by them, there has been hardly any focus on the relationship between testosterone and the performance of these athletes, and instead an obsessive focus on whether these athletes are "real women."

⁵ An interesting complexity in this case is that it is only the sub-class of hyperandrogenist athletes that is actually marked as inferior, while the Regulations causally contribute to the subordination of the broader class of *all* female athletes from the global south. This suggests that an act or policy can causally contribute to the subordination of a broader class even if it only marks out as inferior a sub-class within that broader class.

occurs in part because the policy divides that underprivileged group into a higher and a lower class—here, those athletes who are “real” women and those who are not.

But interestingly, when we hear Dutee Chand’s and Caster Semenya’s descriptions of the impact of the Regulations on their own lives, subordination does not loom so large. There is something else that they are fighting for. And it is not simply their chance to compete.

Semenya, who has been called “a man” to her face by other athletes at international events, has said, “I don’t want to be someone I don’t want to be. I don’t want to be someone people want me to be. I just want to be me.”⁶ In other words, she doesn’t want to be defined by other people’s assumptions about what her gender is. She doesn’t want other people’s assumptions about what her body means to govern fundamental choices that she makes about things such as medical treatment, particularly the kinds of medical treatment that are likely to have profound effects on her body and health. And Semenya thinks that other people’s assumptions about the gender she has, and about the physical characteristics that she should have if she is to be a “real” woman, should not be the kinds of things that figure as impediments or costs in her life.

Similarly, when Dutee Chand refused to undergo the requested medical treatment to make her eligible to compete as a woman, she said: “I am who I am.”⁷ She has spoken out in public about the effects of these Regulations on her life, noting that “[m]ost of my relatives dismissed themselves from me” and “I’m scared to ask [my female friends] to meet me since parents don’t want their daughters to be with me.” She has had to move to a different town away from her parents, and she tries not to call them often, for fear of upsetting them. People frequently come up to her and ask her “Are you an *andirachandi* [tomboy]? Are you not going to get married?”⁸ Her response is just: “I am who I am.” So Dutee Chand, too, is objecting to having to consider other people’s assumptions about her gender when deciding how to live her life and who she is and should become. In her case, these assumptions have been especially intrusive: she has not been free to decide where to live, how often to phone her parents, and which of her friends she should try to associate with, without considering what her gender is or should be, and what other people are saying about her gender.

So the Regulations do not just work to subordinate athletes from the global south, and to mark out Dutee Chand and Caster Semenya as inferior, as less than real women. They also deprive them of freedoms. And it is these freedoms that Dutee Chand and Caster

⁶ Quoted in Sharda Ugra and Susan Ninan, “Castor, Dutee and the Obstacle Race of Their Lives,” *ESPN* (16 August 2016), http://www.espn.com/athletics/story/_/id/17306451/caster-dutee-obstacle-race-their-lives.

⁷ Quoted in Silvia Camporesi, “Why Caster Semenya and Dutee Chand Deserve to Compete (And Win) at Rio 2016,” *The Conversation* (9 August 2016), <https://theconversation.com/why-caster-semenya-and-dutee-chand-deserve-to-compete-and-win-at-rio-2016-63727>.

⁸ Quoted in Debabrata Mohanty, Jonathan Selvaraj, and Nihal Koshie, “I Am Who I Am: Dutee Chand,” *The Indian Express* (29 September 2014), <http://indianexpress.com/article/sports/sport-others/big-picture-i-am-who-i-am/>.

Semenya seem to care most about. They are fighting for the freedom to run a race without being burdened by other people's assumptions about their gender, the freedom to live near their parents and go out with their friends, without fearing aggressive responses from people who think they are too masculine.

These freedoms are what, in earlier writings, I called "deliberative freedoms." I used the term "deliberative freedom" to highlight the fact that these freedoms are important to us because we care about having the opportunity to shape our lives through our own deliberations and choices. But, as I emphasized in my earlier writings, and as I shall explain later in this chapter, deliberative freedoms include not just freedoms of thought but freedoms of action as well. I shall shortly turn to a more detailed exploration of what a "deliberative freedom" is and why we should think that people sometimes have a right to it. I shall also discuss how, on my view, we are to think about the interests of *other* people, which are also often at stake in many cases of apparently wrongful discrimination. For of course, when we think about such cases, there are other people whose interests we need also to factor into our assessment of whether the relevant practice is wrongfully discriminatory. In the case of the Regulations, the I.A.A.F. argued on behalf of other female competitors that it was unfair that athletes such as Chand and Semanya should use the category of "female athlete" when their extra testosterone gave them such an advantage in this category: they claimed that women without such high levels of testosterone would not have a fair chance to compete unless those with high testosterone levels were excluded, and that consequently, the Regulations were not wrongfully discriminatory. I shall return to this argument at a later stage. For now, I want to say a little more about the complaints of those who allege wrongful discrimination of a kind that seems to involve a deprivation of a certain deliberative freedom.

I want to consider one more example in which the victim's complaint seems to be at least in part about an infringement of freedom, rather than just about social subordination. In this next example, subordination seems to play even less of a role than it did in the case of the Hyperandrogenism Regulations, and so it is even clearer that we need to appeal to something else in order to explain why the discrimination here is wrong.

The example I have in mind is the American case of *Craig v. Masterpiece Cakeshop*.⁹ Two men, Craig and Mullins, went to a local cake shop in Lakewood, Colorado to purchase a cake for their wedding reception. The owner of the cake shop, Jack Phillips, refused to bake them a cake, on the grounds that theirs was a same-sex marriage and his religion forbade him from acknowledging or celebrating same sex marriages. Phillips informed them that he would happily provide baked goods of any other sort for them on any other occasions: his objection, he said, was not to them or their homosexuality. But he argued that, if the law were to require him to bake a wedding cake for them, it would in effect be forcing him

⁹ *Craig v. Masterpiece Cakeshop, Inc.*, 370 P.3d 272 (Colorado Court of Appeals, 2015), reversed by the U.S. Supreme Court in *Masterpiece Cakeshop, Ltd. et al v Colorado Civil Rights Commission et al.*, 584 U. S. ____ (2018) [*Masterpiece Cakeshop*].

to acknowledge that a marriage had taken place and forcing him to celebrate that marriage. And this was something his religion forbade him to do.

Whether the law would indeed be requiring Phillips to celebrate a same sex marriage if it required him to bake a cake for Craig and Mullins is a fascinating question, and one I shall come back to in Chapter 7, when I discuss the obligations of individuals such as Phillips in more detail.¹⁰ For now, my main interest in *Masterpiece Cakeshop* lies not in the question of how to conceptualize Phillips' obligations, but in the question of how to conceptualize the complaint of Craig and Mullins, when they were denied the cake. It is not so easy, in this case, to think of the complaint primarily as a complaint about social subordination. For one thing, unless other bakeries shared this policy of refusing wedding cakes to same-sex couples, or unless the policy reinforced negative attitudes towards or stereotypes about gay couples that were accepted in the local neighbourhood, it would be unlikely causally to contribute to the subordination of gay couples. After all, the Masterpiece cake shop is one bake shop among scores of others, and it is providing a relatively trivial thing: sweets. It is not providing an essential social service or a set of opportunities that will make a large difference to the power or *de facto* authority enjoyed by gay couples.

Does the policy nevertheless subordinate gay couples in the sense that it marks them out as inferior? Phillips emphasized in court that he serves gay couples all of his other baked goods. And he argued that because of this, his refusal to bake them a wedding cake did not amount implicitly to the claim that they are unworthy of being served, the way Woolworth's lunch counter implicitly claimed of blacks when it refused to serve them point blank in the 1950s. He was, he argued, simply trying to practise his religion. I think Phillips is right that his policy does not mark out same-sex couples as inferior in as obvious or straightforward a way as does the Woolworth's example or, for that matter, the Hyperandrogenism Regulations. Nevertheless, the particular religious doctrine that he is following does treat same-sex couples as inherently unable to enter the sacred institution of marriage—and it is a rather small step from “unable to enter into this sacred institution” to “unworthy of entering into it.” Since Phillips' refusal to bake them a cake stems from, and expresses, his endorsement of this particular religious doctrine, it seems to come very close to marking out same-sex couples as less worthy than heterosexual couples.

This case may, then, involve subordination, although the subordination is not as obvious or as considerable as it is in the Hyperandrogenism Regulation case. But I want now to explore the way in which, just as in the case of the Hyperandrogenism Regulations, social subordination does not seem fully to capture the unfairness here. We can start by noting that Craig and Mullins regarded it as irrelevant that they could go somewhere else to get their wedding cake. In fact, when they made their complaint public, another cake shop provided them with a beautiful wedding cake at no charge at all, a wedding cake with a rainbow across the top that celebrated their sexual orientation. Yet this act did not, in their

¹⁰ See Chapter 7, Section 7.6.

minds, reverse what had occurred. It did not reverse what had occurred, even though we might think of the beautiful free cake as an expression of what I in Chapter One called “consideration,” a kind of special respect being paid to this gay couple because of their sexual orientation, and hence as the converse of marking them out as inferior. This did not reverse what had occurred, even though it lessened their financial costs and went some way towards challenging stereotypes about gay marriage and empowering gay couples. One explanation of why is that Phillips’ refusal to bake them a cake had made their sexuality *an issue* during their wedding planning. It had made their sexuality, and other people’s assumptions about it, something that loomed before their eyes, something they now had to think about when going to purchase a wedding cake; when in fact part of what same-sex couples had fought for, in fighting for so many years to access the institution of marriage, was the freedom to marry without having to treat their sexual orientation as a liability, and indeed without having to think of themselves and their partner as anything other than “partners in marriage.” Of course, many gay couples enjoy publicly celebrating their sexual orientation, and many make a conscious choice to foreground it in their wedding planning. But what Craig and Mullins were insisting was that they *should not have to* have their sexual orientation, and other people’s assumptions about what it means, constantly before their eyes; nor should they have to bear the costs of other people’s assumptions about it or about what it renders them fit for or unfit for. And it does not matter, from this standpoint, whether these assumptions stem from a recognized religion or from prejudice or dislike. A person should not have to bear the social costs of their sexual orientation, or even to think about it, when buying a cake—even if that cake is a wedding cake. When even one bakery refuses them a cake on the grounds of their sexual orientation, they are suddenly placed in a position where they have to do this. And this is not erased by the ease with which they can find another cake, or by the fact that another bakery celebrates their sexuality.

If this is what Craig and Mullins were objecting to, when they said they had been wrongfully discriminated against, then they were objecting to an infringement of what I have called “deliberative freedom.” And I think we need to appeal to some such idea of freedom in order to explain the unfairness of at least some cases of discrimination. In Chapter Two, we considered the way in which some acts of discrimination create or perpetuate a state of affairs in which there are two classes of citizens, a superior and an inferior one. But the two cases of discrimination we have examined in this chapter reveal that people’s objections to discrimination go beyond not wanting to be treated as second-class citizens. In cases such as the Hyperandrogenism Regulations and *Masterpiece Cakeshop*, people want the freedom to make choices about their lives—from relatively small choices about what cake to have at their wedding ceremony to very profound choices such as what gender they are and what their body should look like—without having to consider other people’s assumptions about their gender or their sexual orientation, or other people’s assumptions about what roles these render them fit or unfit for.

This is what I have called “deliberative freedom.” In the next section of this chapter,

I shall turn to the task of clarifying the idea of deliberative freedom and its role in an account of why discrimination is sometimes wrong.

3.2 What is Deliberative Freedom, and When Do We Have a Right to It?

Deliberative freedom is the freedom to deliberate about one's life, and to decide what to do in light of those deliberations, without having to treat certain personal traits, or other people's assumptions about them, as costs, and without having to live one's life with these traits always before one's eyes. So understood, deliberative freedom seems to consist in a number of different but related freedoms.

On the one hand, it involves certain freedoms of thought. One of these is the freedom to deliberate about one's options without having to treat certain traits *as costs*. Kasper Lippert-Rasmussen has quite rightly noted that the costs that are relevant to this particular aspect of deliberative freedom are *opportunity costs*: a trait is a cost in this sense if it makes it more difficult, or more expensive, for me to pursue a certain option.¹¹

But there is another aspect to the idea of deliberative freedom, which is concerned not with opportunity costs but with what we might call "fixed costs" or "burdens." For deliberative freedom, as I understand it, also involves the freedom not to *have* to think about a certain trait. And there can be circumstances in which a certain trait is made into an issue, and one is forced to have it always before one's eyes, even though there is nothing one can do about it and no option one can pursue under which one will be free from the negative effects of that trait or others' perceptions of it. Suppose (as is generally the case) that female athletes draw salaries that are considerably less than male counterparts who are no less talented, that their sporting events are not as widely televised or publicized, and that their leagues are, in the public eye, regarded for the most part as second-best. It is true that the gender of these athletes is not an opportunity cost in the sense that it makes it more expensive for them to pursue one sporting option rather than another. For whatever sport they enter and however hard they work at it, they will end up drawing a lower salary and attracting less attention than their male counterparts. But their gender has been "made an issue" in the sense that I am concerned with. So they do lack the ability to deliberate "freely" in my sense, without having this feature of themselves and their eyes constantly before them. So, pace Lippert-Rasmussen, it is *not* true that those who are subjected to fixed costs on the basis of a certain trait will have full deliberative freedom on my view.¹² In many cases, they will *lack* deliberative freedom, because that trait will still loom before their eyes as they live their lives. They will still have to bear the burden of it, even though there is nothing they can do to lift that burden.

¹¹ Kasper Lippert-Rasmussen, *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination* (Oxford: Oxford University Press, 2014) at pp. 186–187.

¹² Lippert-Rasmussen, *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*, *ibid.* at p. 187.

This particular aspect of deliberative freedom—that is, the freedom not to have a certain trait looming before one’s eyes, as one lives one’s life—is a very important part of the idea of deliberative freedom. In fact, part of the reason that I am drawn to include denials of deliberative freedom in an account of what makes discrimination wrongful is that the loss of this kind of freedom is a salient feature of the lives of people who suffer from systemic discrimination. It is something they mention very often, when describing their experiences. It may even be *the* salient feature of the oppression that marks their lives. For instance, if you are African-American, you can never enjoy the luxury of forgetting about your race. You carry the burden of other people’s assumptions about your race wherever you go. If you are late for their job interview, your employer will assume you are scattered and lazy --rather than assuming, as they will about your white counterparts, that you simply got caught in traffic. If you have biracial children who look “white,” their teachers will assume, when you come to pick them up from school, that you must be a paid caregiver, because you look black. And even when you are doing something as innocuous as driving your car to work, you will be aware of the fact that you could get pulled over by police, and the police will likely interpret your every move in light of stereotypes about black aggressiveness and criminality.

This is a very real and a very significant lack of freedom. It affects the way in which many African-Americans make their decisions and the options available to them. But it also looms over them, even if they are deliberating in a context in which there is nothing they can do to compensate for other people’s perceptions about them based on their race, and no choice they could make that will not result in their being penalized for being black, or for being perceived as black.¹³

Part of what is so interesting about this particular aspect of deliberative freedom is that it is noticed predominantly by those who *lack* deliberative freedom. Those of us who do have it tend not to be aware that we have it. This is a point that is familiar from discussions of white privilege. Sylvia Law explains it well when she writes that:

Black people invariably note their race and white people almost never do. Surveys tell us that virtually all Black people notice the importance of race several times a day. White people rarely contemplate the fact of our whiteness—it is the norm, the given. It is a privilege to not have to think about race.¹⁴

¹³ I have discussed this aspect of deliberative freedom further in “In Defense of a Liberty-based Account of Discrimination,” in Deborah Hellman and Sophia Moreau (eds.), *Philosophical Foundations of Discrimination Law* (Oxford: Oxford University Press, 2013), pp. 71–86.

¹⁴ Sylvia Law, “White Privilege and Affirmative Action,” *Akron Law Review* 32(3) (1999), pp. 1–23. See also Angela P. Harris, “Race and Essentialism in Feminist Legal Theory,” *Stanford Law Review* 42(3) (1990), pp. 581–616 at p. 604; and Martha A. Mahoney, “Segregation, Whiteness, and Transformation,” *University of Pennsylvania Law Review* 143(5) (1995), pp. 1659–1684 at pp. 1662–67.

This of course makes it harder to convince sceptics that deliberative freedom is a real form of freedom. For those of us who have it tend not to notice that we do. It is more clearly identifiable through the burdens that are imposed on those who lack it than it is through the actual benefits accruing to those who enjoy it, who tend not to notice that they have something that others lack.

Although many scholars speak of this aspect of deliberative freedom as a “privilege”, as Sylvia Law does in the quotation above, I worry that the term “privilege” gives rise to two mistaken ideas. The first is that a deliberative freedom is an isolated opportunity that some people lack, the way others might lack the opportunity to buy a car or the opportunity to join a club. The second is that those who have a particular deliberative freedom are enjoying something over and above what anyone is entitled to –a “privilege”—and that therefore, those who lack that deliberative freedom cannot claim that they lack anything to which they are entitled. Neither of these ideas is correct. A particular deliberative freedom is not an isolated opportunity. It is a power not to be bound by or burdened by certain assumptions and certain costs --and those who lack such powers are in a very real sense unfree, usually across many social contexts. Nor is deliberative freedom a luxury, to which no one can ever claim an entitlement. As I shall argue later in this chapter, those who lack it are often entitled to it, and what they lack is of enormous significance to them. It is not just an opportunity to do something frivolous, like buying a car or joining a club. What they lack is the space to become the people whom they want to be.

Up to this point, I have emphasized the freedoms of thought that are involved in deliberative freedoms, such as freedom from having to treat certain traits as opportunity costs, and freedom from having these traits figure constantly in one’s thoughts as burdens. But in addition to these freedoms of thought, deliberative freedom involves certain freedoms of *action*. A person is not genuinely free to deliberate without considering certain traits or their costs if he is only under the illusion that he is free. I am not free to deliberate about whether to live in Forest Glen if I don’t in fact have the opportunity to do so, because all landlords there deny tenancies to people of my ethnicity, or because all banks deny mortgages to people of my ethnicity. So a necessary condition of a person’s having a certain deliberative freedom is that he or she really does have the opportunity to *do* the thing that she may decide to do. The reason I have highlighted the deliberative aspect of these freedoms by calling them “deliberative freedoms” is not that they are freedoms of thought divorced from their associated freedoms of action, but rather that these freedoms of action matter to us because it matters that we have the opportunity to shape our lives in our own way, through our own deliberations and decisions.¹⁵

I have argued that some claimants see the wrongness of discrimination as consisting partly in a denial of deliberative freedom. But to say this, and to acknowledge

¹⁵ For further discussion of this point, see Sophia Moreau, “What is Discrimination?”, *Philosophy & Public Affairs* 38(2) (2010), pp. 143–179, esp. at pp. 143–153.

that deliberative freedom is important to us, is not to say that we have an interest in full or maximal deliberative freedom or that discrimination is unfair whenever it interferes with our deliberative freedom. Such a view would be implausible. There are many ways in which our deliberations are constrained by the acts and choices of others—and constrained legitimately. For instance, other people’s preferences and choices influence the cost of the products I want to buy and the cost of the activities I want to engage in. If my religion requires me to undertake a pilgrimage to a remote holy site, and there are relatively few people who share my religion and so few people wanting access to this remote site, it will be more expensive for me to travel there. So my religion will therefore impose costs on me; and in this sense, my deliberative freedom will be lessened. But no one would think that this particular deprivation of freedom is unfair, or that others are thereby required to subsidize my religious pilgrimages. So not all infringements of deliberative freedom are problematic. Nor do we require full and complete deliberative freedom. In fact, it is arguable that full and complete deliberative freedom would be an incoherent ideal. As Dworkin has argued, part of what it means for me to take responsibility for my life as an autonomous individual is to make my choices within a framework that is in part defined by your choices and your preferences—for the cost of any product or activity for me will always be partly a function of others’ preferences.¹⁶

We do not, then, have an interest in full and complete deliberative freedom; nor does just any interference with our deliberative freedom count as unfair.¹⁷ Rather, it seems that in certain circumstances we have a right to deliberative freedom, and in other circumstances, we do not. But when exactly does someone have a right to a certain deliberative freedom, and why?

In my earlier work, I suggested that there may be no single, principled explanation of why we have a right to deliberative freedoms in certain cases but not others. We can only ask, on a case-by-case basis, whether the particular deliberative freedom at issue in a given case seems important enough, relative to the other interests of other people affected, that the person in question can be said to have a right to that particular deliberative freedom. And in different cases, we might appeal to quite different considerations, depending on the context and on the particular trait in relation to which deliberative freedom is denied to us.¹⁸ For instance, a right to certain deliberative freedoms about matters of religion may stem in part from the importance of a person’s religion to their life as a whole. By contrast, in the case of race, we may look not to the importance of a person’s race to their life as a whole, but to the fact that membership in certain races in our

¹⁶ See Ronald Dworkin, “Equality of Resources,” in *Sovereign Virtue: Equality in Theory and Practice* (Cambridge, MA: Harvard University Press, 2000), pp. 65–119.

¹⁷ Colin Campbell and Dale Smith have suggested that my view assumes a right to complete deliberative freedom: “the aspiration behind the deliberative freedoms account is that we be completely, and not just sufficiently, free in this respect.” But, as I hope to have shown above, the account is not committed to this. See Colin Campbell and Dale Smith, “Deliberative Freedoms and the Asymmetric Features of Anti-Discrimination Law,” *University of Toronto Law Journal* 67(3) (2017), pp. 247–287 at p. 285, note 100.

¹⁸ Moreau, “What is Discrimination?,” *supra* note 15. at pp. 156–157.

societies—such as indigenous groups—unfairly carries extra costs, costs that are due to other people’s mistaken assumptions about these races, which should not be allowed to thwart these people’s own choices.

But I now think that this is only partly correct. It is true that a diverse array of considerations is relevant to whether someone has a right to a certain deliberative freedom. However, as I mentioned earlier, we value deliberative freedom because we value autonomy. Perhaps, then, we can look to the idea of autonomy to help explain both why we have a right to certain deliberative freedoms in some cases, and why we do not in other cases. That is, it can help to explain both the entitlement and the limits of that very entitlement. Perhaps we can say that someone has a right to a certain deliberative freedom if denying that freedom to her would amount to failing to respect her as a being capable of autonomy. I am using the term “autonomy” in a relatively thin sense here, to mean simply “deciding what is important for you, and living your life as far as is possible in accordance with those decisions.” So this understanding of our entitlement to deliberative freedom does not presuppose that only choices made in certain special ways count as autonomous; and as I shall later argue in much more detail, it is important that the idea of autonomy to which we appeal in an account of discrimination be neutral as between different ways of living one’s life, if it is to further the goals of anti-discrimination law and to enable underrepresented and misunderstood groups to be heard.¹⁹ Note also that what is relevant here is not respect for people who *are* autonomous, but respect for them as beings who are *capable* of being autonomous. And this means we must both show respect for the actual choices that they have made—such as choices about their religion and their job— and must act in such a way that they can continue to make choices about their lives and continue to live their lives in accordance with these choices, without being unfairly hindered by our assumptions about who they are or who they ought to be.

It may seem that to try to shed light on the right to deliberative freedoms by appealing to what it is to ‘respect someone as a person capable of autonomy’ is to gloss one vague idea by appealing to another very similar and equally vague idea. But I think the idea of respecting someone as a person capable of autonomy points us in the direction of a number of relevant considerations, when assessing whether someone has a right to a certain deliberative freedom. One of the most important of these considerations is whether the costs that a discriminatee is being asked to bear reflect her own personal choices, or whether they reflect other people’s assumptions about who she is and what roles she ought to occupy. But of course it also depends on how extensive and pervasive these costs are; whether they affect goals or choices that are particularly important to the discriminatee’s own conception of herself and of her life; and whether most people in the discriminatee’s society, too, face the kind of deliberative burden that she is facing.

Let us turn to a few examples, to see how the idea of respecting someone as a person capable of autonomy might help us to distinguish cases in which someone has a right to

¹⁹ For further defence of this claim, see my arguments in the next section.

deliberative freedom from cases in which she does not. If a Muslim employee is required by her employer to work at the times when her religion requires her to pray, then she will have to choose between continuing in her job and continuing to practice her religion. Given that a person's job and their religion usually reflect choices that are very important to them, forcing someone to choose between their job and their religion is arguably failing to respect them as a being capable of autonomy. By contrast, imagine that this same employee sets up a GoFundMe page in order to finance her religious pilgrimage, but few people donate and she is left without sufficient funds. Although this may have just as great an impact on her deliberative freedom, it is not a failure to respect her as a person capable of autonomy. On the contrary, bearing the costs of her own pilgrimage is, we would say, part of what it is for her to take responsibility for her own life and the costs of some of her chosen activities. As another example, consider the situation of two aspiring hockey players. One happens to live in Canada, where hockey is so popular and so prized that getting a job as a Major League hockey player requires years of training and is highly competitive. This imposes deliberative burdens; but they are burdens that each of us can legitimately be asked to bear themselves, since these are just the burdens that *everyone* must bear in a society that prizes hockey, and they are based on everyone's self-regarding preferences. But it is different for a hockey player who is banned from getting a job in the Major League because she is a woman. This ban seems to fail to show respect for her as a person capable of autonomy, for several reasons. For one thing, it is based on other people's views about who women are and what they can do, which are being used to close off certain options from her. For another, the ban prevents her from making a choice that is open to many other people of the same ability. This is not, of course, a full explanation of the arguments that would need to be made in any of these cases. We might appeal to difficulty or impossibility of changing the trait in question, or to the extent to which the discriminatory practice leaves these claimants' religion or gender always before their eyes. And, as I mentioned above, other considerations will likely be relevant in other cases. But I hope this gives some indication of the ways in which the idea of respect for autonomy provides at least a little guidance.

In determining whether a practice respects someone as a being capable of autonomy –and ultimately, in determining whether this person has a right to a particular deliberative freedom-- it is important to consider not only the situation of that person, but also *the interests of other people*. When we looked, earlier in this chapter, at the case of the Hyperandrogenism Regulations, I promised I would say more at a later point about how we are to factor in the interests of other people. It is time to consider this now.

As I have said, part of what it means for me to take responsibility for my life as an autonomous individual is to make my choices within a framework that is in part defined by your choices and your preferences. We live our lives, not just as beings capable of autonomy, but as beings capable of autonomy *who live among other such beings*. And so whether someone has a right to a particular deliberative freedom in a particular context must depend, not only on facts about that person, but also on facts about the interests of

the other people who are affected by the particular practice that she is challenging as discriminatory. Consider again the Hyperandrogenism Regulations. I argued earlier that we need the idea of deliberative freedom in order to make sense of the complaints of sprinters such as Chand and Semenya: they are alleging, I suggested, that they should not have to bear the costs of other people's assumptions about their gender, or about whether they are "real" women or not. They are alleging that they should be able to run the races they want to race, live where they want to live, without having to face these costs, and without having their gender, or other people's assumptions about it, constantly before their eyes. But theirs are not the only interests at issue, and the mere fact that the Regulations impede their deliberative freedom does not settle the question of whether they have a right to that deliberative freedom. To ascertain that, we need to ask whether excluding them from running in the category of female athletes amounts to failing to respect them as beings capable of autonomy, given the interests of all of the other women that are also at stake. I am inclined to say that it does. There is very little scientific evidence that higher levels of testosterone give these athletes any advantage –and assuming that it does not, then permitting athletes such as Chand and Semenya to run as female does not deprive others of a fair opportunity to compete. Moreover, even if there were some evidence showing that the extra testosterone gave such athletes a minor advantage, it is unclear why we should think that *this* particular advantage takes away a fair opportunity to compete from others, when we do not suppose that *other* natural advantages, similarly, take away anybody's fair opportunity to compete. We do not ban people from the class of female athletes when they have unusually long legs, or genetic variations that provide better blood flow to their muscles. Nor do we ban athletes from the global North, on the grounds that they have unfairly won the birth lottery and get to grow up in countries that can afford to give their athletes better food, better facilities, and better training. So I would conclude that the interests of other women in fair competition are not affected by allowing hyperandrogenist athletes to run as women, and that, given the huge impact of the Regulations on these women's deliberative freedom, we can therefore say that they have a right to deliberative freedom in this context, and that it is infringed by the Regulations. The Regulations therefore amount to wrongful discrimination.

But what about a type of case in which the interests of other people *are* clearly engaged, and make it more difficult for us to conclude that the discriminatee has a right to deliberative freedom? There have been a number of cases recently, in a number of countries, in which visually impaired clients with guide dogs have been denied a ride by Muslim taxi drivers.²⁰ The visually impaired clients have alleged that this amounts to wrongful discrimination. However, the Muslim taxi drivers have argued that it is not wrongful. They have said that dogs are unclean according to their religion; that they can

²⁰ For Canadian cases, see *Gilmour v. North Shore Taxi*, 2006 BCHRT 529; *Malone v. City Wide Taxi* (2016), 83 C.H.R.R. D/99; *Dewdney v. Bluebird Cabs Ltd.*, 2003 BCHRT 7. For American cases, see *National Federation of the Blind of California v. Uber Technologies, Inc.*, 103 F.Supp.3d 1073 (2015) [N.F.B.C. v. Uber]; and *Kauders v. Uber Technologies, Inc.*, 2017 WL 1652551 (2017).

pray only in clean locations and must pray multiple times a day; and that therefore, requiring them to give rides to visually impaired clients would be unfairly burdensome on them, as they would lose considerable time, and hence, considerable amounts of money. The taxi drivers have therefore argued that this does not amount to wrongful discrimination. How, on my view, should we go about thinking through such a case? The question we need to start with is whether the visually impaired clients have a right to this particular deliberative freedom –that is, a right to be free from treating their visual impairment as a cost when they think about where they want to go and how long it will take to get there. This in turn depends on whether we would be treating them as beings capable of autonomy if we denied them this freedom. And to assess this, we have to consider not only the importance for visually impaired people of being able to use cabs as transportation, but also the impact on taxi drivers of having to give them a ride.

This seems to me a genuinely difficult case. On the one hand, visually impaired people are more dependent than others on taxis, since they usually cannot drive by themselves. Moreover, many taxi drivers regularly drive right past people with guide dogs –even drivers who do not have religious reasons for doing so. So we cannot simply say that this has minimal impact on the deliberations of visually impaired people, or that this is an unimportant freedom for them. They cannot usually just go and get another cab; they are reliant on this form of transportation; and for many of them, who are marginalized and live relatively isolated lives, this is the only way in which they can get out and socialize. To have to view your visual impairment as a cost whenever you try to go out and socialize is a considerable burden. On the other hand, it can also be a considerable burden for a Muslim taxi driver to have to drive a visually impaired client. Some Muslim drivers believe that if a dog rides in their taxi, they must clean out their cab seven times, and once with dirt, before praying again in it. One might argue that this is ultimately just a financial consideration for them: what they are really objecting to is the loss of other fares that they will sustain during the time that they are cleaning out their taxi. If this is right, then it seems that their interests ought to be given less weight, in our thought, than those of the visually impaired, who have suffered a loss of deliberative freedom. But perhaps these Muslim taxi drivers could argue that requiring them to give rides to guide dogs would lessen *their* deliberative freedom: they would then be unable to do their job all day, without having before their eyes the very large costs of practising their religion. However, I am inclined to think it matters, as we puzzle through this, that the costs being imposed on these taxi drivers are not the result of other people’s assumptions about them, but are simply the result of their own religion and its dictates. So I am more inclined to think that these costs for them are just the kind of lessening of deliberative freedom that we require ordinary people to bear themselves. By contrast, the lessening of visually impaired people’s deliberative freedoms is due to the Muslim drivers’, and other people’s assumptions about them and their animals. And it seems to me arguable that forcing visually impaired people to suffer the consequences of these people’s assumptions about them is to fail to treat them as beings capable of autonomy. So I am inclined to conclude that, given both the magnitude of the

impact on visually impaired people's deliberative freedom, and the fact that this impact occurs because of other people's assumptions about them, they do have a right to this particular deliberative freedom, whereas the Muslim drivers do not have a right to the freedom to be able to drive their cabs without cleaning them out. But it is not a clear-cut case; and there is room, within my theory, to argue both sides. What is important for our purposes here is not to settle the issue, but to see that whether the discrimination against visually impaired clients in this case is wrongful depends both on the importance to them of the particular deliberative freedom that is at issue, on the nature of the interference with it –that is, the fact that it stems from other people's assumptions about them—and also, on the relative importance of other people's interests, in this case, the taxi drivers' interests.²¹

What about *Masterpiece Cake Shop*? What should we say about the role of other people's interests –and in particular, the role of Phillips the baker's interests—in determining whether Craig and Mullins have a right to the freedom not to have to think about their sexual orientation, and not to bear the costs of Phillips' assumptions about it, when they purchase a wedding cake? Earlier, I tried to make sense of Craig and Mullins' complaint, and I said I thought it was best understood in terms of deliberative freedom. But I did not try, at that earlier stage, to settle whether they had a right to this deliberative freedom –partly because this depends on Phillips' interests also. Do they have a right to this particular deliberative freedom? I think that they do. The decision to get married and to celebrate their wedding is a deeply important one for them, and when they are denied a wedding cake –even if only from one baker-- they are forced to bear the costs of the baker's assumptions about their sexual orientation. Moreover, these assumptions, as I argued earlier, are perilously close to assumptions of lack of worth: the baker's view is that given their sexual orientation, they are not fit to be married, not *worthy* of the institution of marriage. By contrast, the baker's costs are, rather like the taxi drivers' costs in the case we just considered, due simply to his own religion, and not due to the assumptions of others. So if you were persuaded by my argument about the taxi driver case, then you will likely be even more persuaded in this case: we do not treat Craig and Mullins as beings capable of autonomy if we permit them to bear the costs of these assumptions about them, particularly given that these are assumptions about their lack of worthiness to enter an important institution. Moreover, one might argue, as the European Court of Human Rights held in the *Ladele* case, that a person's beliefs about whether *other people* can marry are not at the "core" of their practice of their religion –in the sense that they do not affect the ways in which they, in particular, go about worshipping and living their lives as believers. So one might argue that the baker's position here is even less compelling than the taxi drivers' position, in the case of the client with a guide dog: whereas, in that case, a requirement to drive guide dogs would interfere with a key element of their own practice

²¹ A different way of analyzing this case, which is also consistent with my view, is to say that it is a case in which we wrong someone no matter what we do: perhaps permitting the taxi drivers not to give rides to the visually impaired passengers amounts to wrongful discrimination, but at the same time, requiring them to give rides to those with visual impairments also wrongs them. I discuss cases in which we wrong someone no matter what we do in Chapter 5, Section 5.6.

of their religion –the need to be sure that they pray in a clean space-- a law requiring the baker to bake wedding cakes for same sex couples would not, similarly, interfere with an important part of his practice of his religion.

So far in this chapter, I have argued that discrimination is sometimes wrongful because it violates a right to someone’s deliberative freedom, and I have suggested that a number of factors are relevant in considering whether this person does indeed have a right to that freedom, including the interests of *others*. But I have not yet explained the connection between infringing someone’s right to deliberative freedom and failing to treat that person as an equal. And it is important that I explain this connection. For this chapter, like the previous one, is an attempt to answer the question of inequality: “When we treat some people differently on the basis of certain traits, when and why do we wrong them by failing to treat them as equals?” We saw in Chapter Two that in some cases, we wrong people by failing to treat them as equals because we unfairly subordinate them to others. It is obvious that acts and policies that unfairly subordinate a person fail to treat her as equal. But it not so obvious why we should think that practices that infringe someone’s right to deliberative freedom fail to treat them as an equal. It becomes clearer when we note the importance of autonomy in most of the societies that value non-discrimination.²² Most of these societies hold up, as a social and a political ideal, the idea that each individual ought, as far as possible, to be treated as though they were capable of autonomy.²³ Failing to respect someone as a person capable of autonomy is, in these societies, failing to treat her as an equal. Hence, the value of autonomy helps to explain not only when we have a right to a particular deliberative freedom, but also why, when this right is violated, we fail to treat someone as an equal.

Up to this point, I have defined deliberative freedoms in a largely negative way. They are freedoms not to have to consider certain traits and others’ assumptions about them, freedoms from certain costs and burdens. And I have indicated that these freedoms matter because they enable us to live a certain kind of life, a life that is shaped by our own choices—whatever those choices may be. I have not, however, claimed that living this kind

²² Recall that we are discussing autonomy in the thin sense of “deciding what is important for you and living your life in accordance with it.” One can be autonomous in this sense even if one has chosen to follow the roles scripted for one by a particular religion or a particular culture, as long as there is some scope for personal choice. So I think it plausible to claim that most of the countries that value anti-discrimination and that have anti-discrimination laws value autonomy in this broad sense.

²³ I say “as far as possible” because of course there are people who, because of very severe disabilities, are not capable of autonomy; and it may not be possible to treat them as people capable of autonomy. (Though I am skeptical of how great a number of people really fall into this category: many people with even quite severe disabilities have enough cognitive capacity to make a choice and to conceive of themselves, in some sense, as a being capable of directing their life through their own choices. And I am skeptical, too, of whether we really do give up on the ideal of autonomy even in the case of those very severely disabled persons who really do lack this capacity. It is arguable that we still feel ourselves under an obligation to treat them *as if* they were people who had this capacity, partly as a matter of respect. And so, for instance, when we assess what is in their best interests, we don’t just ask what would be best for this body or this bundle of desires; we ask what is in *their* best interests, imagining that they are a person capable of making a choice.)

of life is valuable only because, and only insofar as, it tracks what is objectively valuable. In his recent book, Tarunabh Khaitan makes just this claim.²⁴ He argues that it is the ultimate goal of anti-discrimination law to secure for all of us the conditions under which we can live lives that are valuable, in an objective sense—that is, lives spent in pursuit of options that are not just valued by their pursuer, but also objectively valuable. In the next section of this chapter, I shall argue that this perfectionist conception of freedom in fact sits uncomfortably with the structure and aims of anti-discrimination law.²⁵ We can better understand both the unfairness of discrimination and the aims of anti-discrimination law by appealing to a non-perfectionist conception of freedom.

3.3 *The Perfectionist Challenge*

We should note at the outset that Khaitan's particular perfectionist conception of freedom is broad-minded and liberal.²⁶ Khaitan assumes the truth of value pluralism—that is, the idea that there are many very different and even incompatible ways of living a successful life. And he allows that a successful life need not involve full self-realization or the maximal pursuit of certain values: it can count as a successful life, in his view, as long as it involves the pursuit of some valuable relationships and goals. So my objection to Khaitan's perfectionist conception of freedom is not that it is overly demanding or implausibly narrow. Rather, it is that, contrary to Khaitan's arguments, it sits uneasily with the aims and structure of anti-discrimination law, and with some of our basic moral intuitions about what we are trying to protect, in cases of wrongful discrimination.

One of Khaitan's central arguments for the relevance of his perfectionist conception of freedom to anti-discrimination law is that we must invoke such a perfectionist conception if we are to understand the function of the prohibited grounds of discrimination.²⁷ Khaitan notes that most courts and tribunals, when considering whether some trait should be protected as a prohibited ground of discrimination, have used a

²⁴ Tarunabh Khaitan, *A Theory of Discrimination Law* (Oxford: Oxford University Press, 2015) at pp. 92–93.

²⁵ Khaitan's account differs from my own in one further way, which I shall not discuss but which it is important to note, in order to understand his view. Whereas I argue that one reason why certain discriminatory acts are wrongful is that they infringe some people's freedoms, Khaitan appeals to freedom only to justify the *rules* of discrimination law, rather than to explain the unfairness of particular *acts* of discrimination. What justifies these rules, on his view, is that they eliminate the kinds of persistent relative disadvantages between groups that prevent members of the underprivileged groups from having the conditions necessary for freedom, conceived as a life spent in choosing and pursuing valuable activities. So Khaitan is not claiming that particular acts of discrimination are wrongful when, or because, they deny certain freedoms to people; rather, his claim is only about what justifies the rules of discrimination law. (He has a different account of why we might consider particular discriminatory acts to amount to personal wrongs, and I have argued elsewhere that this account seems to sit uncomfortably with his main account of why the rules of discrimination law are justified: see Moreau, "Discrimination and the Freedom to Live a Good Life: A Review of *A Theory of Discrimination Law* by Tarunabh Khaitan," *Law and Philosophy* 35(5) (2016), pp. 511–527.

²⁶ Khaitan, *A Theory of Discrimination Law*, *supra* note 25 at pp. 93–94.

²⁷ Khaitan, *A Theory of Discrimination Law*, *ibid.* at pp. 56–60, 134–137.

combination of two criteria: (i) either the ground must be immutable, in the sense that it is not within our immediate voluntary control to change it, or (ii) the ground must demarcate something that we as a society think is a matter of fundamental personal choice. These two criteria have seemed to many scholars to sit rather uncomfortably together, for the first seems to suggest that we care about protecting people from the costs of traits that they have *not* chosen, whereas the second seems to suggest that we care about protecting people from the costs of certain things that they *have* chosen. Khaitan argues that this apparent inconsistency is based on a mistaken reading of these criteria. In his view, the point of the two criteria is to identify which traits or groups are “not immoral” and which are “positively valuable.” He suggests that “when a ground is immutable, possessing it is generally not immoral” and “when a ground represents a valuable fundamental choice, it is positively valuable rather than merely not-immoral.”²⁸ And so he sees the assessment of protected grounds as “an objective moral assessment.”²⁹

There are, however, several potential problems with this view of the assessment of protected grounds. First, although it is true that “valuable fundamental choices” are valuable, criterion (ii) doesn’t actually mention objectively valuable fundamental choices. It just appeals to “fundamental” choices. Khaitan often adds in the word “valuable” or uses the terms “fundamental” and “valuable” interchangeably; but a “fundamental choice,” in the sense that courts have been concerned with, is just a choice that we think people ought to make for themselves, or that people hold particularly dear to themselves. As Khaitan himself acknowledges, “the choice in question is important because it is fundamental *to the person whose choice it is*.”³⁰ But this means that criterion (ii) is entirely neutral on whether the choice is actually objectively valuable or not. It might not be; and yet it could still constitute a prohibited ground of discrimination. Indeed, it is arguable that many if not most of the religions that anti-discrimination law protects are actually not valuable for their members at all—or at least, they limit the roles that their members can adopt and the beliefs their members can hold in pretty substantial ways, and in this sense would be deemed a “harm” on any conception of value that one could plausibly invoke as part of a perfectionist account of freedom. But they are duly protected by anti-discrimination law. If anti-discrimination law really did care about giving us access *only* to valuable opportunities, it seems likely that criterion (ii) would not be neutrally worded and would appeal to some objective conception of value. But it does not.

A different problem arises with Khaitan’s interpretation of criterion (i), as a way of allowing in only those traits possession of which is “not immoral.” For recall that courts have allowed a trait to constitute a protected ground of discrimination if it satisfies *either* (i) or (ii). This means that something can constitute a protected ground as long as it is, on Khaitan’s view, “not immoral.” But if we genuinely cared about giving disadvantaged groups secure access to a range of objectively valuable options, it wouldn’t be enough to

²⁸ Khaitan, *A Theory of Discrimination Law*, *ibid.* at p. 60.

²⁹ Khaitan, *A Theory of Discrimination Law*, *ibid.* at p. 137.

³⁰ Khaitan, *A Theory of Discrimination Law*, *ibid.* at p. 60, Khaitan’s own italics.

ensure that they were not discriminated against on the basis of traits that were not immoral. We would have to take steps to give them options that were positively valuable.

There is a way to reconcile criteria (i) and (ii) without appealing to the relevance of choice, but also without assuming, as Khaitan does, that the assessment of protected grounds is a moral exercise. Both those traits that are out of our immediate voluntary control and those that are matters of fundamental personal choice may be *traits that we think people ought not to bear the costs of having*, or at least not in certain circumstances. Such facts—the fact that, as a society, we think certain groups ought not to have to shoulder the costs within the workplace of, for instance, observing a holy day on a day that is to others a work day, or being the one who carries a child and gives birth to it—these may be what tie together the different protected grounds of discrimination. This is what my deliberative freedom account implies. It allows us to reconcile criteria (i) and (ii) without presupposing a perfectionist conception of freedom.

I have now argued that a perfectionist conception of freedom does not in fact, as Khaitan asserts, make best sense of our criteria for recognizing something as a prohibited ground of discrimination. But there is a further difficulty with Khaitan's perfectionism in the context of anti-discrimination law. This is that anti-discrimination law makes no effort to investigate whether the particular resource or opportunity that is being denied to a given group is itself valuable or not valuable. Within the private sector, anyone offering goods and services to the public is under a duty not to discriminate, regardless of how useless or even morally questionable the good or service they are offering is. A candy store or a soda pop store is obliged to admit everyone, regardless of their race or gender, even though the only immediate benefit this gives to these groups is an equal opportunity to rot their teeth and develop diabetes. A night-club that features degrading live strip shows is similarly obliged not to deny people entry because of their race or gender, even though there is ample evidence that such clubs reproduce and reinforce gendered and racial inequalities and that therefore, what this really accomplishes is to give everyone an equal opportunity to be included in the reproduction of inequalities. If the aim of anti-discrimination law really were to give us all a sufficient range of objectively *valuable* options, one would expect at some stage to have some inquiry into whether the good or service offered really is valuable, even if the inquiry were not decisive and other considerations, such as the impact of continued exclusion on the group's self-respect and negative freedom, were also relevant.

There is also something about "objective moral assessments" of people's traits, or moral assessments of the opportunities that they want access to, that seems to me to run very deeply against the goals of anti-discrimination laws, at least on a lay understanding of these goals—and relatedly, to stand in some tension with what outrages us about discrimination from a moral standpoint. Consider the legal point first. Whatever theory of anti-discrimination law we endorse, I think we cannot deny that part of the point of such laws is to avoid placing some people in a position where they are making pronouncements

about the moral value of belonging to a certain group or having access to a certain opportunity. So a regime of anti-discrimination law that did require judges or tribunals to ask questions like “Is this proposed protected trait really objectively valuable?” or “Would eliminating this disadvantage between group X and group Y really increase the valuable options available to group Y?” would be problematic, at best. It would require judges and tribunals to adopt a paternalistic stance towards the very groups who have, historically, been undervalued and spoken for by others, and whose conceptions of the good life have been ridiculed and misunderstood. It would require officials to make evaluative judgments about the disadvantaged groups that so desperately need a chance to speak for themselves and need to have their own conceptions of value taken seriously.

Khaitan might say that, far from being a problem for his theory, this shows that there is good reason why anti-discrimination law would not include an explicit legal stage of inquiring into the value of the good that is being denied to a particular claimant in a particular case. But it also shows, I think, that it would be problematic for anti-discrimination law to engage in any kind of objective moral assessment of the protected grounds. And there is a moral analogue to this objection. I argued in Chapter One that discrimination subordinates certain social groups, turning them into second-class citizens. As we saw there, subordination does not just involve differences in power and *de facto* authority. It also involves a pattern of giving members of this group lesser consideration or censure—implying that they are less worthy than we are. And it involves persistent structural accommodations that function to render them invisible and to silence them. If our aim is to increase the visibility of these groups and to lift them up to a position of equal status, ensuring that they are treated as just as worthy as the rest of us, it seems an odd way to begin, to propose that we ought to interrogate them to determine whether membership in these groups really is valuable or whether their values really are correct. This approach seems much more likely to perpetuate inequalities than to eliminate them.

Khaitan might reply that this is a pragmatic problem or a problem of institutional design, rather than one that demonstrates that his perfectionist conception of freedom cannot explain why discrimination is unfair. But I think it is not just a pragmatic problem or a problem of institutional design. If the arguments of these first two chapters are right, then the unfairness involved in discrimination involves failing to treat other people as equals. And perhaps we cannot treat them as equals as long as we assume an evaluative stance towards them, taking ourselves to be in possession of an objective conception of the good and evaluating what we think are their values against this idea.³¹ If this is right, then we can only achieve a state of affairs in which all of us are equals, and in which we all have the

³¹ For concerns about patronization in the context of luck egalitarianism, see Elizabeth Anderson, “What is the Point of Equality?”, *Ethics* 109(2) (1999), pp. 287–337, especially at pp. 289, 306–307, 311–313. See also Elizabeth Anderson, “The Fundamental Disagreement between Luck Egalitarians and Relational Egalitarians,” *Canadian Journal of Philosophy: Supplementary Volume on Justice and Equality* 36 (2010), pp. 1–23.

social bases of self-respect, if we refrain from imposing an objective conception of the good on different social groups.³²

3.4 *Why Indirect Discrimination, Too, Can Interfere with Deliberative Freedom*

I have now explained what deliberative freedom consists in, why it is important, and what sorts of considerations we need to take account of when determining whether a particular claimant has a right to it. I have also tried to show that we need the idea of deliberative freedom in order to capture the claimants' complaints in at least some prominent cases of wrongful discrimination. But all of the cases that I have discussed at length in this chapter –the Hyperandrogenism Regulations, *Masterpiece Cakeshop*, and the case of the Muslim taxi driver and the visually impaired client—have been cases of direct discrimination.³³ That is, they are cases in which people are denied a benefit on the basis of a trait that is a prohibited ground of discrimination, such as gender or sexual orientation – or, as in the Muslim taxi driver case, in which people are denied a benefit on the basis of a trait (having a guide dog) that is very closely linked to a prohibited ground (disability). But what about cases of apparently wrongful indirect discrimination, cases in which a practice disproportionately disadvantages a particular group of people who have a certain trait, but it does not explicitly single them out because of this trait or some closely related trait? Kasper Lippert-Rasmussen has recently argued that it is “much less plausible to think that we have a deliberative freedom which is violated in cases of indirect discrimination.”³⁴ I shall offer both some case analyses, and some more general reflections, which together will aim to show that deliberative freedoms can also help us to explain the wrongfulness of some cases of indirect discrimination.

Consider a core case of indirect discrimination that seems wrongful: the Royal Canadian Mounted Police (RCMP)'s longstanding policy of requiring that all police wear the “Stetson hat” at all formal ceremonies. This policy was abandoned in the 1990s because it was realized that it discriminated indirectly against Sikhs, who could not wear the Stetson hat at the same time as wearing their turbans, and who were therefore unable to join the RCMP. Let us compare this case with its easy analogue in direct discrimination: the recent Quebec ban of turbans on the soccer field. Certainly the ban on turbans on the field requires Sikh players to treat their religion as a cost, and constantly to keep their religion before their eyes when considering whether and where to play soccer. So it has an impact

³² Khaitan might argue that this worry is attenuated by his value pluralism: he recognizes that there can be many incompatible but equally valuable ways of living one's life. But this does not seem to me to do away with the difficulty here. Value pluralism does not claim that all ways of life are valuable; it just claims that *many* are. Those who adhere to ways of life that are perceived to be lacking in objective value will still feel as though they have not been treated as equals; and the mere act of checking whether a particular group is engaged in valuable activities is problematic, even if the answer we give ends up being a positive one.

³³ See my definition of direct discrimination on p. , Chapter 1.

³⁴ Lippert-Rasmussen, *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*, *supra* note 11 at p. 188.

on their deliberative freedom. But the Stetson hat policy seems to have no less of an impact on Sikhs' deliberative freedom. This policy made an issue of their religion, and turned it into a serious cost for them: they could not, under this policy, become RCMP officers. Moreover, the policy presupposed, and in turn reinforced, an image of RCMP officers in the public mind as strong and brave, as pillars of our society, and as white and Christian and arguably an image of these officers as strong and brave and pillars of our society *because* they were Christian and white. In this way, too, the policy imposed costs on being Sikh and forced Sikhs to think about their religion, and about others' assumptions about who they were and who they could aspire to be because of their religion—not just when they decided which careers to pursue, but whenever they thought of themselves and their relative place in Canadian society. It seems quite plausible to say therefore both that their deliberative freedom was affected by the policy and that they had a right to this particular deliberative freedom—that is, the policy interfered with choices that they regarded as fundamental to their lives, on the basis of other people's "other-regarding" preferences, and so failed to respect them as people capable of autonomy.

One might object here, however, that the Stetson hat policy did not reflect any other-regarding preferences. Surely it was simply a policy adopted on the basis of *self-regarding* preferences: the RCMP as an organization, preferred to follow their tradition of using this hat, because they themselves wanted to wear it. What is other-regarding about this? Even if it is true, as I have claimed, that the tradition of using this particular hat taps into stereotypes about ideal police officers being white, this is a stereotype about what these RCMP officers wanted *themselves* to be, not a stereotype about what they wanted or thought about others.

What this objection ignores, however, is the fact that many apparently self-regarding preferences define the "self" in relation to an "other"—and therefore implicitly, but just as loudly, assume certain things about that other. Indeed, these apparently self-regarding preferences are so damaging precisely because, although they purport to be about only one group, they subtly alter our perceptions of other groups, and impose further, silent and unacknowledged costs on those other groups, costs that we are then inclined to attribute to these own groups' inner failings or lack of ability.³⁵ This is exactly how the stereotype of the ideal RCMP officer wearing the Stetson hat works. Our imagined officer is strong and brave in part because he is implicitly seen against a backdrop of other cultures who are assumed not to be civilized, respectable, and dependable. It may sometimes be difficult to determine which self-regarding preferences are indeed largely self-regarding, and which are so closely dependent on certain other-regarding preferences that we ought to treat them as effectively other-regarding for the purposes of assessing their impact on other people's deliberative freedom. And it may be true that more cases of indirect discrimination raise this difficulty than do cases of direct discrimination, because

³⁵ See, for instance, Martha Minow, *Not Only for Myself: Identity, Politics and Law* (New York: New Press, 1997); and Nitya Iyer, "Categorical Denials: Equality Rights and the Shaping of Social Identity," 19 *Queen's LJ* 179 (1993), pp. 179-208.

practices that constitute indirect discrimination are often adopted for just such apparently innocuous, allegedly self-regarding reasons, whereas directly discriminatory policies explicitly mention the group that is excluded. But this is not a problem for my account. It simply reflects the very real complexities that arise in cases of indirect discrimination—complexities that I shall have more to say about in Chapter Six, when I discuss indirect discrimination in more detail.

As further examples of indirect discrimination, consider some cases that we examined in Chapter Two: the employers that use tests for promotion that are passed in far greater proportions by Caucasian employees than by blacks and Latinos, and the restaurants that stock no uniforms that are easy for people with certain muscular disabilities to put on. In Chapter Two, our focus was subordination, and so our interest in these examples lay in how these forms of discrimination contributed to the subordination of these racial minorities and of persons with disabilities. But we can see now that these policies also deny members of these groups deliberative freedoms, in circumstances where they seem to have a right to them. If I am black or Latino, and year after year I see blacks and Latinos writing the test and yet the relative percentage of black and Latino supervisors never increases, this does make race an issue in the workplace. It will affect who I choose to associate with, how I view myself, and what value I think I have to my employer. It will affect how I interpret other people's actions, and how I choose to act given the interpretations they are likely to place on my actions, as a member of this racial minority. The same is true of the restaurants that stock only conventional uniforms and no uniforms for persons with disabilities. It's likely that most of these restaurants never explicitly decided to exclude people with disabilities. They just bought an array of easily available, cheap uniforms for the kinds of people who normally apply for jobs as waiters and waitresses, and these do not include people with disabilities. But even so, the absence of these uniforms forces people with disabilities to consider their disability as a cost, and to have it before their eyes, before they fill out a job application or attend an interview.

Lippert-Rasmussen does not argue in any detail for his claim that in cases of indirect discrimination, a person's deliberative freedom is not infringed. He does, however, state that on my account "in cases of indirect discrimination, the sense in which one is denied an opportunity because one has a certain extraneous trait is very different from the sense in which this is the case in instances of direct discrimination."³⁶ By this he seems to mean that in cases of indirect discrimination, the exclusion of a particular group on the basis of a certain prohibited ground is often not a direct effect of the policy but rather connected to the policy through a complex causal chain. So perhaps his worry is that because this causal chain is longer and more mediated in cases of indirect discrimination, the deliberative freedom of the victim is less likely to be affected. But this reasoning is problematic. First, it is not invariably true of cases of indirect discrimination that the causal

³⁶ Lippert-Rasmussen, *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*, *supra* note 11 at p. 188, note 77.

chain connecting the policy and the eventual exclusion is longer or more mediated by other factors than it is in cases of direct discrimination—as the two parallel cases excluding Sikhs show. In the Stetson hat case, the exclusion of Sikhs is no less direct an effect of the policy than is the exclusion of Sikhs from the soccer field by the explicit ban. And secondly, the impact on the deliberative freedom of the excluded group does not seem to depend, for its severity or its significance, on the length or complexity of the causal chain.³⁷ As I argued above, even in a case such as the promotions test that disproportionately disadvantages racial minorities, where the causal chain between the test and the exclusion of racial minorities is quite complex and mediated by a variety of factors, it can still be the case that the policy has a significant impact on the deliberative freedom of those who are excluded. Whether they have a right to this deliberative freedom will depend on whether, under the circumstances, they are still shown respect as beings capable of autonomy. But this does not seem to depend on the length or complexity of the causal chain.

A person's deliberative freedom, then, can be lessened in cases of indirect discrimination just as it can in cases of direct discrimination; and they can have a right to that deliberative freedom in both cases, though of course whether they have that right depends on a complex assessment of the impact of the policy on them and on other people affected by the policy, as I noted above.

3.5 Deliberative Freedom and the Role of the Prohibited Grounds

I have now argued that in some cases of wrongful discrimination, the source of the wrong lies not primarily in unfair subordination but in an infringement of someone's right to deliberative freedom. It may seem that in cases of this type, there is no necessary role for the prohibited grounds of discrimination. Surely what matters is simply whether a particular discriminatory practice denies someone a deliberative freedom to which they have a right. Is it necessary to ask, in addition, whether the practice has treated them differently on the basis of some recognized prohibited ground?

Although this is not a necessary further step, the prohibited grounds of discrimination do play an important heuristic role in cases involving infringements of a right to deliberative freedom. Or rather, they play a number of heuristic roles. I emphasized earlier, in my discussion of white privilege, that when a person lacks a particular deliberative freedom, they are forced to have a certain trait of theirs—whether it represents an actual part of them or only a part of them that others presume is definitive of them—always before their eyes. If you are a black driver in the United States, you are never allowed to forget your blackness. If you are a female professional in a male-

³⁷ In my view, the length of the causal chain is relevant to the responsibility of the discriminator only—and in the passage from my earlier article that Lippert-Rasmussen cites in support of his worry about indirect discrimination, I am discussing the causal chain and its relevance to the *responsibility* of the discriminator, not its relevance to the deliberative freedom of the victim.

dominated profession, you are never allowed to forget that you are a woman playing what is still essentially a man's part. You can play that part wearing trousers or in a "feminine" way; but even when you are positively appraised, it is always in relation to how well you have assimilated into the role that you were not supposed to occupy, or how well you have retained your feminine character while still doing the job. In relation to this kind of privilege—or rather, the lack of it—we can see the prohibited grounds of discrimination marking out the particular traits that often deny people deliberative freedom. It is indeed most often traits such as race, gender, sexual orientation, religion, and disability that people are forced to have before their eyes, in a way that leaves them without the space to become who they want to be. And these same traits are often the ones that we are forced to bear the costs of, in situations in which we ought not to bear those costs.

Of course, as I acknowledged earlier, there are some contexts in which it is quite appropriate for people to bear the costs of having each of these traits: I can legitimately be asked to bear the costs of my religious pilgrimage, of my pregnancy (assuming I have consented to continuing it), of the Black Lives Matter rally that I attend. So it is never enough simply to show that someone has been treated differently, or even disadvantaged, on the basis of a prohibited ground, in order to show that they have had a right to deliberative freedom infringed. The question is always whether the person had a right to that particular deliberative freedom—and the answer, as I argued earlier, depends on whether or not recognizing the right in that context would amount to failing to treat her as a being capable of autonomy. But in answering this question, the prohibited grounds can play a heuristic role, pointing us towards the right questions: Is this person being asked to bear costs that she ought not to bear? Is she being asked to bear the costs of other people's other-regarding preferences, and so not being treated as a being capable of autonomy in her own right? Is she being asked always to have a certain trait before her eyes, when she ought not to?

I have now tried to show that some cases of wrongful discrimination are best understood as wrongful because they infringe some people's right to deliberative freedom. In the Appendix to this chapter, I shall turn to some objections that have recently been made to this approach, and I shall offer some replies to these objections.

Appendix to Chapter Three: Replies to Critics

In this Appendix, I reply to several critics who have engaged with my work on deliberative freedom and who have offered counter-examples designed to show that we cannot explain wrongful discrimination by appealing to the idea of deliberative freedom.

Reply to Lippert-Rasmussen

Kasper Lippert-Rasmussen has argued that lessening someone's deliberative freedom is neither necessary nor sufficient for wrongful discrimination.³⁸ He raises a number of purported counter-examples that are designed to show this. Three purport to show that a lessening of deliberative freedom is not necessary for wrongful discrimination, by presenting scenarios in which there is unfair discrimination but no lessening of a person's deliberative freedom. And three purport to show that a lessening of a person's deliberative freedom is not sufficient for wrongful discrimination, by presenting scenarios in which there is an infringement of someone's deliberative freedom but no wrongful discrimination. In this reply, I shall consider each of these examples in turn. My aim is partly to contest the conclusions that Lippert-Rasmussen draws from them. But I also think that his analysis of the examples misconstrues the nature of deliberative freedom and the role that it plays in my account of the unfairness of discrimination. So my other aim is to clarify the relevant features of my view.

In this spirit of clarification, before turning to the six purported counter-examples, I should note one important feature of my view that is overlooked in Lippert-Rasmussen's discussion. He presents counter-examples to the view that the wrongfulness of discrimination results from its "restriction of deliberative freedom" or its "infringement of deliberative freedom." But my claim is not that each and every lessening of someone's deliberative freedom constitutes wrongful discrimination. It would, for some of the reasons I have already examined, be quite implausible to claim that we are wronged whenever our deliberative freedom is lessened. My claim is rather that in certain circumstances people have a right to deliberative freedom, and that discrimination is wrongful when it infringes this right.

Bearing this in mind, let us consider first the three examples that Lippert-Rasmussen gives to show that an infringement of a right to deliberative freedom is not a necessary condition of wrongful discrimination. It might be thought that, because the pluralist theory defended in this book invokes deliberative freedom only as one of the sources of wrongful discrimination, it is not necessary to respond to these three examples at all. For the pluralist view developed in this book does not claim that infringements of someone's right to deliberative freedom are *necessary* conditions of unfair discrimination. I

³⁸ Lippert-Rasmussen, *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*, *supra* note 11 at pp. 185–189.

do, however, still claim that such infringements play an *important* role in explaining the unfairness of many core cases of discrimination. So I think there is still value in considering these examples.

The first example concerns a case of fixed costs, in which whatever a woman does, her gender will be a liability. Lippert-Rasmussen asks us to imagine that “everyone faces the same opportunity sets, except that for the fact that the expected value of any available opportunity is 10% higher for a male individual than for a female individual.”³⁹ He notes that the women in this example are “free from” the costs associated with, and the deliberative pressures imposed by, their gender, since “as a matter of fact they can do nothing to avoid them.”⁴⁰ I considered a version of this example earlier in this chapter, when I looked at female athletes and their inability to generate salaries equal to their male counterparts. I noted there that these athletes do have their deliberative freedom lessened, precisely because their gender has been made into an issue. Their gender is not deliberatively irrelevant, simply because they can do nothing about it: on the contrary, because it is the cause of their difficulties and they know it, it will hang over them, constituting what I earlier called a “burden” on them. So this first example is not, in fact, a case in which there is unfair discrimination but no lessening of deliberative freedom. Lippert-Rasmussen is incorrect in suggesting that fixed costs are irrelevant to a person’s deliberative freedom on my view.

Lippert-Rasmussen’s second example concerns an employer who “directly discriminates against women in hiring, but also indirectly discriminates against men in a way that exactly counterbalances her direct discrimination against women.” This employer “makes successful, good faith efforts to make applicants think that they can decide whether to apply for a job with her independently of their sex.” It is not clear from this example whether the reason the female applicants’ deliberative freedom is purportedly not affected is that these applicants are deceived, or whether the reason why it is not affected is that the “counterbalancing” effects of the indirect discrimination against men result in a situation in which the female applicants are not disadvantaged at all. If the former is what Lippert-Rasmussen has in mind, then I would deny that the applicants’ deliberative freedom is unaffected. For recall that deliberative freedoms involve freedoms of action as well as freedom of thought; and in such cases of deception, a person still lacks the relevant freedom of action. If the latter is what he has in mind, then I have to say that it is unclear to me what it would mean for indirect discrimination to “counterbalance” direct discrimination, or how it could be possible for a single policy to discriminate against two cognate groups in these two different ways at the same time. But if “counterbalance” means that the female applicants are not in fact disadvantaged, then it seems likely that we would deny that they faced wrongful discrimination: this might be a case in which they were

³⁹ Lippert-Rasmussen, *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*, *ibid.* at p. 187.

⁴⁰ Lippert-Rasmussen, *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*, *ibid.* at p. 187; subsequent quotations are also from this page, unless otherwise indicated.

discriminated against, but the discrimination was not wrong. So this example similarly fails to constitute an example in which there is wrongful discrimination but no negative effects on a person's deliberative freedom.

Lippert-Rasmussen's third example involves an employer who is an "incompetent sexist." This employer wants to exclude women and writes job advertisements saying "Men only need apply"—but, through a series of bumbling errors, he ends up hiring women rather than men. Lippert-Rasmussen states that this employer does not "restrict anyone's deliberative freedom"—yet, he states, we would think of this as a case of wrongful discrimination. However, what makes this a case of wrongful discrimination is the *advertisement* that explicitly excludes women. And this advertisement does lessen women's deliberative freedom, in ways that are not erased by the fact that women end up being hired by this employer at the end of the day. Indeed, this example reveals that whether someone is denied a certain deliberative freedom, and whether they had a right to that deliberative freedom, do not depend only on the overall balance of advantages and disadvantages—on whether, at the end of the day, this person ends up better off than they were before. It is, as I have argued, a question of whether she has been shown respect as a person capable of autonomy. And it is arguable that the advertisement in this example does not show such respect for her; even if she ultimately gets the job. So this third example does not show that a right to deliberative freedom is not a necessary condition for wrongful discrimination. (Though, as I have mentioned, I am not claiming that it is a *necessary* condition for wrongful discrimination—only that it is an important source of the wrongfulness of many cases of discrimination).

What about Lippert-Rasmussen's other examples, which aim to show that lessening deliberative freedom—or, as I would prefer to say, "infringing a person's right to deliberative freedom"—is not a sufficient condition for unfair discrimination? These examples are all presented as instances in which a discriminatory act *does* lessen someone's deliberative freedom, and yet the discrimination is not wrongful or unfair. It is worth noting once again that my view is not that merely lessening a person's deliberative freedom is sufficient to constitute an unfairness—the view is rather than denying someone's right to deliberative freedom wrongs them. With that in mind, let us consider these three remaining purported counter-examples.

One involves a scenario in which "the only way to give everyone better options is through reducing people's deliberative freedom."⁴¹ Lippert-Rasmussen imagines two options: one "yielding 10 and 12" for men and women respectively, and one "yielding 15 and 17" in which although everyone is better off, men face discrimination and so have their deliberative freedom lessened. Lippert-Rasmussen does not specify what exactly these numbers are supposed to represent, but I assume they are supposed to represent the results of some kind of overall cost-benefit analysis that takes into consideration

⁴¹ Lippert-Rasmussen, *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*, *ibid.* at p. 188.

everything that is valuable or disvaluable in these people's lives. I shall later question whether such a cost-benefit analysis is possible. But first, we need to amend this example to suppose that on the second scenario, as my view would require, men do not just have their deliberative freedom lessened, but they have a right to deliberative freedom and it is violated; for unless their right is violated, there is no wrongful discrimination on my view. If this is the case, however, then on the second scenario, the men must not be shown respect as beings capable of autonomy, and consequently, they are not treated as equals. But it is then unclear to me how this is consistent with their being at level "15" in this scenario but only at level "10" in the first scenario. Whatever disvalue we assign to not having one's autonomy respected, and to not being treated as an equal, this disvalue would surely, in the eyes of many moral theorists, be large enough to outweigh even a very large gain in welfare. Moreover, I think many moral theorists would balk at the assumption that we can weigh on some single metric things as different as gains in welfare, on the one hand, and such disvalues as disrespect for one's autonomy and not being treated as an equal. So many would hold that there is no basis for saying in such an example that "everyone benefits overall from the discrimination"—for there is no sound basis on which we could make this overall assessment.⁴² Finally, even if we could coherently suppose that men were at level 15 in the second scenario but only level 10 in the first, and even if there were some single metric for weighing all value and disvalue, it does not seem to me to follow without further argument that there is no wrongful discrimination in this case. It is certainly a coherent moral position to hold that an act of discrimination could still wrong the discriminatee, even if it is all things considered the best act for the discriminator to perform under the circumstances.⁴³

Lippert-Rasmussen's next objection compares two cases: a case of an employer who engages in objectionable nepotism and excludes "non-family" applicants on the grounds of their family status, and the case of an employer who excludes people on the basis of their green eyes. He uses these two cases to make two claims. First, he suggests that there is no difference in the effects on a person's deliberative freedom in these two cases and in a case of ordinary discrimination on the basis of race, and yet we would think of racial discrimination as being unfair in a different and special way from nepotism and green-eye exclusion. And this suggests that an appeal to deliberative freedom does not capture what is distinctive about the kind of wrongness that is involved in discrimination. And second, he argues that we would normally think of racial discrimination as a more serious wrong. Nepotism seems to me to be a form of discrimination coupled with a set of other, unrelated wrongs: an abuse of power, a selection of candidates on the basis of arbitrary features (such as their relation to you) rather than on the basis of their merit alone, and so on. So we

⁴² Although Lippert-Rasmussen presents his analysis of my view, and his discussion of discrimination, as neutral between different moral theories, he often seems to make assumptions that are most natural for welfare consequentialists—such as the assumption that we can come up with a single figure that represents someone's "level," or the assumption that the values promoted by or realized through a particular act can be weighed or compared on a single metric.

⁴³ See Chapter 7 for further exploration of such cases.

can, I think, look to these other wrongs—however we define them—to explain the difference between nepotism and mere discrimination on the basis of family status. The case of exclusion on the basis of green eyes might seem more difficult. But I think my account offers a ready explanation of why this does not constitute unfair discrimination, and of why the unfairness here seems different from the unfairness involved in discrimination. We do not have a *right* to deliberative freedom in circumstances where we are denied a job because of our green eyes. No one has a right not to have their green eyes held against them. Why? Because this is not the kind of thing that fails to show respect for someone as a being capable of autonomy. To be sure, it leaves one feeling that one's employer has acted arbitrarily and on a whim. But that is different. And it does seem less serious than the cases of wrongful discrimination that Lippert-Rasmussen compares with it. So my account does give us ways of distinguishing wrongful discrimination from nepotism and green-eye exclusion.

Finally, Lippert-Rasmussen imagines two cases, both of which seem to involve a lessening of deliberative freedom, but which seem to differ in their degrees of moral seriousness. He tries to use these to show that decreases in deliberative freedom aren't sufficient to explain the wrongfulness of discrimination. He supposes that one employer "gives men and women the same opportunities but gives the impression that she does not." We are, I think, supposed to imagine that this employer gives such a strong impression of excluding one group that this group does, in their deliberations, have to treat their gender as a cost, or at least has to have it constantly before their eyes when applying. By contrast, the other employer gives the false impression that she does *not* discriminate, and does this so successfully that in their deliberations, no one feels they need to treat their gender as a cost or have it before their eyes. But they still have their deliberative freedom lessened, because they lose the associated freedom of action. So both groups suffer from a decrease in deliberative freedom. But, he argues, we would view the second employer's act as "more wrong" than the first, and that obviously an appeal to deliberative freedom can't explain why. As I have argued, the relevant question on my view is not whether a particular group has some decrease in their freedom, but whether they have a right to deliberative freedom in these circumstances, and whether that right has been violated. If the answer in both cases is yes, then it is true that my view implies that both are wrongful.

Does my view imply that they are equally wrongful? Given that my account is now a pluralist one, I think it does not imply this: it is open to me to argue, and seems quite plausible, that the employer who actually does give different opportunities to men and women subordinates the one gender, whereas the one who only gives a false impression of doing so doesn't contribute as substantially to social subordination, and doesn't explicitly mark out one group as inferior. So we could appeal to subordination to explain why the one case might seem unfair in an additional or further way. But *pace* Lippert-Rasmussen, I do not think it is obvious that the one act is "more wrong" than the other. If we are inclined to think that the latter is "more wrong," I wonder whether this is really because we are confusing a judgment about wrongness with a judgment about blameworthiness. The

employer who actually discriminates may seem more blameworthy than the one who gives the impression that she discriminates but in fact does not—particularly if we are tacitly supposing that the latter gives this impression mistakenly or out of carelessness, rather than deliberately. But this does not show that the former employer’s act constitutes a more serious wrong—it only shows that she is more culpable for having performed it.

Reply to Campbell and Smith

Colin Campbell and Dale Smith have recently argued that an attempt to explain the unfairness of discrimination by appealing to deliberative freedom will have implausible implications in cases of affirmative action. So it seems worthwhile clarifying what my view implies about affirmative action.

Consider first the kinds of affirmative action policies that deny people a benefit simply and straightforwardly because of a certain protected trait—such as the policies at some Canadian universities of opening jobs in Indigenous Studies only to applicants who are indigenous, and not to applicants of any other race. Suppose one of the people who would like to apply for such a job is a Canadian scholar of European heritage—and suppose she has devoted years of her life to research in Indigenous Studies and to developing deep and meaningful ties with indigenous communities. It does seem that this affirmative action policy lessens her deliberative freedom. Since the policy refuses even to consider her for a job in her field because of her race, it turns her race into quite a serious cost for her. It also means that she cannot work (or more accurately live, since this work is her life) without having before her eyes the fact that no matter what she does, no matter how many scholarly books she writes or personal contributions she makes to particular indigenous communities, she will never be truly indigenous, will never truly belong in the field.

But does it follow that, according to my theory, this affirmative action policy obviously wrongs her, or is unjustified? I do not think so. My theory is consistent with a number of different ways of analyzing this case. On the one hand, one might argue that this scholar does not actually have a right to this particular deliberative freedom. For recall that whether a person has a right to a particular deliberative freedom depends, on my view, on a complex set of factors, including the interests of those whom the policy is intended to protect or benefit. We might argue that this kind of affirmative action policy is needed to increase the visibility of indigenous academics, to give indigenous communities a voice within academia, and to help give indigenous communities more power and authority, relative to other communities. We might note that it is implicit in such policies that they are temporary and remedial, and that this makes a difference to the kind of impediment that they present to our scholar’s deliberative freedom, and to what it says about her. It is not equivalent to a policy banning a person from a certain kind of job because of their race, because it is temporary and remedial and therefore doesn’t impugn her or her race as such. And perhaps, for this reason, the policy doesn’t imply that our scholar will never be “truly indigenous.” Instead, it sends a signal that, for the present, there is a more urgent need for others to be given these jobs.

However, there are other interesting features of this case that might lead us to conclude that our scholar does actually have a right to this deliberative freedom, and that she is wronged by the affirmative action policy. Perhaps, for instance, the interests of indigenous communities could be served just as well through a quota system at this particular university, a system which requires that a certain percentage of hires in Indigenous Studies be from indigenous groups but which does not require that *all* such hires be from these groups. If such an alternative were available and reasonably likely to be effective, then we might be inclined to argue that the current affirmative action policy isn't strictly necessary from the standpoint of the underprivileged group, and that therefore, we fail to show respect for this scholar as someone capable of autonomy if we adopt this policy. So the scholar would have a right in these circumstances, and it would be violated by the affirmative action policy.

There is also a third way in which we might analyze this case, which is consistent with my view as well, and which Campbell and Smith might find more attractive. We might be inclined to say that this scholar has a right to deliberative freedom, and that it is violated; but that even though the affirmative action policy wrongs her, it is nevertheless justified "all things considered," simply because of the urgency of giving indigenous communities a voice in academia and of placing them in positions of power after years of unjust marginalization. Perhaps this urgency means that the University does not need to take the time to find the perfect alternative—the one that would enable our scholar to be hired while also promoting indigenous scholars. We can still, on my view, acknowledge that she is wronged by the policy. But that wrong may be justified, all things considered.⁴⁴

Consider now a different case, such as *Fisher II*, which Campbell and Smith discuss.⁴⁵ In this case, race was merely one among a number of factors treated as relevant to whether applicants were admitted to the University. It was never treated as sufficient, on its own, to deny someone a place. This might make a difference to the policy's impact on the deliberative freedom of members of the privileged group. Since this policy does not deny people a place solely on the basis of their race, members of this group are not, in the same way as our hypothetical scholar of Indigenous Studies, forced always to worry about their race or always to treat it as a cost. And because there are many other ways that white students can gain entry to the University—through high marks, through exceptional extracurricular achievements, etc.—those white students who are denied admission are not implicitly being told that they are not "black enough" to matter to the University, the way our scholar was implicitly told that she was not truly indigenous. So in this case, we might want to say that the policy does not really have an impact on the more privileged group's deliberative freedom at all. My account, then, seems to give us ample resources to make different kinds of judgments about different cases of affirmative action, just as we

⁴⁴ For further elaboration of this argument, see Chapter 5, Section 6.

⁴⁵ *Fisher v. University of Texas*, 579 U.S. 2016, discussed by Campbell and Smith, "Deliberative Freedoms and the Asymmetric Features of Anti-discrimination Law," *supra* note 17 at pp. 259, 261–263.

would want an adequate account of discrimination to do.

Chapter Four:

Access to Basic Goods

4.1 A Third Form of Wrongful Discrimination

I have, so far, discussed two ways in which discriminatory practices can wrong people by failing to treat them as equals. I argued in Chapter Two that some discriminatory practices wrongly subordinate some people to others, by marking them out as inferior or by contributing in some other way to their social subordination. And I tried to show in Chapter Three that some discriminatory practices wrongly deny to some people a deliberative freedom that they have a right to have. There are, however, some discriminatory practices that fail to treat people as equals, but not primarily for either of the two reasons we have already examined.

Consider a situation that I mentioned at the start of this book: the lack of safe drinking water on reserves for indigenous populations in countries such as Canada. There are now over 70 indigenous communities in Canada whose reserves have water advisories, ranging from “boil water” advisories to “do not use in any capacity” advisories. Almost half of these advisories have been in existence longer than ten years; and more than half of them are in response to what the UN deems a “moderate” to “high” health risk posed by contaminated water supplies.¹

One reason why we might find this situation troubling is that the various governmental policies that allow this situation to persist seem to violate a basic human right, the right of each person to a sufficient amount of safe drinking water for personal and domestic use.² As I mentioned at the start of the book, when we think of the indigenous water crisis in this way, we are not focusing on it as a problem of discrimination. That is, we are not suggesting that it is wrong to not to provide clean water to these indigenous peoples because this fails to treat them as equals. Rather, our objection is that they have

¹ See, for instance, “Glass Half Empty: Year 1 Progress Toward Resolving Drinking Water Advisories in Nine First Nations in Ontario”, Feb. 2017, David Suzuki Foundation; see also <https://www.ecojustice.ca/world-water-day-the-state-of-drinking-water-in-indigenous-communities>.

² As even Canada acknowledged, in its response to the UN Conference on Sustainable Development in 2012: Canada acknowledged that a human right to “a sufficient quantity and safe quality of reasonably affordable and accessible water for personal and domestic uses” is implicit in Article 11 of the International Covenant on Economic, Social and Cultural Rights: Guillermo E. Rishchynski, *Annex to the Letter Dated 22 June 2012 from the Permanent Representative of Canada to the United Nations Addressed to the Secretary-General of the United Nations Conference on Sustainable Development*, UN Doc A/CONF.216/12 (17 July 2012), http://www.un.org/ga/search/view_doc.asp?symbol=A/CONF.216/12.

not been given something they are owed, owed by virtue of certain fundamental human needs.

However, the current water crisis also seems to be troubling as an instance of wrongful discrimination, as a failure to treat members of these indigenous communities as equals. For most Canadians have constant easy access to clean water; and even other remote communities—those that are not located on reserves, and that are not indigenous—have experienced only a few periods of contamination, which are quickly resolved. And this is partly because, although the different levels of government in Canada cooperate to ensure a high quality of water in off-reserve contexts, the federal government has provided unpredictable and insufficient funding for on-reserve water issues and is only now starting to investigate whether the particular indigenous communities are themselves in a fiscal position to provide the additional funding that is required to ensure reliable access to safe water.³

Before I go on to consider what our equality-based concern is in this case, I want to dispel a certain objection, an objection that is often made to cases of this type. It is that this is not really discrimination by the government at all, but merely a case in which the government, through its inaction, is allowing an unfortunate situation to persist. After all, one might say, what causes the water contamination on reserves is not actually the action of any government, but rather the pollution from nearby industries and the sewage generated by the reserves themselves. Consequently, our objector might conclude, this is not actually a case of discrimination by the government at all—even though there is no denying that indigenous peoples are left disadvantaged.

But there are two responses we can make to this objection. First, government actions *are* a cause of the indigenous water crisis. The government has given unpredictable and low levels of funding to indigenous communities for water sanitation, while providing more funding for, and oversight of, safe water sanitation practices in off-reserve locations. So the government is just as much a cause of the indigenous water crisis as are nearby industries and faulty sanitation systems on the reserves. Second, and more importantly, the objection assumes that whether a government's behaviour counts as a cause of the indigenous water crisis is a factual question. But in fact, it is a normative question: it depends not only on empirical facts about what the government has done or left undone, but on facts about what the government's responsibilities are. This idea is familiar to us from tort law, where it is a basic legal doctrine that a public authority can be held liable for what would otherwise be regarded as an omission if that authority has a duty of care to particular individuals to see that a certain thing is done, and nevertheless fails to have it

³ See the 2016 Human Rights Watch Report, "Make It Safe: Canada's Obligation to End the First Nations Water Crisis," Human Rights Watch (7 June 2016), <https://www.hrw.org/report/2016/06/07/make-it-safe/canadas-obligation-end-first-nations-water-crisis>.

done.⁴ If, for instance, it is a government's responsibility to fund frequent highway inspections and carry them out, then when it fails to do so and a rockslide injures passengers on a highway, the government cannot turn around and say that the injuries were only caused by the rockslide and not caused by its own actions.⁵ So, given that the federal government in Canada has legal responsibility for funding sanitation on reserves, its funding practices—including its failure to provide consistent and adequate funding—can certainly be thought of as a cause of the water crisis, and as the kind of thing that can be evaluated as wrongfully discriminatory.⁶

And it certainly seems to be wrongfully discriminatory. Why? Not just for reasons of subordination. It is, of course, true that the contaminated water on reserves contributes to the social subordination of indigenous peoples. Trying to find alternative sources of clean water in order to avoid disease is time-consuming and energy-sapping, and so indirectly contributes to indigenous people lacking the social power and authority that others have. Moreover, the persistence of the water crisis on reserves also reinforces public stereotypes of indigenous peoples as unclean and as incompetent, unable to maintain the most basic of facilities; and these stereotypes support public habits of censure towards indigenous peoples. But the links in these causal chains are very long, and mediated by many other factors. And this might explain why, when we think about the water crisis as discriminatory, its contribution to patterns of social subordination seems to be only one part of the story, and not the part that is in the forefront of our minds. Similarly, the impact on indigenous people's deliberative freedom—though undoubtedly severe—does not seem to tell the whole story. At least intuitively, there is a further

⁴ See, for instance, *Stovin v. Wise*, [1996] UKHL 15 at p. 2 (per Lord Nicholls); and *Home Office v. Dorset Yacht Co Ltd*, [1970] UKHL 2 at p. 29 (per Lord Diplock).

⁵ *Just v. British Columbia*, [1989] 2 SCR 1228, 1989 CanLII 16 (S.C.C.).

⁶ A different, and more subtle, objection to thinking of the water crisis as a genuine case of discrimination is that the failure to give adequate funding to reserve communities, and the efforts to give these things to non-reserve communities, technically involve different *levels* of government. It is municipalities and provinces that are technically responsible for funding the water treatment off-reserves, while it is the federal government that is responsible for funding water treatment on reserves; so it can look as though there is no single agent who is giving to one group while withholding from another group. However, there is both an easy way out of this objection, and a deeper response. The easy way out is to note that provincial governments and municipalities are only able to do their jobs because of the cooperation of, and extra funding from, the federal government; and it is exactly these things—cooperation and extra funding—that the federal government is not providing to indigenous communities. So there is a single agent here, giving to one group and not giving to another. The deeper response is that, on this conception of wrongful discrimination, it actually *does not matter*, when someone is wrongfully discriminated against by being denied access to a basic good, whether the agent denying them that access is the same one as the agent who has given it to others. One wrongs someone by denying them access to a basic good, when it is in one's power to provide it; and it is irrelevant, for the purposes of answering this question, whether it is the same organization or individual who has given the good to others.

problem here. Indigenous peoples are being wrongfully discriminated against, we want to say, because they are being denied access to *something so basic*: clean water.

Of course, we have to be careful here. If by “something so basic” we mean “something to which they have a basic human right,” then we are right back where we started, with a wrong that does not seem to be a denial of anyone’s status as an equal but instead involves the violation of a prior moral right. So how can we make sense of this wrong as what it seems to be—namely, a wrong that involves, centrally, the failure to treat indigenous peoples as equals, but at the same time, a failure to give them something basic?

I think that a clue to the reasons we are reaching for in this case lies in the Canadian Supreme Court’s insistence, in a number of its early equality rights cases, that sometimes whether a practice is wrongfully discriminatory depends on whether it “restricts access to a fundamental social institution, or affects a basic aspect of full membership in Canadian society.”⁷ This helps to explain what is so troubling about the water crisis on reserves. Without clean, safe drinking water, it is much more difficult to do any of the things that count as participating in Canadian society: working at a job or a vocation and making a meaningful contribution to society; raising children; practicing a religion or a culture. The water crisis does not just deny indigenous peoples something basic to survival, to which they have a human right. In the process, it prevents them from participating fully and as an equal in Canadian society. And it also denies them the ability to be seen as full and equal participants, and to see themselves as such.

I shall call this “denying someone access to a basic good.”⁸ In what follows, I shall sometimes shorten this to “denying someone a basic good”. But what matters in all such cases is that the discriminatee has *access* to the good in question, which I take to mean a real opportunity to obtain that good, one that they can take advantage of with their current resources and current abilities, not an opportunity that is formally open to them but in practice impossible for them to take advantage of. Access matters rather than actual possession of the good for the purposes of wrongful discrimination, because in order to be treated as equals, people also always need the opportunity to determine for themselves whether they want to make use of these goods or not.

Although it happens that the good in this case is necessary for our survival and well-being, this is not a necessary condition for a good’s constituting a “basic good” in the sense

⁷ *Egan v. Canada*, [1995] 2 SCR 513, 1995 CanLII 98 (S.C.C.) at 556 (L’Heureux-Dubé J., dissenting), cited in *Law v. Canada* (Minister of Employment and Immigration), [1999] 1 SCR 497, 1999 CanLII 675 (S.C.C.) at para. 74 (Iacobucci J.).

⁸ For ease of writing, I shall sometimes refer simply to “the denial of a basic good” instead of “the denial of access to a basic good”; but what matters in all such cases is whether the discriminatee has *access* to the good in question, which I take to mean a genuine opportunity to have that good. Access matters rather than actual possession of the good for the purposes of wrongful discrimination, because in order to be treated as equals, people also always need the opportunity to determine for themselves whether they want to make use of these goods or not.

that I am concerned with. Rather, a good is a “basic good” for a particular person in my sense if and only if the following conditions are satisfied:

- (i) Access to this good is necessary in order for this person *to be* a full and equal participant in her society, and
- (ii) Access to this good is necessary, in order for this person *to be seen by others and by herself as* a full and equal participant in her society.

Both of these conditions are satisfied in the case of indigenous communities denied access to safe drinking water. As I argued above, lack of access to clean drinking water prevents indigenous peoples from participating fully in many of the institutions that comprise Canadian society. And, particularly because of the stereotypes surrounding indigenous peoples—that they are unclean, lazy in their habits and primitive in their practices—lack of access to clean drinking water also prevents them from being *seen* as full and equal participants in Canadian society.

I have presented conditions (i) and (ii) as though they were independent. But they are, of course, related to each other. If access to a certain good really is necessary in order for certain people to be full and equal participants in society, then it seems plausible to think that, if that group of people is left without that good for a long time, this may send the social message that they are not worthy of it. And this in turn may contribute to their actually being seen by others, and also being seen in their own eyes, as less than full or equal participants in their society. However, I think it is important to note that simply because the first condition is satisfied in a particular case, it does not follow that the second will also be satisfied. How particular people are seen by others, and how they see themselves, depends on other facts, such as facts about their social position relative to others in their society, and facts about the particular stereotypes associated with them. As I mentioned above, indigenous peoples in Canada have for many years been stereotyped as unclean, lazy in their habits and primitive in their practices. So the absence of clean drinking water on their reserves will certainly, in light of these stereotypes, prevent them from being seen as full and equal participants in Canadian society. Contrast their case, however, with the case of other remote communities that lack clean drinking water, but that are not indigenous and have no history of being thought of as unclean or incompetent. In most cases, these other communities have simply had the misfortune of being located near the sites of chemical spills or polluting mines. Condition (i) is likely satisfied in their case: they have been denied a good that is necessary if they are to participate in society as equals, and their lives will, at least for a time, be much more difficult, and their other opportunities, fewer. But, because they have not historically been stereotyped as unclean or primitive, this lack of water will likely not lead to their being seen by others or by themselves, as less than full or equal participants—at least, not unless their water crisis persists for some years. So in their case, condition (ii) is not satisfied; and this may explain why we are reluctant to say that these other remote communities have been wrongfully discriminated against when they are left for a time without access to clean drinking water.

And other factors, beyond a group's social position and the stereotypes surrounding them, may also be relevant to whether condition (ii) is satisfied. Consider a remote community comprised of a group of scientists, who have chosen to work in the Arctic but discover that the water near their site is contaminated. Particularly given that, unlike indigenous communities, they have a choice as to whether to stay in their location, it seems that even if their water crisis persisted for years and made their scientific work much more laborious and their lives, more difficult, they would likely not be seen as less than full and equal participants in society. However, again, the background social facts matter: if they were a group of scientists investigating climate change and the government's refusal to provide proper water treatment facilities were part of a concerted program to deny credibility to proponents of climate change, then this too might, over time, affect how others saw them or how they saw themselves.

I have suggested that if condition (ii) is not satisfied, we are reluctant to see the case as a case of wrongful discrimination. But what about condition (i)? Is it really necessary? Perhaps the only thing that is relevant to whether these different communities that all lack clean water have really been wrongfully discriminated against is whether their members can be *seen* as full and equal participants in their society. Why insist, in addition, that access to the good in question must be a precondition of their *being* full and equal participants?

I think we need to insist on this first condition because we need to leave room for error. Although I have been emphasizing throughout this book the importance of taking the discriminatee's perspective seriously, it is nevertheless true that people can be mistaken about what is necessary for them, or others, to be a full and equal participant in society. Simply being unable to see yourself as an equal, or having others unable to see you as an equal, does not in and of itself make you unequal. And I think that it is particularly important for us to be able to allow for such mistakes, if we are to offer plausible responses to claimants in some of the cases of apparent discrimination that are not, in fact, wrongful. Consider a certain sub-set of these cases, in which a certain privileged group of individuals is denied access to some special program designed for an underprivileged privileged group, and the more privileged group challenges that program as wrongfully discriminatory, on the grounds that it denies them access to some basic good. Although in some of these cases, the programs are indeed discriminatory, there are others to which the correct response seems to be that the claimants are mistaken. For instance, in the Canadian case of *R. v. Kapp*, a group of non-aboriginal Canadian fishermen alleged that they were being treated as second-class citizens because they were denied a special commercial fishing license issued to aboriginal fishermen.⁹ This special fishing license was given to aboriginal communities by the local government as a way of increasing the self-sufficiency and economic viability of these aboriginal communities, whose members did not have many other opportunities.

⁹ *R. v. Kapp*, 2008 SCC 41, [2008] 2 S.C.R. 483.

The Canadian Supreme Court held that the exclusion of non-aboriginal fishermen from the special license program did not amount to wrongful discrimination, because of the ameliorative purpose of that program. The Court did not contest that non-aboriginal fishermen collectively *felt* inferior, because they did not have the special licenses. Perhaps it was true that, in this particular location, special fishing licenses were so coveted that any fisherman who did not have one would see himself as, and be seen by others as, a less than full or equal member of society. But nevertheless, what the Court seemed to be suggesting in its judgment was that these fishermen were not *in fact* missing out on a good that was necessary in order for them to be full and equal participants in society.

You may disagree with my analysis of this particular case. But we surely want to allow that there *could* be such a case. And if so, we need some way of recognizing that people can be mistaken about the opportunities and resources that are necessary in order for themselves, and others, to be full and equal participants in society. And consequently, when defining a “basic good” for the purposes of discrimination, we need to invoke condition (i). We need to maintain that there is an independent truth of the matter as to whether access to a particular resource really is necessary for someone’s being a full and equal participant in a society, and that people’s own assessments can fail to accord with this truth. At the same time, we can allow that one important part of being an equal to others is being recognized by them as an equal, and being able to see yourself as an equal. This is why condition (ii) is necessary as well.

I have now introduced the idea that in some cases of discrimination, what is wrongful is that some people have been left without access to basic goods, goods that they need to have access to if they are to be, and to be seen as, full and equal members of their society. I think that this idea lies at the heart of a number of prominent cases of wrongful discrimination—not just the case of the indigenous water crisis.

For instance, the push to recognize same-sex marriage was, in large part, motivated by the belief that same-sex couples lacked access to a fundamental institution in society, the institution of marriage, and for this reason could not truly participate in their societies as equals. In countries such as Canada and the U.K. there were, at the time the initial court challenges were brought, alternative ways in which same-sex couples could attain the same fiscal and material benefits as married couples. So the couples who brought these challenges were not seeking these particular material benefits. Rather, they saw marriage very much as a “basic good” in my sense—that is, as the kind of institution that they needed at least to have the opportunity to belong to, because it was only if they were officially granted that opportunity that their relationships would be deemed equal in commitment and maturity to the relationships of married couples. For instance, one of the applicants in the case of *Halpern v. Canada*, Julie Erbland, testified that: “I want the family that Dawn and I have created to be understood by all of the people in our lives and by society. If we had the freedom to marry, society would grow to understand our

commitment and love for each other.”¹⁰ Another, Carolyn Rowe, said: “We would like the public recognition of our union as a ‘valid’ relationship and would like to be known officially as more than just roommates.”¹¹ These applicants felt that, until they were officially recognized as eligible to marry, they would not be recognized in public as capable of making the kind of long-term commitment to another person that each member of a married couple makes to the other. And without such public recognition, they could not be or be seen as full and equal participants in their societies.

Certain cases involving discrimination against people with disabilities also seem best conceptualized as a denial of a basic good to certain people. Consider, for instance, the case of *Eldridge v. British Columbia*, which involved a challenge by hearing-impaired individuals to legislation that failed to insure sign-language interpreters in hospitals.¹² The claimants argued that because they were denied sign-language interpreters, they were unable properly to communicate with their doctors. Hand-written notes, they argued, were insufficient: not only are they impractical during emergencies, but, more importantly, many hearing-impaired individuals are unable to read or write at a sophisticated level, so cannot communicate effectively through writing. Two of the claimants, John and Linda Warren, had no sign-language interpreter during the premature birth of their twin daughters. The staff was reduced to using random hand gestures to inform them of difficulties during the birth; and although the staff provided a hastily handwritten note that said “fine” as they whisked the babies away to the NICU, the couple was left with no understanding of what their daughters faced and no opportunity to be a part of the decision-making process.

The Canadian Supreme Court accepted that this amounted to wrongful discrimination; and much of their judgment is consistent with seeing the central problem in this case as a denial of access to a basic good.¹³ For instance, the Court noted that, without sign-language interpretation, hearing-impaired people are unable to communicate with their doctors, and so are effectively left out of the normal conversation between doctor and patient. And the Court further emphasized that, if we are to understand the full impact of this situation on hearing-impaired people, we need to think of the background social context: the history of marginalization of people with such disabilities, the fact that they have been systematically “excluded from the labour force,” “denied access to opportunities for social interaction and advancement,” and silenced “in a world that assumes that most people can hear.” We might add that, as Denise Réaume has argued, without sign-language

¹⁰ See *Halpern v. Canada (Attorney General)*, [2003] 65 O.R. (3d) 161, O.J. No. 2268 (Ont. C.A.) at para 9.

¹¹ *Ibid.*

¹² *Eldridge v. British Columbia (Attorney General)*, [1997] 3 S.C.R. 624, 1997 CanLII 327 (S.C.C.).

¹³ Though not all of them. The Court suggests at times that the wrong in question is simply failing to give hearing-impaired people “effective medical care” when all other Canadians have it. But this cannot be right, since not all other Canadians have effective medical care. Nor is it really consistent with the other claims made by the Court in this judgment about the importance of attending to the marginalization of hearing-impaired people in determining whether they have faced wrongful discrimination. For if the problem here were simply that hearing-impaired people lacked the same quality of medical care that others have, then their social position and marginalization would be irrelevant.

interpretation, these claimants were denied a meaningful opportunity to consent to their own and their children's treatment, and so were effectively treated like children themselves.¹⁴ Access to sign-language interpretation in hospitals is, for all of these reasons, a basic good, and denying it to these people prevented them from being, and being seen as, full and equal participants in Canadian society.

I have tried to show that in a number of cases of wrongful discrimination, the wrong seems to stem from the denial of what I have called "access to a basic good." But there is still much that remains murky in the idea of a "basic good", and much that is potentially problematic about the claim that this is why certain discriminatory practices are wrong. In the rest of this chapter, I want to clarify the idea of a basic good, and to defend the claim that this is a distinctive and important reason why certain discriminatory practices are wrong.

4.2 *Basic Goods: Further Clarification*

4.2.a *To identify a good as "basic" is not to claim it is objectively good*

The basic goods that I have discussed at greatest length so far—clean water and sign-language interpretation in hospitals—are things that many would identify as objectively good. But in order to count as a "basic good," it is not necessary that a particular resource or opportunity should be actually or objectively good. All that must be true is that, given the practices and beliefs of people in a particular society, access to that resource or opportunity is necessary for this person, if she is to participate fully and as an equal in her society, and to be seen as an equal. Marriage is a good example. Many people believe that it is on balance good, allowing for public recognition of a long-term commitment to another adult. However, a significant number of people see marriage as oppressive, a social institution that has historically relegated women to the position of men's property and that still works to undermine women's autonomy. Even if they are right, marriage can still count as a "basic good" for the claimants in same-sex marriage cases. All that must be true is that, given the society in which these claimants live and people's shared assumptions in that society, these claimants will not be, and will not be regarded as, full and equal participants in their society until they too are given the opportunity to marry the people of their choosing, regardless of their sex.

That an opportunity or resource can count as a "basic good" for someone even if it is not objectively good is not a problem for my view. Rather, it reflects the fact that claims of wrongful discrimination of this kind are different from claims to a certain resource or institution that are grounded in its objective value. And indeed, as we saw in Chapter Three, most countries' anti-discrimination laws protect our right even to some things that are not good for us. As I mentioned in that chapter, most countries' domestic anti-

¹⁴ Denise Réaume, "Discrimination and Dignity," *Louisiana Law Review* 63(3) (2002), pp. 1–51 at p. 44.

discrimination laws protect people from discrimination in a very broad array of contexts, without any qualification concerning the goodness of the opportunity or the situation. They protect us from discrimination in the provision of any kind of good or service—from candy stores to casinos—and in the provision of any kind of accommodation, whether it is beneficial for us or not, and in negotiations over membership in any kind of trade union, whether this will help us or not. This is because what anti-discrimination laws are protecting is not only access to objectively valuable resources, but access to the resources and opportunities that we need if we are to be treated as equals in our society.

4.2.b Some basic goods are privately appropriable; others are public

What, then, are some other examples of basic goods? Some are privately appropriable goods. Among these, some are preconditions for the claimants' survival, such as clean drinking water, sufficient food, enough basic clothing that they can be warm, and a shelter that will keep them dry. Other privately appropriable goods are not preconditions for survival, but are nevertheless preconditions for the claimants' functioning as equals in society. You may die without a name. But unless you have one, you will not be able to exercise any of the other rights that your society accords to its members; and this is partly why the right to a name is recognized in the United Nations *Convention on the Rights of the Child*.¹⁵

Another privately appropriable good that is not required for survival but is arguably a "basic good" in my sense is a *home*. I am, here, using the term "home" to mean something different than the term "shelter." We need shelter as a matter of survival, as I noted above. But a "home" in the sense I have in mind is a place in which you have some say over who enters and exits, and in which you cannot yourself be asked to leave. Chris Essert has argued intriguingly that, since everything we do must be done somewhere, those who do not have a home in this special sense—a place where they can do what they wish to without being told to leave or to curtail their activities by others—are in a significant sense unfree.¹⁶ I do not need quite as strong a claim for my purposes here. All that I need to note is that, given the significant number of people who do have a "home" in this sense, and the number of social activities that depend on one's having a home, those who lack a home cannot be, or be seen as, full and equal participants in our societies.

I have been talking so far about privately appropriable basic goods. But most of the basic goods that seem to be the focus in prominent cases of discrimination concern shared public institutions. Sometimes, the basic good at issue seems best described as access to a status or a resource made possible by certain public institutions: for instance, access to the

¹⁵ U.N. General Assembly, *Convention on the Rights of the Child* (20 November 1989), *Treaty Series* 1577 at p. 3. See Article 7: "The child shall be registered immediately after birth and shall have the right from birth to a name . . ."

¹⁶ See Chris Essert, "Property and Homelessness," *Philosophy & Public Affairs* 44(4) (2016), pp. 266–95.

status of marriage, access to effective health care through sign-language interpreters, access to the funds available from pension plans, access to public transport. Sometimes, the basic good is better described as the right to perform certain social or political acts without having to change some aspect of your appearance, such as the right to vote without having to remove your headscarf, or the right to be a waiter, and wear the uniform of a waiter, without having to be clean-shaven. Sometimes, the basic good claimed is a right to be in certain public places while doing certain things, such as the right to breastfeed in public shopping malls. And some claimants have argued—and some courts have accepted—that it is a basic good in my sense for certain social groups to have access to institutions that are specially necessary for them given their histories and needs, such as an Indigenous Child Protection Service that is better funded than any Child Protection Service available to non-indigenous groups, and offers different programs, specially tailored to indigenous groups.¹⁷

4.2.c “Basic” in relation to particular people in a particular society

Although I have spoken of “basic goods” as though it is the goods that are basic, I do not mean to imply that we can decide whether a certain good is basic by looking at the good in isolation from particular people within a particular society. On the contrary, I have spoken throughout of whether a particular good is a “basic good” *for* a particular person or group *in* a particular situation. There are two important points to note here.

First, whether a certain good counts as a “basic good” for the purposes of wrongful discrimination depends not just on facts about that good, but also on facts about the particular people that claim to have been denied this good. Something can be a basic good for some people but not for others. Sometimes, this is for the simple reason that some people do not need, or could never use, a particular opportunity, and so having it is not necessary for them to be equals. For instance, non-hearing impaired individuals do not need sign-language interpretation; and men could not make use of the opportunity to breastfeed in public. But in other cases, the reason a certain good is not a basic good for a particular group of people is that, even though they could use it, its availability does not affect whether they can be *seen* as equals in their society—as we saw, for instance, in the case of non-indigenous remote communities that suffer from temporary water crises.

Second, whether a good counts as a basic good for a particular group depends on the particular society in which they live. Access to the institution of marriage may not be a basic good for any social group several centuries from now, if fewer and fewer couples seek to marry and the institution declines drastically in its social importance. But it likely is in our own society, here and now. How we define what counts as the relevant “society,” in

¹⁷ *First Nations Child and Family Caring Society of Canada (FNCFCFS) et al. v. Attorney General of Canada (for the Minister of Indian and Northern Affairs Canada)*, 2016 CHRT 2 [FNCFCFS v AG of Canada].

determining whether a particular good is basic for a particular person in a “particular society” is an important question here. Most of us live concurrently in a number of different social groups. We are a part of a particular country and its practices, which could be called a “society”; but we are also a part of a particular city and a neighbourhood within that city that has a certain character, which are also societies; and we may also be a member of a certain religion or of a certain racial group with particular traditions and beliefs. We also have online presences, where we feel the pull of different online cultures. Which of these is the relevant “society” in relation to which we ought to evaluate whether the claimants in a particular case have been denied a basic good? When dealing with a case of discrimination against indigenous peoples, for instance, should we look solely within the particular indigenous group at issue, or should we look at the country as a whole and how other members of that country perceive the group in question? This is not a question that I think can be answered in the abstract. How wide a net we cast when we define the relevant “society” in a given case, and which social circles we include within it, will depend on the claimants and on the good in question. We can draw a helpful parallel here to nuisance law. In Anglo-American nuisance law, whether something amounts to an “unreasonable interference” with someone else’s use of their land depends on what is called “the standard of the locality”—that is, the practices and expectations of people in the local area.¹⁸ But there is no fixed rule for determining what counts as the local area, or how large a circle we must draw when delimiting this local area. Rather, nuisance law recognizes that the relevant area will sometimes be as large as a town, and sometimes as small as just one street or two, depending on the kind of complaint that is at issue. I am making the same suggestion here.¹⁹

4.2.d Importance of the discriminatee’s perspective

I have just explained that any basic good needs to be identified as such in relation to some specific group of people, within a specific society. While in the case of many basic goods, anyone can understand a particular person’s need for them without looking too deeply into the beliefs and circumstances of that person, other basic goods can only be understood as basic from the perspective of that particular person or group. When I laid

¹⁸ See, for instance, *Colls v. Home and Colonial Stores, Limited*, [1904] AC 179; *Rushmer v. Polsue and Alfieri, Limited*, [1906] 1 Ch. 234, [1907] AC 121. See also *Coventry & Ors v. Lawrence & Anor*, [2014] UKSC 13; *Sturges v. Bridgman*, (1879) 11 Ch. D; and *Campbell v. Seaman*, 63 N.Y. 568, 577 (1876).

¹⁹ I shall go on to argue in the next section that we need to pay particular attention to the discriminatee’s perspective when we assess whether a certain good is a basic good for that person. I think we also need to attend to the discriminatee’s perspective when we decide what the relevant “society” is, for the purposes of assessing whether a particular good is indeed basic for her (that is, necessary if she is to be and be seen as an equal in her society). This is because we cannot determine what the relevant society is unless we have a sense of the role that this particular good plays in her life, and to understand this role, we will most often need to consider her beliefs and values.

out the example of the water crisis earlier in this chapter, I presented the problem as though it involved only a lack of clean drinking water. I did this because I wanted to introduce the idea of a “basic good” in a way that was easy to understand, and most of us can readily appreciate many of the reasons for which a lack of clean drinking water might amount to a lack of access to something basic to full and equal participation in our societies. But actually the indigenous water crisis is more complicated than I first suggested; and to present it only as a problem of contaminated drinking water is to under-describe the good in question, in relation to indigenous peoples. Water is the medium in which many of their cultural activities, such as fishing, are practiced. More importantly, to most indigenous peoples, water is sacred. It has a spiritual force, connecting them to the earth and to their ancestors, and it plays a crucial role in many of their cultural practices. So when they lack access to clean water, they do not just lack access to a consumable commodity and a precondition for health. They lose the ability to live in their traditional ways. And the lack of clean water has a particularly strong impact on many indigenous women. In many indigenous cultures within Canada, women are believed to have a sacred connection to the earth and its water. The earth is perceived as female and water is the earth’s blood. Women give birth to children just as the earth gives birth to vegetation; and because of this connection, women are the ones who, in many indigenous communities, are responsible for keeping the earth’s blood pure. They are called “Keepers of the Water” or “Carriers of the Water.” When others pollute their water and offer them no infrastructure to clean it, these women are unable to fulfil their cultural responsibilities, unable to be the people whom their culture says they must be.²⁰

I hope that even this brief description makes it clear just how rich and complex the basic good at issue in the indigenous water crisis is, and how little of that good will actually be visible to us if we look at it without a full appreciation of its place in indigenous culture. Of course, not all basic goods are like this. But many can be fully comprehended only from the perspective of the person or group who has been denied the good. Ask any woman who claims the right to breastfeed in public, and she will tell you that the good at issue here is not simply a matter of convenience or enjoyment, not simply the opportunity to enjoy the benefits of a particular public place and to avoid the inconvenience of going somewhere else. What is at stake for these women is also the opportunity to have their bodies publicly acknowledged as *theirs* to use, theirs to use to nurture their child with when they see fit, rather than treated as a body that is defined by *others’* feelings of embarrassment, or *others’* assumptions about what a breast is and where it belongs. I think that many basic goods—more than we might at first think—are like this. That is, in order to understand their significance for the discriminatee, we need to look at them from the perspective of

²⁰ Kate Cave and Shianne McKay, “Water Song: Indigenous Women and Water”, *The Solutions Journal* 7(6) (2016), pp. 64–73. See also Kim Anderson, “Aboriginal Women, Water and Health: Reflections from Eleven First Nations, Inuit, and Métis Grandmothers,” Paper commissioned by the Atlantic Centre of Excellence for Women’s Health and the Prairie Women’s Health Centre of Excellence (October 2010).

that person or group.²¹ We need to try to understand, in light of their situation, their needs, and their beliefs, what the real impact of being without a certain good is for them. This is not, of course, to say that a person or group can never be mistaken about *whether* some good is in fact a basic good for them. As I acknowledged above when discussing the non-aboriginal fishermen in *Kapp*, claimants can certainly be mistaken about this. But when we try to define *what* the good in question is, we need to do so from the discriminatee's perspective, taking into consideration her needs and the practices and history of the relevant social group or groups. Only then will we see, for instance, that the good is not just "clean drinking water and sanitation" but also "water needed for ritualistic purposes, so that indigenous women can continue to fulfil their cultural roles as purifiers of the water."

4.2.e *Something can be a basic good for some people even if no others need it*

I have now argued that a basic good is "basic" only in relation to certain people in a certain society and that in some, and perhaps many, cases we will only be able to understand what the basic good is if we consider the practices, beliefs, and history of the claimants. But can something be a basic good for certain people even if no other group of people that society needs it, or needs it to such a great extent? I think that it can, and to show this, I want to consider another recent Canadian case involving indigenous communities, *First Nations Child and Family Caring Society of Canada et al. v. Attorney General of Canada (for the Minister of Indian and Northern Affairs Canada)*.²² This case concerned whether the federal government of Canada discriminates against members of indigenous communities living on reserves on the grounds of race, by failing to provide a high enough level of funding for family and child protection services on reserves for these families to have a greater chance of remaining together, with their children staying on the reserves instead of being removed to foster care in locations remote from their own communities.²³ The Canadian Human Rights Tribunal accepted that this constituted unjustifiable racial discrimination. Some aspects of the basic good at issue in this case are not specific to the indigenous communities in question: for instance, the claimants argued, and the Tribunal accepted, that indigenous children deserve at least the same level of funding as is given to those Child and Family Services programs that target non-indigenous children. But the claimants went on to assert that, given the history of abuse of indigenous peoples in Canada—in particular, the legacy of residential schools, through which families

²¹ This is what some scholars working on women's reproductive rights in the context of international human rights law have done: see Joanna N. Erdman and Rebecca J. Cook, "Women's Rights to Reproductive and Sexual Health in a Global Context," *Journal of Obstetrics and Gynaecology Canada* 28(11) (2006), pp. 991–997; and Rebecca J. Cook, Bernard M. Dickens, and Mahmoud F. Fathalla, *Reproductive Health and Human Rights: Integrating Medicine, Ethics, and Law* (Oxford: Clarendon Press, 2003).

²² *FNCFC v AG of Canada*, *supra* note 17.

²³ There are currently an estimated 27,000 First Nations children in welfare care, and this accounts for 30 to 40% of all children in child welfare care, even though they represent less than 5% of the child population in Canada: Pamela Gough, Nico Trocmé, et al., "Pathways to the Overrepresentation of Aboriginal Children in Care," Centre of Excellence for Child Welfare Information (2005), pp. 1–3 at p. 1.

were torn apart, and children, sexually and emotionally abused—indigenous families often require more assistance, and assistance of a special kind, in order to ensure that things do not reach a point where children need to be removed from homes. The Tribunal held that the combination of parents who were themselves victims of abuse in residential schools; inadequate housing on reserves; widespread poverty on reserves; and substance abuse together form a special set of circumstances that uniquely characterize many indigenous communities. Using the language of basic goods that I have developed in this chapter, we might say that a special set of child and family service programs, of a kind that is not required elsewhere, and that necessitates funding to a level that is not given elsewhere, is a “basic good” for these indigenous communities. So something can be a basic good for one group even if no other group needs it, or needs it to the same extent.

4.2.f What counts as “denying” someone a basic good?

Lastly, I want to raise a question that I have so far left open, and that can be answered in a number of ways. This is: what exactly counts as “denying” someone a basic good? When I introduced the idea of a basic good earlier in this chapter, I used the example of the indigenous water crisis, and I said that we can certainly treat the Canadian government as having denied indigenous peoples this good, in part because they stand under a duty to provide adequate funding for sanitation on reserves. This is probably the clearest type of case in which an action counts as a denial of a basic good—namely, where the agent already has a duty or a responsibility to provide the good in question. Some may argue that this is the only type of case in which a mere failure or omission to provide a good can count as a denial of it, and hence as an instance of wrongful discrimination. But on a more expansive version of my view, one denies others a basic good whenever it is *in one’s power* to give them access to that good, and one does not do so. I am more sympathetic to this broader view, for reasons that I shall set out in Chapter Seven. I shall argue there that we can only create a society of equals if each of us takes ourselves to stand under a duty to treat others as equals, and that this includes doing what we can to give others access to basic goods. I shall explain in Chapter Seven why I do not feel this is overly demanding. But those who disagree could adopt the more limited version of the view: that we deny others a basic good only if we have a responsibility or duty to provide it, and we do not.

4.3 Why this is a problem of inequality, and a distinctive form of wrongful discrimination

I have now clarified a number of features of basic goods. In this next section of the chapter, I want, first, to explain why the denial to someone of a basic good is genuinely a problem of inequality. I shall then defend the claim that this is a distinctive reason why discriminatory practices can be wrong, a reason that is different from the reasons of social subordination that we examined in Chapter Two and from the infringements of deliberative freedom we considered in Chapter Three.

The first of these tasks—that is, explaining why the denial to someone of a basic good is a genuine problem of equality—is relatively easy. Since a basic good is “basic” for a particular person if she needs it in order to be, and to be seen as, a full and equal participant in her society, it follows that if this person is left without this particular good, then she is not a full and equal participant in her society. Basic goods are basic not by virtue of their objective value or their connection to our survival, but by virtue of their impact on a particular person’s ability to participate as an equal in their society. So when someone is left without one, they are unable to be, or unable to see themselves as, an equal. In this particular sense, then, they are not treated as an equal.

But this explanation, though helpful in laying out why the denial to someone of a basic good is genuinely a problem of inequality, might cause one to wonder whether this reason for certain discriminatory practices being wrongful is really so distinctive, so different from the reasons of subordination we examined in Chapter Two. Are these really two different reasons why discriminatory acts can be wrong? Or are they, at bottom, the same reason?

There are, it seems to me, several respects in which the denial of a basic good is different from the “social subordination” that was discussed in Chapter Two. First, as we saw in Chapter Two, social subordination is concerned with the unequal social status of a *group* of people, all of whom share a trait that I described as “socially salient,” in the sense that others in society take that trait to have implications for the character and behaviour of members of the group. By contrast, our main focus, in assessing whether someone is denied a basic good, is on the status of particular *individual* claimants: are they able to participate in their society as equals and to be seen as equals? So the focus of the two inquiries, and the locus of the wrong in each case, is different. In the one case, our focus is on the group and the group’s standing, relative to some other group; in the other case, the focus is on the individual and whether that individual is missing one of the necessary conditions of their participating fully and equally in society. A second difference concerns the ways in which the two sorts of judgments—about social subordination and about individuals being denied a basic good—are comparative. The judgment that some practice contributes to social subordination is what we might call directly comparative: it always depends on comparisons about the relative amounts of power, authority, deference, and structural accommodations enjoyed by different social groups. By contrast, the judgment that some individuals are denied a basic good seems in only an indirect way to depend on comparisons. It is primarily a judgment about what that individual lacks. And although, in order to assess whether an individual has been denied a basic good, we often look to what other social groups have, we do so only in order to understand the opportunities that *this* individual is now lacking. And there are cases in which the best way to understand this is to focus, not on a comparison with other groups’ resources or opportunities, but on the claimant’s own situation and the history of her own social group. The case of inadequate child and family service support to indigenous families is a case of this type. As we saw, the judgment that they lack this basic good was based primarily on information about their

own special history and situation, which has left them with unique needs. Thirdly, it is possible for the members of a social group that is, in some or many contexts, not socially subordinate to others, nevertheless to lack a certain basic good. Otherwise put, you can lack one of the necessary conditions for participating fully and equally in society even if, overall, your social group is much better off than others, and has a much higher social standing than certain other social groups.

To see this, it may help to consider the situation of those heterosexual couples in the U.K. who claim that they are wrongfully discriminated against if they are not, like same-sex couples, allowed the option of entering into a civil partnership.²⁴ Civil partnerships were first recognized in the U.K. in 2004, as a way of granting the same rights and privileges to same-sex couples that were available to heterosexual couples through the institution of marriage. But although the U.K. permitted same-sex couples to marry in 2013, it did not at that time abolish the institution of civil partnerships. Instead, the government chose to wait, apparently to investigate whether the best course of action was to abolish civil partnerships or not. This interim period therefore gave same-sex couples a choice that was not open to heterosexual couples: they could choose whether to enter a civil partnership or a marriage, whereas heterosexual couples had to choose either marriage or no marriage. Some heterosexual couples brought lawsuits, alleging that this was wrongfully discriminatory. They claimed that they too ought to have the opportunity to be civil partners, primarily because they viewed marriage as an oppressive institution and felt that they would rather not be a part of an institution that has, historically, enabled men to have a degree of power over women.

This is a good case for us to use in testing the differences between claims based on social subordination and claims based on the denial of a basic good. For it seems implausible to suggest that the exclusion of heterosexual couples from civil partnerships contributes to their social subordination, even though it may deny them a basic good. Why should we think that it does not contribute to their subordination? For one thing, heterosexual couples are not normally thought of as standing in a subordinate position to any other kind of couple: it is same-sex couples who occupy a subordinate position relative to heterosexual couples. One might object that the relevant group here, the group that may be socially subordinated, is not heterosexual couples but rather “women who have a male partner.” However, even if we accept that this is the relevant group and that it is a group that is, in certain respects, subordinated, it is not clear that the mere absence of a choice to enter into a civil partnership contributes to the social subordination of this group. It seems likely that, within this group, it is only those women whose partners support their full autonomy and wish to distance themselves from the kind of power had by traditional

²⁴ See, for instance the facts and background given in *R. (on the application of Steinfeld and Keidan) v. Secretary of State for International Development (in substitution for the Home Secretary and Education Secretary)*, [2018] UKSC 32.

husbands who would agree to civil partnerships if such a choice were available. In other words, it is only those women who are already not subordinate to their partners who would be able to take advantage of the choice to enter a civil partnership. For these reasons, I do not think we can plausibly claim that the absence of this choice contributes to women's social subordination. But does it nevertheless deny heterosexual couples access to a basic good? Or is the position of heterosexual couples in this case akin to the position of the non-aboriginal fishermen in the case of *R. v. Kapp*?²⁵ I argued earlier that the Canadian Supreme Court suggested that, although these fishermen felt they were not treated as equals and lacked a basic good, they actually did not. Is this what we ought to say about the heterosexual couples who claim that they, too, ought to be able to enter civil partnerships?

I am not sure. In *Kapp*, the aboriginal license program was necessarily limited, and its ameliorative function would be entirely undermined if everyone had a special license. There would be no advantage to aboriginal fishermen, and hence no increase in their communities' welfare, if every non-aboriginal fisherman also had such a license. And it is partly because of this that it seemed less plausible for the non-aboriginal fishermen to claim they had been denied a basic good, in being refused a special license. If something can necessarily only be had by a special few, as part of an ameliorative program, it seems implausible to claim that everyone else must have it too, as a precondition for equal standing. By contrast, although it is true that heterosexual couples are in many respects more privileged than same-sex couples, it is not true that the entire purpose of granting civil partnerships to same-sex couples would be undermined if the institution were opened to all. A civil partnership is arguably the kind of institution that can be open to all, without in any way sacrificing the benefits that accrue from it to same-sex couples, and without changing its social meaning as a way of recognizing a life-long but not patriarchal commitment to another person.

But is it really necessary for heterosexual couples to have the opportunity to choose to become civil partners, if they are to have an equal social standing? Is this choice, in other words, a basic good for them? On the one hand, now that there is a social institution available for having one's long-term commitment to another person publicly recognized in a way that is disassociated from marriage's patriarchal history, it does seem that there is a meaningful opportunity that heterosexual couples lack. But is it a basic good—a necessary condition for their being, and being seen as, equals in their society? I am not sure. The institution of civil partnerships is so very young, and the number of actual civil partners, so relatively few, that the institution itself does not have the kind of widely understood social meaning or symbolic force that the institution of marriage does. So whereas it did seem plausible for same-sex couples to claim that they were denied a basic good by being excluded from the institution of marriage, it seems much less obvious that heterosexual couples are denied access to a basic good when they are denied access to civil partnerships.

²⁵ *R. v. Kapp*, *supra* note 9.

But if, over the next ten years, more same-sex couples opted for civil partnerships, and the institution came in the public eye and the eye of the media to symbolize the ideal domestic partnership between equals, then perhaps we would be more likely to think that the choice to enter this institution *is* one that heterosexual couples too must have, if they are to be full and equal participants in society.

I do not need, for the purposes of my argument, to settle this question. What is important for my purposes is just to note that, even though the exclusion of heterosexual couples from the institution of civil partnerships does not seem to contribute to their social subordination—nor to the social subordination of the female members of heterosexual couples—it is conceivable that it could nevertheless constitute a denial of a basic good.

Another case which sheds some light on the difference between wrongs grounded in social subordination and wrongs involving a denial of basic goods is the case of *Manuel Wackenheim*, brought before the U.N. Human Rights Committee.²⁶ Wackenheim, who lives with the condition known as “dwarfism,” challenged bans on the sport of dwarf-tossing imposed by several municipalities in France. He argued that these bans violated his right to non-discrimination under Article 26 of the International Covenant on Civil and Political Rights. Dwarf-tossing is a form of entertainment offered at some bars and public events in certain European towns. People with dwarfism don protective clothing and are thrown by the competitors onto air mattresses, with the winning competitor being the one who can throw the dwarf the farthest. Understandably, the towns who banned dwarf-tossing did so because they felt that it was degrading for people with dwarfism to be treated as projectiles: this practice, in their view, was “an affront to human dignity.” However, Wackenheim argued that, as a person living with dwarfism, he had so few employment opportunities that dwarf-tossing was his one hope of having a steady job and a steady income, and that “dignity consists in having a job.” In other words, put into my language of basic goods, Wackenheim’s argument was that even if it is true that the practice of dwarf-tossing encourages people to ridicule people with his condition and to treat them as objects, and even if it thereby contributes to the social subordination of people with dwarfism, it is nevertheless *also* true that in French society at the moment, dwarf-tossing is one of the only jobs available to people with dwarfism. So the opportunity to be employed in the sport of dwarf-tossing is, right now, a basic good for him. Without this opportunity, he cannot participate fully in French society; and so he cannot be, or be seen as, an equal.

The Wackenheim case is helpful for us to consider for several reasons. First, it gives us a very clear example of the difference between claims of wrongful discrimination based on the social subordination of a particular social group (in this case, people with dwarfism), and claims of wrongful discrimination based on the denial to an individual of a particular basic good (in this case, the denial to Wackenheim of employment through dwarf-tossing). The towns’ argument that dwarf-tossing should be banned can be seen as based

²⁶*Manuel Wackenheim v. France*, Communication No 854/1999, U.N. Doc. CCPR/C/75/D/854/1999 (2002).

on a claim about the practice's contribution to the social subordination of all those living with dwarfism. By contrast, Wackenheim's challenge of the ban seems to be appealing to something different, even though it is still a claim based upon inequality. His claim, I am suggesting, is helpfully understood as based on an appeal not to the social subordination of people with dwarfism, but to the basic good of employment. Given the structure of French society at the moment and the limited opportunities available for employment for people living with dwarfism, Wackenheim can only participate fully in French society if he is given the opportunity to seek employment in the sport of dwarf-tossing.

Second, the Wackenheim case also helps us to see that the framework I am proposing for thinking about discrimination—as wrongful for a number of very different reasons—can provide more clarity in helping us think through the different positions in different cases than does an appeal to a single value such as dignity. The Wackenheim case was argued before the Human Rights Committee not as a matter of basic goods or social subordination, but as a question of what violated dignity. And in this case, *both* sides claimed an affront to dignity. The towns viewed dwarf-tossing as an affront to the dignity of all people living with dwarfism, and claimed that their bans restored dignity to these people. But Wackenheim viewed the bans, and the resulting lack of employment, as an affront to his own dignity, and claimed that removing it was necessary to restore his dignity. The Human Rights Committee sided with the towns, finding that the bans were reasonably justified and concluding that they could therefore not be an affront to dignity. One problem with seeing the disagreement in this particular way—as a disagreement over what infringes dignity—is that, if the towns win, as they did, then Wackenheim is left with no residual moral objection to the bans, at least on grounds of discrimination. That is, either the bans are, or they are not, discriminatory as an infringement of his dignity; and if the towns are correct that the bans are not an infringement of dignity, then it seems to follow that Wackenheim has no objection to them on the grounds of discrimination. But we may want to allow instead that even if the towns are correct and the bans are justified, there is a meaningful sense in which Wackenheim still has a residual moral objection to them, an objection that is grounded in considerations of discrimination. And we can say this if we see the case not as a disagreement over what dignity requires, but as a disagreement over how to prioritize the towns' need to eliminate social subordination, on the one hand, and Wackenheim's own need for a job as a precondition of his being, and being seen as, an equal participant in society. Whereas in most of the cases we have considered so far, it is the same practice that contributes to social subordination and denies someone a basic good, so these two different reasons for thinking a practice wrongfully discriminatory point us in the same direction, the tragedy of the Wackenheim case is that the very practice that seems necessary for eliminating the social subordination of a certain group (the ban on dwarf-tossing) denies some members of that group a basic good, the good of employment. If we see the case not as a case about what dignity means, but as a case where, unusually, two different reasons for something's constituting wrongful discrimination pull us in two different directions, this opens the possibility of recognizing

that even if we ultimately conclude that the towns are all things considered justified in imposing these bans, we can still maintain that there is a very real sense in which Wackenheim has not been treated as an equal. He has been denied a basic good.

The Wackenheim case also leads us to another interesting and important set of questions, which I shall explore further in the next Chapter of this book. This is: how ought we to go about reasoning through those difficult cases in which a practice seems required if we are to eliminate one form of wrongful discrimination, but also seems wrongful, in light of another of the reasons why discrimination can be wrong? How ought we to reason through such cases? Do considerations of social subordination, for instance, trump claims that a basic good has been denied to someone? And if they do, what are we to say about what I have been calling the “residual moral objection” of those who are denied basic goods? Does it make sense for us to say that, although they have been treated in a way that is all things considered justifiable, they have nevertheless been wronged? I shall discuss this in detail in Chapter Five, when I consider the ways in which these different reasons for wrongful discrimination relate to each other.

I have now explained why we need to think about denials of basic goods as different from claims about wrongful social subordination. But what about infringements of a right to deliberative freedom? Are these really distinct from the wrongs that I have been calling a denial of basic goods? Why shouldn't we think of deliberative freedom as one *type* of basic good—so that an infringement of deliberative freedom is really just a denial of a basic good? If this is right, then these are not really two different kinds of wrongs; rather, infringements of deliberative freedom are a sub-class within the broader class of denials of basic goods.

I am reluctant to treat deliberative freedom as just another basic good, however, for two related reasons. First, the basic goods we have been discussing in this chapter involve resources and opportunities, such as access to the institution of marriage, access to sign-language interpreters, and access to a robust child and family service program. But deliberative freedoms are not such resources or opportunities. They are, as I argued in the previous chapters, best thought of as freedoms, including freedom from the fixed and opportunity costs of having a certain trait, and freedom from having that trait always before your eyes, whether you wish to or not. But this is only a partial reply. For it seems simply to invite a follow-up question: why not expand our list of basic goods to include not only the resources and opportunities discussed in this chapter, but also the freedoms discussed in the last chapter? The reason for not doing this is my second reason for thinking that the wrong of infringing someone's right to deliberative freedom is different from the wrong of leaving them without access to a basic good. This is that the structure of the two wrongs is different. The wrong of infringing someone's right to deliberative freedom is a wrong that depends upon the value of autonomy. And it is only an instance of failing to treat someone as an equal because we live in societies that so value autonomy that failing to treat someone as a person capable of autonomy amounts to failing to treat

them as an equal. By contrast, the wrong of leaving someone without a basic good does not depend on the value of autonomy or its role in our society, and it directly engages with the value of equality. Leaving a particular person without such a good is wrong simply because these are goods that this person must have if they are to be, or to see themselves as, an equal in our society. So it seems to me that, structurally, these are two different wrongs. They are still, to be sure, both ways of failing to treat someone as an equal. But they are different ways; one is not an instance of the other.

I have now argued that the claim that a practice denies someone a basic good differs both from the claim that it infringes their right to deliberative freedom and from the claim that it subordinates them. But there are still a number of puzzles presented by the idea that some discriminatory practices are wrongful because they deny people basic goods. In the last section of this chapter, I shall tackle what I believe to be the two most important ones.

4.4 Basic Goods, Prohibited Grounds, and Responsibility

One might still feel uneasy at the thought that denying someone a basic good can be sufficient to ground a claim of wrongful discrimination. One source of unease might be the legal requirement that claimants must prove that their wrongful discrimination has occurred on the basis of a prohibited ground. For it is not clear that this requirement serves any helpful function in cases where discrimination denies someone a basic good. In such cases, what matters is simply whether the opportunity or resource in question is genuinely a “basic good” for the discriminatee in my sense, and whether the allegedly discriminatory practice is one of the causes of the discriminatee lacking that good. If the good is genuinely a basic good for a particular person, then without it, he cannot be, or be seen as, a full and equal participant in his society. So, provided that the allegedly discriminatory practice is one of the causes of his lacking this good, then the practice will have wrongfully discriminated against him. But we can determine all of this without knowing whether the discrimination has occurred on the basis of a trait that is, or ought to be, on our list of prohibited grounds of discrimination. So it may look as though the prohibited grounds of discrimination have no important role to play in these cases.

However, although it is true that it is not necessary for wrongful discrimination to have occurred in such cases on the basis of a prohibited ground, it does not follow that the prohibited grounds have no role at all to play in these cases, or that there is no way to justify the common legal requirement that discrimination must occur on the basis of a recognized prohibited ground. We saw in Chapter Two that, in cases where discrimination is wrongful because it contributes to social subordination, the prohibited grounds help us to identify those social groups who most often stand in relations of subordination to other groups. That is, the grounds play a kind of heuristic role, directing us in those cases towards the social groups that are most likely to be victims of wrongful discrimination.

Here too, in cases involving a denial to someone of a basic good, I think we can see the common lists of prohibited grounds as heuristic devices—that is, as attempts to mark out those individuals who are most likely to be unable to be, or to be seen as, equals in their particular society. For instance, recall my earlier discussion of the indigenous water crisis. In the earlier part of that discussion, I compared the indigenous communities who were left without safe drinking water to other remote communities, who also lack safe drinking water but who are not indigenous. I argued that although safe drinking water is, for members of both communities, a precondition for their being equal participants in society, the lack of safe drinking water particularly affects how indigenous communities are seen, and it does so because they are indigenous. Given the stereotypes surrounding indigenous peoples and cleanliness, it is much more likely that their water crisis will leave them unable to be seen as equals in Canadian society than that it will leave a community of non-indigenous Canadians unable to be seen as equals. So even in this case, prohibited grounds do seem to play a role. They point us towards those individuals who are more likely, as a result of lacking a certain resource or opportunity, to be either unable to be, or unable to be seen as, equals in their society. They also point us towards some of the reasons why the lack of these resources or opportunities will have a distinctive impact on the social standing of these particular individuals. What makes the indigenous communities more vulnerable to the lack of clean drinking water, more likely than others to have their social standing affected by the lack of such water, is precisely that they are indigenous. Similarly, what makes a person living with dwarfism such as Manuel Wackenheim particularly affected by the ban on dwarf-tossing is precisely his disability. So, although there is no requisite extra step in our reasoning in such cases, in which we must make sure that the claimant lacks the requisite good because of a trait that amounts to a prohibited ground, we can see the legal requirement that the claimant refer to a prohibited ground as a way of honing in on those situations in which it is most likely that a claimant's lack of some resource or opportunity really does have an impact on whether they can be, or be seen as, an equal in their society.

Even if one accepts this explanation of the role of prohibited grounds in such cases, one might still find it difficult to accept that certain cases of discrimination are wrongful because they deny people a basic good. This might be because of concerns about responsibility. In some of the cases we have discussed in this chapter, the claimant's lack of an opportunity is directly due to the allegedly discriminatory agent. Governments, for instance, have control over how they define marriage, just as municipalities have the power to ban dwarf-tossing. But in other cases that we have discussed, the claimant's lack of a certain resource or opportunity is due to the concurrent actions of many other agents, and also to the operation of non-agential forces. As we saw earlier, for instance, the reasons many indigenous communities in Canada lack clean drinking water are complex, and have to do not just with government fiscal policies, but also with the remoteness of the communities, the polluting activities of a variety of industries and mining companies, the prevailing winds, the absence of easy alternative local water sources, and a myriad of

relevant geographical conditions. This raises an important question of responsibility. Is it fair to hold the alleged discriminator responsible for providing a basic good, in situations where the claimant's lack of that good is also due to so many other concurrent factors? This is a particularly worrisome issue in cases where, like the indigenous water crisis, the costs of providing the basic good in question are enormous. When the costs are so large, and the relative contribution of the alleged discriminator is only partial, is it fair to hold the discriminator responsible for eliminating the wrongful discrimination?

This set of concerns is helpful and important. But I think it blurs together a number of quite different questions. One of these is the question I have been trying to answer in this chapter and the two previous chapters: when does a discriminatory policy wrong someone by failing to treat them as an equal? I have argued in this chapter that it is wrongful as long as it does indeed result in someone's being denied a basic good. But there is a set of further questions, whose answers we cannot just read off of our answer to this question about wrongfulness. These include questions about all things considered wrongness. They also include questions about culpability: How far should the discriminator be held culpable for the wrong that he has committed, or that his practice perpetuates? And they also include questions about responsibility for cost: How much of the cost of eliminating wrongful discrimination is it fair to require the discriminator to bear? As I shall argue in Chapters Six and Seven, these are separate questions, and we need not assume that they will always be answered in the same way. Our answers may depend in part on who the discriminator is—or rather, on what type of agent the discriminator is, and what that agent's other responsibilities are. Is the discriminator the state? Is it a private individual who has stepped into the public sphere and offered certain things to the public? Is it a private individual, fulfilling what are traditionally regarded as more private familial responsibilities? Just because a particular practice amounts to wrongful discrimination, it does not follow that the discriminator is culpable, nor that he or his organization must bear the full costs of eliminating it.

If we separate out these questions, I think it becomes easier to accept that a denial of a basic good can indeed lead to wrongful discrimination. To claim this is not yet to draw any conclusions about which costs the government, or any other agent of discrimination, can fairly be asked to shoulder. It is simply to acknowledge that in some cases, leaving people without a resource or an opportunity can leave them unable to participate in society as an equal. So, just like social subordination and infringement of a right to deliberative freedom, leaving people without a basic good, too, can wrong them.

Chapter Five

A Pluralist Answer to the Question of Inequality

5.1 Why the Theory is both Pluralist and Unified

I have now explored three different ways in which discriminatory practices can wrong people by failing to treat them as equals. In Chapter Two, I argued that many discriminatory practices subordinate some people to others, sustaining the conditions of social subordination and either constituting an expression of censure of members of a certain group or rendering them invisible. In Chapter Three, I looked at the ways in which certain discriminatory practices infringe a person's right to deliberative freedom. And in Chapter Four, I tried to show that some discriminatory practices leave people without access to basic goods.

Each of these chapters aimed to demonstrate, not just that some discriminatory practices have these harmful effects, but that they thereby wrong people by failing to treat them as equals. This is most evident in the case of subordination. When a practice marks someone out as inferior or contributes to their having a lower social status than others, it clearly fails to treat them as equals. But we also saw that when a practice denies someone a deliberative freedom to which they have a right, it fails to treat them as an equal. For when someone's right to deliberative freedom is infringed, they are not treated as a person capable of autonomy. And given that our societies hold up, as a social and political ideal, the idea that each individual ought, as far as possible, to be treated as though they were capable of autonomy, it follows that when we fail to respect someone as a person capable of autonomy, we fail to treat them as an equal. Finally, we saw that leaving someone without access to a basic good is also a way of failing to treat them as an equal. For a "basic good" just is the kind of good that a particular person needs if she is to be, and to be seen as, an equal in her society.

All of the wrongs that we examined in these three chapters, then, can be seen as ways of failing to treat others as equals. But, as I hope the discussions in these different chapters showed, what does much of the work, in explaining why discrimination is wrongful, are the particular explanations of *why* people are not treated as equals: namely, because they are subordinated, or because their right to deliberative freedom is infringed, or because they have been denied a basic good. And, as we have seen, these explanations are genuinely different from each other. My theory of wrongful discrimination, therefore, is "pluralist" in the sense that we examined in Chapter One. It gives us a number of quite different interpretations of what it is to fail to treat people as equals, and it does not claim that all of these are reducible to some single issue. As I hope my analysis of different cases of wrongful discrimination in these three Chapters has shown, this pluralist theory is able to capture the rich, multi-faceted nature of discriminatees' complaints precisely because it

does not try to reduce the wrongness of discrimination to some single set of harmful effects, some single way of conceiving of what it is to fail to “treat someone as an equal.” At the same time, however, the theory does offer us a unified account of wrongful discrimination. It is capable of explaining why all of these different wrongs are all instances of wrongful *discrimination*, as opposed to diverse wrongs that have nothing to do with each other. They are all instances of wrongful discrimination because they are all cases in which someone is treated differently from others on the basis of a certain trait and thereby not treated as an equal. But when we ask *why* exactly this person or this group was not treated as an equal, our answer may be different in different cases. In some cases, it will appeal to unfair social subordination. In some cases, it will appeal to the infringement of a right to deliberative freedom. In some cases, it will involve a denial of a basic good. And as we have seen, some cases may involve multiple wrongs, simultaneously.

At the start of this book, I noted a number of worries about arbitrariness that any pluralist theory faces, and I promised to address the legitimate worries at a later point, after we had laid out the different components of my pluralist theory. We are now in a position to do this. Section 2 of this Chapter aims to alleviate these worries about arbitrariness and explanatory power. Section 3 briefly explains that on this theory, although all wrongful discrimination fails to treat others as equals, other kinds of acts—that is, non-discriminatory acts—could also do this. So there is nothing distinctively or uniquely wrong with discrimination. I explain why this is not a problem. I then turn in subsequent sections of this Chapter to a number of questions that are raised by the different components of my pluralist theory. Why should we think that each component, on its own, is sufficient for wrongful discrimination? Can there be different sets of victims and different kinds of obligations, depending on *which* wrong we are concerned with? What weight do these different wrongs have, relative to each other—and how should we reason through cases such as the Wackenheim case, in which it seems that we must continue to commit one of these wrongs if we are to take the necessary steps to eliminate another one? Finally, in the last section of the Chapter, I argue for some further advantages of this pluralist theory, over and above its capacity to offer a nuanced account of the wrongs at issue in different cases. I suggest that the theory helps to explain a number of persistent disagreements between legal scholars over discrimination—disagreements over whether assessments of wrongful discrimination require comparative judgments, disagreements over the role of individuals and groups, and disagreements over the role of the prohibited grounds. These disagreements have persisted, I argue, because *there is no one answer* to any of these questions. It depends on the particular kind of wrong that is at issue in a certain case, and certain cases may involve more than one wrong. So a pluralist theory such as mine can help to explain both why these disagreements have persisted for so long, and how we might address them.

5.2 Resolving Worries about Arbitrariness

In Chapter One, I noted that pluralist theories of discrimination give rise to certain special concerns about arbitrariness.¹ I suggested that we could distinguish several different objections among these concerns. We are now in a position to answer these objections.

One objection is that if a theory of wrongful discrimination appeals to several different ways of failing to treat someone as an equal, it risks being arbitrary in the sense that we have no greater reason to appeal to these ways of failing to treat someone as an equal than we have for appealing to any others.² When I first laid out this objection, I noted that if I could show in Chapters Two, Three, and Four that my own theory makes good sense of the complaints of real victims of discrimination, and if I could show that the theory is consistent with certain basic features of anti-discrimination law, then I would have provided an answer to this particular concern. For the fact that a theory offers nuanced explanations of our lived experiences of discrimination, and the fact that it explains certain basic features of our laws, together give us good reason for thinking that it tracks something correct about the moral phenomenon in question, and that these really are at least some of the reasons for thinking discrimination wrong in certain cases. In each of these chapters, I tried to derive my understanding of the relevant way of failing to treat someone as an equal from the complaints of discriminatees, looking in each case in detail at the structure of their complaint and the claims they made about it. Moreover, I looked mostly at cases that most of us would agree are wrongful (with a few exceptions, where there was a special reason for looking at a more controversial case). I also tried to show that my theory is consistent with the idea that claimants must bring their claim of wrongful discrimination on the basis of certain prohibited grounds, and with the idea that there is a distinction to be drawn, at least in certain contexts, between direct and indirect discrimination. Because my theory can make sense of these legal doctrines, and because it is rooted in the real concerns of complainants in cases of discrimination, it gives us good reason to think that subordination, infringements of rights to deliberative freedom, and denials of basic goods *are* some of the important reasons why discrimination is wrong, when it is.

But the arbitrariness worry may take a different form. It may instead be the worry that if a theory tries to explain a certain moral concept with reference to a number of irreducibly different ideas, then it will not really be *explaining* this moral concept. It will simply be giving us a list of items that are in some way related to it. So it will not really be a theory at all.³ This is a familiar concern about pluralist theories of moral phenomena.⁴ We

¹ See Chapter 1, Section 1.5.

² See Lawrence Blum, "Racial and other Asymmetries: A Problem for the Protected Categories Framework for Anti-discrimination Thought," in Deborah Hellman and Sophia Moreau (eds.), *Philosophical Foundations of Discrimination Law* (Oxford: Oxford University Press, 2013), pp. 182–202.

³ See, for instance, Patrick Shin, "Is There a Unitary Concept of Discrimination?" in Deborah Hellman and Sophia Moreau (eds.), *Philosophical Foundations of Discrimination Law* (Oxford: Oxford University Press, 2013), pp. 163–181.

⁴ Dating back to Plato: see Socrates' discussions in the *Meno* (at 71d–77a) and in the *Euthyphro* at 6d–e. *Plato: Complete Works*, ed. John Cooper (Indianapolis: Hackett, 1997).

can see it voiced, for instance, in relation to “objective list” theories of well-being: scholars have argued that, rather than being theories of what well-being is, objective list theories are really just lists of its different components.⁵ Without a single underlying thread to tie the items on the list together, such theories appear to give us no real explanation of what “well-being” is. The same worry might be expressed of a pluralist theory of wrongful discrimination. It risks being a mere list of some of the circumstances in which discrimination is allegedly wrong. And if it is just a list, without offering an underlying explanation for why certain items are on the list and certain items, off it, then how is it really a theory of wrongful discrimination?

But are the different wrongs done by discrimination unconnected, on my theory? No—in fact, they are linked by two features. The first is a feature of all cases of discrimination *qua* discrimination: one or more people are treated differently from others on the basis of certain traits. The second is a feature of all cases of *wrongful* discrimination: they fail to treat someone as an equal. So all instances of wrongful discrimination, on my view, share two features. First, they treat certain people differently on the basis of certain traits; and second, under the circumstances, they thereby fail to treat these people as the equals of others. We can fail to treat someone as an equal because we subordinate them to others; we can fail to treat them as an equal because we infringe their right to deliberative freedom; and we can fail to treat them as an equal because we leave them without access to a basic good. But these are all ways of failing to treat someone as an equal. So the three items in my pluralist theory are not unconnected items on a list. They are not analogous to the items on an objective list theory of well-being. Objective list theories do not offer us any explanation of *why* these items belong on the list, other than the claim that they contribute to our well-being. But my theory does offer us an explanation of why subordination, infringements of a right to deliberative freedom, and denials of basic goods belong on our list of the wrongs done by discrimination. They are all reasons why, when one treats someone differently on the basis of certain traits, one can fail to treat them as an equal. Moreover, I did not, in earlier chapters, simply take it for granted that these were all ways of failing to treat others as equals. Both in these previous chapters and at the start of this current chapter, I explained *why* each of these, in its own right, constitutes a failure to treat others as equals.

But although this may help to satisfy you that the theory’s components are connected, you may still wonder whether the theory has sufficient explanatory power. I have said that what unifies the different wrongs on this pluralist theory is the fact that all of them involve a failure to treat others as equals, in the process of distinguishing between people on the basis of certain traits. But I have also said that the abstract moral idea of failing to treat others as equals does not, on its own, explain why discrimination is wrong. Much of the explanatory work is done by the particular reasons why a practice fails to treat others as an equal, such as the fact that it subordinates people, or infringes their right to deliberative

⁵ This criticism has been made by L.W. Sumner, *Welfare, Happiness, and Morality* (Oxford: Clarendon Press, 1996) at p. 45; and Mark Murphy, *Natural Law and Practical Rationality* (Cambridge: Cambridge University Press, 2001) at p. 95. Others, however, have argued that so-called “enumerative theories” can be legitimate: see, for instance, Roger Crisp, *Reasons and the Good* (Oxford: Clarendon Press, 2006) at pp. 102–103.

freedom, or denies them a basic good. Is this not a problem? I do not think so. Each of the wrongs that I have discussed in detail in Chapters Two, Three, and Four provides a different interpretation of what it is to fail to treat others as an equal. They are, to borrow a distinction from Rawls, different “conceptions” of this basic “concept” of differentiating between people on the basis of certain traits in such a way as to fail to treat some people as equals.⁶ When we ask: “Why is discrimination wrong?” we could say “Because it fails to treat people as equals.” But that does not explain what it is for a discriminatory practice to fail to treat some people as equals, and so it is an incomplete answer. If we are to answer what I called “the question of inequality,” we need to go on to explain in detail *why*, when we distinguish between people on the basis of certain traits, we fail to treat them as equals. The fact that our more complete answer appeals to a diverse array of considerations, and not simply to the ideal of treating people as equals, is not a problem: it is a proper response to the explanatory task at hand. The adequacy of each of my explanations in Chapters Two, Three, and Four needs to be judged on the basis of such considerations as whether it seems accurately to capture the legitimate complaints of discriminatees and whether it accords with basic features of legal doctrine. The mere fact that these explanations differ does not, on its own, cast doubt on their explanatory power.

I have now tried to address concerns about arbitrariness. But perhaps lurking under these concerns is a different worry, a worry about the lack of distinctiveness of the wrongs involved in cases of discrimination, on my pluralist account. I shall address this concern in the next section.

5.3 Nothing Distinctively Wrong With Discrimination?

My theory of wrongful discrimination recognizes that not all discriminatory practices are wrongful. They become wrongful when they fail to treat some people as equals; and they can fail to treat some people as equals for a number of different reasons. However, at least some of these reasons are also reasons why *other* kinds of acts can be wrongful, acts that are not acts of discrimination because they do not involve distinguishing between different people on the basis of the kinds of personal traits that would normally appear on a list of prohibited grounds of discrimination—nor indeed, distinguishing between people on the basis of anything that could accurately be called a “personal trait.” For instance, I can mark people out as inferior simply by following certain social conventions in our society, such as pushing them to their knees as they approach me or spitting in their direction. If I do this, not on the basis of any personal trait, but just randomly, I am still marking these people out as inferior, and my behavior is likely still wrongful; but I am not discriminating against them. Similarly, suppose a government fails to provide a certain good to some of its citizens over a number of years, a good that we would deem a basic good for these citizens in this society. But suppose the lack of this good cannot be traced to any personal feature of these citizens—it is not on the basis of any personal trait that their government has left them without this good, nor do they independently lack the good because of some trait such as a disability or their race. Perhaps they lack it simply for reasons of geography. This

⁶ John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971), Ch. 1, s. 2.

would not be recognizable as a case of wrongful discrimination, but it would nevertheless be a failure to treat them as equals.

My account, then, has the implication that at least some of the reasons why certain discriminatory practices are wrongful are not reasons that are unique to cases of discrimination. Although all cases of wrongful discrimination involve practices that differentiate between people on the basis of certain traits in a way that fails to treat them as equals, it turns out that some of the reasons why these practices are wrong are also reasons for thinking other practices are wrong.

Is this a problem? I do not think so. Many moral theories imply that acts of different kinds are wrong for the same reason. So it cannot be that this conclusion is problematic in the case of discrimination because it is, as a general rule, implausible to think that different kinds of acts could all be wrong for the same reason. If it seems problematic for a theory of discrimination to claim that all of the reasons why acts of discrimination can be wrong are also reasons why other sorts of acts can be wrong, I think this must be because we are assuming that wrongful discrimination is somehow especially heinous, as compared with other kinds of wrongdoing—and therefore, that acts of wrongful discrimination must somehow all be wrong for the same special reason. Certainly within the popular media, charges of “discrimination” carry with them a peculiar kind of stigma: it is often assumed that they are particularly serious, and that the agent is especially blameworthy. We might therefore think that only an account of wrongful discrimination that traces the wrongfulness of discrimination to some unique feature of discriminatory acts, shared by all and only these acts, could explain this stigma.

But I think we should pause before accepting the idea that discriminatory acts or practices are, as a group, especially heinous. As the various examples we have discussed so far within this book indicate, the practices that appear to be wrongfully discriminatory vary enormously in their underlying motivation and in the level of the agents’ awareness of the impact on the excluded group. Within the class of direct discrimination, some acts are maliciously done; others reflect a patronizing but well-meaning prejudice; and still others are done with regret and what is felt by the agent as an unfortunate financial necessity. Within the class of indirect discrimination are cases in which the agent is fully aware of the impact of a certain practice on certain social groups, other cases in which the agent is unaware of the impact but chooses not to investigate, and still other cases in which the agent is unaware that there is even an issue that they could consider investigating. Do all of these seem equally heinous? Do all of these agents seem equally deserving of blame? I shall look in much more detail both at the question of the moral seriousness of discrimination and at the different but related question of the blameworthiness of the agent in Chapter Six. But for now, I think it suffices for us to note that these differences between cases of wrongful discrimination cast some doubt on the view that all wrongfully discriminatory practices are equally heinous. And if only some wrongfully discriminatory acts are particularly heinous, then this cannot be because of the wrong at issue—for in all cases, it is the wrong of failing to treat some people as an equal. The special heinousness of *some* acts of wrongful discrimination must instead have to do with certain special facts about the agent in these cases—perhaps, facts about their motives, or about the knowledge they had available to them, or about the special roles or responsibilities they had under the

circumstances. And if this is true, then the heinousness of these particular acts of wrongful discrimination cannot provide us with a reason for thinking that we need an account of the wrongness of discrimination that traces it back to some feature that is common to all and only discriminatory acts.

I shall now turn to a number of questions that concern the different components of this pluralist theory, and their relationship. Some of these are questions about the nature and weight of the different reasons that we have to avoid these different wrongs. Others are questions about who is wronged by each of these different ways of failing to treat people as equals: that is, who the victims are. For they are not the same across all wrongs. And this leads to a complication, which I shall in the last section of this Chapter argue is in fact an advantage of this pluralist theory. It means that in the same case of wrongful discrimination, there can be different victims, depending on what wrong we are focused on, and different obligations of different kinds placed on the agent. And this of course means that, when tribunals or courts analyze cases of discrimination, it matters that they figure out exactly which way of “failing to treat others as equals” is at issue. For the consequences, both for discriminators and for discriminatees, will be different, depending on the particular wrong that is at stake.

5.4 Each Wrong Sufficient for Wrongful Discrimination

Although I separated out the different ways of failing to treat people as equals and devoted a separate chapter to each of them, most of the cases we examined in these chapters involved practices that failed to treat people as equals in more than one of these ways. For instance, I noted in Chapter Three that the Hyperandrogenism Regulations both contribute to the social subordination of women from the global south, while at the same time infringing the right to deliberative freedom of those female athletes whose hormones test at higher than acceptable levels. Similarly, we saw in Chapter Four that the indigenous water crisis in Canada both contributes to the social subordination of members of indigenous communities and denies them a basic good. But I have claimed that each of these is, on its own, sufficient for wrongful discrimination; and I have been calling each of them, on its own, a “wrong.” Why should we think this is true, rather than thinking that each works in tandem with the others, and would be insufficient on its own? I shall give two quite different arguments for this claim. The first is a more theoretical argument; the other involves an appeal to particular cases.

The theoretical argument for the claim that each of these is, on its own, sufficient for wrongful discrimination appeals to the fact that, as I have argued, each of these is a way of failing to treat people as equals. I began this book by noting, in Chapter One, that the law assumes that wrongful discrimination is wrongful not just because it differentiates between people, but because, in doing so, it fails to treat certain people as equals. It treats them as inferiors, instead. I then argued, in each of Chapters Two, Three, and Four, that there were certain distinctive ways of failing to treat others as equals: namely, by subordinating them, by infringing their right to deliberative freedom, and by denying them a basic good. If my arguments in these chapters are sound, then each of these really is, on

its own, a way of failing to treat others as an equal. It then follows that each is sufficient to constitute wrongful discrimination, even in the absence of the others.

But in case you are not persuaded by this, there is another argument we can turn to. It appeals to several cases in which, unlike the bulk of cases of wrongful discrimination, the wrongful practice seems to fail to treat people as equals in only one of these ways that I have discussed. It is not a problem for my argument that these are unusual cases, rather than representative ones. For as long as it is *possible* for there to be some such cases, then it follows that each of these can, on its own, be sufficient for wrongful discrimination.

Consider first the “Sketching the Line” program that I discussed in Chapter Two, in which sketches of allegedly representative Toronto transit riders are posted up on transit vehicles across the city. I mentioned this program in that Chapter to make a quite specific point: I was arguing that practices can mark groups out as inferior even in the very act of rendering them invisible. But this program also seems to be a good example of a program that discriminates only in the sense of subordinating visible minorities, and not also in the sense that it denies them a deliberative freedom or leaves them without access to a basic good. Precisely because the program renders such riders invisible, we cannot claim that it forces them always to have their race before their eyes when riding the subway; nor does it place any special cost or burden on them. And having a picture of riders such as yourself posted on public transit is in no sense a basic good. I do not even think one could plausibly claim that being recognized as the typical rider of public transit is a basic good, since our society places no particular importance on riding public transit. (If anything, having to ride public transit rather than having your own private vehicle is, in some social circles, an indication that you lack a certain prestige.) But it is nevertheless true that, through this program, riders of visible minorities are rendered invisible. And, given that the images in public transit are seen by so many people, and are often taken as a microcosm of society at large, the invisibility of these minorities on public transit seems to contribute in a significant way to their invisibility in society at large. So this is an example in which members of visible minorities are wronged, because they are subordinated; but they are not wronged in either of the other two senses that we have discussed. Hence, social subordination is sufficient for wrongful discrimination, even in the absence of the other two ways of failing to treat people as equals.

What about a case that only seems to involve a denial of deliberative freedom, but not social subordination, and not a denial of a basic good? A recent case of employment discrimination against Caucasian employees at a resort seems to be a good example of this.⁷ The resort was sold, and its new owner made repeated comments about how he would prefer ethnically Chinese employees over the resort’s current Caucasian employees, because he believed that ethnically Chinese employees would not demand overtime pay or pay on statutory holidays. The Human Rights Tribunal hearing the case found that this attitude was in large part responsible for the subsequent firing of some of the employees and for the resignation of the others. In this case, the employees were left with their race always before their eyes; and it became a very real cost, both within the workplace while they were there and then subsequently, when they were forced to leave. But they were all

⁷ *Eva and Others v. Spruce Hill Resort and Another*, 2018 BCHRT 238.

Caucasian, members of a racially privileged group; and the owner's attitude was in fact more insulting of the ethnically Chinese employees whom he hoped to replace them with. So it does not seem plausible to suggest that his discriminatory practices socially subordinated the Caucasian employees or marked them out as inferior. It simply marked them out as more expensive. And although the Caucasian employees lost their jobs—and this is a significant loss, particularly in a community with limited employment opportunities—I do not think we can claim that they were denied a basic good. It may be true that everyone in our society needs to have access to *a* job in order to be, and be seen as, an equal (though even that might be too strong, and the correct claim might be only that everyone needs to have the opportunity to apply for, and be taken seriously as a candidate for, some job). But, barring very special circumstances, it would not be plausible for us to claim that having a particular job at a particular place of employment is a basic good. So this, too, is a case in which an act of discrimination is wrongful for one of the reasons we have examined, without being wrongful for the others.

Finally, is there a case in which a person or group is denied a basic good, but the denial of that good does not contribute to their social subordination, and their right to deliberative freedom is not infringed? Consider again the dwarf-tossing case from Chapter Four.⁸ I mentioned there that the main argument of Wackenheim, the complainant, could be seen as based on a denial of a basic good: because the towns together banned dwarf-tossing, and because dwarf-tossing was the only employment available to people with dwarfism in this area, he was left without any form of employment. Although assessing the bans' effect on the social subordination of people with dwarfism is a complicated task—it would, for instance, be naïve to suggest that the bans were entirely beneficial, in part because they carry the patronizing implication that these people, like children, require protection from certain kinds of consensual activity—nevertheless, let us assume that the bans do more to combat the subordination of members of this group than they do to perpetuate it. So they do not, all things considered, contribute to the social subordination of people living with dwarfism. Do they infringe the right to deliberative freedom of people with dwarfism? They are, at least in intent, supposed to liberate them: people with dwarfism will no longer have to think of themselves as objects of ridicule, or, indeed, as beings akin to objects. So it is possible that this is a case of being denied a basic good but not being wronged in either of the other ways we have examined.

I have now tried to show that each of these is, on its own, sufficient to wrong someone; though, as we have seen throughout the book, many cases of discrimination involve more than one of these wrongs. But, as I shall now go on to explain, there are important differences between these wrongs—differences in their scope, in the people who are wronged, and in the kinds of reasons they provide us for rectification or restitution. So it matters that we examine any given case of discrimination closely, to see which particular wrong or wrongs are at issue.

5.5 Personal Wrongs and Group Wrongs

⁸ *Manuel Wackenheim v. France*, Communication No 854/1999, U.N. Doc. CCPR/C/75/D/854/1999 (2002).

The three wrongs that we have explored are all ways of failing to treat people as equals. But these wrongs differ in a number of important ways. To see this, it will help us to distinguish between what I shall call a “personal wrong” and what I shall call a “group wrong.” A personal wrong is a violation of an obligation to a particular person, a violation that generates a claim on that person’s part to some form of restitution: a claim to a job, for instance, or a claim to have their particular schedule adjusted so that they can pray at the requisite times of day. By contrast, a “group wrong” is a wrong that involves a failure to treat a group of people properly, but that does not generate any distinct claim on the part of any particular members of this group for any special form of restitution. Rather, the appropriate way to rectify a group wrong is to change a practice so as to ensure that no members of this group are, in the future, disadvantaged in a certain way or denied certain opportunities. Note that this particular usage of the term “group wrong” is quite consistent with the claim that wrongs to a group are reducible to wrongs to the group’s members. So I am not suggesting that a group wrong is a wrong to some separate entity, “a group,” over and above its particular members. On the contrary, when a group is wronged in a case of discrimination, this wrong just consists in the wrongful treatment of each person who is a member of that group, at the time the wrong was committed.

With this distinction in place, we can see that some of the wrongs involved in discrimination are personal wrongs, and others are group wrongs. Moreover, the same case can involve both personal and group wrongs. Think back to the Hyperandrogenism Regulations, for example. We saw that these regulations infringed the right to deliberative freedom of particular athletes such as Dutee Chand and Castor Semenya. This infringement constitutes a personal wrong towards Chand and Semenya. Each can claim that their deliberative freedom was interfered with, and that they have a right not to have it interfered with. We also saw that the Regulations mark out these athletes as inferior to other women—as not “real” women. This wrong too seems a personal one: each female athlete who has naturally high levels of these hormone levels has a claim, we want to say, not to be censured in this way, as less than a real woman, and as a matter of restitution, she is owed the opportunity to run in the women’s races without having to take hormone supplements. Finally, we saw that these regulations play a causal role in sustaining the social subordination of a broader class—namely, women from the global south, regardless of whether they are athletes or non-athletes. So they also subordinate this broader class of women. In so doing, however, they do not generate a further personal claim on the part of any one woman from the global south to any particular good or opportunity. Rather, this last wrong is what I called a “group wrong.” It is true that each member of the group “women from the global south” has not been treated as an equal. But this does not generate a special claim on their part to personal restitution. It is not a personal wrong; it is a group wrong. And the way to rectify this particular group wrong is to change the regulations, so that in future no female athlete from the global south will encounter these barriers. This will not, of course, eliminate all of the barriers faced by women from the global south. But it will eliminate *this* cause of their marginalization and subordination.

As this example suggests, the wrong of infringing someone’s deliberative freedom is a personal one: it generates a personal claim for redress. By contrast, the wrong involved

in causally contributing to subordination is a group wrong: it does not generate any particular claim on any one person's part to any special kind of restitution, over and above the measures that need to be taken in order to ensure that the group does not face this discrimination in the future. In the Hyperandrogenism Regulations example, we also saw that marking out female athletes with higher hormones as inferior to other women generates a personal claim on their behalf. But I think that not all instances of marking a person or a group out as inferior, or censuring them, are personal wrongs. So whether this particular kind of wrong—that is, marking out certain people as inferior—is personal or group-based may vary from case to case. In some cases, such as the “Sketching the Line” program that excluded visible minorities, it is the group that is made invisible, as a group, and in these cases, we may be inclined to say that it is a group wrong, which generates no particular claims for restitution on the part of individual members. But in other cases, an individual may be censured for being a member of a certain group, and we may feel that this person has had a personal wrong committed against them.

I have not yet discussed the wrong of leaving people without access to basic goods. This seems to be a personal wrong, since it generates a claim on the part of each member of the group to be given the basic good in question. So, for instance, the members of the indigenous communities who lack water for drinking and for symbolic purposes, and who are therefore unable to participate in Canadian society as equals, are entitled to clean water. They have a personal claim to clean water. Similarly, the gay couples who want access to the institution of marriage have a personal claim to be given such access. As both these cases suggest, a personal claim does not need to be a claim to some individually divisible or privately appropriable good: clean water is something that is provided to the community as a whole if it is provided to anyone, and marriage is not privately appropriable.

I have been focusing on the different *kinds* of claims that these different wrongs generate. But the same examples that demonstrate that they generate different kinds of claims also show that they generate claims *by different groups*. And so it is particularly important that we focus separately on each of the wrongs that is at issue in a given case of discrimination, so that we can be sure we are thinking of the right discriminatee or claimant. In the case of the Hyperandrogenism Regulations, the infringement of the athletes' right to deliberative freedom and the censuring of them as less than real women wrong the group of female athletes with higher than acceptable natural levels of certain hormones. But they contribute to the social subordination of a number of broader groups: women, women from the global south (regardless of whether they are athletes or not), women athletes, and women athletes from the global south.

That a single discriminatory practice can wrongfully discriminate against different people in different ways is not a problem for my theory. It simply shows us that we need to be careful when thinking through and adjudicating cases of wrongful discrimination. For a single case may involve multiple wrongs and multiple discriminatees. It will be true of all of the discriminatees that they have not been treated as equals, in relation to others. But

some of them will have suffered a personal wrong, whereas others may have suffered a group wrong. And this leads to an interesting complexity. It means that the reasons that we have to eliminate these forms of wrongful discrimination can't be assumed always to be reasons for implementing the same solution. For instance, in some cases, the policy that we would need to adopt in order to counteract the social subordination of one group of people may result in our being unable, for a time, to give a smaller sub-class within this group a particular basic good, or a deliberative freedom to which they have a right. The Wackenheim case that I discussed in the last chapter involves just such a conflict. I shall now turn to it, and to the broader question that it raises concerning the relative weight of the different reasons that we have to rectify these different wrongs.

5.6 Cases of Conflict and the Relative Weight of Different Reasons

There can be difficult cases, in which we cannot eliminate a practice that wrongfully discriminates against certain people unless we take measures that wrongfully discriminate against a certain sub-group of this broader group, either for a different reason or for that same reason. As we saw in the last chapter in our discussion of the Wackenheim case, dwarf-tossing is a practice that discriminates against those living with dwarfism, in ways that fail to treat them as equals. Most notably, the practice contributes to social subordination by encouraging members of the public to view people with dwarfism as sources of amusement and as toys that can be thrown rather than as subjects with as much agency as the rest of us. But, as Manuel Wackenheim argued before the U.N. Human Rights Committee, there is such deep prejudice against people with dwarfism in French society that one of the only sources of employment for people with dwarfism is the sport of dwarf-tossing; and Wackenheim was adamant that having a job was necessary to his self-esteem and to his ability to view himself as an equal to others. So this may be a situation in which the French municipalities in question will wrongfully discriminate against someone no matter what they do. If one of these municipalities bans dwarf-tossing, then it will prevent those people who depend on dwarf-tossing for their employment from getting any job at all. So it will leave these particular people without access to a basic good. But if the municipality does not ban dwarf-tossing, it will perpetuate the social subordination of people with dwarfism, and so will fail to treat this group as equals.

That there can be such cases, in which we wrong someone no matter what we do, is a familiar idea from moral philosophy. Bernard Williams once argued, even more strongly, that there can be cases in which we act wrongly no matter what we do—that is, cases in which every act available to us is what I earlier called “all things considered wrong.”⁹ He called these cases “moral tragedies.” I am not making that strong claim here; but I am making an analogous weaker claim. I am suggesting that in some cases of discrimination, all of the acts that are open to us will *wrong someone*. They may not all be “all things

⁹ Bernard Williams, “Moral Luck” in *Moral Luck* (Cambridge: Cambridge University Press, 1981).

considered wrong.” Perhaps there is one act that is, all things considered, preferable to others, and so we ought to perform it. But even when we do this, we will be wronging someone by failing to treat that person as an equal.

This is, I think, a helpful way of understanding our ambivalence about cases of affirmative action. When a firm adopts a quota for female employees or employees from racial minorities, it is singling out members of this group and implying that they need special help securing a proportionate number of jobs. A standard objection to such quotas is that they invite us to see members of these groups as unable to secure a position on the basis of merit alone, and therefore as less talented than others. So even when quotas are beneficial over the long term, they still contribute in an unfortunate way towards temporarily re-entrenching stereotypes about the inadequacy of the very groups they aim to protect, and temporarily facilitating their social subordination.¹⁰ Quotas also temporarily lessen the deliberative freedom of those members of these groups who are actually hired, as the members of these groups must go about their work aware that others think of them as “charity cases.” But of course the aim of quotas is to eliminate social subordination; and they do seem, in at least some circumstances, to be a necessary means of achieving this goal. Perhaps the correct way to think about such cases is that these are, like the *Wackenheim* case, instances in which we wrong people no matter what we do. In order to substantially reduce the social subordination of a particular group over the long term, and thereby not wrong them, we adopt measures that, for a short time, wrong either this group as a whole or a certain sub-group within it, by temporarily contributing to their social subordination and denying them deliberative freedoms to which they have a right.

I have argued that in some affirmative action cases, and in *Wackenheim*, we wrong someone—by failing to treat them as an equal—no matter what we do. Knowing this may help us to make sense of why such cases seem so difficult. But what does my theory tell us about what we ought to do all things considered, in such cases? What are the relative weights that we ought to assign to these different wrongs? Is it, for instance, more important or more urgent to stop social subordination, or more important to give a particular member of the subordinated group a deliberative freedom or a certain basic good?

Before I respond to these questions, I want to note that they take us past the question of equality and into a very different stage of reasoning about discrimination—the stage of justification, at which we determine whether practices that wrong people are nevertheless justified all things considered, or whether they are instead wrong all things considered. I shall have other things to say about justification in Chapters Six and Seven; but even there, I shall not attempt to offer a complete theory of justification. My aim in this book, as I have said, was simply to answer the question of inequality. I wanted to figure out

¹⁰ See, for instance, Tristin K. Green, “Discomfort at Work: Workplace Assimilation Demands and the Contact Hypothesis,” *North Carolina Law Review* 86 (2008), pp. 379–440 at p. 388.

when discriminatory practices wrong people by failing to treat them as equals. I did this in Chapters Two, Three, and Four. The question of justification, though important from a practical standpoint in helping us figure out what we ought to do in such cases, takes us into a further inquiry. And it is beyond the scope of this book to offer a full theory of justification.

I do not think that we can determine in the abstract what weight or normative force these three wrongs have, relative to each other, because I doubt that they have a single unvarying weight or normative force across all cases. You might think that they would—because after all, they are all ways of failing to treat people as equals. Isn't it just as bad, or just as serious, to fail to treat people as equals in one particular way as it is to fail to treat people as equals in some other way? But of course, as I have indicated, each of the three ways of failing to treat people as equals that I have explored provides us with a very different *conception* of inequality, of what it is to wrong people by failing to treat them as equals. And so the seriousness of failing to treat someone as an equal in each of these ways may well differ, depending on the particular way in which one is not treated as an equal; and even a single one of these wrongs may carry a different weight in different cases. For instance, contributions to subordination clearly come in degrees. A practice can contribute to the subordination of a particular group to a greater or a lesser extent. Compare the practice of not admitting women to law school at all on the grounds of their sex with the various discriminatory practices that female students at such schools faced once the schools started admitting them, such as practices of directing women towards more “feminine” areas of law like family law; and compare this, in turn, with the lingering forms of discrimination that female students in some schools face today, such as being required to dress as a model for their clerkship interviews. It seems unlikely that these different ways of failing to treat women as equals make equally large contributions to the subordination of women. And while infringements of a right to deliberative freedom and denials of a basic good do not come in degrees (either you have a right to a certain deliberative freedom or you don't, and either a certain good is a basic good or it is not), nevertheless, the weight of these two wrongs in different cases may be different because of another variable. This is the *number* of people who have been wronged in these ways—and numbers are also of course a variable factor in cases of wrongful social subordination. Should it matter, in cases where we are forced to choose between wronging people in one of these ways and wronging people in another, how *many* people are wronged in each way? For instance, if it seems plausible in the Wackenheim case that the municipalities should ban the practice of dwarf-tossing because it contributes to the subordination of a large number of people, is this partly because *so many people* are subordinated when we allow dwarf-tossing to occur, whereas relatively few are actually involved in the sport of dwarf-tossing and so relatively few people will be denied a basic good if the sport is banned?

The question of what moral significance we should give to the number of people who are wronged in a particular way is a very complex one, and I do not have the space to

discuss it in any detail here.¹¹ But I do think it is worth noting that, whatever the significance of numbers is, it cannot be that the correct approach is simply to weigh the number affected on the one side with the number affected on the other. For one thing, as I have just noted, contributions to subordination come in degrees; so even if a very large group is subordinated by a particular practice, it may be that the contribution that this particular practice makes to their subordination is relatively small. It might, then, seem more urgent to provide a basic good to a smaller group, even if that meant that the subordinating practice had to persist. Another complication here is that some of the people wronged in one of these ways may also be members of the group that is wronged in one of the other ways. This is the case both in affirmative action cases and in the *Wackenheim* case. In *Wackenheim*, if we deny those people who are seeking employment through dwarf-tossing the basic good of a job in order to combat the social subordination of the much broader class of “all those living with dwarfism in France” then the class that we are failing to treat as equals is a sub-set of the broader class that we are treating as equals. So in one way, we are treating the members of this sub-class as equals. Moreover, we are doing it for their long-term benefit: the hope is that, if enough attitudes change through the elimination of demeaning sports such as dwarf-tossing, prejudices will be lifted and more people with dwarfism will be able to find other sorts of employment. Similarly, in the cases we considered of quotas for women in certain places of employment, the same women whose right to deliberative freedom is temporarily denied and who are forced to endure ongoing stereotypes about women in that workplace are also members of the broader group that stands to benefit from lesser social subordination as a result of more women assuming positions in that workplace. So although we ought to leave open the possibility that it is relevant in some cases that a far greater number of people will be wronged in one way than will be wronged in another, we need to bear in mind that other considerations will also be relevant here, such as who it is that is suffering these different wrongs, what the degree of the subordination in question is in cases of social subordination, and whether the imposition of one wrong on certain people is a necessary step in the elimination of other wrongs to a broader group that includes this smaller group.

I have said that I cannot settle here, in the abstract, the question of how the numbers might matter when we weigh different wrongs against each other and try to determine what is the right thing to do overall. But there is a way in which my view can help us think clearly about this question. It is sometimes tempting to think that either numbers always matter to the moral seriousness of a particular wrong, or they never do. But on the view of wrongful discrimination that I have proposed, there are two stages to our reasoning. First, we ask whether a particular discriminatory practice wrongs people by failing to treat them as equals. And when we engage in this inquiry, it does not matter how many people are

¹¹ For general discussions of the moral significance of the number of people who are wronged or harmed by a particular act, see John M. Taurek, “Should the Numbers Count?” *Philosophy and Public Affairs* 6.4 (1977), pp. 293–316; F.M. Kamm, “Aggregation and Two Moral Methods,” *Utilitas* 17.1 (2005), pp. 1–23; and Veronique Munoz-Darde, “The Distribution of Numbers and the Comprehensiveness of Reasons,” *Proceedings of the Aristotelian Society* 105.2 (2005), pp. 207–33.

affected by a particular practice. As long as *some* people are not treated as equals, it follows that some people have been wronged, and the discrimination is wrongful. But then we ask: is this practice all things considered justified? And at this stage, the numbers may matter. It may, at this second stage, matter how many people suffer each sort of wrong; though, as I have indicated, it is not a simple matter of aggregating the wrongs on one side and the wrongs on the other.

5.7 Advantages of this Pluralist Theory

I have tried to show how the different components of my pluralist theory of wrongful discrimination fit together; and I have defended the theory against the objection that it is arbitrary, and against the related objection that it cannot explain the distinctiveness or peculiar seriousness of wrongful discrimination. But I have not said very much in general terms about the advantages of the theory.

5.7.a It makes possible a nuanced analysis of cases

One reason for my silence about the theory's advantages is that, in my view, one of the main advantages of the theory is that it offers us a rich and nuanced way of understanding what goes wrong in different cases of discrimination. It helps us pry apart the different wrongs that may be involved, even in a single case of discrimination. And, rather than requiring us to focus only on one kind of fact—for instance, the demeaning nature of certain discriminatory practices, or the way they restrict our freedom—my theory enables us to see how a number of different features of these practices could all be relevant to whether they are wrongful. It thereby enables us to explain and validate many claimants' thoughts about the specific ways in which they have been wronged, without oversimplifying their complaints. But of course I cannot prove that it does this in the abstract. This depends on whether the arguments in Chapters Two, Three, and Four are sound. I hope that they are, and that the analyses I have given of the complaints of discriminatees in these chapters seem plausible and persuasive.

But there are also other advantages of my pluralist theory, ones that can be discussed in the abstract, and these are the ones I shall focus on here. Legal and philosophical scholars, and also courts and tribunals deciding cases of discrimination, have for a long time disagreed on a number of quite fundamental questions. One of these is the way in which claims of wrongful discrimination are comparative, and how we are to determine who the relevant comparator group is. Another source of disagreement is whether anti-discrimination law aims to protect individuals or groups. That practitioners and scholars might disagree on fundamental questions within a particular area of the law is unsurprising; but often some incremental progress is made in resolving them, or at least in laying out what is at stake in the dispute. And it is unclear that we have made even this

kind of progress in anti-discrimination law. I shall argue in what follows that my pluralist theory can help us both in understanding why there are such persistent, apparently unresolvable disagreements, and in resolving them.

5.7.b. It resolves “the comparative puzzle”

One persistent puzzle that my pluralist theory can help us to explain is whether and how judgments about wrongful discrimination are comparative.¹² It has seemed to many people that such judgments are necessarily and inherently comparative. That is, in order to make them, we must compare the discriminatee with certain others; and the judgment that someone has wrongfully been discriminated against says something about the discriminatee, relative to those others. But which others? And on what basis are we to compare them? Both scholars and courts have found it difficult to settle on a single answer to these questions.

Interestingly, the disagreements often take a certain form. Some scholars suggest that the relevant comparisons are between the discriminatee and the people who do *not* have the particular trait on the basis of which the discriminatee was treated differently, and who were therefore not disadvantaged in whatever way the discriminatee was disadvantaged.¹³ They then disagree amongst themselves about which “comparator group” is relevant here. For in any given case, many groups and sub-groups will not have the particular trait on the basis of which the discriminatee was treated differently, and each of them will likely have been treated somewhat differently and will stand in a slightly different relationship to the discriminatee. How do we know which group is the relevant one, with which to compare the discriminatee and to assess the kind of treatment that the discriminatee has received? Other scholars claim that the relevant comparison is between the discriminatee and a hypothetical version of this same person, who would have been treated differently under these same circumstances had she not had the trait in question.¹⁴ Those who take this latter view deny that judgments of wrongful discrimination are comparative across different actual people. Such judgments do, on their view, involve comparisons; but the comparison is between how someone was in fact treated and how this “same” person would have been treated if they had not had a particular trait. Now,

¹² See Deborah Hellman, “Two Concepts of Discrimination,” *Virginia Law Review* 102(4) (2016), pp. 895–952; Denise Réaume, “Dignity, Equality, and Comparison,” in Deborah Hellman and Sophia Moreau (eds.), *Philosophical Foundations of Discrimination Law* (Oxford: Oxford University Press, 2013), pp. 7–27; Suzanne B. Goldberg, “Discrimination by Comparison” *Yale Law Journal* 120(4) (2011), pp. 728–812 at p. 731; Timothy Macklem, *Beyond Comparison* (Cambridge: Cambridge University Press, 2003); and Sophia Moreau, “Equality Rights and the Relevance of Comparator Groups,” *Journal of Law & Equality* 5 (2006), pp. 81–96. For some legal judgments in which courts have explicitly discussed whether and in what sense claims of discrimination are comparative, see e.g. *Withler v. Canada (Attorney General)*, 2011 SCC 12 at para. 2; *Law v. Canada (Minister of Employment and Immigration)*, [1999] 1 S.C.R. 497 at para. 56; *Andrews v. Law Society of British Columbia*, [1989] 1 S.C.R. 143 at para. 8; and *Sweatt v. Painter*, 339 U.S. 629 at p. 631 (1950).

¹³ For instance, Kasper Lippert-Rasmussen, *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*, (Oxford: Oxford University Press, 2013).

¹⁴ For instance, John Gardner, “Discrimination: The Good, the Bad, and the Wrongful,” *Proceedings of the Aristotelian Society* 118(1) (2018), pp. 55–81.

whereas those who think that the relevant comparisons are with actual people disagree over who the relevant actual people are, those who think that the relevant comparisons are with hypothetical versions of the discriminatee disagree over which other traits or circumstances of the discriminatee we need to import into the imagined, hypothetical situation in which that person lacks the trait on the basis of which they were discriminated against. These disagreements are, in a sense, disagreements over the same thing. What both groups of scholars are trying to figure out is which circumstances are relevant to the wrongness of the discrimination in question. And of course, the bare idea that discrimination involves treating someone differently from others cannot help us answer this.

My theory can help us get through this impasse because it suggests that the kinds of comparisons that are dispositive, in determining whether wrongful discrimination has occurred, are not the kinds we make in our initial assessment *that* discrimination has occurred. They are, to use the distinction I drew in Chapter One, not the comparisons that we use to determine that there has been an instance of *wrongful differentiation*. So the whole debate that I sketched above, concerning whether it is actual comparisons or hypothetical comparisons that are relevant, is arguably itself irrelevant. That is because the relevant comparisons are those that we need to make in order to determine whether these discriminatees *have been treated as equals*. The judgment that someone has not been treated as an equal is a comparative judgment. But it does not involve a straightforward comparison, either with the group that received whatever immediate benefit the discriminatee was denied, or with whatever the discriminatee would have received under hypothetical circumstances if they had lacked a certain trait. Rather, in order to assess whether someone has not been treated as an equal in any of the three ways that I have considered, we need to make a number of different kinds of comparative judgments, depending on the particular way in which they apparently have not been treated as an equal. If we are concerned with social subordination, a number of comparisons will be relevant, including comparisons of the power and authority and consideration given to members of the allegedly superior social group with those given to the subordinated group, and comparisons between the ways in which certain practices accommodate and normalize the needs of a superior group while ignoring and marginalizing the needs of subordinate groups. If we are assessing whether a certain practice infringes someone's right to a certain deliberative freedom, we will need to ascertain whether she is forced to have a certain trait before her eyes, or to bear its costs, in circumstances where others do not have to bear the costs of other, similar traits—other races, for instance, or other religions. So this will involve comparisons of the costs accruing to bearers of different traits of the same type. It will also matter what other interests are at stake in a given case, as we saw in the case of the Muslim taxi-driver and the visually impaired passenger. And if what is at issue is whether a certain discriminatee has been denied a basic good, it will be relevant whether others in his or her society also enjoy this good. Though, as we saw in Chapter Four, a good can be necessary for a certain person or group to function as equals without being necessary for everyone. So although comparisons with other people's situations are

relevant, they are not dispositive; and their role is simply to help us assess whether, given what other people in that society do and think, this good is necessary for this person to function as an equal.

I have explained that, on my pluralist view, different comparisons are relevant in different cases, depending on which wrong is at issue. So my view does provide an answer to the question of which comparisons are relevant in cases of wrongful discrimination. It also helps us see past the impasse we find ourselves in when we think of the relevant comparisons in relation to the wrongful differentiation question, instead of in relation to the question of inequality.

Moreover, if my pluralist theory is correct, this would also help to explain why these disagreements about the relevance of different kinds of comparisons have persisted. According to my theory, a variety of different comparative judgments, with different actual and hypothetical comparators, is necessary, depending on the wrong that is at issue. But this implies that each of the different positions in these debates is, in a sense, right—though right for the wrong reasons. Sometimes, in order to assess whether and how a discriminatee has been wronged, we need to make actual comparisons between the discriminatee or social groups to which they belong, on the one hand, and other people, who are members of other social groups. But other times, such as when we assess certain opportunity costs of a practice on a particular claimant in order to determine whether their deliberative freedom has been lessened, we need to invoke hypothetical judgments about what would have happened to the claimant under different circumstances, if they had lacked this trait. We require different kinds of comparisons in the case of different wrongs. So there is some truth to each of these views—and this may explain why they have all persisted.

5.7.c. It resolves “the puzzle about groups and individuals”

A further longstanding puzzle about discrimination concerns whether it is primarily a personal wrong, akin to a tort, or primarily an injustice to particular social groups. Anti-discrimination laws have some features that suggest that the wrong to which they are responding is a personal one, and other features that suggest they are aiming to rectify a group-based injustice. For instance, many private sector anti-discrimination law regimes rely on individual claimants to instigate legal proceedings against alleged discriminators, and the claimant is, at least nominally, treated as though she is bringing a personal complaint akin to a tort. Moreover, many of the available remedies in private sector anti-discrimination law are personal ones: discriminatees can seek personal accommodations, reinstatement in their jobs or some equivalent monetary compensation, and special damages for personal injury to their dignity and self-respect. At the same time, however, there are other, much more transformative remedies available—remedies that are designed to fundamentally alter discriminatory practices rather than just to carve out a personal accommodation for the claimant. Remedies can include mandatory educational

programs for discriminators, and quotas and changes in hiring practices that are designed to help a much larger portion of the social group to which the claimant belongs.¹⁵ Remedies can also include orders to replace a particular wrongfully discriminatory practice with one that is inclusive of a social group that has certain needs, and that does not single out members of this group as different—for instance, abandoning “clean-shaven” rules and allowing employees to choose whether to shave, rather than selectively exempting African Americans who have PFB (a condition that makes shaving very painful, and which occurs mainly in people of African descent). These more transformative remedies have suggested to some that anti-discrimination law is addressing a group wrong or injustice, and not, or not only, a personal wrong done to the claimant.

My pluralist theory can allow us to see all of these structural features as reflecting a different aspect of the moral truth about discrimination. For discrimination, on this theory, sometimes involves personal wrongs, and sometimes involves group wrongs, and the same case can involve both kinds of wrongs. Though, as I mentioned earlier in Section 5, the group wrongs recognized in my account are not wrongs to some separate entity, a “group,” over and above its different members; rather they are wrongs done to each of the group’s members, by virtue of their membership in that group. As I explained in Section 5, the wrong of infringing someone’s deliberative freedom and the wrong of denying someone a basic good are both personal wrongs, which generate personal claims for redress. But the wrong of causally contributing to subordination is not a wrong that generates a claim on any one person’s part to any special kind of restitution, over and above the measures that need to be taken in order to ensure that the group does not face this discrimination in the future. And, as I also argued in Section 5, the related wrong of marking out a person or group as inferior or rendering them invisible may sometimes be a personal one, where special personal remedies are necessary in order to end the censure or the invisibility; but it may sometimes be a group wrong, in cases such as the Sketching the Line program, where an entire social group—in this case, visible minorities in Toronto—has been rendered invisible. No one member of the group has a claim to a special benefit, such as the benefit of seeing their own picture on the wall; but the practice of excluding them as a group from the subway posters needs to change, if they are to be treated as equals.

One way in which certain scholars have tried to resolve the apparent tension within anti-discrimination law between measures that seem to presuppose a personal wrong and measures that seem better suited to a group wrong is by suggesting that direct discrimination, when wrongful, is a personal wrong, whereas indirect discrimination, when wrongful, is a group wrong.¹⁶ The fact that a common way to prove indirect discrimination is to show that a certain group as a whole was disproportionately disadvantaged by a certain practice may seem to lend support to this suggestion. But the suggestion is a rather procrustean one; for direct discrimination sometimes seems to wrong a group in the ways

¹⁵ See e.g. *Canadian National Railway v. Canada (Canadian Human Rights Commission)*, [1987] 1 S.C.R. 1114.

¹⁶ See, for instance, Colleen Sheppard, *Inclusive Equality: The Relational Dimensions of Systemic Discrimination in Canada* (Montreal: McGill-Queen’s University Press, 2010).

that I have described, and indirect discrimination can sometimes give rise to personal claims on the part of group members. As we have seen in Chapters Two, Three, and Four, both forms of discrimination, direct and indirect, can fail to treat people as equals in either of the three ways I canvassed—that is, by subordinating them, by infringing their right to deliberative freedom, and by denying them a basic good. So both direct and indirect discrimination can, on my view, impose personal wrongs, and both can impose group wrongs.

I have now tried to show that my pluralist theory can help us understand the kinds of comparisons that judgments of wrongful discrimination require, and can help us understand why attempts to reduce these to a single sort of comparison will not succeed; and I have argued that my theory can also make sense of the fact that discrimination seems to involve both personal and group wrongs.

There is also, however, a third persistent puzzle that besets our thinking, and our legal practices, concerning discrimination. It concerns the relationship between direct and indirect discrimination. Many scholars have questioned whether indirect discrimination is indeed a form of discrimination at all.¹⁷ In their view, it is too different from direct discrimination to be an instance of the same kind of wrong. And some have questioned whether indirect discrimination is a wrong at all, suggesting that it is simply what we might call a misfortune, a harm that we certainly have good reason to try to rectify, but not something that wrongs people if it is allowed to persist. Even among those who treat both direct and indirect discrimination as wrongs, there is often an underlying suspicion that indirect discrimination is generally less serious from a moral standpoint than direct discrimination. Moreover, some legal regimes, such as the U.K., permit justification in the case of indirect discrimination, while they imply that no instance of genuine direct discrimination could be justified. I shall turn to these issues in the next chapter. There, I shall clarify what my theory implies about indirect discrimination. I shall argue that the differences between direct and indirect discrimination are less stark, and less important, than one might think, and that it is largely for pragmatic reasons of proof that they should remain a part of our laws. I shall argue that we can reasonably ask questions about justification—that is, about whether a particular instance of wronging someone is all things considered wrong—in all cases of discrimination, not just in cases of indirect discrimination. And I shall suggest that we need to separate questions about how far the agent is responsible for the costs of rectifying the wrong and how extensive the agent's obligations of rectification are from questions about culpability, or how far and in what sense the agent is to blame.

¹⁷ See e.g. Jeremy Waldron, "Indirect Discrimination" in Stephen Guest and Alan Milne (eds.), *Equality and Discrimination: Essays in Freedom and Justice* (Stuttgart: Franz Steiner, 1985) at p. 93; and Benjamin Eidelson, *Discrimination and Disrespect* (Oxford: Oxford University Press, 2015).

Chapter Six

Indirect Discrimination

I argued in the earlier chapters of this book that indirect discrimination can wrong people by failing to treat them as equals in the same ways that direct discrimination can: by subordinating them, by infringing their right to deliberative freedom, and by leaving them without access to a basic good. It may seem, therefore, that there is no need for a separate chapter on indirect discrimination. However, there are still a number of difficult questions concerning indirect discrimination that I have not yet addressed.

One of these questions is whether there is a morally salient difference between direct and indirect discrimination at the stage of justification, when we ask whether a discriminatory practice that wrongs someone might nevertheless be justified, all things considered. Up until this point in the book, I have only considered the question of whether and why discrimination can wrong someone. But what ought we to say about the further question of all things considered justification? Is there, as some legal jurisdictions posit, a difference in the way in which we ought to treat direct and indirect discrimination at the stage of justification? I shall argue that on my account, there is no reason to treat the two forms of discrimination differently at the stage of all things considered justification. It is true that different justificatory factors are relevant in different cases. But these differences do not line up neatly with the distinction between direct and indirect discrimination.

A further set of difficult questions raised by indirect discrimination concerns the responsibility and culpability of the discriminator. Some of these are questions about what I shall call “responsibility for cost”—that is, responsibility for the cost of altering one’s practices so as to eliminate wrongful indirect discrimination. This is a morally thin sense of responsibility because it need not imply culpability or blameworthiness: it is just the judgment that it is fair, under the circumstances, to make this discriminator bear the cost of eliminating the discrimination. The idea that we ought to hold discriminators responsible for the cost of eliminating discrimination may seem especially problematic in cases of indirect discrimination. This is because in these cases, the disproportionate disadvantages accruing to a particular group are the result of many different factors operating together—not just the wrongfully discriminatory practice, but also the practices of other institutions, the actions of other individuals, general social conventions, tacit assumptions, and in some cases also our natural environment. Moreover, the disadvantages accruing to the relevant social group are often much farther down the causal chain—to borrow a term from tort law, they are more “remote” from the discriminator than are the harmful effects of direct discrimination. And so it can seem unfair to hold the discriminator responsible for indirect discrimination, even in the morally thin sense of responsibility for cost.

Other important questions about indirect discrimination concern responsibility in a morally thicker sense, which I shall call “culpability.” Even if we accept that direct and indirect discrimination wrong people in the same way, by failing to treat them as equals, and even if we accept that discriminators ought to be held responsible for the cost of eliminating indirect discrimination, we may nevertheless feel that those who indirectly discriminate are not culpable, and ought not to be subjected to the same kind of moral criticisms to which we subject those who discriminate directly. I shall argue, however, that this feeling is largely misguided. Although there are certain heinous cases of direct discrimination in which agents are motivated by hate or prejudice, nevertheless, in many cases of direct and indirect discrimination, we can see the culpability of agents as stemming from the same source: their negligence. And I shall explain what this negligence seems to me to consist in.

But before I address these questions, it will be helpful to summarize what conclusions I drew about indirect discrimination in the earlier chapters of this book, when I looked in detail at each of the ways in which discrimination wrongs people by failing to treat them as equals.

6.1 What My Theory Implies About Indirect Discrimination

I have argued that, just like direct discrimination, indirect discrimination can wrong people by failing to treat them as the equal of others. I looked extensively in Chapter Two at the way in which indirectly discriminatory practices can subordinate certain social groups to others, both by causally contributing to the four conditions that characterize persistent and unfair subordination across a number of social contexts, and by rendering the needs and situations of certain groups invisible and thereby marking them as inferior to others. I then argued in Chapter Three that there are a number of indirectly discriminatory practices that infringe certain people’s right to a particular deliberative freedom and thereby fail to treat them as equals. For instance, we considered tests for promotion in the workplace that disproportionately disadvantage certain racial minorities, leaving these people always with their race before their eyes and unfairly bearing the costs of lacking whatever experiences or background are assumed by the test. Finally, I tried to show in Chapter Four that indirectly discriminatory practices can leave people without access to basic goods and thereby fail to treat them as equals, the way the Canadian government’s inadequate provision of clean water to many indigenous communities leaves them without access to the water they need both for their health and for the symbolic rituals that are essential to their cultures and identities.

Indirect discrimination, then, can wrong people in the very same way that direct discrimination does. There is, to be sure, a difference in the *mechanism* through which wrongful subordination occurs, depending on whether the discrimination is direct or indirect. Wrongful direct discrimination, as we saw, marks people as inferior by explicitly

naming a trait that is a prohibited ground of discrimination and thereby branding or stigmatizing the person or group that possesses it. By contrast, as we explored in detail in Chapter Two, indirect discrimination usually works to subordinate people by ignoring their needs and thereby rendering them invisible. But this is a mere difference in the mechanism through which the wrong of failing to treat others as equals comes about, rather than a difference in the kind or degree of seriousness of the resulting wrong. And in cases involving the other two wrongs that I have discussed—that is, where people’s right to deliberative freedom is infringed, or where they are left without access to a basic good—it does not seem to make a significant difference whether this occurs through direct discrimination or through indirect discrimination. So although it was useful, and indeed necessary, to start our investigation of discrimination with this legal distinction, so as to ensure that our theory of wrongful discrimination was capacious enough to capture all of what we consider to be “discrimination,” we have now reached a point where we can question the theoretical usefulness of this distinction. Or rather, we can recognize that it may be helpful in identifying the mechanisms through which some discriminatory practices fail to treat people as equals, while nevertheless questioning whether it marks a difference in the kind, or degree of seriousness, of the moral wrong at issue.

Of course, we can question this while still recognizing the history and evolutionary importance of the distinction between direct and indirect discrimination. As a part of our anti-discrimination laws, this distinction reflects the evolution of our public views about discrimination.¹ And interestingly, the initial evolution of these views after World War II was strikingly similar across a number of different countries such as the United States, Canada, and the U.K.² Discrimination was, in these countries, originally regarded by the law as a form of wrongful treatment that centrally involved some sort of offensive and unwarranted motive—such as hatred of, or prejudice against, a particular group of people, based on some trait of theirs such as their race. Liability was gradually extended to acts that lacked this kind of illicit motive but nevertheless involved *intentional* treatment of members of one group differently from that of others; and from here, in jurisdictions such as Canada and the U.K., the intent requirement was formally dropped and direct

¹ For helpful legal overviews of the development of anti-discrimination law in certain countries, see Denise Réaume, “Harm and Fault in Discrimination Law” *Theoretical Inquiries in Law* 2(1) (2001), pp. 349–385 (detailing the growth of discrimination law in Canada); Alfred W. Blumrosen, “Strangers in Paradise: Griggs v. Duke Power Co. and the Concept of Employment Discrimination” (1972) *Michigan Law Review* 71(1) at pp. 59–110 (detailing American developments); and Christopher McCrudden, “Changing Notions of Discrimination,” in Stephen Guest and Alan Milne (eds.), *Equality and Discrimination: Essays in Freedom and Justice* (Stuttgart: Franz Steiner, 1985) (detailing developments in the U.K.).

² Although of course now indirect discrimination is treated differently by different countries. Some, such as Canada, embracing it as no different from direct discrimination (see *British Columbia (Public Service Employee Relations Commission) v. B.C.G.E.U.*, [1999] 3 S.C.R. 3, 1999 – known as the “Meiorin case,” after the claimant, Tawney Meiorin). Others, such as the U.K. recognize indirect discrimination but allow for justifications in the case of indirect discrimination that are not available for direct discrimination: see *Equality Act 2010*. Still others, such as the United States, limit liability for indirect discrimination and regard it with suspicion: see *Washington v. Davis*, 426 U.S. 229 (1976); and *Wards Cove Packing Co. v. Atonio*, 490 U.S. 642 (1989).

discrimination came to involve simply an explicit or facial distinction.³ At the same time, indirect discrimination was recognized as a form of the same wrong, in which a particular group was neither intentionally treated differently nor even explicitly singled out for different treatment, but nevertheless disproportionately disadvantaged.

On the theory of wrongful discrimination that I have proposed in this book, both direct and indirect discrimination are wrongful when they fail to treat people as equals. And whether they fail to treat people as equals does not depend on the motive or the intent of the discriminator. It depends on the kinds of considerations we examined in the earlier chapters of the book—those relevant to subordination, deliberative freedom, and basic goods. However, we can recognize this and yet still leave room for the possibility that victims of discrimination are *also* wronged in a *further* and different way when the discriminator acts from certain motives, such as malice or prejudice. All that my theory denies is that having such motives is a necessary part of wronging someone by failing to treat that person as an equal. It seems quite plausible to suggest that in the most heinous cases of discrimination, such as the Jim Crow laws, or violence directed at Muslims out of hatred, there is also an additional wrong done to the victim. There are a number of ways in which we might characterize this further wrong: the wrong of acting out of hatred toward another person, and with enjoyment of the harm that comes to them; the wrong of deliberately insulting another person; the wrong of deliberately assigning another person not just a less than equal status, but a sub-human status. What is important for my purposes is that my account is quite consistent with our recognizing that such a further wrong is present in some cases of wrongful discrimination. What my account insists upon, however, is that not all cases of discrimination—and not all or even most cases of *direct* discrimination—involve this further wrong. And it is not a necessary component of the wrong of failing to treat another person as an equal.

The view that direct and indirect discrimination both wrong people by failing to treat them as equals and that these particular wrongs do not depend on the agent's having any particular motive or intent is arguably the view that underlies current Canadian laws on discrimination. Canada no longer recognizes a distinction between direct and indirect discrimination: it applies a single test to any form of discrimination, though the relevant test is different in private sector anti-discrimination law and in the constitutional equality rights provisions under the *Charter*.⁴ When the Supreme Court of Canada harmonized the approaches to direct and indirect discrimination in private sector anti-discrimination law, it noted a number of pragmatic reasons for not treating the two forms of discrimination differently: for instance, that as long as the law treats indirect discrimination as easier to justify, employers and other agents of discrimination may try to re-frame policies that they know are directly discriminatory using neutral language, in the hope that they can bring about exactly the same effect through different means and thereby escape legal sanctions;

³ See Réaume "Harm and Fault in Discrimination Law," *supra* note 1; and Blumrosen, "Strangers in Paradise: *Griggs v. Duke Power Co.* and the Concept of Employment Discrimination," *supra* note 1.

⁴ See the *Meiorin* case, *supra* note 2.

and that it may end up inadvertently legitimizing systemic discrimination.⁵ I hope that, in this book, I have given some theoretical reasons that complement these pragmatic ones. I have tried to show that the two should not be treated differently—at least at the stage of determining when they wrong people—because in fact they involve the same kind of wrong. Both fail to treat some people as equals.

I have argued that there is no difference between the wrongs done by direct and indirect discrimination—that is, no difference in their kind, and hence, no difference in their seriousness or urgency. But recently, some scholars have gone further than this. They have argued that in fact all indirect discrimination just *is* direct discrimination—direct discrimination on the basis of a *different* ground. So, for instance, John Gardner has argued that indirect sex discrimination is just discrimination “against people with some *other* property (people of less height, or with less availability for evening work, or having less upper body strength, or with a record of lower earnings), where that other property is statistically correlated with sex.”⁶ And so, he argues, legislation such as the U.K. *Sex Discrimination Act* “does not regulate only sex discrimination. It also regulates, in a derivative and relatively circumscribed way, height discrimination, strength discrimination and so on.”⁷ There are two claims here, and I think we need to question both of them. The first is that all indirect discrimination is reducible to direct discrimination. The second—and it is an implicit, rather than an explicit claim—is that we can grasp what is morally problematic about indirect discrimination, and accurately form a picture of the particular social group that has not been treated as an equal, under that new description. Both of these claims are, in my view, problematic. First, although it is true that some practices that indirectly discriminate distinguish between people on the basis of some other property, other indirectly discriminatory practices don’t explicitly employ *any* criteria in order to distinguish between different groups. Their aim is not to distinguish between different groups, by applying a certain criterion to those groups. Rather, they simply *ignore* the needs of a particular group, while lavishing resources on particular causes that happen to satisfy the needs of other groups. This is true, for instance, of the governmental practices we considered in Chapter Four, which have resulted in the indigenous water crisis. These are not practices that deliberately or explicitly assign indigenous communities fewer resources on the basis of some *other* criterion, such as remoteness. Rather, the needs of these indigenous communities have simply been ignored, as funding has been directed at other problems that happen to be the problems of non-indigenous communities. It seems to me that it is much more helpful to call this what it is—a set of practices that disproportionately disadvantages a certain group because of a trait that is a prohibited ground of discrimination—than it is to try to re-describe all of these many practices as attempts to distinguish between people on some other basis. And the indigenous water crisis is not an unusual type of case. Many of the cases of indirect discrimination that leave

⁵ *Ibid.*

⁶ John Gardner, “Discrimination: The Good, the Bad, and the Wrongful,” *Proceedings of the Aristotelian Society*, 118(1) (2018), pp. 55–81 at p. 60.

⁷ Gardner, “Discrimination: The Good, the Bad, and the Wrongful,” *ibid.*

people without access to basic goods will also be difficult to re-describe as an attempt to classify people on the basis of some other criterion. So it is not clear that Gardner's re-description can capture all instances of indirect discrimination.

Second, however, even in the case of those indirectly discriminatory practices that *can* be re-described as direct discrimination on the basis of some other property, I worry that we will lose both our ability to see them as wrongful and our ability to pick out the particular social group that is in fact wronged if we see them as direct discrimination on the basis of some other property.⁸ *Pace* Gardner, the Sex Discrimination Act does *not* regulate height discrimination and strength *qua* height and strength discrimination. It regulates them only insofar as they constitute sex discrimination. And what makes them wrongful is not their impact on persons because of their height or strength or lesser availability to work in the evening, or record of lower earnings. What makes them wrongful is their impact on people because of their sex.

Gardner might reply that these remarks are perfectly consistent with his view. He might argue that yes, in order to understand what makes these forms of direct height discrimination and direct strength discrimination *wrongful*, we have to make reference to sex: they are wrong because they disproportionately disadvantage women. But they are nevertheless forms of direct discrimination. But if this is right, then I cannot understand what is to be gained by saying that these are forms of direct discrimination. For in order to capture what is wrongful about them, we will have to say that they are a special kind of direct discrimination. Whereas in ordinary cases of direct discrimination, the discrimination is wrongful because of the particular property that is explicitly used to differentiate some people from others, in cases of indirect discrimination, the discrimination is only wrongful in relation to some other property, which is the prohibited ground, and only wrongful insofar as the criterion that is explicitly used happens to track that other ground. And this seems to me to amount, in effect, to an admission that indirect discrimination is *not* just like wrongful direct discrimination.

Moreover, I think it is important to note that we *learn* something about indirect discrimination when we describe it as "indirect discrimination on the basis of sex" that we do not learn when we re-describe it as "direct discrimination against people of less height" or "direct discrimination against people with less availability for evening work." Even if we could come up with a perfectly accurate set of conditions that pick out all and only the particular people who have been disadvantaged by a certain indirectly discriminatory policy—it would be a conjunction, I think, of many different conditions, such as "discrimination against shorter people, with less lung power, with less evening availability, who did not present as sufficiently aggressive or assertive at their interview, who might shortly need to take a parental leave" and so on—I do not think we would be able to understand why and how all of these conditions contributed to our treating this particular group as less than equals in society unless we thought of their treatment in light of the

⁸ I am very grateful to Andy Yu for suggesting this objection.

prohibited ground of sex. When we conceptualize wrongful discrimination, after all, our aim is not only to pick out the *right* group of people in each case, the group that has not been treated as an equal to others. Our aim is also to understand *why* this has happened to them. And in order to understand why wrongful discrimination against women occurs, it is crucial that we think of it as discrimination on the basis of their *sex*. When we think of how and why women are subordinated, or what it means for them to be denied deliberative freedom—what trait it is that they must always have before their eyes, as they deliberate—or why, in certain circumstances, they are denied basic goods, we will only be able to understand women’s situations, and to see the relevant practices as failures to treat women as equals, if we think of them as failures to treat women as equals on the basis of their sex.

For all of these reasons, I think it is helpful to retain the distinction between direct and indirect discrimination. But we can do this while still maintaining that when they are wrongful, they are wrongful for the same reason: because they fail to treat some people as equals.

6.2 *Is Indirect Discrimination Easier to Justify?*

As I have mentioned, in some legal jurisdictions, direct and indirect discrimination are treated differently at the stage of justification. Wrongful direct discrimination is simply prohibited; whereas wrongful indirect discrimination is treated as justifiable under certain circumstances. And this is sometimes understood as implying, from a moral standpoint, that direct discrimination can never be justified all things considered, whereas indirect discrimination can sometimes be justified, all things considered.

Such an approach to the justification of direct and indirect discrimination would be sound if they involved two different kinds of wrongs, or if direct discrimination always involved a deeply troubling motive, such as the kind of hatred or prejudice that I discussed briefly above. And indeed, this legal approach to justification may, historically, hark back to the time when we did think of direct discrimination as essentially motivated by hostile attitudes. But many countries now treat direct discrimination as requiring no such attitudes, and this is the view I have defended in this book.⁹ So, for instance, if a sports club permits blacks and Latinos to access its facilities only at different times of the day in order to reduce racial tensions, this amounts to direct discrimination on the grounds of race—even if, far from being motivated by racial prejudice or hatred, it is driven simply by an innocent (though perhaps misplaced) desire to use the club’s opening hours to try to reduce racial tensions. Since direct discrimination can occur, as it does in this case, without animus, there does not seem to be any reason for thinking that, in principle, it could not be all things considered justified by whatever considerations might justify certain cases of wrongful indirect discrimination. I am not suggesting here that the direct discrimination in this case is justified—only that it seems to be the kind of practice that could in principle be

⁹ See Chapter 1, Section 3.

justified. Moreover, if the same wrong—the wrong of failing to treat people as an equal—is done by direct and indirect discrimination, it seems reasonable to suppose that similar considerations will be relevant to the question of when wrongful direct and indirect discrimination are nevertheless all things considered justified.

The only difference between direct and indirect discrimination, as I have understood them, is that wrongful direct discrimination explicitly singles out a certain group or person using a prohibited ground of discrimination (or some trait that is closely connected to such a ground), whereas practices that discriminate indirectly do not: the latter are apparently neutral, but nevertheless have a disproportionately disadvantageous effect on a group that shares a trait that is a prohibited ground of discrimination. So the two forms of discrimination both wrong people by failing to treat them as equals; but one explicitly uses prohibited grounds of discrimination, whereas the other does not. It is not clear why this difference should give rise to a difference in the kinds of factors that would *justify* wrongful direct and wrongful indirect discrimination.

My theory, then, implies that there is no difference, as a group, between the justificatory factors relevant in cases of indirect discrimination and the justificatory factors relevant in cases of direct discrimination. This, too, is the approach that has been adopted by Canada. Both in its interpretation of private sector anti-discrimination laws and in its interpretation of constitutional equality rights in the *Charter*, Canada applies the same test to cases of direct and indirect discrimination, to assess whether they can nevertheless be justified.¹⁰ In the private sector, for instance, when assessing whether employers or providers of goods and services can be justified in continuing to discriminate, Canadian tribunals ask whether the discriminatory practice was adopted in good faith for a purpose that is rationally connected to the function being performed, and whether it is “reasonably necessary” in the strong sense that there is no alternative practice that the discriminator could adopt that would accommodate the claimant’s needs without imposing “undue hardship” on the discriminator.¹¹ This test is applied both to direct and to indirect discrimination. As the reference to “undue hardship” on the side of the discriminator suggests, Canada allows that the burden on the discriminator is relevant in determining whether a practice that fails to treat certain people as equals is nevertheless justified, all things considered. But it is relevant *both* in cases of direct discrimination and in cases of indirect discrimination.

I have argued that there is no reason to think that different factors are relevant to the justification of direct and indirect discrimination. Which particular factors are relevant in any given case is a question I shall pursue in more detail in Chapter Seven. For, as I shall argue there, it seems to depend very much on the particular role that is occupied by the

¹⁰ Though, interestingly, one test is applied to both forms of discrimination in the private sector context, and another, to both forms of discrimination under the Constitution. I shall discuss these differences in Chapter 7, as they stem in part from the different demands of inquiries into discrimination by private agents and inquiries into discrimination by the state.

¹¹ See the *Meiorin* case, *supra* note 2 at para. 54.

discriminator, relative to the discriminatee, and on the constitutive responsibilities of that role. “Undue hardship” is the language that Canadian legislation uses in connection with those who hold themselves out to the public as providers of employment, or goods, or services. These are individuals with their own projects to advance and their own lives to live; and the law recognizes that certain costs would make it impossible for them to pursue these projects, or impossible for them to pursue these projects safely. By contrast, government actors do not pursue private projects and do not have the same interest in being allowed to pursue them autonomously. They, however, have other obligations—obligations both to promote the equality of other groups and to respect other rights of these groups. And so the relevant tests for justifications of violations of constitutional equality rights by the government in Canada take account of these other obligations. I shall look at these and other differences between relevant justificatory factors in Chapter Seven. What is important for our purposes now is to note that these differences depend on the status and responsibilities of the agent, and not on whether the discrimination is direct or indirect.

6.3 Responsibility for Cost

Even if you accept that direct and indirect discrimination are wrongful for the same reasons, and even if you accept that both can be justified by the same sorts of considerations, you might still question whether the agents of direct and indirect discrimination are responsible to the same degree.

It is helpful to distinguish two sets of questions here. There are, firstly, questions concerning how far a particular agent can be held *responsible for the cost* of eliminating a particular discriminatory practice and replacing it with a practice that treats the individuals or the group in question as equals. It can be fair to hold someone responsible for the cost of a particular alteration even if they are not culpable in some weighty moral sense—that is, even if we are not justified in blaming them, or in acting towards them in a way that expresses disapproval of them as a person. This morally thin idea of responsibility is familiar to us from tort law. Judges adjudicating cases of negligence commonly distinguish between the judgment that a particular defendant is blameworthy and the judgment that he/she is liable.¹² A judgment of liability, they repeat, is simply the judgment that, of all of the people who are in some way causally connected with a particular loss, it is just and fair to make *this* person cover the cost of that loss. This is what I mean by a morally thin notion of responsibility. It is the idea of responsibility for a certain cost—in our case, the cost of eliminating the discriminatory practice and replacing it with something that treats these people as equals. And although a person can be both responsible for cost and blameworthy, they can also be responsible for a cost even when, like many of those who are found liable in tort law, they are not to blame.

¹² See e.g. *Rylands v. Fletcher*, [1868] UKHL 1; and *Roberts v. Ramsbottom*, [1980] 1 All E.R. 7 (Q.B. 1979).

I want to set aside for now questions about responsibility in any thicker moral sense and focus on questions about responsibility for cost. It may seem that responsibility for cost ought to depend at least in part on how *close* one's actions are on the causal chain to the disadvantageous effects on the discriminatee, and on how many *other* factors are also causes of the discriminatee's situation. And so it may seem that there are some morally relevant differences here between direct and indirect discrimination. It is often assumed that these two forms of discrimination differ in what we might call the directness or the closeness of the effects on the discriminatee to the discriminatory practice itself. Indeed, this assumption is likely what explains the choice of the terms "direct" and "indirect" to describe these two forms of discrimination. The thought seems to be that in cases of indirect discrimination, the discriminatee is disadvantaged, not directly or immediately by the discriminatory practice, but only through a much more complicated causal chain, involving other institutional practices, other agents' actions, and other features of our shared social environment. By contrast, in cases of direct discrimination, it is the discriminatory practice itself that directly or immediately disadvantages the discriminatee.

It may be true that in many cases of indirect discrimination, the disadvantage to the discriminatee occurs by means of a complex causal chain. But I doubt whether the causal chains at issue in direct discrimination are any less complex, or the disadvantageous effects, that much more remote from the policies, than they are in the case of indirect discrimination. Think of the sports club policy I discussed above that constitutes direct discrimination. If restricting blacks and Latinos to different opening hours amounts to wrongful discrimination, this is only because of a complex set of social circumstances, social conventions, a history of racial segregation and so on: the failure to treat these groups as the equals of others is hardly an immediate or "direct" effect of the policy.

More importantly, whether the disadvantage to the claimant is direct or indirect cannot be relevant to our thinking about responsibility for costs on my own theory, because what makes discrimination wrongful according to this theory is not the disadvantages suffered by discriminatees but the failure to treat them as equals. So it is irrelevant how closely connected their disadvantage seems to be to the discriminatory practice, or how many other items appear in the causal chain connecting them to the discriminatory practice. And the fact that a particular practice fails to treat people as equals will always depend, not just on the practice itself, but on the surrounding social context. It will always depend on the history of certain social groups and their interactions, on social conventions and expectations, on the rules of other institutions, and on the built and natural environments in which the discrimination occurs. This will be true regardless of whether the discrimination is direct or indirect.

Nor does it seem correct to suppose that a person is less responsible for a certain outcome simply because it occurs farther along a certain causal chain, or because there are a greater number of relevant background factors. If I negligently drop a broken bottle in the sand and it is tossed about, carried here and there, and finally washes up on another beach many years later and injures a child, I am morally culpable for this injury in spite of

the many factors and the many years intervening between my act and the actual injury—because this is precisely the kind of injury that makes it morally problematic to drop broken bottles on a beach in the first place. It is true that if the causal chain is very long and mediated by many other people's acts, then we do not generally think a particular agent is responsible for the outcome unless he failed to do something that he had a *duty* to do. But on my account, both agents of direct and agents of indirect discrimination have a duty to treat others as equals.

But at this point, one might object that I have misconstrued the nature of the problem here. Perhaps the problem isn't that all agents seem less responsible for indirect discrimination than for direct discrimination. Perhaps the concern is that it is simply unfair to hold *private* agent—that is, individuals, corporations, any non-government entity—responsible for what are really the cumulative effects of many different social institutions, interacting in complex ways to disadvantage certain social groups. Perhaps there are, as I have said, some cases of direct discrimination that, like indirect discrimination, involve complex causal chains. Well then, our objector might say, these cases of direct discrimination, too, are ones in which it is not fair to hold private discriminators responsible. Unlike the government, private agents do not have a constitutive responsibility to create the conditions under which people can relate to each other as equals. So it is *never* fair to hold them responsible for the cost of eliminating what is essentially a shared social problem.

This may sound attractive from a theoretical standpoint. But I think it is important to remember that in practice, somebody will always bear the costs of wrongfully discriminatory practices, regardless of what we do. If we alter these practices—moving our organization's meeting times from 5pm to lunchtime so that more women can attend, giving employees a more flexible work schedule so that they can take breaks as required by their religion or their disability, changing hiring and promotions practices so that more racial minorities are given an equal chance to contribute to the organization—then it is true that the cost will be borne by private parties, such as employers. But if we do nothing, if we do not require employers to alter such practices, then there is also a cost. There is a cost to the people who are treated as inferiors: they do not enjoy a status equal to that of others. It may be easy for us to overlook this cost—after all, as we saw in Chapter Two, in those cases of wrongful discrimination that involve subordination, the group that is treated as an inferior is often rendered invisible. So the burdens they must carry go unnoticed. But of course this does not mean they do not exist. And there are also shared social costs to allowing wrongful discrimination to continue, costs that *all* of us will bear, if we live, as we do now, in a society in which some people are deemed inferior across a number of different social contexts and over a long period of time. We will all be poorer. We will lack the ideas and the perspectives that members of these groups might have shared with us. We will have greater misunderstandings and greater mistrust between social groups—the kind of mistrust that ferments when certain social groups are persistently excluded from important social institutions and from positions of power and privilege. And perhaps most

sadly, a certain kind of life together will not be possible, a life in which we all relate to each other as equals.

So there are costs either way—costs to eliminating discrimination, and costs to allowing it to continue. Because *somebody* will always bear the cost of wrongfully discriminatory practices, and because the alternative to holding discriminators responsible is to let the costs lie where they fall, on the shoulders of the equally innocent discriminatees, it seems less plausible to me to claim that it is unfair to hold the discriminator responsible for the cost of eliminating such practices. Moreover, I think the costs of eliminating discriminatory practices are often less than we imagine. When we think of responsibility for cost, we may have in mind prohibitive costs, and costs for physical objects, such as the costs employers incur in accommodating people with disabilities: retrofitting a historic building with elevators, or purchasing a braille printer. But in very many cases, the costs of eliminating discrimination are better thought of as an investment of time and effort and dialogue with members of subordinated social groups, thinking creatively together with them about how to redefine certain roles or alter certain rules so that they are more inclusive. And it can be as simple as not requiring fire-fighters to meet male aerobic standards, or purchasing a chair for a person with a disability so that they do not have to stand at the cash register.

One might also argue that paying the costs of eliminating a particular discriminatory policy, when you are the person or the organization who is best positioned to eliminate it, is just one of the responsibilities that one takes on when one lives together with others within a democratic society. Democratic societies, as Rawls noted, are “systems of social cooperation”—and, we might explicitly add, systems of social cooperation between people conceived of as equals.¹³ We all share in the benefits of this system, both the economic benefits and also the relational benefits. But the relational benefits are only possible if we all do what we can to ensure that others are treated as equals. And the kinds of private agents whom our anti-discrimination laws place under a duty to eliminate the wrongful discrimination to which they contribute—employers, providers of goods and services, and accommodation—are particularly well positioned to make a difference in eliminating wrongful discrimination.

There are, of course, some cases where the cost of altering a particular discriminatory practice is so great that it would threaten the survival of an otherwise beneficial organization or social institution: a small bookstore, for instance, in a historic building, that simply does not have the funding to install an elevator. This may be one of the considerations that could legitimately justify an agent in continuing with a wrongfully discriminatory practice. Our conclusion in such a case might be that the employer wrongs those whom his practices do not treat as equals—in this case, those in wheelchairs—but

¹³ See, for instance, Elizabeth Anderson, “What is the Point of Equality?”, *Ethics* 109(2) (1999), pp. 287–337; Kasper Lippert-Rasmussen, *Relational Egalitarianism: Living as Equals* (Cambridge: Cambridge University Press, 2018).

because of his difficult circumstances, he is under no all things considered obligation to provide the elevator. But we may also want to say that governments have an obligation to put into place a compensation scheme that might enable at least some of the private agents in these situations to procure additional funding and thereby eliminate the discriminatory practice. In other words, even in cases where private agents are justified in not changing their practices for financial reasons, there may be an overarching governmental obligation to provide funds to address the problem, or funds to address those cases that affect the largest numbers of people in the most serious ways.

6.4 *Culpability in Direct and Indirect Discrimination*

But what about responsibility in the morally thicker sense of “culpability”? We often assume that those who engage in wrongful direct discrimination are especially blameworthy, whereas those who engage in wrongful indirect discrimination are often seen as innocent. Sheila Day and Gwen Brodsky put this point particularly vividly when they noted that direct discrimination is frequently perceived as “loathsome” and “morally repugnant,” whereas indirect discrimination is commonly held to be “innocent, unwitting, accidental, and consequently not morally repugnant.”¹⁴ I want to argue in this section of the chapter that many agents of indirect discrimination are no less culpable than are agents of direct discrimination. I shall proceed by looking at some ordinary cases of indirect discrimination and some parallel cases of direct discrimination, and by analyzing in what respects the agents seem culpable. This will lead us to the arguments of the final section of this chapter, in which I try to show that it is most helpful to think of agents of direct and indirect discrimination as culpable for their *negligence*.

Acts of indirect discrimination frequently occur as part of a whole set of policies, practices, and assumptions that together form what is called “systemic discrimination.” So we can start by considering one common instance of systemic discrimination: the culture of sexual harassment within the military. Recently, an External Review of the Canadian Armed Forces revealed an environment in which harassment and assault of women and LGBTQ members have become so commonplace that they are regarded as normal and natural.¹⁵ Some of the worst aspects of this culture involve direct discrimination: frequent use of sexualized language and sexual jokes targeting women’s body parts; comments and posters proclaiming that a woman enters the army “to find a man, leave a man, or become a man”; and sexual assaults and date rape of younger women by senior ranking officers. But these acts of direct discrimination have been allowed to continue in large part because they are sustained by a whole set of policies that are indirectly discriminatory and that work to

¹⁴ Shelagh Day and Gwen Brodsky, “The Duty to Accommodate: Who Will Benefit?”, *Canadian Bar Review* 75 (1996), pp. 433–473 at p. 457. These authors went on to argue that this view was misguided—but for reasons different from the ones that I want to foreground in what follows.

¹⁵ *External Review into Sexual Misconduct and Sexual Harassment in the Canadian Armed Forces*, conducted by The Honourable Marie Deschamps, March 27, 2015, available at: <http://www.forces.gc.ca>.

silence women and LGBTQ members. These include: a practice of ostracizing recruits who speak up about any kind of problem; a complaints process that has no provision for confidentiality; a policy of documenting only serious physical injuries and no “lesser” injuries; and a training program that does not focus on appropriate behaviour towards others. These policies amount to indirect discrimination because, even though they are neutral on their face, they have a disproportionate impact on women and LGBTQ members in a culture in which these people are the most frequent targets of sexual abuse.

If we look at these cases of indirect discrimination within their context—that is, within the culture of sexual harassment that exists in the military, in which everyone is aware that such acts are occurring even if they think this is normal and natural—it is difficult to view the members of the Armed Forces as less than seriously culpable. They have failed to do certain crucial things to stop the subordination of women and LGBTQ members, such as develop a proper training program, encourage victims of abuse to come forward, cultivate a culture of openness and honesty, and implement a confidential complaints process. And they have failed to do these things, and failed to see the importance of doing them, presumably because they have failed to see women and LGBTQ members as equals, as beings whose interests are just as important, and deserve just as much weight in their deliberations, as the interests of straight men.

Indeed, when we look closely at this example, the moral failings involved in the indirectly discriminatory policies do not seem so very different from the moral failings involved in the acts of direct discrimination—the sexual jokes, the assaults, and the harassment. Those who engage in such acts of direct discrimination are likely either trying to put victims “in their place” because they think of them as inferior and want their victims to know it, or they are just “having a bit of fun” on the assumption that having fun at the expense of these groups is perfectly acceptable because women and LGBTQ members aren’t “real” soldiers anyway. Either way, these agents, too, are failing to take the harms suffered by these people as a reason to act differently, and they are failing to treat these people’s aims and ambitions as seriously as they treat their own. So both those engaged in direct discrimination and those engaged in indirect discrimination in this case are failing to see others as equals. They are failing to give others the moral significance that they should be given in their deliberations. Of course, those who engage in direct discrimination must, in addition, know that they are directly causing physical or emotional harm to the people whom they assault or harass. But indirect discrimination also harms these groups, and the members of the Armed Forces who continue to support the indirectly discriminatory policies must be aware that they are contributing to the harm that is suffered by these groups. They are just contributing to it in a less direct way, with the causal chain being somewhat longer and mediated by other factors—such as other policies, and other people’s words and actions. I argued earlier that this fact should not make much of a difference to our judgments about responsibility in the sense of “responsibility for costs.” For the same reasons, one might doubt whether it should make much difference to our judgments about culpability.

One might object that this is an unhelpful type of example to use when trying to assess the moral status of indirect discrimination, because the indirect discrimination in this case is so closely bound up with direct discrimination. The policies that amount to indirect discrimination in this example do so only because they help to condone and so to perpetuate direct discrimination against these same groups. So it might seem that in this type of case, if the agents of indirect discrimination seem as culpable as the agents of direct discrimination, this is only because the practices in question are so deeply bound up with directly discriminatory practices.

What we require, then, is an example of indirect discrimination by agents who are not themselves engaged in direct discrimination, and where the indirectly discriminatory policy works to impose disadvantage by some means other than encouraging or permitting agents to engage in acts of direct discrimination against these groups. So consider the physical fitness tests used for hiring in occupations that require considerable strength and stamina—such as fire fighters, forest fire fighters, or security guards. Some of the fitness tests used for these occupations have faced legal challenges in the U.S. and Canada, on the grounds that they hold everyone to standards that were originally based on male aerobic capacity and male fitness targets and are therefore much harder for most women to succeed at.¹⁶ The tests do not amount to direct discrimination: there is no reference to gender in the application of the test, the tests are open to both men and women, and some women do pass them. However, as a group, women find it disproportionately harder to pass the tests than men, and it seems that this is because of their physique and aerobic capacities as women.

A second, and similarly structured example of “independent indirect discrimination” involves written tests for aptitude or intelligence that are used by some employers for purposes of promotion, which I considered in earlier chapters of this book. As I noted in those earlier chapters, some of these tests have been found to be very difficult for certain racial minorities to pass: the percentage of blacks or Hispanics that pass the tests, out of all of those who attempt it, is a much smaller percentage than the percentage of Caucasians who succeed, relative to the number who attempt it. Often, this occurs in part because the questions on the test presuppose knowledge of certain kinds of life experiences and certain sorts of social interactions, of a sort that are more commonly had by Caucasian families than by these racial minorities. In some cases, the disparity in success rates results also from direct discrimination: white employees are part of a social network from which minority employees are excluded, and senior employees within this network are happy to coach friends and family members but not minority candidates. So that this will remain an example of “independent indirect discrimination”, let us suppose that this is not occurring.

Most countries’ laws would deem these tests unjustified wrongful discrimination only if there were alternative tests available that could successfully track aptitude for the job, while at the same time increasing the number of minority candidates who pass the test.

¹⁶ See e.g. the *Meorin* case, *supra* note 2.

And the availability of these alternative tests is important, because it makes a difference to what agents of wrongful indirect discrimination are doing and failing to do when they persist in applying the current tests. They are continuing to use their original tests in circumstances where there are alternatives available that would harm the minority groups less, while disadvantaging the employer in only a relatively small way. In some of these cases, the employer presumably realizes that there are alternative tests available but decides not to implement an alternative test, either for reasons of cost or simply out of laziness. In other cases, the employer does not know that there are alternative tests available, but has a vague suspicion that there might be, and avoids looking into this because it is easier to turn a blind eye. And in still other cases, it may never have occurred to the employer that the original test poses difficulties for certain minorities, because the employer doesn't often bother to think about minority employees as the kind of people who deserve to be promoted. You may think there are significant differences between these different employers' actions. The first involves knowledge that a harm is avoidable; the second involves willful blindness; whereas the third involves a complete lack of awareness. But *all* of these actions seem to manifest exactly the same failure to see others as an equal that we saw in the example of indirect discrimination in the armed forces. Here, it is a failure to see other people's interests as significant enough to outweigh the relatively small trouble or cost that would be involved in looking into a particular test's effects on this group, in searching for a viable alternative, or in changing the test once an alternative is found.

So in cases of wrongful indirect discrimination that are not all things considered justified, the agents do seem culpable. And I wonder if they seem even more culpable when we reflect that in many of these cases, part of the reason why the organization in question has not tried to look for or develop alternative tests has to do with a lingering stereotype. Perhaps it is the stereotype that women don't really belong in "rough" professions such as fire-fighting: they are too delicate, too emotionally fragile, and too distracting to men. Or the stereotype that racial minorities couldn't really cope with managerial positions: they lack initiative, they don't have their lives together, and anyway, they probably have an enormous extended family at home that would take their attention away from their job. I suggested earlier that the cases of indirect discrimination that we examined all involved a failure to see certain groups as equals. I think we often fail to see these groups as equals because we see them through the lens of a stereotype—sometimes the same stereotypes that are used to rationalize direct discrimination. By "stereotype", I mean, as I suggested in Chapter Two, a generalization about a trait that is allegedly possessed by some or all members of a particular social group, which is used as a justification for seeing members of that group as different from ourselves and often as less than fully capable. There may certainly be circumstances in which reliance on stereotypes is necessary and unproblematic; but when an agent is responsible for attending to the real needs and circumstances of the people affected by his policies and, instead of making an effort to engage with these people and to inquire about their real needs, instead relies upon a stereotype, he may seem more culpable than someone who was simply oblivious.

If I am right about the way in which stereotypes often figure in the reasoning of agents who engage in wrongful indirect discrimination, this means that what I have called “independent indirect discrimination” is not completely independent of direct discrimination. Both can be rationalized by stereotypes, and the same stereotypes that were once given as explicit justifications for particular instances of direct discrimination can be cited to try to avoid having to search for alternatives to policies that disproportionately disadvantage certain groups. This does not pose a problem for my argument: independent indirect discrimination is still “independent” in the sense that it does not impose disadvantage on minority groups by encouraging other agents to engage in separate acts of direct discrimination towards this group. And so my examples of indirect discrimination still serve the purpose of helping to demonstrate that agents of indirect discrimination can be just as culpable as agents of direct discrimination, and in much the same way. Both, I have argued, often fail to see members of certain other groups as equals—and this is a significant failing, regardless of whether it results from deliberate neglect, or from willful blindness, or from ignorance.

6.5 *The Negligent Discriminator*

I have argued that agents of direct and indirect discrimination often share a single moral failing: they have failed to think of others as their equals. I think we can see this as a form of negligence—or rather, as involving two concurrent sorts of negligence. It is negligent in a sense akin to the negligence of tort law, which has to do with *treating* someone in an unreasonable way. And it is negligent in a moral sense as well, the sense of unreasonably failing to think of something that one ought to have thought of—or more accurately, in this case, failing to think of *someone* as one ought to have thought of them, as the equal of others.

I should emphasize here that in suggesting that both direct and indirect discrimination involve negligence, I am not offering the concept of negligence as a test or instruction manual for determining when an alleged discriminator has violated a duty towards a discriminatee. Their duty, as I have explained earlier in the book, is to treat the discriminatee as an equal. They act negligently when they fail to treat her as an equal, in circumstances where their failure is not all things considered justified. And they show moral negligence when they fail to think of her as a person whose interests merit a certain weight in their deliberations, and whose real needs and circumstances ought to be taken into consideration when deciding on relevant policies.

There are at least two objections one might make to my suggestion that both direct and indirect discrimination involve negligence. First, one might argue, as some legal scholars have done, that in fact prohibitions on indirect discrimination are much more akin

to *strict liability* than they are to prohibitions on negligence.¹⁷ It is no defence to a claim of wrongful indirect discrimination that one took all of the precautions that one reasonably could have taken to avoid disproportionately harming a particular protected group or that one did one's best to look into alternative policies. And presumably, even though in many cases we do think that a reasonable person in the agent's position would have been aware of the disproportionate effects of their policy on a particular group and would have located a viable alternative, nevertheless there will be some cases in which agents, through no fault of their own, fail to notice either the availability of alternative policies or the negative effects of their existing policy. Would we really want to say, as I seem to have done above, that these agents are negligent? Wouldn't we want to say, instead, that although they are not negligent and are not in any way at fault, there are nevertheless sound policy reasons for holding that they too should bear the costs of fixing their policies, to eliminate these harmful effects on protected groups?

It is true, of course, that absence of fault is generally no defence in law to a claim of indirect discrimination. But this does not seem to me to show that it is unhelpful to think about the kind of failing that is involved in cases of discrimination as a form of negligence. I think we can view the absence of such a defence as reflecting the diverse nature of the aims of anti-discrimination laws. One of the aims of anti-discrimination law is clearly to encourage governments, corporations, employers, providers of goods and services—any agent who is in control of significant resources or is in the position of offering significant opportunities to members of the public—to consider the impact of his actions on groups that have historically been treated as second-class citizens and significantly disadvantaged. But another aim of anti-discrimination law, and particularly of prohibitions on indirect discrimination, is to try to rectify or reduce some of these disadvantages. In other words, anti-discrimination law focuses not just on the agents of discrimination but on the effects of their policies on protected groups. Presumably, if lack of fault were a defence, this would impede the goal of improving the prospects of these groups. It might also function as a disincentive to employers or other agents of indirect discrimination, who might, under an explicit fault standard, be less likely to stretch themselves, less likely to take that extra step to try to figure out alternatives that would cause less harm to members of protected groups. So we can explain the absence of this defence in a way that is entirely consistent with my theory.

Having said this, I do not think it is true that many actual cases of wrongful indirect discrimination involve agents who have made perfectly reasonable assumptions and investigations but were simply unable, through no fault of their own, to grasp that their policies have had disproportionate effects on protected groups or to locate better alternatives. Most people are aware of the history of exclusion of certain social groups in our societies, and we are bombarded by reports of discrimination from the media—so although it certainly does not follow that every organization will be aware of every

¹⁷ See e.g. David Benjamin Oppenheimer, "Negligent Discrimination," *University of Pennsylvania Law Review* 141 (1992), pp. 899–972.

discriminatory aspect of their policies, it does seem to me reasonable to expect most people to look into their practices and policies and to consider their impact on members of different social groups. And it seems to me that most of us are likely already to have access to much of the relevant information we need, in order to assess the impact of our policies on these groups. Most employers and providers of goods and services are a part of many social networks of similar employers or providers of similar goods. The idea that there could, under these circumstances, be many sincere employers who are simply unable to figure out that their tests have unfair adverse effects on ethnic minorities or oblivious to the fact that their policies unfairly disadvantage women—seems to me a convenient fiction, one that some agents of indirect discrimination might like us to believe, but not one that has much basis in actual fact. So I think we need to be honest that such cases arise rarely. When they do, we can say, as I did above, that there are nevertheless sound reasons for holding these agents responsible for cost, even if from a moral standpoint they are not culpable.

There is also a second objection one might make to my suggestion that both direct and indirect discrimination involve negligence. One might object that it is only negligent to fail to give other people's interests a certain weight in one's own deliberations, and to fail to act accordingly, if we are actually *obliged* to give others' interests that weight. And one might claim that the existence of such obligations is precisely what is contested by at least some of those people who think that indirect discrimination is less morally problematic than direct discrimination. For instance, scholars such as John Gardner and Richard Arneson would argue that the agent who discriminates indirectly does not inappropriately elevate her own interests above theirs because she stands, in the first place, under no obligation to give their interests any particular weight in her own deliberations. Of course, both Gardner and Arneson allow that there could be beneficial effects to prohibiting indirect discrimination. For instance, such prohibitions likely result in a redistribution of opportunities from the privileged to the underprivileged, and this will increase the well-being of underprivileged groups. But these are just beneficial consequences of a certain policy choice; and they do not, for Gardner and Arneson, track any kind of prior moral duty that we have to members of these groups. And so agents who fail to give such weight to the interests of others are not, on their views, negligent.

It is true that, on the view I have been defending in this book, we do have a moral duty to treat others as equals. This is, as I said in the section on responsibility for cost, is a responsibility we take on, as members of democratic societies, where we share the benefits and burdens of living together in society as equals.

But what about in our private lives? One might accept that, when one holds oneself out to the public as an employer, or as a provider of goods and services, then one stands to benefit from social cooperation, and so one must take on the corresponding burdens of helping to ensure that others are treated as equals. But is the same true in the private context—for instance, within the family, or among friends? Do I really have a duty to treat everyone as equals, even when I am choosing whom to invite to a dinner party in the

privacy of my home? Haven't many philosophers written about the importance of special relationships in our lives, and about how such relationships require us to prioritize certain people over others, giving special preference to our children, our parents, and our friends? How is my account consistent with the recognition of such relationships? I shall turn to these, and other related questions about the difference that the roles and responsibilities of the discriminator make to the scope of his duties of non-discrimination and the availability of certain justifications, in the next and final chapter of the book.

Chapter Seven

The Duty to Treat Others as Equals: Who Stands Under It?

7.1 Situating the Question

I have argued in this book that when we treat people differently on the basis of a certain kind of trait—their race, their religion, or any other trait that ought to be part of a list of prohibited grounds of discrimination—we may wrong them in one or more of a number of ways. We may unfairly subordinate them to others, perhaps by marking them out as inferior to others, or rendering them invisible in a certain context, or contributing to the unfair subordination of a social group to which they belong. Or we may infringe their right to a particular deliberative freedom, their right not to have to think about traits such as their gender, or other people’s assumptions about these traits. Or we may deny them what I have called a “basic good”—that is, a good that these people need to have access to, if they are to be and to be seen as, full and equal participants in their society. I have argued that a detailed explanation of why such cases of discrimination are wrongful needs to refer to such facts as these—the fact that the agent unfairly subordinates some people to others, or infringes their right to a particular deliberative freedom, or denies them a basic good. But I have also suggested that these different ways of understanding why discrimination is wrongful can be viewed as three different conceptions of what it is to “fail to treat someone as an equal.” So when discrimination is wrongful, it wrongs people by failing to treat them as the equal of others; but what, in particular, this means—what exactly is involved in “failing to treat someone as an equal”—can be different, in different circumstances.

I have not yet said anything, however, about who stands under a duty to treat people as equals, in the first place. Governments? Individuals acting in what we might call a “public” capacity, such as employers or providers of goods and services? What about individuals when they make more personal decisions? I have been able to postpone consideration of these questions until this point, because I have so far confined my examples to two kinds. Most of the cases that I have used in order to explore what makes discrimination wrongful have been cases in which, although we might disagree over whether the discrimination in question is wrongful, it is nevertheless clear that the discriminator is the sort of body or individual that stands under a duty to treat people as equals when making decisions of that type—for instance, governments making decisions about funding water treatment on and off reserves, and employers adopting dress codes for their employees. By looking at these sorts of examples, and taking it for granted that governments and employers stand under a duty to treat people as equals in these contexts, we were able to focus our attention instead on the question of how best to understand the complaints of those who argued that they

had been wrongfully discriminated against. Second, though I did discuss a few cases in which some may doubt whether the agent has a duty to treat everyone as equals—such as *Masterpiece Cake Shop*, in which Phillips the baker argued that to force him to sell a wedding cake to Craig and Mullins, a gay couple, was tantamount to failing to respect his freedom of speech and freedom of religion¹—nevertheless, I used this second type of example mainly in order to explore how discriminatees such as Craig and Mullins experience discrimination that seems wrongful.

So up until now, I have not said anything about why we might be justified in supposing, for instance, that the state has a duty to treat those whom they deal with as equals. Nor have I said anything yet about the obligations of non-discrimination that we might have as individuals, when we make personal or familial decisions—decisions about whom to date or pursue friendships with, which babysitter to hire for our children, or how to educate our daughters and sons. Do we have a moral obligation to treat everyone as equals when making such personal decisions? If so, why? And what about businesses, that seem in some respects akin to private individuals making a personal decision, and in some respects akin to the state, exercising significant amounts of control over people and distributing important resources or benefits? Or the individuals who work for such businesses—the employers and the employees, the bakers, the flower arrangers, who are serving the public but doing this as part of a life that they are trying to live in accordance with their own beliefs? What is the extent of their obligations of non-discrimination? These are the questions I shall pursue in this chapter.

Before I turn to them, however, there are two important things to note. First, when I speak in this chapter of a “duty” to treat others as equals, I am referring to a duty that we may have, independently of whether the state chooses to recognize it or chooses to attach sanctions to its violation, as a matter of positive law. I shall sometimes call this duty a “moral duty.” But this is only to distinguish it from legal duties, or duties that are recognized by the law. My arguments do not presuppose any particular view about the nature or strength of moral duties, or their relation to other duties that we have.\

Second, a reminder that, as I have understood it in this book, the duty to treat others as equals is broader than a duty of non-discrimination. Back in Chapter One, I explained that there were different ways of failing to treat people as equals; and I explained, also, that the main concern of my book would be those ways that involve wrongful discrimination. In this chapter, I shall focus on the three forms of wrongful discrimination that I have been discussing throughout the book. But, as with the rest of the book, my arguments here are consistent with the recognition that one can fail to treat others as equals in certain other ways as well, some of which do not involve discrimination.

¹ *Craig v. Masterpiece Cakeshop, Inc.*, 370 P.3d 272 (Colorado Court of Appeals, 2015), reversed by the U.S. Supreme Court in *Masterpiece Cakeshop, Ltd. et al v Colorado Civil Rights Commission et al.*, 584 U. S. ____ (2018) [*Masterpiece*].

7.2 A Seemingly Plausible Answer

Who, then, has an obligation to treat others as equals? One seemingly plausible answer, endorsed by some legal philosophers writing on discrimination and also suggested by the arguments of some moral philosophers, is that although governments have a duty to treat everyone whom they govern as equals, as do individuals who have stepped into the public sphere and occupy institutional roles that render them in certain respects like the state—employers, for instance, or providers of goods, services or accommodation to the public—nevertheless, individuals making personal decisions generally do not have such a duty. Many have argued that we have a very strong interest in freedom of association and freedom of contract, at least when making personal decisions about our families and friends.² And this suggests that we cannot stand under a duty to treat everyone as equals when making these personal decisions.³ Moreover, moral philosophers have argued that it would be too demanding if people were required to give everyone's interests equal weight in their personal decision-making, rather than being permitted to favour the needs and preferences of those they love, and that it might even make certain kinds of deep personal relationships impossible.⁴

One question that proponents of this common view need to answer is: how it could be that individuals have no duty to treat others as equals when they are acting in a more personal capacity, and yet acquire such a duty when they occupy certain more public institutional roles? Most countries that have anti-discrimination laws treat legal duties of non-discrimination as being owed, not just by the state to those whom it governs, and not just by government employees or agents, but also by ordinary individuals, when they occupy certain institutional roles: for instance, employers, in their treatment of employees, and providers of goods and services and accommodation, when they offer these things for

² See Matt Zwolinski, "Why Not Regulate Private Discrimination?" *San Diego L Rev* 43.3 (2006) pp. 1043–61 at p. 1043, and Michael Blake, "The Discriminating Shopper," 43 *San Diego L. Rev.* 1017 (2006) pp. 1017–34 at pp. 1017–18, describing what he calls "a settled point for liberals." (Note, however, that Zwolinski goes on to argue that we have a similarly strong interest in freedom of contract *even* in commercial contexts, when we are acting as employers, or providers of goods, services or accommodations).

³ See, in addition to the works cited above, Hugh Lazenby and Paul Butterfield, "Discrimination and the Personal Sphere," Ch. 31 of the *Routledge Handbook of the Ethics of Discrimination* (New York: Routledge, 2018) pp. 369–78 at p. 372.

⁴ These particular claims of moral philosophers have been made in a somewhat different context – that is, not as part of discussions of discrimination, but within debates over the soundness of utilitarianism and consequentialism. But their plausibility and their centrality within our moral thought seems to me to explain some of the reticence of those working on discrimination to hold that we stand under a duty to treat others as equals in our personal decision-making. See Bernard Williams, "A Critique of Utilitarianism", in J. J. C. Smart and B. Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge UP 1973) and 'Persons, Character, and Morality', repr. in Williams, *Moral Luck* (Cambridge: Cambridge University Press, 1981); Samuel Scheffler, *The Rejection of Consequentialism* (Oxford: OUP, 1994) esp. Chs. 1–3; Samuel Scheffler, *Human Morality* (New York: Oxford University Press, 1992) esp. Chs. 6–7; and David Brink "Impartiality and Associative Duties", *Utilitas* 13.2 (2001) pp. 152–72.

sale to the public. And we assume that the law is justified in imposing these legal duties on individuals in these contexts, because they really do have such duties when they occupy these particular roles. In other words, this feature of our laws seems to reflect something about the moral obligations that individuals stand under, when they occupy certain institutional roles.

One way to answer this question is to suggest, as Gardner has done, that when individuals occupy these institutional roles, then the state can justifiably impose a legal duty on them to treat others as equals. But that legal duty is not an attempt to recognize a pre-existing moral duty: individuals have no such moral duty.⁵ The state can choose to impose a legal duty on certain individuals—for instance, employers, or providers of goods and services—not to discriminate against certain people in certain contexts. And imposing such a duty is justifiable if doing so would serve important social goals, such as incentivizing behaviour that the state views as desirable, or transferring the costs of certain disadvantaged people's needs onto the shoulders of those who, like the large employer, are better able to bear these costs. So the state may have good reasons to impose such a legal duty on individuals in certain circumstances. But importantly, this is not because these individuals have a prior moral duty to treat everyone as an equal. And if a particular government decides not to impose such a legal duty on these individuals, it is not making a mistake. It is just making choices different from the ones made by societies with anti-discrimination laws that apply to the private sector.

This answer seems to me to sit uncomfortably with our ordinary beliefs about the duties of individuals who occupy such institutional roles. Most of us believe that when the law places employers or providers of goods and services under such obligations, it is justified in doing so because these people really do have an obligation to treat others as equals. They have such an obligation, whether or not the law chooses to recognize it. And so a state that failed to recognize such obligations under similar social conditions to ours would not just be doing things differently, but making a mistake. Of course, Gardner denies this. But it seems to me that our sense that these duties are not just the law's way of turning discrimination into a *malum prohibitum*, but the law's way of recognizing a *malum in se*, runs very deep. And so we ought to see if there is a coherent account of the duty to treat others as equals that can make sense of this appearance.

Perhaps Khaitan's somewhat different view could help here.⁶ When discussing the duties owed by individuals who occupy certain institutional roles, Khaitan proposes that what distinguishes these individuals from individuals engaged in more personal deliberations is the fact that they occupy roles that "have a sufficiently public character." This in turn is relevant, he says, because when a person occupies a role with a sufficiently public character, she has a much weaker claim to negative liberty. And this means that the

⁵ John Gardner, "Discrimination: The Good, the Bad and the Wrongful," *Proceedings of the Aristotelian Society*, 118(1) (2018), pp. 55–81.

⁶ Tarunabh Khaitan, *A Theory of Discrimination Law* (Oxford: Oxford University Press, 2015) at Ch. 7 s. 1. All further quotations from Khaitan in this part of the paper are taken from this section.

kinds of reasons that might weigh against that person's having a duty to treat others as equals are simply not present, or not as strong, in the case of employers, service-providers, and others who occupy such institutional roles.⁷ But why, and in what sense, do employers and service-providers have a "public character"? Khaitan has a two-fold answer to this. An employer's public character is, he says, "based on the institutional power she enjoys": employers wield a great deal of power in our society. By contrast, providers of goods and services have "assumed a degree of public-ness by offering to serve the public generally." In both cases, however, the public character of these roles means that the individuals who occupy them have a reduced interest in negative liberty; and so, when this is weighed against the state's very significant interest in ensuring that people in such public roles treat others as equals, the latter outweighs the former, and the individuals have a duty to treat others as equals.

This reasoning seems intuitively plausible. We do think of institutional roles such as that of employer or provider of goods and services as having something of a "public character," and it seems plausible that this public character, whatever it is, is in some way relevant to their duty to treat others as equals. But if we take a closer look at the particular claims in Khaitan's argument, I think we will see that many of them are problematic, and cannot do the work that the argument needs them to do.

It is true that large employers wield a great deal of power in our society. But does this make them "public" in the right sense, the sense that Khaitan needs to support his claim that the individuals occupying these roles have a reduced interest in negative liberty? Surely I wield just as much, if not more power, over my small daughter than any employer wields over his employees—and this power is just as much a function of our social institutions as is any employer's power. Yet we do not think that this particular state-like aspect of my parental role reduces my interest in negative liberty. On the contrary, the parental role is usually assumed to be a paradigmatically "private" role, in the sense that its bearers are thought to be entitled to a significant amount of freedom from state interference with their decisions about how to raise their children.

With respect to providers of goods and services, it seems to me that the claim that they are "public" because they have voluntarily undertaken to serve the public is problematic, for at least two reasons. First, many providers of goods and services would argue that they have not undertaken to serve the public at large: they have only set out to serve a sub-group of the public, those who accept their mission as they define it, or those whom they can serve in a manner that is consistent with their religious beliefs—in the way, for instance, that Phillips the baker argued that his bakery was able to serve wedding cakes only to heterosexual couples. It seems question-begging to claim that in setting up shop as

⁷ This, of course, is not on its own sufficient to show that such individuals *do* have a duty to treat others as equals—for that, we need a positive reason for supposing that they stand under such a duty. Khaitan's account of how these particular individuals can help to achieve the goal of eliminating group disadvantage provides this positive reason. I am interested at this point only in the part of his argument that I have included in the main text, so I shall not address the rest of it here.

a baker, Phillips has implicitly undertaken to serve everyone: this is exactly what he is contesting. Secondly, however, the implied undertaking to serve the public is, on Khaitan's argument, supposed to make the baker more state-like because the state's job, too, is to serve the public; and it is supposed to be what leaves such providers of goods and services with less of an interest in negative freedom. But it seems to me that the sense in which the state "serves" the public and therefore has no interest in personal freedom is completely different from the sense in which the baker serves the public. The state serves the public in the sense that its *raison d'être* is to promote the interests of the public. It is acting in the service of their interests. Indeed, it has no interests of its own, apart from the individual interests and the collective interest of its members. And this is precisely why we do not speak of the state having a personal interest in negative liberty. But the same is not true of the baker. He may literally "serve" the public in the sense that he serves up his cookies and cakes. But his purpose in opening up his shop is not to promote the public interest: it is to promote his own interests, possibly those of his employees, and possibly those of the people to whom he wishes to sell his baked goods. And so it seems reasonable to suppose that he still has as much of an interest as ever in his own negative liberty, even while he is serving the public.

More generally, it is not obvious that when people occupy such institutional roles as the role of employer or the role of a provider of goods to the public, their interest in negative liberty weakens. We do not stop living our lives as private individuals the minute we arrive at work: underneath the baker's hat and inside the employer's suit are people who are still trying to live out their lives in the ways they think best. Indeed, it is often through our jobs that we realize some of our most important personal aspirations. This is what makes cases such as *Masterpiece Cake Shop* so difficult to think about: we cannot simply say that once Phillips dons his baker's hat for the day, he assumes a public role and straightforwardly acquires the kinds of obligations that the state is normally thought to have and loses all or most of the interests in freedom that individuals have when deliberating in more personal contexts.⁸ And this is why it is so difficult for us to come to an answer about whether discrimination is wrongful in such cases. We need an explanation of the duty to treat others as equals that will allow such difficulties to be represented and will show us how to conceptualize them, rather than an explanation that implies that these difficulties do not exist because such individuals have lost their interest in negative liberty when they step into certain institutional roles.

So far, I have tried to show that there are some problems with the ways in which scholars have tried to justify the common view that we have no duty to treat others as equals when we make personal or familial decisions in our private lives, but then acquire

⁸ For a more detailed analysis of the way in which people realize their personal goals through their work, see Zwolinski *supra* note 2. For an argument that public employees cannot be asked to set aside their personal values when taking up their public roles, see Christopher McCrudden, "Marriage Registrars, Same-Sex Relationships, and Religious Discrimination in the European Court of Human Rights," Ch. 16 of *The Conscience Wars: Rethinking the Balance between Religion, Identity, and Equality*, ed. Susanna Mancini and Michael Rosenfeld (Cambridge: Cambridge University Press, 2018), esp. section 16.5.1.

such a duty when we step into certain institutional roles. I have focussed so far on the explanations that have been given for why we acquire such a duty when we step into certain institutional roles. But it seems to me that there is a further problem with the common view. It is a problem with the way in which the view portrays our personal lives. It seems to me to take much too thin a view of our obligations to others in the context of family and friendship --and relatedly, to underestimate the ways in which our personal lives are always in a sense "public," always lived through and in relation to a variety of institutional roles. I shall argue that, if we are committed to living together in a society of equals, then we must be committed not only to recognizing a duty owed by the state to treat those whom it governs as equals, but to recognizing a duty owed by each of us to every other member of our society, to treat them as equals as well. However, I shall urge that this obligation is actually less onerous than one might expect, and certainly less onerous than the common view supposes. It can seem implausible that individuals stand under a duty to treat others as equals in their own personal lives, if we suppose that this must involve giving every other person's interests equal weight in one's deliberations at all times, and never favouring some people's interests over others. But of course this is not how, in this book, I have understood what it is to treat others as an equal. I have articulated three distinct conceptions of treating others as an equal, and I shall appeal to these three conceptions in the rest of this chapter to try to show that the duty to treat others as equals does not impose unreasonable demands on us, or demands that are inconsistent with recognizing that we have interests in freedom of association and freedom of contract. I shall also show how, on my view, we can reason through cases such as *Masterpiece Cake Shop* without explaining away what is difficult about them. And I shall argue that, even though we do have a moral obligation to treat others as equals even when making more personal decisions about our family and our friends, there are nevertheless sound reasons for not extending anti-discrimination laws to these contexts --that is, for not recognizing a parallel legal obligation in these contexts.

Before I turn to these arguments, however, I want to consider the state and its obligations to treat others as equals. After we have done that, we will be in a better position to understand the obligations of individuals, and of those individuals and organizations who seem to be straddling the line between public and private.

7.3 The State's Duty to Treat Those It Governs as Equals

There are at least three different kinds of arguments we might give to show that the state has a duty to treat those whom it governs as equals.

On the one hand, we might start from a pre-existing commitment to creating what relational egalitarians have called "a society of equals." As Elizabeth Anderson has noted,

such a society can be defined both negatively and positively.⁹ Negatively, it is a society that is not characterized by the oppression or marginalization of some social groups by others. Positively, it is a society that treats all adults as equal and independent agents, giving them the opportunity to participate equally in important social institutions and political governance, and also giving them the opportunity to live out their lives in accordance with their own personal aspirations, at least insofar as these are compatible with recognizing others as equals. It seems plausible that a pre-condition for establishing and maintaining such a society of equals is that the government must treat the various members of this society as equals, at least in the three senses that I have discussed in this book. First, it cannot subordinate some people to others, by marking them out as inferior or rendering them invisible in certain contexts. Second, it cannot deny certain people what I have called a “basic good” —that is, a good that, given the needs and circumstances of these particular people and the significance of that good in that society, they must have access to if they are to be, and to be seen as, equals in that society. And lastly, the state cannot deny, to any person whom it governs, a deliberative freedom to which that person has a right. This last claim may seem less obvious: why should we suppose that the denial of a deliberative freedom—even in circumstances where one has a right to it—would affect a person’s equal status in society? But of course, a person’s equal status in a society of equals does not consist only in their equal social and political status: it also involves being recognized as a certain kind of agent, one who is trying to live out her vision of a valuable life. And even if depriving someone of a particular deliberative freedom to which they had a right did not lower this person’s social or political status, it would nevertheless fail to show respect for her status as an agent.

I have suggested that, if we are committed to creating a society of equals, then we must assume that the state is under a duty to treat those whom it governs as equals. But some have argued that this is too strong a claim, and that, in order to create a society of equals, the state would only have to treat us as equals *for the most part*—that is, in most of its decisions, but not necessarily in all of them. Lippert-Rasmussen, for instance, has suggested that a certain social group might face wrongful discrimination, and yet might nevertheless enjoy equal status in their society, “because they enjoy offsetting advantages relative to those fellow citizens who are not subjected to discrimination.”¹⁰ So it is a contingent empirical matter, he says, whether the state needs to treat any particular group of people as the equal of others with respect to any particular state decision: if those affected could enjoy offsetting advantages elsewhere, then the state would not have a duty

⁹ See Elizabeth Anderson, “What is the Point of Equality?”, *Ethics* 109(2), pp. 287–337. For other discussions of what relational egalitarians mean by a ‘society of equals,’ see Elizabeth Anderson, “The Fundamental Disagreement between Luck Egalitarians and Relational Egalitarians,” *Canadian Journal of Philosophy, Supplementary Volume* 40 (2010) pp. 1–23; Samuel Scheffler, “What Is Egalitarianism?” *Philosophy and Public Affairs* 31.1 (2003) pp. 5–39; and Kasper Lippert-Rasmussen, *Relational Egalitarianism* (Cambridge: Cambridge University Press, 2018).

¹⁰ Lippert-Rasmussen, *Relational Egalitarianism*, *ibid.* at Ch. 2, footnote 24.

to treat them as equals, or at least, it would not have a duty that derived from the need to maintain a society of equals.

I do not find this line of argument persuasive, because I do not think that equal status within a society of equals is something that admits of this kind of offsetting of certain inferiorizing acts by other privileges. The kind of equal status that relational egalitarians care about seems to me to involve certain claims of inviolability, rather than a certain quantum of benefits. A society of equals is not a society in which we are all equally well off when one weighs the humiliations each of us has to endure in certain contexts against the privileges we enjoy in other contexts. Rather, it is a society in which no one has to endure certain kinds of humiliations, even if they enjoy huge privileges in other contexts. So I do not think the right way to qualify our claim that the state is under a duty to treat those whom it governs as equals is to say that the state is only under this duty sometimes, as a contingent matter, when the disadvantages of being treated as an inferior can be offset by advantages elsewhere.

Nevertheless, it seems right that the claim that the state must treat us as equals requires some qualification. I think the qualification we need has to do with justification: the state is always under a duty to treat us as equals, but it may sometimes be justified in violating that duty. However, it is only certain kinds of considerations that can count as adequate justifications. This is because, assuming that we are committed to creating a society of equals, the state's duty to treat those whom it governs as equals is what we might call a "constitutive duty." It is a duty that derives from the very purpose of the state. The state has a duty to treat those whom it governs as equals because this is one of its central, or constitutive purposes—at least in a country whose people are committed to living as equals. But of course, the state also has other constitutive duties, such as taking steps to maintain the health and safety of the population, in order to safeguard its own existence. And in some cases, it may be impossible for the state to fulfil all of its constitutive duties simultaneously. Consequently, it cannot be the case that the state's duty to treat those whom it governs as equals is absolute: it must be justifiable for the state sometimes to violate this duty. But it is arguable that it can only be justifiable for the state to violate this duty in cases where it can appeal to the need to fulfil some *other* constitutive duty—some other duty that, like the duty to treat everyone as equals, grows out of the very purpose of having a state. So it is not quite true, then, that the state must always treat us as equals, in order to maintain a society of equals. It always stands under such a duty. But it can sometimes be justified in violating this duty, in cases where this violation is necessary in order to fulfil some other constitutive duty. And it can also be justified in violating this duty in cases where there is a conflict between duties of non-discrimination—that is, where the state cannot treat one person as an equal without temporarily violating another person's claim to be treated as an equal, as we saw in Chapter 5.

So far, I have been exploring one argument for the claim that the state has a duty to treat those whom it governs as equals. That argument started from a commitment to a society of equals. But what about those who are not sure whether they, or we collectively,

are committed to creating a society of equals? How could we persuade them that the state has a duty to treat those whom it governs as equals?

We might borrow an argument made by democratic theorists. Some have argued recently that what justifies democracy is not that it is useful instrumentally in achieving individual liberty or promoting welfare, or that it helps us collectively to govern ourselves, but rather that it is a constituent part of a society of equals.¹¹ That is, regardless of whether democracy serves as a means to achieving any other goals, it is important because it is a necessary condition for, and indeed a constituent part of, a society of equals. Why is this? Because—or so these democratic theorists argue—democracy is the system of government that enables each of us to be ruled by ourselves rather than formally and persistently ruled by the will of others. Such democratic mechanisms as a guarantee of universal suffrage, a guarantee of equal opportunity for political influence, and a fair distribution of political power and authority across all members of society—these are all necessary to ensure that some members of society are not dominated by others. Indeed, Kolodny has gone even further than this, and has argued that insofar as we care about such democratic mechanisms, the best way of understanding our concern is ultimately as a desire for a society in which no one has a superior status to anyone else.¹² We care about giving each person an equal influence in the political sphere, and about ensuring that political decisions are justifiable to all, precisely because we care that no one should be ruled by anyone else. So it is not just the case that democratic mechanisms are a constituent part of a society of equals: it is also true that insofar as we value democratic mechanisms, this is because we already care about living in a society of equals.

So this second argument may take us somewhat further than my first argument. It does not provide a further reason for caring about living in a society of equals. But it suggests that many of us are already committed to this ideal, simply by virtue of our commitment to democracy. And perhaps more powerfully, it suggests that we in a collective sense—that is, we as groups of people who live within democratic states, or states that aspire to be democratic—are already committed to creating a society of equals.¹³

¹¹ See Niko Kolodny, "Rule over None I: What Justifies Democracy?" *Philosophy and Public Affairs* 42.3 (2014) pp. 195-229 and Daniel Viehoff, "Democratic Equality and Political Authority," *Philosophy and Public Affairs* 42.4 (2014) pp. 337-75. See also Elizabeth Anderson, *The Imperative of Integration* (Princeton: Princeton University Press, 2010) and Elizabeth Anderson, "Equality," *The Oxford Handbook of Political Philosophy*, Ed. D. Estlund, (Oxford: Oxford University Press 2012).

¹² Kolodny, "Rule over None I: What Justifies Democracy?" *ibid.* at pp. 224-25.

¹³ Seana Shiffrin has argued, more strongly, that *any* full and proper legal system *must be* democratic, because the very function of a legal system is to execute our collective moral duties through shared, communicative means. If any full and proper legal system must be democratic, and if democratic legal systems presuppose that we are all equals, then it follows that, if we are committed to the idea of a legal system (or, more accurately, a full and proper legal system, in Shiffrin's sense), then we are committed to the idea that we are one another's equals. So this argument, if it succeeds, provides an even stronger justification for the claim that we are equals—for, whereas one might say, in response to the arguments of democratic theorists that I have given above, "I don't endorse democracy for that reason," or "I don't endorse democracy at all," Shiffrin's

One might object that, since both of my arguments so far have appealed to a pre-existing commitment to creating a society of equals, they do not accomplish enough. We need a reason to think that people actually *are* the equals of others. If we could locate some fact about people, as moral agents, that would show that each of us really is the equal of each other person, we could then ground our claims about the state's duties to treat us as equals in these prior claims about our nature as moral agents. And this might seem to be a more secure foundation for our arguments about the duties of the state.

Jeremy Waldron has recently tried to provide such an argument, to locate what he calls "some basis for human worth and human dignity that constitutes us all as one another's equals."¹⁴ However, even Waldron notes that we need to be careful when we think about what exactly such an argument will show. This is because any property of people that we might seize upon—for instance, their potential for rationality and moral agency—will not literally *entail* that people with this property ought to be treated as equals; for there will always be a gap between empirical facts about us and moral facts about how we ought to be treated. Rather, the most such facts can do, he notes, is help us "make sense of an inclusive understanding of human equality."¹⁵ So perhaps there is less of a difference between this strategy and my first two arguments than there might initially seem to be.

Waldron's nuanced and complex attempt to make sense of our equal status seems to me to reveal a problem with this approach. To notice the problem, we need to start from a lesser problem, one which Waldron quite openly admits. This is that whatever property of people we pinpoint as the one that grounds their claim to equal status, there will always be some people who do not possess that property. And yet most of us would be deeply unwilling to say that for this reason, these people are not entitled to be treated as equals: as Waldron emphasizes in his discussion of the profoundly disabled, "we are determined to include them as humans and as our equals—grimly determined . . ."¹⁶ The conclusion Waldron draws from this problem is that we need to think differently in the case of the profoundly disabled, appealing possibly to an unrealized potential, or possibly to a tragic brokenness that links such people to us because the possibility of it is always present in our own lives, as well. In my view, however, this is the wrong conclusion to draw from this problem—and this is why I think that this lesser problem points us to a deeper problem with this approach. It is true that there will always be people who do not possess whatever property we might invoke as the basis for treating people as equals. And it is true that we are deeply unwilling to cast any person away, as ineligible for equal status, simply because they lack the property or properties that we have chosen. But this seems to me to show, not that we need to locate a *different* property of the profoundly disabled that might link them to us and salvage their claim to equal status, but rather that we do not need to locate

argument implies that we cannot but endorse democracy, and endorse it for this reason, if we care about having a full and proper legal system. See Shiffrin's Tanner Lectures, forthcoming [insert citation].

¹⁴ Jeremy Waldron, *One Another's Equals: The Basis of Human Equality* (Cambridge, MA: Harvard University Press, 2017), at p. 215.

¹⁵ Waldron, *One Another's Equals: The Basis of Human Equality*, *ibid.* at p. 248; see also p. 57.

¹⁶ Waldron, *One Another's Equals: The Basis of Human Equality*, *ibid.* at p. 252.

any such property *in anyone at all*, because our belief in each person's equal status is foundational. It is not a belief that we are willing to abandon. So in my view, this difficulty suggests that it is a mistake to search for a deeper foundation for our belief in each person's equal status. Any argument that tries to locate such a foundation will have to appeal to claims that we are less certain about, and more readily willing to abandon, than our conviction that we are all each other's equals. And it looks rather as though these claims will simply serve to rationalize a conviction that we are unwilling to give up in any case, rather than pointing to what really justifies that conviction. So why not just start, as I have done in my first two arguments, with our commitment to creating a society of equals?¹⁷

There is also a further problem with the strategy of trying to locate some human capacity or property that might ground an obligation to treat others as equals. It seems to me to misunderstand the nature of our commitment to creating a society of equals. At least as expressed by relational egalitarians, this commitment seems to me to be a commitment to *creating* a community in which everyone has a certain status. It is a commitment to living together in a certain kind of way, and to governing ourselves in a certain way—so, in this community, no one is treated as though their life matters more or less than anyone else's.¹⁸ If I am right about this, then the status of being the equal of others does not depend on our each having some independent property which makes each of us, separately, deserving of recognition as the equal of others. It does not depend on our having any such property, because our commitment to treating others as equals is not a recognition of some prior fact about each person, but a commitment that we make going forwards, a commitment to treat everyone within our society in certain ways, so that no one is treated as the superior, or the inferior, of anyone else. But then we do not need the kind of argument that appeals to some property of ours as human beings. We do not need to search for a property that could ground a claim to equal status, because each person's claim to equal status derives from their membership in a society that is committed to treating them as equals, not from some prior and independent property of theirs.

Waldron gives rather short shrift to a version of this view. He imagines someone objecting that equality "need not be predicated on any descriptive property of human nature" because, "by political convention, we *hold* ourselves to be one another's equals."¹⁹ His response is that the view is "slightly mad, as though we could just decide to hold trees, tigers, teapots and teenagers as one another's equals."²⁰ But this seems to me to misunderstand the idea that our commitment to treat others as equals is a practical one

¹⁷ Others, of course, have also argued that our belief in the equal status of all members of society is foundational: see, for instance, Joel Feinberg, *Social Philosophy* (Englewood Cliffs, N.J.: Prentice Hall, 1973).

¹⁸ This is how I understand the views, for instance, of Elizabeth Anderson, in "What is the Point of Equality?" *supra* note 9, and Carina Fourie, in "What is Social Equality? An Analysis of Status Equality as a Strongly Egalitarian Ideal" *Res Publica* 18, pp. 107–126.

¹⁹ Waldron, *One Another's Equals: The Basis of Human Equality*, *supra* note 13, at p. 58. He proposes initially that this was Arendt's view, but later he argues that her views were more complex and that she did take human equality to be grounded in some further property: natality, or the freedom to do or be new things.

²⁰ Waldron, *One Another's Equals: The Basis of Human Equality*, *ibid.* at p. 59.

rather than a theoretical one. The idea is not that we should, or ever could, get together and arbitrarily dictate that certain things are to be treated as equals. The suggestion is, rather, that we have already found ourselves with a commitment to treating each other as equals. Certain Kantians would argue that this is one of the basic commitments that underlies, or makes possible, our various acts of willing. I have not gone quite so far—I have suggested only that it is a foundational commitment that many of us in fact do have, and that it is implicit in our endorsement of democracy. I do not mean to suggest that the case for this commitment is watertight: indeed, one has only to look at the rise of the far right in many countries, and the upsurges in racism and religious tensions even in democratic countries, to doubt whether we do in fact have such a shared commitment to taking others as equals. But the idea of it is not, in itself, ridiculous, or an appeal to an arbitrary decision-making process, or a flight into a fantasy world of alliteration.

One might object that there is a much simpler and more powerful argument for the claim that we are all equals than any of the three arguments I have explored so far. Waldron's strategy assumes that, if our equal moral status is to have a foundation, then this foundation must be provided by certain empirical facts about us. But why should we assume this? One might argue, on the contrary, that the only necessary foundation for our equality is another *moral* fact about us: namely, the fact that our lives matter, and each person's life matters just as much as, and no more than, every other person's life. Why isn't this enough to generate a duty, on the part of the state, to treat us as equals? I am not sure that these reflections provide a foundation for the claim that we are all equals: they seem to me to be a way of spelling out part of what it means for us to be recognized as equals. To say that we are equals and that no one is superior to, or inferior to, any other person, is in part to say that our lives matter just as much as each other's, and that no one ought to be treated as though their life mattered more than others. So I accept these claims, but I am not sure they take us any farther than we were before. It is still open to someone to deny them --to say that we don't matter equally. And to this, I am not sure that we have an answer: there is no further fact that we could point to, to explain *why* we matter equally, that isn't more controversial than the claim *that* we matter equally.

But this brings us to a further objection. One might argue that this view of our equal status --that it is a commitment we make to each other, which involves recognizing that each person's life matters just as much as any other person's life-- makes our equal status seem too fragile.²¹ I think this objection, too, is in a sense, correct, but that it is an accurate statement of the nature of our commitment rather than an objection to my view. In one sense, of course, our commitment to treating others as equals is not at all fragile. It is what we might call "moral bedrock": as I have argued, it is an assumption that we are *less* willing to abandon than any claim we might make about the properties of human beings that allegedly ground this entitlement. And it arguably underlies many people's commitment to democracy. But it is not clear that there is anything further that it rests upon, that could in

²¹ I am very grateful to Larry Sager for pressing this objection.

turn be invoked to justify it to those who dig in their heels and deny that others are their equals. And so in this sense, the commitment is fragile. But this is not a fault of our arguments: it is just the nature of the commitment. And perhaps we need to realize its fragility, in order to be spurred on to taking greater measures, both individually and collectively, to ensure that people do treat others as equals.

7.4 The Individual's Duty to Treat Others as Equals

I want now to argue that if we are committed to living in a society of equals, then we must suppose not only that the state owes a duty to treat people as equals, but also that each of us, as individuals, owes a duty to every other member of society, to treat them as everyone else's equal. And I shall try to show that we have this duty not just when we occupy certain institutional roles, such as employer or purveyor of some good or service to the public, but even in our private lives, when we make more personal decisions.

This may seem implausible. But recall that I am appealing here to three quite specific conceptions of what is required, in order to treat someone as an equal: not subordinating them to others by marking them out as inferior or rendering their needs invisible, or contributing to their ongoing social subordination; not infringing their right to a particular deliberative freedom; and not denying them access to a certain basic good, in circumstances where you have the power to give them such access. One can fulfil these requirements without having to give everyone's interests equal weight in one's deliberations. So the view that I am going to defend does not have the implausible implication that we cannot prioritize the needs of those we love or care for, in our personal lives. Nor does it follow, simply because we have a moral duty to treat others as equals in these senses, that the state is justified in creating a parallel legal obligation and sanctioning us whenever we violate it.²² Indeed, I shall argue in the next section of the chapter that the state has good reason not to place sanctions on individuals' failure to treat others as equals in many personal contexts; though the state ought nevertheless to take other measures, of a more indirect and non-coercive kind, to assist individuals in complying with their moral duty to treat others as equals.

Why, though, should we think that all of us stand under this duty to treat others as equals, even when we make more personal decisions? Partly because these decisions—decisions about how to raise our children, whom to have as friends, what social and political causes to support—have significant effects on the power relations between different social groups in our society, and play a large role in perpetuating stereotypes of the kind that result in certain people being regarded as inferior to others, or less worthy of deference. They have such effects not just because they are decisions about matters that are very important to most of us, but also because they are not purely “personal” decisions,

²² When I refer to sanctions in this chapter, I have in mind any unpleasant consequence, whether a penalty or a requirement that one compensate the victims of wrongful discrimination.

even though we often think of them this way. They too are decisions we make when we occupy certain institutional roles—the role of a parent, the role of an adherent of a certain religion, the role of a host or a guest—and these institutional roles are structured by shared social expectations, and by shared social assumptions. So the actions we perform when we occupy these institutional roles have the power to perpetuate a variety of stereotypes about the people we are dealing with, to perpetuate habits of deference to some, and ignorance or censure of others. And this means that even the private or personal realm is a realm in which my actions have significant effects on the power, authority, and freedoms enjoyed by others. Eleanor Roosevelt once commented that equality needs to be respected:

in small places, close to home - so close and so small that they cannot be seen on any maps of the world. Yet they are the world of the individual person; the neighborhood he lives in Such are the places where every man, woman, and child seeks equal justice, equal opportunity, equal dignity without discrimination. Unless these rights have meaning there, they have little meaning anywhere.²³

I am making the same argument about the duty to treat others as equals. It is not only large organizations such as the state that have the power to change our situations and our social status: many of the actions that determine how we stand, relative to others in our society, are performed, as Roosevelt said, “close to home.”

And I think we already do take ourselves and others to be under duties to treat people as equals in our personal lives. For instance, few would doubt that I have a duty to treat my children as equals, in the sense that I cannot justifiably mark some out as inferior to others, or act in ways that contribute to their social subordination, either within my family or in our broader social circles. This means that I am making a mistake, for instance, if I pay to send my son to an expensive private school while insisting that the overcrowded, underfunded public school is good enough for my daughter, or if I quietly allow my son to behave like a slob in the house, while insisting that my daughter tidy up after herself and him. Similarly, I think many of us already believe that we have a duty to treat strangers as equals, and not to infringe their right to a particular deliberative freedom, when they have one. For instance, most of us do not think ourselves entitled to make cat-calls at women as they walk along the street, and we feel anger at those who do precisely because this is a way in which complete strangers try to assert that someone else is not their equal, while veiling their demonstration of their own greater power as a compliment. We hold ourselves to be under an obligation to our guests to find out about their allergies, so that no one is left with a constant reminder of their allergies or a feeling of being second-class because of them. And if we find out that someone in our neighbourhood lacks what, in

²³ Eleanor Roosevelt, “In Our Hands” (Speech delivered on 27 March 1958 on the tenth anniversary of the proclamation of the *Universal Declaration of Human Rights*), reprinted in *ABC: Teaching Human Rights – Practical Activities for Primary and Secondary schools* (United Nations: New York and Geneva, 2004) at p.11.

Chapter Four, I called a “basic good” —a good that a particular person must have access to if he is to be, and be seen as, an equal in his society—and if we know that we have the power to give this person access to that good, we generally take ourselves to be required to do so. We think we ought to help the elderly man next door who lives in social isolation, or the children at our neighbourhood school who will be left hungry during the school vacation because they normally rely on school breakfasts and lunches. Of course, our reactions to such cases can be explained in other ways as well, by appealing to other reasons we have for reaching out to help these individuals. But it seems quite plausible to suppose that one of the explanations is that we think it is not just the state who has a duty to treat people as equals, but also each of us, in our personal lives.

When legal academics and philosophers deny that we have such a duty, as private individuals, they standardly invoke two examples: the example of someone deciding whom to date, and the example of a host deciding whom to invite to a party. They argue that it is implausible to suppose that we have a duty to treat others as equals when making these decisions.²⁴ If we want not to date a certain person because of their race, this is our prerogative, just as if we don’t wish to invite a particular person to our party because of their sexual orientation, we should be given the freedom to do this. But I think we need to be careful here. First, I am not contesting that each of us should be able to date or party with whomever we want, without state interference. As I have noted, and will discuss further in the next section, one can consistently hold that we stand under a moral duty to treat others as equals and yet deny that it would be a good thing for antidiscrimination laws to apply to these personal decisions, or for the state to attach any kind of sanction to a failure, in personal contexts, to treat others as equals. It is also a separate question whether other people are morally obliged, or even morally permitted, to intervene when they see someone failing to treat another person as an equal. So we can accept that people ought to have considerable freedom to make personal decisions as they see fit, without interference from the state and without pressure from other people, quite consistently with recognizing that each of us nevertheless has a moral duty to treat people as equals, and that we exhibit some kind of moral failing when we do not treat others as equals.

Second, on the three conceptions of treating people as equals that I have explored and defended in this book, merely declining to invite someone on a date or to a party because of their race or their sexual orientation is not, in and of itself, a failure to treat this person as an equal. In order to know whether this decision amounts to a failure to treat people as equals, we need to know more. We need to know, for instance, whether the potential date or guest had a right to some deliberative freedom which this decision breached, such as the freedom to be considered as a date regardless of her race—and I think most of us would say that no, my prospective date had no right to this particular freedom. We also need to

²⁴ See, for instance, Zwolinski, “Why Not Regulate Private Discrimination?” *supra* note 2; Lazenby and Butterfield, “Discrimination and the Personal Sphere,” *supra* note 3; and Carina Fourie, “Wrongful Private Discrimination and the Egalitarian Ethos,” Ch. 35 of the *Routledge Handbook of the Ethics of Discrimination* (Routledge: New York, 2018) pp. 421–32.

know, whether my decision amounts, in the context, to marking out this person as inferior or treating her as though she does not exist, simply because of her race—that will depend very much on the context. And we will need to know whether she is thereby denied access to a basic good: this seems unlikely, given that all the agent is offering is a date or a dinner. Finally, we will need to know whether this decision contributes to the social subordination of the group of people who share the trait on the basis of which I rejected this person.

And this brings us to a further complexity. Of course, one decision about whom to take on a date tonight or whom to invite to tomorrow's party is very unlikely, on its own, to make much of a difference to the social status of anybody. But the cumulative effect of many dating decisions and many decisions not to invite people of certain races or certain sexual orientations to parties is clearly a subordinating one: these practices contribute in a very large way to the ongoing social subordination of members of these groups. And this suggests to me that the duty to treat people as equals is a complicated one. Sometimes it is clear that I must do a particular thing on a particular occasion in order to treat certain people as equals. For instance, having just ordered my daughter to pick her socks up off the floor, I must order my son to pick up his socks, too; having invited a group of friends over for dinner, I must check with each of them that they do not have an egg allergy before deciding that the sole item on our menu will be an omelette. In these cases, it is quite clear in advance what I have to do, in order to treat each person as the equal of others, and everyone in a similar situation will have similar obligations. But in other cases, such as our decisions about whom to invite to parties in the first place, it is not clear in advance what precisely our obligations are, and not clear that we will all be required to do the same things. How far exactly I must extend myself, in order to treat people as equals, will often depend on many things, including my other obligations, my projects, and my institutional role. So the duty to treat people as equals looks more like an imperfect duty, and perhaps even a collective imperfect duty. That is, it gives all of us together a required end, and that end cannot be satisfied unless, for the most part, we all act in certain ways; but it does not always require that each of us does a particular thing on a particular occasion, and we cannot always tell on the basis of an isolated case whether the duty has been complied with or not.

It is common to think of imperfect duties as duties that leave room for the agent's choice: the agent, it is sometimes said, gets to choose what counts as compliance with that duty. This is not, however, the conception of an imperfect duty that sits most comfortably with my analysis here. If we do have a duty to treat others as equals, it surely cannot be the case that the agent gets to decide what counts as treating others as an equal. So, in suggesting that we think of the duty to treat others as equals as an imperfect duty, I have in mind a different conception of imperfect duties, such as the conception defended by Barbara Herman. Herman has argued that what makes a duty imperfect is not that it leaves more room for the agent's choice than do perfect duties, but rather the fact that the precise content of an imperfect duty, for a given agent, depends very much on the other demands that this agent stands under, on facts about her other activities and projects, and also her

institutional role. Imperfect duties also complement public or institutional duties --and in this respect, too, the duty to treat others as equals seems to fit the model of an imperfect duty. Just as our duties of beneficence increase the more the state neglects to look after the welfare of its members, so our duty to treat others as equals will be more demanding, the more the state fails to create the conditions under which we can relate to others as equals.

Because the duty to treat others as equals is an imperfect duty, it raises difficult questions. How far does each person have to extend themselves, in order to treat others as equals? To what extent am I responsible for alleviating or eliminating the unfair subordination of certain groups in our society? What counts as it being “within my power” to provide another person with a basic good? Our answers to these questions will vary, depending on the agent’s other obligations, attachments, and projects. But this is also true of other imperfect duties such as duties of beneficence or duties of gratitude. It is not evidence that, as individuals, we stand under no such duty.

I have now argued that, even in our personal lives, we stand under a duty to treat others as equals, and I have suggested that it looks, in our individual cases, like an imperfect duty. But, like my argument for the state’s duty to treat others as equals, this argument has depended on our shared social commitment to creating a society of equals. It is because we live in societies with certain aspirations that we have this duty. So Gardner is in a sense correct: our duty not to engage in wrongful discrimination depends on the commitments made by our society. But this does not make wrongful discrimination into a *malum prohibitum* rather than a *malum in se*. We have a duty to treat others as equals, regardless of whether the state chooses to recognize this duty or chooses to use coercion to ensure that we comply with it.

But now my argument may seem to run into the following problem. As I noted earlier, most countries’ anti-discrimination laws impose certain special duties of non-discrimination upon people who occupy certain public roles—such as employers, and providers of goods and services. We think of these legal duties as legitimate, insofar as we suppose that people who occupy these roles really do have such duties. But I have argued that we *all* have such moral duties, even when we do not occupy these particular institutional roles. So what changes, when we occupy these particular institutional roles? And does the “public” nature of these roles make no difference to our moral obligations?

I shall argue in the next section of the chapter that what changes when we occupy these roles is that we lose the reasons that we have, in more personal contexts, for not having the state prohibit wrongful discrimination and attach some kind of sanction to acts and practices that wrongfully discriminate. In my view, the “public” nature of these institutional roles is relevant, not to the existence of a moral obligation to treat others as equals, but to the absence of certain kinds of reasons for not applying antidiscrimination laws to the actions of those who occupy these roles. In other words, on my view, the relevant question is not “Why do we have obligations to treat others as equals when we assume certain public roles, but not otherwise?” but “Given that we *always* have such

obligations, why should the state *not* apply antidiscrimination laws to us in certain more personal contexts, and which contexts, exactly, are these?" I shall turn to this inquiry now.

*7.5 Reasons for Not Legally Prohibiting Discrimination in Private Contexts*²⁵

There are a variety of different personal contexts that many countries treat as "private" in relation to discrimination, in the sense that they assume the state is not justified in interfering with people's choices in these contexts, even when some people are not treated as equals. Which contexts are these, and what might be some good reasons for not legally prohibiting wrongful discrimination in these contexts?

We can start with the most personal. Most countries that have anti-discrimination laws do not generally impose legal obligations of non-discrimination on families, spouses, or between friends: we permit people to decide for themselves how to relate to the members of their families and to their friends, and importantly, we let them decide for themselves how to allocate authority and power between family members and friends. One plausible explanation of why this is so is that part of what is valuable about these relationships is the fact that they grow naturally out of their members' own desires and aspirations. Of course, we have to be careful here: even our family lives and our friendships are already subject to a considerable amount of state regulation. They are bounded, and structured, by the rules we have for marriage, by rules requiring us to provide necessities to our children, by the rules of negligence law and property law. So the problem here is not that it would be difficult to have deep and meaningful personal relationships with state interference: we already have a considerable amount of state interference in these relationships, and much of it is arguably necessary for the flourishing of these personal relationships. The kind of state interference that would be problematic is the kind that would interfere with what is valuable in such relationships—and, as I have suggested, part of this value appears to inhere in the fact that spouses have the chance to choose each other and to choose, together, the kind of life they are going to live and the kind of family they want to create for their child, just as friends have a chance to choose each other and choose the kind of friendship they want to have. Anti-discrimination laws that governed personal relationships might prevent us from making these choices on our own, in our own way. The same sort of reasoning seems to underlie legal exemptions for private clubs: we want, similarly, to allow people to form recreational associations and pursue their passions together, in the company of the people they choose.

Note, importantly, that this argument presupposes that these relationships are valuable insofar as they reflect the shared desires and aspirations, and the free choices, of spouses and friends. So there is room, consistently with this argument, for us to suggest

²⁵ For ease of writing, I shall often in this section refer simply to "legally prohibiting discrimination," instead of "legally prohibiting wrongful discrimination," but of course what is at issue are prohibitions on *wrongful* discrimination.

that where marriages are forced, or where the family is a site of male authority and female oppression, then we lose the reasons we might otherwise have had for thinking that the state should not regulate discrimination within families. Indeed, the 2015 “UN Working Group on the Issue of Discrimination against Women in Law and In Practice” recommended specifically that states ought to prohibit discrimination within the family, and to take appropriate measures to enforce such prohibitions.²⁶ As the Working Group noted, in many countries, women are forced into marriage; do not have equal decision-making power within the family; are denied education by family members; and are denied, by their family members, the privilege of engaging in economic or social activities outside of the family without the supervision of a male family member. In countries where the family is characterized by this kind of asymmetry in power and authority, there may be no justification for the state not intervening coercively to enforce individuals’ obligations to treat their family members as equals.

At this point, one might object that many commercial enterprises are valued by their owners, too, because they provide the means through which these people can freely shape a life in accordance with their own beliefs. Think of Wholefoods, which markets itself not just as a profit-making enterprise, but as a way “to nourish people and the planet.”²⁷ And yet we do think that the state *can* justifiably intervene to prohibit discrimination in the hiring of employees by companies such as Whole Foods, and in their dealings with customers. So why do we treat employers and providers of goods and services differently from the way we treat private individuals making personal decisions? One reason may be that we think of the relationship between employer and employee, and between business owner and customer, as predominantly commercial relationships. The parties to these relationships may share a vision of what they are doing; but they do not have to, and they have chosen to enter these relationships primarily in order to turn a profit. For all of Whole Foods’ rhetoric, its CEO’s and its shareholders’ main aims are to nourish their profits, and they have found a way to do so by appealing to people’s desires to nourish their health and the planet’s. Of course, the same may not be true of small, artisanal businesses—the Haida art store, for instance, that aims to promote awareness of Haida art and enable a new generation of indigenous artists to learn and in turn develop the art of their ancestors, or the willow basket-maker, who carries on a heritage trade, or the cheese-monger, who has made a career out of creating new artisanal cheeses.²⁸ But such smaller, artisanal businesses are sometimes exempted from anti-discrimination laws, and perhaps this explains why.

But why, exactly, should the profit motive of larger companies make a difference here? I think it makes a difference because it is not clear, then, that what we value in these

²⁶ Report of the U.N. Working Group on the Issue of Discrimination Against Women in Law and in Practice, submitted on 2 April 2015, pursuant to Human Rights Council resolutions 15/23 and 26/5.

²⁷ See the Whole Foods Mission Statement: <https://www.wholefoodsmarket.com/our-mission-values>.

²⁸ For an argument that this is untrue even of larger businesses, and that private sector discrimination law overestimates the importance of profit and undervalues the need to protect the autonomy of employers and providers of goods and services, see Zwolinski, “Why Not Regulate Private Discrimination?”, *supra* note 2.

commercial relationships would be threatened by state regulation of discrimination, in the way that what we value about personal relationships would be threatened by the state regulation of discrimination in those contexts. Consider, for instance, clothing stores that cater to the tastes of wealthy white clients. If the state enacts anti-discrimination laws preventing all stores from wrongfully discriminating against indigenous people when they hire their sales staff, then such stores may well lose some profits because of lingering prejudices on the part of its clients, who equate trendiness with whatever white people want to sell them. But the financial losses of such stores will generally not be so large that it would be impossible for them to continue in business—particularly if all of their competitors are *also* under a legal requirement not to discriminate against indigenous people in their hiring decisions. Moreover, over time, the presence of indigenous sales staff in the stores selling the latest fashions will presumably help to combat the prejudices that cause financial losses: the store's clients will learn from experience that indigenous peoples can sell trendy fashions, and as these prejudices are lost, there will be even less of an impact on the store's profits.

I have argued so far that part of what we most value about certain relationships – including, most prominently, relationships of friendship, and familial relationships-- would be threatened by state regulation of discrimination in these contexts, and that consequently, the state has a strong reason not to create a legal duty of non-discrimination in these contexts, notwithstanding the fact that individuals do still stand under a moral duty not to engage in wrongful discrimination even in these contexts.

A second reason for not applying anti-discrimination law to these more personal contexts is provided by the fact that, as I noted in the previous section, the duty to treat others as equals is an imperfect duty. It will take a different shape for different people, depending on their other obligations and their other attachments, and their various projects and activities. As I noted earlier, this is not to say that it is up to each of us to decide how and when to treat others as equals: there is a fact of the matter about what we are required to do, whether we like it or not. But there is no single rule that each of us can follow that will tell us all what to do, in order to treat others as equals in our personal lives. And this makes it difficult to articulate a general legal standard that could be known in advance, and fairly applied to everyone, in the context of our personal and familial decision-making. But the situation seems to me to be different in the context of employment and commercial relationships. Here, there *are* specific steps that all employers can take, which will go at least some way towards treating their employees as equals, just as there are rules that all providers of goods or services or accommodation can follow, which will similarly take them at least part of the way towards treating their clients as equals. Employers can ensure that hiring and promotions decisions are made in a way that does not wrongfully exclude or disadvantage people on the basis of traits that mark out subordinated groups, or groups that have been denied deliberative freedom, or groups that have been denied access to basic goods. The same is true of providers of goods or services or accommodations, when they make decisions about whom to contract with.

There is of course much more that each of them could *also* do, to treat others as equals. But we can know in advance that they must at least do this much. So a second reason why antidiscrimination law justifiably marks out duties of nondiscrimination within such commercial contexts, but not in the context of friends or family, is that in the former case, we can know in advance that all those who occupy the institutional role of employer or provider of certain goods are in a position to treat others as equals in these ways, whereas it is more difficult to know what is required of individuals in more personal contexts.

There are also a third and a fourth reason that I want to discuss for not extending antidiscrimination laws into the more personal contexts of family and friendship, but for nevertheless enforcing them in commercial contexts. The third concerns a practical difficulty: it is much more difficult to monitor the decisions we make within families and between friends than it is to monitor companies' decisions about hiring, promotions, and sales. Corporations are already under a variety of obligations to keep records of such decisions and of their reasons. But familial decisions and decisions among friends are not usually recorded. And it can be very difficult to determine who has decided what within the context of a family or a friendship, or for what reason. This is of course not a decisive reason against imposing such obligations in personal contexts, as it pertains only to the practical difficulties of enforcing them; but since these difficulties would be considerable, this does seem to carry some weight.

A fourth reason for not extending antidiscrimination laws to more personal contexts relates specifically to one way of failing to treat others as equals—namely, by contributing to the unfair subordination of particular social groups. Many of the discriminatory acts that we commit in our personal relationships make a difference to the unfair subordination of particular social groups only cumulatively, over time, and because they are repeated in many friendships and many families. Each individual act may seem, on its own, to have almost no impact at all, and to be quite innocuous; and it may be almost impossible, at a later time, to figure out which acts together made a difference, and which did not. But I wonder if the situation is different in employment, and in the provision of goods and services and accommodation—in other words, the contexts in which we do impose legal obligations of non-discrimination on individuals. Being granted or denied a particular job or a promotion can have a huge impact on an individual's social status, and derivatively, on the social status of the group to which they belong. The same is true of accommodation: being denied accommodation in a certain area of town can, similarly, result in the ghettoization and marginalization of particular social groups, as has happened to blacks in many urban areas of the United States.²⁹ And although most often, denials of particular goods and services seem to work in very small increments to make a difference to the status of particular individuals and groups, nevertheless, there are at any given time particular goods and services that become status symbols, with the result that being denied

²⁹ For a history of this marginalization through denials of accommodation, see Richard Rothstein, *The Colour of Law: A Forgotten History of How Our Government Segregated America* (New York: Liveright Publishing Corp, 2017).

these particular goods and services can, on its own, have a large impact on the social position of an individual or a group. Perhaps this is a further reason for legally enforcing obligations not to engage in wrongful discrimination in commercial contexts, but not in more private contexts. It might also help to explain why these legal duties are imposed on employers but not on employees, and on providers of goods and services but not on purchasers of these goods and services: it is the decisions of employers, and of providers of goods and services, that have the potential to have a large single impact on a particular individual or group, because of the importance of certain jobs and certain goods; whereas the decisions of employees and consumers tend to have an impact only cumulatively, over time.

I have argued that there are good reasons for not extending antidiscrimination law to personal decisions made within the family or between friends, but that these reasons largely do not apply when people are acting as employers or providers of such things as goods, services and accommodation. But it is worth emphasizing that even if it is true that the state should not apply antidiscrimination laws within these more personal contexts, it does not follow that the state cannot legitimately take other measures to assist individuals in complying with their moral obligation to treat others as equals. There are a great many things the state could do to help us —and many things that, if we are genuinely committed to creating a society of equals, the state *ought* to do. For instance, through educational policies and programs, governments can foster attitudes of respect for diversity and an understanding of different cultures and different identities among children and teens in school; through school districting rules, governments can increase the likelihood that students of different racial and cultural backgrounds and their families will mix with each other as members of the same school community, and so come to understand and respect each other. Governments can provide public spaces open to all, such as public parks and community centres, where people from different backgrounds can come together and share recreational pursuits and gradually learn more about each other. And governments can enact generous parental leave policies encouraging fathers to take parental leave, which studies have shown results in fathers sharing the tasks involved in child-rearing more equitably with mothers throughout the family's child-rearing years. Of course, these are only a few examples; but it does not take much imagination to think of many more. This seems to me an area that is ripe for new work by legal scholars writing on discrimination. Relatively little has been written on discrimination in more personal contexts. And the little that has been written tends to focus exclusively on whether the state can legitimately prohibit wrongful discrimination in these contexts, as though the state's role in these contexts must be limited to either prohibiting wrongful discrimination and imposing sanctions on those who engage in it, or standing out of the way and doing nothing at all about it. But there is surely room for us to think creatively about how the law might be used, in an indirect and more supportive way, to foster the kinds of relationships and the kinds of attitudes that will help us to treat each other as equals in our own personal lives. Discrimination law is only one way of addressing failures to treat others as equals; and it

does not follow, from the fact that it is not the best means to use in personal contexts, that there are no other means at the state's disposal.

7.6 The Value of Freedom and the Duty to Treat Others as Equals

I have not said much yet about the value of freedom, and the role or roles that are left for it to play, assuming that we all have an obligation to treat others as equals in our own personal lives. I argued earlier that, as long as the state and others do not interfere with our decisions, then we can quite consistently grant each person a sphere of negative freedom while still supposing that they stand under a moral duty to treat others as equals. But this answer might seem unhelpful, for two reasons.

First, if the only kind of freedom that my view makes room for is the freedom not to be interfered with when we do the wrong thing, this may seem to be a hollow victory for freedom. For this seems a rather unimportant kind of freedom. What we really care about, one might argue, isn't just the freedom to make moral mistakes, but the freedom to decide what we care about, to be, to a certain extent, the masters of our own moral lives. And if the duty to treat others as equals is conceived of in a capacious enough sense, it may threaten to engulf our entire personal lives, leaving us no room to decide for ourselves how we want to live, no room to be the masters of our own lives. However, as I have argued both in this chapter and earlier in the book, I am primarily concerned with the duty to treat others as equals in the three specific senses we have looked at: not subordinating them to others by marking them out as inferior or rendering their needs invisible, or contributing to their ongoing social subordination; not infringing their right to a particular deliberative freedom; and not denying them access to a certain basic good, in circumstances where you have the power to give them such access. And a duty to treat people as equals in these senses does not seem to me to be so demanding as to rob us of the power of shaping our own lives in accordance with our own ideals. On the contrary, it is arguably a *precondition* for the more subordinate groups among us to have the power to shape their own lives that the rest of us take ourselves to be under a duty to treat them as equals. That is to say, it is only if we suppose we are all under such a duty that we will all actually be able to have the freedoms that we care about.

This first worry concerned the value of freedom and the demands of equality, as they relate to each other within a single person's life. But the second, and more serious worry that I want to respond to concerns apparent conflicts between one person's claim to certain freedoms and another person's claim to be treated as an equal. We normally think of cases such as *Masterpiece Cake Shop* as involving such conflicts. And it may seem that my view leaves no room for such conflicts. Earlier, I criticized Khaitan's view on the grounds that it seems to explain the conflicts away rather than explaining why they exist and how we ought to deal with them: for his view implies that, once the baker enters the public sphere as a commercial baker, he loses his most of his interest in negative liberty. But does my

view, too, explain away the conflict? I do not think so, and I shall try to explain why in what follows.

So let us turn back now to *Masterpiece Cake Shop*. Recall that Craig and Mullins, the same sex couple, had argued that Phillips the baker was wrongfully discriminating against them by denying them a wedding cake, contrary to Colorado's public accommodations law. By contrast, Phillips had argued that forcing him to provide them with a cake would violate his rights to freedom of speech and freedom of religion. On my view, if we are assessing Craig and Mullins' charge of wrongful discrimination, we need to start by asking which of the three forms of wrongful discrimination that I have discussed are at issue here. I argued back in Chapter Three that Craig and Mullins' complaint is partly a complaint about social subordination, but also partly –and more significantly—a complaint about an infringement of their right to deliberative freedom. In particular, they are objecting to having to consider, and to bear the costs of, the baker's assumptions about their sexual orientation and what roles it makes them fit or unfit for, when buying their wedding cake. So, as I argued in Chapter Three, the question that we need to focus on is: do Craig and Mullins have a *right* to this particular deliberative freedom? We have seen that whether a discriminatee has a right to a particular deliberative freedom in a given context depends, among other things, on the countervailing interests of the discriminator. So, in deciding whether Craig and Mullins have a right to deliberative freedom, we need to consider not only their interest in deliberative freedom, but also any relevant interests of Phillips the baker, including his interests in freedom of speech, freedom of religion, and freedom of contract.

Phillips of course claimed that, if he were forced to bake this couple a cake for their marriage celebration, he would be forced implicitly to affirm that their marriage was a real marriage. And he argued that this would amount to compelled speech, contrary to the First Amendment. He also argued that it was an infringement of his right to practice his religion, for his religion forbade him from celebrating same sex marriages. In my opinion, neither of these two arguments succeeds. The freedom of speech argument seems to me dubious: surely in selling a person a product or service, one is not thereby compelled to endorse whatever purpose that person uses the product or service for. The florist who arranges the flowers for the wedding, and the limousine driver who drives the couple to the wedding venue, are not thereby implicitly celebrating the marriage –and neither is the baker. The argument from freedom of religion also seems problematic. As I noted in Chapter Three, the European Court of Human Rights has held that a person's beliefs about whether *other people* can marry are not at the "core" of their practice of their religion –that is, they do not affect the ways in which *they*, in particular, go about worshipping and living their lives as believers.³⁰ If this is right, then not only is it a mistake to think that the baker is required to celebrate or endorse gay marriage when he bakes a cake for a gay couple, but it is also a mistake to think that baking a cake for a gay couple interferes in a significant way with his

³⁰ See Ladele [citation].

own practice of his religion. As I argued in Chapter Three, the baker's position here is even less compelling than the taxi drivers' position, in the case of the Muslim driver who refuses to give a lift to a client with a guide dog: whereas, in that case, a requirement to drive guide dogs would interfere with a key element of the taxi driver's own practice of their religion – the need to be sure that they pray in a clean space-- a law requiring the baker to bake wedding cakes for same sex couples would not, similarly, interfere with an important part of his practice of his religion.

But even if Phillips' freedom of speech and freedom of religion are not at issue in this case, there is a clear sense in which his freedom is lessened when he is required to bake Craig and Mullins a cake. I am not sure that the type of freedom that he loses is helpfully described as "freedom of association": in selling Craig and Mullins a cake, he need not interact with them personally for more than a minute, and so it is unclear to me that he is, in any deep sense, being required to "associate" with them. But he clearly loses some freedom of contract, if he is required to sell them a cake contrary to his own wishes. My view leaves room for us to recognize this loss of freedom of contract. But my view asks us to think about the moral relevance of this loss within the broader inquiry into whether Craig and Mullins have a right to deliberative freedom in this case. And, importantly, my view gives us a way of distinguishing between the kind of freedom that Phillips is losing, and the kind of freedom that is at issue for Craig and Mullins. For Craig and Mullins face not just a loss of freedom of contract, but a threat to their deliberative freedom. And it is a particularly significant deliberative freedom. The decision to get married and to celebrate their wedding is a deeply important one for them, and when they are denied a wedding cake –even if only from one baker-- they are forced to bear the costs of the baker's assumptions about their sexual orientation. Moreover, these assumptions, as I argued earlier, are perilously close to assumptions of lack of worth: the baker's view is that given their sexual orientation, they are not fit to be married, not *worthy* of the institution of marriage. By contrast, as I suggested in Chapter Three, the baker's costs are, rather like the taxi drivers' costs in the case we just considered, due simply to *his own beliefs*, and not due to *the assumptions of others*. So the baker's deliberative freedom –in my special sense of the term—is not actually engaged at all in this case. What Phillips loses, when he is required to bake the cake for Craig and Mullins, is only some freedom of contract. But Craig and Mullins risk losing both that freedom of contract and a particularly important deliberative freedom, if Phillips is permitted not to sell them the cake. It seems to me that, partly because the freedoms of the baker at issue in this case are of lesser significance than the freedoms of Craig and Mullins, we can conclude that Craig and Mullins do have a right to deliberative freedom in this case.

For my purposes here, however, this conclusion is less important than is the fact that my view does not, like Khaitan's, explain away the conflict in this case. My view gives us a way of representing the freedoms of Phillips that are at stake, and a way of representing the freedoms of Craig and Mullins that are at stake, and it offers a plausible explanation of

why, even though Phillips still loses some freedom when he is required to bake the cake for Craig and Mullins, it is nevertheless true that he has an obligation to bake it.

Interestingly, in this case, the conflict between the discriminator's freedom and the discriminatee's right to be treated as an equal emerges *within* our discussion of what it is to treat someone as an equal, rather than as a conflict *between* the value of freedom and the value of equality. There is no question, on my view, that Phillips is under a duty to treat Craig and Mullins as an equal. But his freedoms are nevertheless relevant, because we need to consider them in determining *what is required* of him, in order to treat Craig and Mullins as equals. This depends, on my view, not only on facts about Craig and Mullins, but also facts about the freedoms of Phillips that would be restricted if he had to provide them with the cake.

I have argued in this chapter that it is not only the state, but also we ourselves as individuals, who have a duty to treat others as equals. I have tried to show that we can acknowledge that individuals have this duty consistently with recognizing that they have many freedoms that also matter. I have also argued that, although individuals do have this moral duty to treat others as equals, it does not follow that the state should extend antidiscrimination law to the more personal contexts of family and friendship: on the contrary, there are good reasons for antidiscrimination law not to apply in these contexts. However, I have noted that there are nevertheless many measures that the state could take in order to help us comply with our duty to treat others as equals, even in these more personal contexts. And I have urged that we need to think more creatively about what sorts of measures these must include, if we are to create a true society of equals.

Although this chapter has emphasized that both the state and individuals have duties to treat others as equals, it is not an implication of my arguments that these duties are alike in all respects. So I want to end the chapter by briefly mentioning three important differences.

There are, firstly, differences in the content of the duty. Both the state and individuals must ensure that they treat each person as the equal of every other person, in the sense that they do not unfairly subordinate some to others, or infringe their right to deliberative freedom, or deny them access to basic goods. But the state's duty is broader than this: it must also create the background conditions necessary for us to relate to each other as equals in our personal lives, not just by enacting and enforcing antidiscrimination laws where necessary, but also through the kinds of indirect legal and political measures that I discussed at the end of Section 7.5, such as creating an education system that fosters tolerance and celebrates diversity, and creating public spaces in which people can come together and pursue important goals together. The content of our duty, as individuals, depends on how well state fulfils its duty. If it sets up many of the necessary background conditions for us to relate to each other as equals, then each of us will need to do less, on our own, to ensure that others are not unfairly subordinated, or denied deliberative freedoms that they have a right to, or denied access to a basic good. However, if the state

does too little to establish the necessary legal rules, the relevant political and social institutions, and the right sorts of social expectations, then our duty as individuals becomes more demanding, and we must, through our own individual and collective actions, try to do our best to fill the gaps.

Secondly, the duty of the state to treat others as equals and the duty of each of us, as individuals, differ in what we might call their concomitant duties. The state also has a duty to make transparent the ways in which it is treating us as equals, and to announce that it is doing so --and perhaps to compensate those whom it cannot treat as equals, due to exceptional circumstances. As private individuals, we do not usually have a duty to announce each time we are treating others as equals; although, when we occupy certain institutional roles, the demands of that role, and the kind of authority we exercise over others, may require us to announce that we are treating others as equals and to be more transparent about the way that we do so. As a professor, it is important that I not only treat my students as equals but let them know, in various ways, that and how I am doing so; the same is true of private clubs, when they run sports competitions. So the extent of the concomitant duties depends on the particular institutional role we occupy and the expectations attendant on it.

Thirdly, there are important differences in the kinds of factors that can occasionally justify the state or an individual in continuing to engage in wrongful discrimination. I have, throughout the book, left open the possibility that in certain special cases, we may wrong someone by wrongfully discriminating against them, and yet our act may nevertheless be justified all things considered. While I have not had the space in this book to elaborate a detailed theory of justification, I suggested earlier in this chapter that certain factors may be relevant in the state's justification of ongoing wrongful discrimination. I argued there that the state's duty to treat others as equals is a constitutive duty --that is, it is part of the very purpose of the state to treat others as equals and to create the conditions under which we are able to treat others as equals. And I suggested that, because it is a constitutive duty, breaches of this duty can only be justified by the need to fulfil some *other* constitutive duty. So the fact that a majority might wish not to treat others as equals, or might have some shared preference that can only be satisfied if one group is unfairly subordinated, is not a sufficient reason. But if the government is facing an emergency, and must take certain measures to protect the health of part of the population, and a necessary side-effect of these measures is that certain other people are denied a basic good, then the fact that protecting the health of the population is also a constitutive duty may justify the government in wrongfully discriminating against this other group. What in particular count as "constitutive duties," other than the state's duty to treat others as equals and the state's duty to protect the health of its members, and how tight the connection must be between the measures necessary to fulfil these other duties and the acts that wrongfully discriminate, are large questions, which I shall not pursue further here; but they would need to be answered as part of a complete theory of what justifies the state, occasionally, in continuing to engage in wrongful discrimination.

By contrast, the duty to treat others as equals is not, in the same way, a constitutive duty for individuals –though I have tried to argue that it is a deep commitment of ours, as members of democratic societies. But there are nevertheless circumstances in which we, too, may be justified in wronging others. One such set of circumstances I discussed in Chapter Five, when I considered cases in which we wrongfully discriminate against someone no matter what we do: if we do not adopt a particular policy, we unfairly subordinate a particular group, but when we do adopt that policy, certain members of that group are denied a deliberative freedom to which they had a right –or perhaps certain other people are denied access to a basic good. But there may also be other circumstances. Most antidiscrimination laws allow that something analogous to what Canadian laws call “undue hardship” on the part of the discriminator can sometimes justify what would otherwise amount to impermissible wrongful discrimination. Canadian law employs a particularly demanding interpretation of undue hardship, allowing that only considerations of health and safety, and not a mere loss of profits or minor inconvenience, can count as “undue hardship.” I shall not take a stand here on the question of how we should best interpret this idea, other than to note that what the correct interpretation is for a particular country –that is, which considerations we allow individuals to invoke, to justify continuing to engage in wrongful discrimination—may depend on how well the state is fulfilling its duty to treat others as equals and to create the background conditions for equal social relations. Perhaps, if the state is doing a great deal, then there may be more space for individuals to appeal to the importance of their various projects and activities as a justification for continuing to engage in wrongful discrimination; whereas, if the state is doing very little, our duty may be more stringent, and there may be less room for us to claim undue hardship. But of course we need to be careful here: too generous an interpretation of undue hardship will risk draining the duty to treat others as equals of much of its content.

Even when the state or an individual is justified in continuing to engage in wrongful discrimination, however, it is important to note that the discrimination is still, on my view, *wrongful*. That is, it wrongs someone, even when it is all things considered justified. I think it is important that my view gives us the resources to acknowledge this. The fact that some other constitutive duty of the state must be pursued, or that the individual discriminator faces undue hardship if they abandon the practice that causes wrongful discrimination, does not eliminate the wrong to the discriminatee. My view gives us a way of acknowledging this. We can say that in such cases, the discrimination is all things considered justified, but nevertheless, wrongful to a particular person or group.

Conclusion

In this book, I have laid out and defended a pluralist theory of when and why discrimination wrongs people. I started from actual legal cases, in which claimants have alleged wrongful discrimination by other people or by the state. I suggested that we can understand these people's complaints best by thinking of them as complaints about different ways in which they have not been treated as equals in their societies –in particular, through unfair subordination, through the violation of their right to a particular deliberative freedom, or through the denial to them of access to a basic good, that is, a good access to which is necessary if they are to be, and to be seen as, an equal in their society. I argued that each of these wrongs is distinctive, but that they are all ways of failing to treat others as equals. And I tried to show that both the state and we as individuals have a duty to treat others as equals, in these three specific senses.

This pluralist theory of wrongful discrimination has a number of advantages, which I have also tried to draw out. Rather than treating only one of the many harms resulting from discrimination as the source of its wrongness, my theory suggests that different features of discriminatory acts and practices are relevant to different wrongs. My theory thereby enables us to explain and validate many claimants' thoughts about the specific ways in which they have been wronged, and it offers us a rich and nuanced way of understanding what goes wrong in different cases of discrimination. Moreover, the fact that there are many reasons why discrimination can be wrong helps us also to understand our ambivalence about certain special cases of wrongful discrimination: in some cases of affirmative action, and in cases such Wackenheim's challenge to the ban on dwarf-tossing, it seems as though we wrong someone no matter what we do. My theory can explain why this is so: these are cases, I have suggested, in which, if we adopt a certain policy, we discriminate wrongfully against one individual or group, but if we do not adopt that policy, we risk wrongfully discriminating in a different way against another individual or group. I also tried to show in Chapter Five that my theory provides us with the resources to address a number of puzzles that have beset theories of discrimination –puzzles about the comparative nature of claims of wrongful discrimination, and about whether the wrong in question is a personal wrong or a group wrong. Lastly, as I tried to show in Chapter Six, the theory paints a compelling picture of why indirect discrimination is wrongful, and it gives us the resources to explain why it is often just as wrongful as direct discrimination, and to see both as forms of negligence.

If my theory is correct, then there are a number of questions we need to think further about. I raised these questions in earlier chapters and offered some thoughts about them; but I have not tried to give complete answers to them in this book. They are

questions for future study. For instance, my theory leaves open the possibility that there may be other ways in which discrimination wrongs us by failing to treat others as equals – though, as I have suggested, I think that the three ways that I have discussed in this book are among the most important. Further work is needed to think through other ways in which discrimination wrongs us by failing to treat us as the equal of others. I have also suggested that, although practices that amount to wrongful discrimination are most often, for this reason, wrong all things considered, there are nevertheless certain special situations in which either the state or an individual can wrong someone through discrimination, and yet be justified in doing so, all things considered. I offered some thoughts in the later chapters of the book about what might count as relevant justifying factors for the state and for individuals, and about the differences between the factors that justify the state in continuing to engage in wrongful discrimination, and the factors that justify individuals in continuing to engage in it. But there is room for further work on which factors exactly these are. And lastly, and perhaps most importantly, I suggested in Chapter Seven that we need to think at much greater length about the ways in which the state can support individuals, in discharging their duty to treat others as equals. Anti-discrimination law, though important, is only one of these ways –and, as I argued in Chapter Seven, it is not always helpful or appropriate for the state to intervene directly to ensure that we comply with our moral duty to treat others as equals. But there is nevertheless a great deal that the state can do to help create the conditions under which we are able to relate to others as equals, and a great deal more academic work that needs to be done, in thinking through other ways in which the state can help, outside of anti-discrimination law.

Most of the ideas in this book have been presented through philosophical arguments. But of course, the arguments began as attempts to make sense of the complaints of people who have suffered from discrimination that they believe is wrongful – people such as Dutee Chand, Manuel Wackenheim, people affected by the water crisis on indigenous reserves in Canada, and many others. So, in a sense, this is a book about their stories. It seems fitting, then, for the book to end by relating their stories to another set of stories involving wrongful discrimination - the stories behind the faces in Robert Davidson's serigraph on the cover of this book.

These faces are quite literally “faces of inequality.” They are adaptations of traditional Haida depictions of characters in their legends, drawn by the Canadian Haida artist Robert Davidson. For many years after colonization, the Haida people faced systemic discrimination. Their lands were taken from them; their children were sent away to residential schools where they could not speak their language or learn their stories; and many of the practices and rituals that were integral to their culture were made illegal, including those that kept alive the characters depicted in this serigraph. As a result, many of the stories associated with these faces were lost. This is particularly true of Mouse Woman, *kuugan jaad*, whose features appear in some of these faces. Her history is, like the history of indigenous peoples in Canada, a history of second-class citizenship, and of loss.

And yet, these faces also carry a message of hope. The serigraph is entitled “I am You and You Are Me,” and the artist, Robert Davidson, has said that this title is based on the Haida saying: “I am you, that is also you.”³¹ This means, among other things, that there are echoes of each of us in every other person, and that we must therefore be careful of how we treat others. But if it is true that there are echoes of each of us in every other person, then it must also be true that we are capable of understanding each other, capable of working towards a society in which no one is a second class citizen, and in which no one’s stories are left to be forgotten. And indeed, the few stories of Mouse Woman that have survived tell us that she is a guide who leads people through transformations, and that she helps to restore equality between beings. So, in addition to telling a story about loss, the serigraph also symbolizes this hope for our future. The faces in the two circles are not exact reflections of each other. They are different, but also related. They are each a ‘you,’ that isn’t exactly the same as the other ‘you,’ but is “also you.” They point, not just backwards, to stories of loss and disenfranchisement, but forwards, towards a possible future in which different faces, with different colours and different backgrounds, can stand together on the same page, as equals.

Whether this will one day come to pass –whether our future will be a story of treating others as equals, or a continuation of our past, where some are treated as inferiors-- is up to us.

³¹ Davidson has written that: “This print is based on a line in an oral history of Raven travelling: “I am you, that is also you”, and the wisdom in Haida culture that when you point out a fault in someone, you’re actually pointing out the fault in yourself. The two central design elements are reflections of each other, but they’re actually variations of each other. The design echoes the designs found in the ends of bent-wood chests and boxes, where the design elements are usually “non-conforming” and do not follow the conventional symmetrical nature of designs found on the fronts and backs. There is not a single being depicted in the print, but *kuugaan jaad* (Mouse Woman) is prominent, as she is in bentwood boxes and chests.” See http://www.spiritwrestler.com/catalog/index.php?products_id=3494.