Kasper Lippert-Rasmussen

BORN FREE AND EQUAL?

A philosophical inquiry into the nature of discrimination



Born Free and Equal?

This page intentionally left blank

Born Free and Equal?

A PHILOSOPHICAL INQUIRY INTO THE NATURE OF DISCRIMINATION

Kasper Lippert-Rasmussen

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide.

Oxford New York
Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in
Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trade mark of Oxford University Press
in the UK and certain other countries.

Published in the United States of America by
Oxford University Press
198 Madison Avenue, New York, NY 10016

© Oxford University Press 2014

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, by license, or under terms agreed with the
appropriate reproduction rights organization. Inquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above.

You must not circulate this work in any other form
and you must impose this same condition on any acquirer.

Library of Congress Cataloging-in-Publication Data
Lippert-Rasmussen, Kasper, 1964–
Born free and equal? : a philosophical inquiry into the nature
of discrimination / Kasper Lippert-Rasmussen.
p. cm.
Includes bibliographical references and index.
ISBN 978–0–19–979611–3 (hardback : alk. paper) 1. Discrimination—Philosophy.
2. Discrimination—Moral and ethical aspects. 3. Equality—Philosophy. I. Title.
HM821.L57 2014
305—dc23
2013014484
9780199796113

1 3 5 7 9 8 6 4 2

Printed in the United States of America on acid-free paper

{ CONTENTS }

Acknowledgments ix
Sources xi

Introduction **1**
1. The Questions 1
2. The Approach 4
3. Overview of the Book 6

PART I **The Concept of Discrimination**

1. What Is Discrimination? **13**
1. Introduction 13
2. Discrimination in the Generic Sense 14
3. Irrelevance Discrimination 22
4. The Moralized Concept of Discrimination 24
5. Group Discrimination 26
6. Social Salience 30
7. Because 36
8. Treatment 40
9. Summary 45

Appendix 1: Methodology 47
Appendix 2: Discrimination Skeptics: Oppression and Dominance 49

2. Indirect Discrimination **54**
1. The Distinction between Direct and Indirect Discrimination 54
2. Altman's Definition 56
3. The No-Intention Condition 59
4. The Disadvantage Condition 61
5. The Disproportionateness Condition 65
6. Sufficient for Indirect Discrimination? 68
7. Direct vs. Indirect Discrimination 72
8. Conclusion 73

Appendix 1: Some Other Definitions of Indirect Discrimination 74
Appendix 2: Institutional and Structural Discrimination 77

3. Statistical Discrimination **79**
1. Introduction 79
2. Statistical Discrimination vs. Nonstatistical Discrimination 81
3. Direct vs. Indirect, Statistical Discrimination 87

4. What Statistical Discrimination Is Not 89
5. Conclusion 94
Appendix: Genetic Discrimination and Social Salience 95

PART II **The Wrongness of Discrimination**

4. Mental-State-Based Accounts **103**
1. Introduction 103
2. Some Common Accounts 105
3. Mental States and Permissibility 111
4. Different Mental-State Accounts 113
5. Alexander on Disrespect and Discrimination: The Falsehood Account 113
6. Alexander on Disrespect and Discrimination: The Comparative Falsehood Account 121
7. Alexander on Disrespect and Discrimination: The Irrational Comparative Falsehood Account 122
8. Conclusion 127

5. Objective-Meaning Accounts **129**
1. Introduction 129
2. Hellman's Account: Demeaning Others 131
3. Some Challenges to Hellman's Account 134
4. Scanlon on Racial Discrimination and the Meaning of Actions 139
5. An Important Ambiguity 144
6. Some Worries About Scanlon's Account 146
7. The Moral Distinctiveness of Discrimination Based on Judgments of Inferiority 149
8. Conclusion 151

6. Harm-Based Accounts **153**
1. Introduction 153
2. The Essentials of the Harm-Based Account 154
3. The Baseline Issue 157
4. The Metric of Harm 160
5. Some Challenges to the Harm-Based Account 163
6. A Desert-Prioritarian Account 165
7. Some Objections 170
8. A Test Case: Moral Wrongness of Indirect Discrimination 177
9. Conclusion 183
Appendix: Moreau on Deliberative Freedom and Discrimination 185

PART III **Neutralizing Discrimination**

7. Discrimination and the Aim of Proportional Representation **193**
1. Introduction 193
2. The Simple View and Ambition-Sensitivity 196
3. The Counterfactual, Holistic View 202
4. Which Counterfactual Scenario? 204

5. Is Absence of Discrimination Necessary for Suitable Representation? 211
6. Second-Best Representational Aims 213
7. Conclusion 215

8. Discrimination in Punishment **217**
1. Introduction 217
2. Loci of Legal Discrimination 219
3. Criteria vs. Indicators of Discrimination 221
4. The Pure Discrimination Case 224
5. The No-Complaint Argument 231
6. Conclusion 234

9. Reaction Qualifications **235**
1. Introduction 235
2. Discounting Qualifications Based on Illegitimate Preferences 240
3. Refining Meritocracy 249
4. Illegitimate Preferences Not Disadvantaging Targeted Groups 253
5. Respect and Reaction Qualifications 257
6. Conclusion 260

10. Discrimination in the Private Sphere **261**
1. Introduction 261
2. A Legal Duty to Engage in Wrongful Private Discrimination 266
3. A Legal Right to Engage in Wrongful Private Discrimination 266
4. A Legal Duty Not to Engage in Wrongful Private Discrimination 267
5. A Legal Duty or Permission to Engage in Private Discrimination That Is Not Wrongful 268
6. A Legal Duty Not to Engage in Private Discrimination That Is Not Wrongful 269
7. Conclusion 269

11. Racial Profiling **272**
1. Introduction 272
2. A Right to Be Treated as an Individual 275
3. Unequal Treatment 278
4. Unfairness 279
5. The Making of Statistical Facts and the Justifiability of Statistical Discrimination 283
6. Putting the Argument to the Interpersonal Test 285
7. Noncomprehensively Justified? 293
8. Challenges 297
9. Conclusion 299

Bibliography 303
Index 311

This page intentionally left blank

{ ACKNOWLEDGMENTS }

Over the years, many people have commented on (predecessors to) different parts of this book, all parts of which at some stage have been commented on by several friends and colleagues. In particular I would like to thank the following persons for helpful suggestions and criticisms: Larry Alexander, Martin Marchman Andersen, Elizabeth Anderson, Richard Arneson, Gustaf Arrhenius, John Baker, Linda Barclay, Anders Berg Sørensen, Theresé Björkholm, Lene Bomann-Larsen, Alexander W. Cappelen, Paula Casal, Matthew Clayton, Mathew Coakley, Jakob Elster, David Estlund, Nir Eyal, Eli Feiring, Katrin Flickshuh, Mikkel Gerken, Bob Goodin, Gina Gustavsson, Edward Hall, Lena Halldenius, Allan Dreyer Hansen, Rasmus Sommer Hansen, Dan Hausman, Heine A. Holmen, Nils Holtug, Brad Hooker, Robert Huseby, Magnus Jiborn, Pia Justesen, Signe Kjær Jørgensen, Chandran Kukathas, Sigurd Lauridsen, Carsten Bagge Laustsen, Annabelle Lever, Jason Lindsey, Christian List, Mats Lundström, Sune Lægaard, Raino Malnes, Andrew Mason, Jeff McMahan, Søren Flinch Midtgaard, David Miller, Morten Ebbe Juul Nielsen, Henrik Palmer Olsen, Michael Otsuka, Ingmar Persson, Anne Phillips, Paolo Picquadio, Roland Pierik, Morten Raffnsøe-Møller, Soran Reader, Jeffrey H. Reiman, Jouni Reinikainen, Mathias Risse, Paul Robinson, Hans Ingvar Roth, Joakim Sandberg, Jens Saugstad, Robin May Schott, George Sher, Kai Spiekermann, Marie Tandrup, Laurence Thomas, Frej Klem Thomsen, Dag Einar Thorsen, Kristian H. Toft, Torbjörn Tännsjö, Laura Valentini, Peter Vallentyne, Robert van der Veen, Philippe van Parijs, Kira Vrist Rønn, Nicholas Vrousalis, Leif Wenar, Alan Wertheimer, Jonathan Wolff, Lea Ypi, and two anonymous Oxford University Press readers. I am particularly indebted to Dick Arneson and Shlomi Segall for detailed, critical comments on a previous version of the entire manuscript. Also, I am very grateful for having had the opportunity to present different parts of the manuscript at various places: London School of Economics, Université catholique Louvain, University of Aarhus, University College London, University of Copenhagen, University of Hull, University of Manchester, University of San Diego, and at various meetings at the Nordic Network for Political Theory. I thank the Danish Research Council (FSE) for financial support for completing this monograph. Finally, I dedicate this book to my children—Cæcilie, Hannah, and William—and to Kira. I hope there are many more chapters to come.

This page intentionally left blank

{ SOURCES }

This book grows out of several articles etc. on discrimination that I have written over the years:

1. Chapter 1: "Discrimination: What Is It and What Makes It Morally Wrong" (2007), in Thomas Petersen, Jesper Ryberg, and Clark Wolf (eds.), *New Waves in Philosophy: Applied Ethics* (Basingstoke: Palgrave Macmillan), pp. 51–72. Reproduced with the permission of Palgrave Macmillan: http://us.macmillan.com/newwavesinappliedethics/JesperRyberg.
2. Chapters 1 and 6: With kind permission from Springer Science+Business Media: "The Badness of Discrimination" (2006), *Ethical Theory and Moral Practice* 9.2, 167–188.
3. Chapters 3 and 11: "Nothing Personal: On Statistical Discrimination" (2007), *Journal of Political Philosophy* (John Wiley and Sons), 15.4, pp. 385–403.
4. Chapters 4 and 5: "Equality and Discrimination" (2012), in Andrei Marmor (ed.), *Routledge Companion to Philosophy of Law* (London: Routledge), pp. 569–583: http://www.routledge.com/books/details/9780415878180/.
5. Chapter 5: "Intentions and Discrimination in Hiring" (2012), *Journal of Moral Theory* (Brill) 9.1, pp. 55–74.
6. Chapters 6 and 10: "Private Discrimination: A Prioritarian, Desert-Accommodating Account" (2007), *San Diego Law Review* 43, pp. 817–856.
7. Chapter 7: "Discrimination and the Aim of Proportional Representation" (2008), *Politics, Philosophy, and Economics* (SAGE), 7.2, pp. 159–182.
8. Chapter 8: "Punishment and Discrimination" (2010) in Angelo J. Corlett and Jesper Ryberg (eds.), *Punishment and Ethics: New Perspectives* (London: Palgrave Macmillan), pp. 169–188. Reproduced with the permission of Palgrave Macmillan: http://us.macmillan.com/punishmentandethics/JesperRyberg.
9. Chapter 9: "Reaction Qualifications Revisited" (2009), *Social Theory and Practice* 35.3, pp. 413–439.

10. Chapter 10: With kind permission from Springer Science+Business Media: "'We are all Different': Statistical Discrimination and the Right to be Treated as an Individual" (2011), *Journal of Ethics* 15.1–2, pp. 47–59.
11. Chapter 11: "Racial Profiling versus Community" (2006), *Journal of Applied Philosophy* (John Wiley and Sons) 23.2, pp. 191–205.

I thank the above-mentioned journals and publishers for permission to reuse materials published in their journals or books.

Introduction

1. The Questions

What is discrimination? There are certain instances of differential treatment that almost anyone would describe as discriminatory. Among such paradigmatic cases are the denial of women's right to vote and the way in which nonwhite South Africans were treated under apartheid. Yet, when one looks deeper into the issue it seems that such near-unanimity about certain core cases is simply the tip of an iceberg of disagreement or difference. There are many nonparadigmatic cases in which it is unclear whether "differential treatment" is, in fact, discriminatory. For instance, is it discrimination against smokers when hospitals hire nonsmokers only?;[1] religious discrimination when churches and religious organizations reject applicants belonging to a different faith for a job as a secretary or, for that matter, as a priest?;[2] discrimination against atheists when people who refuse a job at a munitions factory on religious grounds are entitled to employment benefits, whereas people who do so on secular, pacifist grounds are not?;[3] sex discrimination when swinger clubs charge male, unlike female, members a fee, or when pensioners are admitted to a council's swimming pool at a concessionary rate and where, due to gender differences in the state pensionable age, a 61-year-old man must pay the full rate, while his similarly aged wife may pay the concessionary rate?;[4] age discrimination when people who are older than 65 can be excluded from unemployment insurance

[1] http://www.nytimes.com/2011/02/11/us/11smoking.html?_r=0 (accessed October 17, 2012).

[2] http://www.ecfa.org/Content/TopicReligiousDiscrim (accessed October 17, 2012).

[3] Christopher L. Eisgruber and Lawrence G. Sager (1994), "The Vulnerability of Conscience: The Religious Basis for Protecting Religious Conduct," *University of Chicago Law Review* 61.4, 1245–1315, p. 1294.

[4] http://www.dr.dk/Nyheder/Indland/2012/02/17/133930.htm (accessed October 17, 2012); *James v. Eastleigh Borough Council,* described in Hugh Collins (2003), "Discrimination, Equality and Social Inclusion," *Modern Law Review*, 66.1: 16–43, p. 35.

schemes simply on the basis of their age?;[5] or racial discrimination against whites when, in view of the fact that only 600 out of 550,000 of its theater-goers in the 2005/2006 season were Danish citizens with a non-Danish ethnicity, the Royal Danish Theatre planned to sell tickets at significantly reduced prices to Danish citizens of non-Danish ethnicity?[6] Ideally, we would like to have an account of what discrimination is that enables us to distinguish between discriminatory and nondiscriminatory treatment, and to explain why we consider cases such as those just mentioned one or the other.[7]

Not only do people differ over which cases of differential treatment they see as discriminatory, they also disagree about when discrimination is morally wrong and about what makes discrimination morally wrong.[8] Some believe that the moral objection to discrimination is that it undermines equality of opportunity. Others believe that discrimination is wrong because it is disrespectful. And yet others believe that the wrongness of discrimination is to be explained by the harm it does to the victims of discrimination. These are different accounts—presumably not all forms of inequality involve disrespect of or harm to those with worse opportunities, and some forms of disrespect may not harm the person who is disrespected; and while paradigmatic forms of discrimination are wrong according to all three accounts, there are other cases that are not. For instance, racial discrimination against members of a well-off, racial minority may well reduce inequality of opportunity, while still being disrespectful of and harmful to members of this minority.[9] About such cases, friends of the three accounts of the wrongness of discrimination mentioned may disagree. Moreover, there are cases of differential treatment that arguably are discriminatory but not clearly wrong, perhaps even morally required. Some think that giving priority to young people in the distribution of scarce organs

[5] http://bm.dk/da/Aktuelt/Nyheder/Arkiv/2012/10/Dom%20i%20sag%20om%20aldersdiskrimination.aspx (accessed October 17, 2012).

[6] http://avisen.dk/ulovlig-indvandrerrabat-i-det-kongelige-teater_86357.aspx (accessed October 17, 2012). Slightly more than 6% of the Danish population has a non-Danish, non-Western ethnic background.

[7] In the sense of "discrimination," which is introduced in chapter 1 ("group discrimination") and employed in most of this book, all but the first case (or, possibly, the second example) are cases of discrimination, albeit some of them are cases of nonwrongful discrimination.

[8] By saying that an act is morally wrong in virtue of a certain feature, I mean that this feature counts toward its being morally impermissible or, to put it differently, one that, morally speaking, ought not to be done, all things considered. Hence, as I use the term in this book, it does not follow from the fact that an act is wrong in virtue of, say, harming the discriminatee that it is morally impermissible, since it may have other right-making features, e.g., that it benefits many others, that, morally speaking, outweigh its wrong-making features. Also, in my sense, an act can be more or less wrong because the act's wrong-making features can count more or less strongly toward its being an act that ought not to be done, all things considered: cf. Jeff McMahan (2002), *The Ethics of Killing: Problems at the Margins of Life* (Oxford: Oxford University Press), p. 190.

[9] Thomas Nagel (1997), "Justice and Nature," *Oxford Journal of Legal Studies* 17.2, 303–321, p. 310.

available for transplantation amounts to ageist discrimination.[10] Yet, it is unclear whether this form of ageist discrimination, assuming it is one, is wrong, given that old patients have enjoyed more life than young ones. Giving priority to younger patients might be what justice requires.[11]

Finally, some kinds of discrimination are morally wrong and very much so. This raises the question of what should be done about it. Again there is considerable disagreement. Some believe that, in the case of many discriminatory acts, even though they may be morally wrong, there is little the state can permissibly do to prevent them, e.g., because much discrimination takes place when people exercise their rights in a way that may be morally wrong, but in contexts in which others do not have the right to force them not to. For instance, it may be morally wrong for people to engage in racial discrimination in their love life, but impermissible for the state to interfere with citizens' love lives. Others believe that the state can permissibly go to great lengths in trying to eliminate discrimination or to counteract its effects. For instance, it may permissibly encourage interracial marriages by making sure that people mix across races. Moreover, in some areas, state agencies pursue policies that are widely seen as discriminatory, e.g., racial profiling, in pursuit of aims that are widely seen as worthy and legitimate, e.g., preventing crime. Understandably, such policies are controversial. Finally, since people disagree about which kinds of wrongful discrimination exist— for instance, some will reject the view that one can engage in wrongful racial discrimination in one's love life—they are also bound to disagree about which policies will actually reduce the amount of discrimination present in society.

This book addresses these three issues, to wit: What is discrimination?; What makes it wrong?; and What should be done about wrongful discrimination?[12] First, it argues that there are different concepts of discrimination and that different purposes pertaining to different contexts determine which one is the most useful. It gives special attention to a concept of discrimination that ties discrimination to differential treatment of people on the basis of their membership in socially salient groups. Second, it argues that discrimination is wrong first and foremost because of its harmful effects, and not because of the discriminator's disrespectful mental state or because discrimination is objectively demeaning. Third, it takes issue with some of the standard methods used to counteract discrimination, e.g., proportional representation of different groups, and it submits that combating discrimination requires more than state actions.

---

[10] John Harris (1987), "QALYfying the Value of Life," *Journal of Medical Ethics* 13.3, 117–123.

[11] Klemens Kappel and Peter Sandøe (1992), "QALYs, Age and Fairness," *Bioethics* 6.4, 297–316.

[12] For a brief sketch of some of the issues related to these three main questions, see Kasper Lippert-Rasmussen (2013), "Discrimination," in Hugh LaFolette (ed.), *International Encyclopedia of Ethics* (Oxford: Wiley-Blackwell), 1405–1415.

2. The Approach

As Richard Epstein observes: "A broad principle of antidiscrimination lies at the core of American political and intellectual understandings of a just and proper society."[13] A similar claim is true of many other states. Yet, not much philosophical literature has been written on the concept and morality of discrimination. The 1970s saw a surge in writings on affirmative action, or, as it was also called then, "positive" or "reverse discrimination," but with few exceptions conceptual and moral issues in relation to discrimination were mostly mentioned in passing. This state of relative neglect is regrettable. Recently, the situation has begun to change. Seminal here is Larry Alexander's rich 1992 article—"What Makes Wrongful Discrimination Wrong?"—and Deborah Hellman's 2008 book—*When Is Discrimination Wrong?*[14] Philosophy offers valuable tools to enhance our understanding of the nature and morality of discrimination, and it is puzzling why more philosophers have not attended to the topic. Other topics in applied ethics, e.g., abortion, have certainly received much more attention than discrimination despite the fact that discrimination hardly has a less prominent place in ordinary moral and political thinking than abortion. As this example also shows, it cannot be the fact that discrimination is an emotionally and politically charged topic that has kept philosophers from addressing it to any great extent.

In this book I treat discrimination in the way that analytic philosophers treat any other topic. For instance, I draw sharp distinctions; I separate various objectionable features of cases of discrimination; and to explore the limits of the concept of discrimination I sketch hypothetical cases of discrimination that might appear very different from the kinds of cases people typically have in mind when they object to discrimination. Comparing simple, carefully crafted, sometimes extraordinary hypothetical cases that only vary in terms of one particular factor whose moral relevance that we want to test is useful. This is because any moral difference between the two cases must derive from the variation in the relevant moral factor, and because moral equivalence between the two cases suggests that variation in the relevant factor is morally irrelevant. We can hardly ever draw such conclusions when comparing actual cases, because they are bound to differ in terms of a lot of different factors, thus preventing us from isolating the influence of any one of them. Also, hypothetical examples are useful in testing universal, unconditional moral principles—e.g., "Killing

[13] Richard A. Epstein (1995), *Forbidden Grounds: The Case against Employment Discrimination Laws* (Cambridge, MA: Harvard University Press), p. 1. Epstein finds the centrality in American understandings of a just and proper society that the antidiscrimination principle has, or at least its result in the form of antidiscrimination legislation, regrettable.

[14] Larry Alexander (1992), "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies," *University of Pennsylvania Law Review* 141.1, 149–219; Deborah Hellman (2008), *When is Discrimination Wrong?* (Cambridge, MA: Harvard University Press).

the innocent is always wrong" or "All forms of discrimination are morally wrong"—because these are often adhered to on the basis of a rather narrow range of typical cases, and constructing atypical cases forces us to reflect on whether we really endorse the implications of the relevant moral principles in the full range of cases to which they apply. When presenting papers on discrimination, I have sometimes encountered the reaction that this approach takes us too far away from real-life discrimination, and that doing what I propose to do here is in some way not taking discrimination seriously, or that it manifests insensitivity to the plight of those who suffer from discrimination.

Admittedly, in one sense there is a great distance between the present inquiry and the pressing issues of real-life discrimination. So in response to the reaction just described let me offer the following rejoinders. First and meekly: I have partners in guilt. Worries of the present kind should be had about many issues in political philosophy, if they apply to discrimination.[15] Consider the following case much discussed in the "Equality of What?" debate—Ronald Dworkin's case of someone who has an expensive taste for pre-phylloxera claret and plovers' eggs.[16] Dworkin introduces this case to cast light on what egalitarian justice requires that individuals have equal amounts of at the level of fundamental principles of justice. The example is meant to show that welfare cannot be the relevant metric, because if it were, justice would require, implausibly, that persons who cultivate the relevant expensive tastes should be provided with more resources than others to prevent them from having less welfare due to the frustration of their expensive tastes. This case will attract little attention from egalitarian politicians, and understandably so. From the point of view of activists concerned with real-life inequalities, there are other inequalities that need to be addressed much more urgently.[17] Still, Dworkin's example of expensive tastes legitimately plays a large role in theoretical discussions of the nature of equality. In my view, the existence of such a gap between political philosophy and political reality raises no real objection to the former.

---

[15] In fact, many philosophers outside the branch of political philosophy are partners in guilt. For instance, when exploring what matters in survival, philosophers imagine sci-fi, split-brain cases, where the connection between the two hemispheres of a person's brain is cut. One hemisphere is then transplanted into another body with an empty skull, resulting in two persons coming into existence, both of whom are psychologically connected and continuous with the first person, whose brain has been divided; e.g., Derek Parfit (1984), *Reasons and Persons* (Oxford: Clarendon Press), pp. 253–261. Many philosophers believe that such cases show that identity is not what matters in survival. While the option of having one's brain split may not merit a significant place in a person's assessment of her situation from the point of view of prudence, it may nevertheless yield significant insights into what does merit such a concern.

[16] Ronald Dworkin (2000), *Sovereign Virtue* (Cambridge, MA: Harvard University Press), p. 49.

[17] Cf. Elizabeth Anderson (1999), "What Is the Point of Equality?" *Ethics* 109.2, 287–337, p. 287. She complains, polemically, that had neoconservatives penned recent luck-egalitarian writings, the result could not have been more damaging.

Second, there is a division of intellectual labor. I write this book not as an activist engaged in the struggle against discrimination, but as a philosopher trying to understand the phenomenon and its moral qualities. Generally speaking, philosophers are not well qualified to do the work of politicians or activists (and I for one certainly am not.) Hence, when I consider, for instance, whether people can discriminate against members of their own group, I do this to clarify the concept of discrimination, not to propose that self-discrimination, as it were, should be a significant part of any political agenda concerning discrimination.

Finally, generally speaking, a division of labor makes participants therein better off, and something similar is true about the present inquiry in relation to antidiscrimination activism. As noted, it is often unclear when a case of differential treatment constitutes discrimination as well as what makes it morally wrong. In such cases, activism might be assisted by philosophical reflection on the nature and morality of discrimination. Accordingly, the philosopher's attention to hypothetical, and even weird hypothetical cases of discrimination, does not suggest that the philosopher has no interest in saying something that speaks to any concerns that people might have about real-life issues. As the so-called "Equality of What?" debate within egalitarianism shows, careful consideration of contrived and marginal cases—here of discrimination—can illuminate and help us to sensibly assess actual cases of inequality and discrimination. In particular, they can help us to identify which features of discrimination are morally objectionable. The philosophy of discrimination does have a contribution to make to antidiscrimination activism, as I hope to be able to show as we go along, despite its indirect nature.

3. Overview of the Book

This book is divided into three parts. Part 1 analyzes the concept of discrimination. Chapter 1 defines discrimination in the generic sense and then distinguishes between three different and more specific senses of discrimination, all of which are prima facie, or even necessarily, morally wrong kinds of differential treatment. One of these senses—group discrimination, i.e., treating people differently on the basis of their membership in different socially salient groups—is then explained in greater detail, as this is the sense of discrimination that I employ in the rest of the book.[18]

[18] Throughout this book I shall use the phrase "membership of a socially salient group" (see chapter 1). Obviously, membership in a racial, religious, or gender group is very different from membership in a political party, a family, or a study group. When I use the term I do not mean to suggest that a member of a socially salient group necessarily identifies with his or her group, or that much, if anything, about him or her can be explained through membership of the relevant group.

Chapter 2 explores the distinction between direct and indirect discrimination. The chapter takes its point of departure from a definition proposed by Andrew Altman and uses it as a springboard for a better one, which, among other things, does not involve the requirement that indirectly discriminatory acts have effects on discriminatees, i.e., those who are being discriminated against, that are in some sense disproportionate. The chapter also seeks to distinguish between inequalities that reflect indirect discrimination and inequalities that do not. It does so by tying indirect discrimination to the perpetuation of disadvantages resulting from past direct discrimination. More generally, I argue that indirect discrimination is asymmetrically parasitic on direct discrimination in that the former requires the (past) presence of the latter, but not the other way round.

Chapter 3 defines statistical discrimination and pries it apart from nonstatistical discrimination. It then distinguishes between direct and indirect statistical discrimination, arguing that the latter is a real possibility. Statistical discrimination often appears together with other morally objectionable features, which the chapter tries to isolate from "pure" statistical discrimination. This exercise is crucial to a later assessment of why statistical discrimination is wrong when it is.

Whereas part 1 concentrates on the concept of discrimination and its various subcategories, part 2 explores what makes discrimination morally wrong when it is. Chapter 4 focuses on mental-state accounts of the moral wrongness of discrimination, i.e., accounts according to which the wrongness of discrimination derives from the fact that the discriminator has an objectionable mental state. It takes a close look at Larry Alexander's influential account, according to which certain forms of discrimination are wrong for noninstrumental reasons, because they reflect that the discriminator falsely believes that the discriminatee has a lower moral status. Various challenges to this account—and to mental state based accounts in general—are presented.

Chapter 5 addresses accounts that locate the wrongness of discriminatory acts in the objective meaning of acts of discrimination. In particular, it focuses on a recent account by Deborah Hellman, according to which the distinctive wrongness of discriminatory acts lies in their being objectively demeaning whereby they violate a bedrock principle of equal human worth, and one by Thomas Scanlon, according to which acts of racial discrimination are wrong in part because of the offensive meaning discriminatees can reasonably ascribe to them. The chapter first presents a number of challenges to Hellman's account, arguing that, at best, it explains the wrongness of a very narrow range of cases of discrimination. Even if we take a broad view of which acts demean others, many kinds of discrimination do not demean. The chapter then takes a critical look at Scanlon's account, arguing that it is ambiguous on whether the offensive meaning of acts of racial discrimination in itself renders them wrong. Once this ambiguity is uncovered, Scanlon's account appears quite problematic.

Chapter 6 presents the account of the wrongness of discrimination that I think is the most promising one, i.e., the harm-based account, even if, as I concede, it is not without problems of its own and may need to be supplemented with other accounts. According to this account, discriminatory acts are wrong by virtue of the harm they cause. However, since there can be "harmless" acts of discrimination, the harm-based account implies that discrimination is not always wrong. Some might see this as an objection to the harm-based account even if they keep firmly in mind the relevant concept of discrimination. If so, however, it is an objection that with equal right can be directed against mental-states and objective-meaning accounts. In the second half of the chapter, I propose a specific, consequentialist version of the harm-based account, i.e., the desert-prioritarian account. This account gives greater moral weight to avoiding harms to the worse-off—which makes it able to explain one way in which affirmative action is different from paradigm cases of discrimination—and it gives greater moral weight to benefiting the more deserving—which makes it able to explain why harms to discriminatees typically are morally worse than harms to discriminators.

Part 3 explores various issues that arise in connection with eliminating discrimination or with neutralizing the bad effects of discrimination. Using a term introduced by John Rawls, we might say that this part of the book addresses various issues in nonideal moral theory, i.e., it seeks to answer the question of what agents, including the state, should do given other agents' noncompliance with moral requirements, in particular a moral requirement not to discriminate.[19] Chapter 7 takes a critical view at what should be the proper aim of proportional representation. It is often assumed that if members of a certain socially salient group are less well represented within a certain attractive category, say CEOs, than it is within the population as a whole, this representational gap results from discrimination and, accordingly, aiming at neutralizing the effects of discrimination implies aiming at making sure that this group's representation corresponds to its proportion of the population as a whole. The chapter questions this idea and then shows that it is surprisingly difficult to know what should then be the aim, in view of the fact that discrimination may affect not only how many members of a certain group succeed in becoming CEOs, but also how many people are members of the relevant group.

Chapter 8 explores a different issue of proportional representation, i.e., one that arises in connection with punishment. Suppose, say, that murderers belonging to a certain racial group are much more likely to receive capital

[19] It would, however, be misleading to say that part 2 addresses issues in ideal normative theory, where ideal theory is understood to assume full compliance with the demands of justice and morality; see John Rawls (1971), *A Theory of Justice* (Oxford: Oxford University Press), pp. 245–246. None of the accounts of the wrongness of discrimination discussed in part 2 makes any assumptions about whether people actually comply with the demands of justice and morality.

punishment than murderers from other racial groups and that the difference reflects discrimination in punishment. Assuming for the sake of argument, that a murderer deserves capital punishment, might a convicted murderer belonging to the former group complain, on grounds of justice, that he would not have been sentenced to death had he belonged to a different racial group (assuming that to be true, empirically speaking)?[20] Some argue that given that a murderer deserves capital punishment, he is in no position to complain. In response, I defend the view that this contention ignores the nature of what being in a position to complain involves and the nature of comparative justice.

Chapter 9 returns to discrimination on the labor market. More specifically, it explores the issue of the morality of cases where employers discriminate, not because they themselves harbor any discriminatory attitudes, but on the basis of reaction qualifications, i.e., employers know that, say, customers harbor sexist, discriminatory preferences against women as a result of which male salespersons have better qualifications in that they outsell their female colleagues. The existence of reaction qualifications seems to imply that despite the fact that all hiring is done on a purely meritocratic basis, members of some groups may do much worse than members of other groups. To avoid this implication, some theorists have argued that employment justice, including the ideal of meritocracy, requires that reaction qualifications rooted in discriminatory responses should be disregarded. The chapter challenges this view and argues that it fails to consider a sufficiently wide range of cases where reaction qualifications matter.

Chapter 10 explores discrimination outside the state and labor market, e.g., in relation to people's preferences regarding with whom they share their neighborhood, their religious places, or their lives. Generally, discrimination in the private sphere has not received the same amount of attention as, say, labor-market discrimination. In part this may reflect the belief that people have a moral right to act in a morally wrong way in matters that pertain to their private lives. However, most of the reasons that render discrimination wrong in the labor market apply equally well to private discrimination. While state policies may often be ineffective or even counterproductive in eliminating private discrimination, this does not establish that discrimination in the private sphere may *not* be seriously morally wrong. I approach these issues from the point of view of the desert-prioritarian account expounded in Chapter 6 and show how it implies that there can be a gap between the morally desirable legal status and the moral status of discriminatory acts.

[20] Given that this is a book on discrimination (and, possibly, even if it were not) I should make a note on my use of gendered personal pronouns. I alternate (unsystematically) between using "he/his" and "she/her" in cases such as this one where, due to the hypothetical nature of the situation, I could use either gender.

Chapter 11 considers the controversial issue of racial profiling. Suppose that being a member of a certain racial group is a useful proxy for, say, an above-average probability that one engages in drug smuggling on transatlantic flights. Is it permissible for customs officers to discriminate against members of this racial group, i.e., by being more inclined to stop and search travelers with this particular racial identity? Obviously, most find racial profiling morally objectionable and perhaps even noninstrumentally wrong. The chapter argues that racial profiling is a form of statistical discrimination; that, like other forms of statistical discrimination, racial profiling is not morally wrong per se; but that it is often wrong, and for interesting reasons. One such putative reason is that the statistical facts that supply an instrumental justification for racial profiling result from an unjust social structure reproduced and supported by many of those people who favor and benefit from racial profiling and who, over time, could undo this structure.

{ PART I }

The Concept of Discrimination

This page intentionally left blank

{ 1 }

What Is Discrimination?

1. Introduction

We are all familiar with lists of the usual victims and the usual perpetrators of discrimination. So, for instance, we would all say that women were (and, in a few places, still are) discriminated against because, unlike men, they do not have the right to vote or to run for public office; that whites discriminated against blacks under apartheid; and that, unlike members of the majority religion, members of a minority religion are discriminated against, when they are not legally permitted to practice their religion publicly. While such lists offer some understanding of what discrimination is, this understanding is incomplete. This becomes apparent, when we come across new, putative cases of discrimination and need to determine whether they extend the list. So, for instance, when smokers complain that they are being discriminated against when not hired for jobs in the health sector, or obese people think that they are being discriminated against because the public space is designed on the basis of needs and requirements of nonobese people, are they right?

In the light of the limitations of the current understandings of discrimination embodied in lists of usual victims and usual perpetrators of discrimination, we should seek to formulate explicit criteria for discrimination.[1] One advantage of such criteria is that they tell us when to delete items from, or add items to, the well-known lists of discriminators, discriminatees, and discriminatory acts. Another advantage is that in order to know what makes discrimination

[1] Some might say that there is "*a complicated network of similarities overlapping and criss-crossing*" (Ludvig Wittgenstein [1953], *Philosophical Investigations* [Oxford: Basil Blackwell], p. 32 [§66b]) between different kinds of discrimination, but that there is no set of necessary and sufficient conditions for discrimination. However, this view is compatible with the desirability of an explicit definition. After all, it is often regrettably impossible to settle which members (and nonmembers) of a family resemble one another. Also, if different kinds of discrimination are merely related to one another by way of family resemblances, different kinds of discrimination are bound to be morally wrong for different reasons.

morally wrong—a question I return to in part 2 of this book—we need to know which nonmoral properties discrimination have. A definition of discrimination supplies this knowledge. A similar point applies to measures taken to counteract discrimination and its effects—the topic of part 3. To know what such means are and to assess their efficiency, we must first know what discrimination is.

This chapter proposes various definitions of discrimination (sections 2 to 5). It claims that the term "discrimination" is used to signify different things and, thus, that it refers to several concepts of discrimination. The chapter gives special weight to a definition that ties discrimination to differential treatment on the basis of membership of socially salient groups. This concept of discrimination is expounded in sections 6 to 8 and it is the one that will be employed in later chapters of the book.

2. Discrimination in the Generic Sense

An individual can be discriminated *in favor of* as well as *against*. Ignoring cases where initially the discriminatee is much better off than others, if the discriminator treats the discriminatee worse (or better) than others, we have a case of discrimination against (or in favor of).[2] In cases where an agent treats people differently but treats no one worse than others, she discriminates *between* them. If an institution has separate showers for men and women, presumably it discriminates *between* men and women without discriminating *against* either (but perhaps against transsexuals).

This book focuses on discrimination *against* individuals. The reason for doing so is not that discrimination between or in favor of individuals raises no interesting moral issues. Discriminating between often involves constraining social scripts, and sometimes people complain of being treated differently, though not worse. However, discrimination *against* raises these issues in addition to other and no less important ones.

So what does it take to discriminate against someone? Here is a very broad definition of discrimination:

[2] Suppose X treats Y disadvantageously compared to Z and Y is much better off than Z. Suppose also that X still treats Y better than X ought to do, and Z worse than X ought to do given the fact that Y is so much better off than Z. Say, Y is a white billionaire and Z a nonwhite impoverished person, and yet X gives significant amounts of money to both of them though somewhat more to the impoverished nonwhite who, however, ought to have received all the money given by X. Here one might think that X discriminates against Z despite the fact that, in one sense, X treats Z better than Y. However, there is also a sense in which X treats Z worse than Y, e.g., by giving less weight, comparatively speaking, to the needs of Z than to the needs of Y.

Generic discrimination
An agent, X, discriminates against someone, Y, in relation to another, Z, by Φ-ing (e.g., hiring Z rather than Y) if, and only if:

(i) There is a property, P, such that Y has P or X believes that Y has P, and Z does not have P or X believes that Z does not have P,
(ii) X treats Y worse than he treats or would treat Z by Φ-ing, and
(iii) It is because (X believes that) Y has P and (X believes that) Z does not have P that X treats Y worse than Z by Φ-ing.[3]

In this basic sense, to discriminate against someone is to treat her disadvantageously relative to others because she has or is believed to have some particular feature that those others do not have. In short, discrimination against someone simply is disadvantageous differential treatment.

The proposed definition is obviously very broad.[4] We might say that the definition captures discrimination in the generic sense and certainly "discrimination" in this sense has no automatic condemnatory force. Indeed, there is not even a presumption that someone who engages in generic discrimination acts wrongly. Discrimination in the generic sense is what—in a way that distracts attention from what is his real topic—J. R. Lucas has in mind when he "confesses" to being guilty of discrimination:

> In an article I once said I should discriminate against black men when casting the part of Iago, and in another article against women when considering potential brothers-in-law; and I have often advocated in public the exclusion of undergraduates from making various sorts of decisions in College and University. I have received much stick in consequence, but am impenitent. To decide is to discriminate. In choosing one line of action we exclude another. And often decisions involve the interests of others, in particular when we are appointing people to posts.[5]

A judge who convicts the guilty and releases the innocent is not for that reason engaged in presumptively wrongful behavior. Similarly, a person who

---

[3] Jan Narveson (2002), *Respecting Persons in Theory and Practice* (New York: Rowman and Littlefield), p. 205. On the definition proposed here, X might discriminate against Y even if he believes that both Y and Z have P provided that X treats Y worse than Z because, as a matter of fact, Y, but not Z, has P. I return to what "because" might mean in section 7, this chapter. "[B]ecause (X believes that) Y" in the definition means "because Y has P or because X believes that Y has P."

[4] Arguably this definition is not broad enough to cover all uses of "discrimination" and its derivatives. A wine connoisseur might be said to have a discriminating taste, even though there is no Y and Z of the relevant kind whom he or she treats differently. In my sense, however, discrimination is a way of treating individuals with moral status, not inanimate things.

[5] J. R. Lucas (1985), "Discrimination and Irrelevance," *Proceedings of the Aristotelian Society* 86, 307–324, p. 307. Lucas's confession is a distraction because it misleadingly portrays opponents of discrimination as loonies who oppose choosing one course of action over another.

refuses "to sit next to people who haven't bathed recently" engages in generic discrimination, but she is hardly doing something that is presumptively wrongful, or at least, if it is, it is wrong for very different reasons why, say, sex discrimination is morally wrong.[6] It is doubtful that we could avoid engaging in discrimination in the generic sense. Yet, when people complain about discrimination, they have in mind behavior that is at least presumptively wrong and generally avoidable. Hence, they do not have generic discrimination in mind. So not all instances of generic discrimination constitute discrimination in the sense we are after, which I take to be, prima facie, morally wrong differential treatment. Yet, all such more specific cases of discrimination amount to discrimination in the generic sense. Before I get to which additional conditions disadvantageous differential treatment must satisfy to amount to discrimination in the morally relevant sense—call the challenge of specifying these conditions *the definition challenge*—I will point to some interesting features of generic discrimination, and thus of discrimination in these more specific senses as well. Doing so will serve to highlight some important structural features of discrimination.

The first feature of discrimination that the basic characterization brings out is that discrimination is *essentially comparative with respect to individuals*, i.e., a matter of how an agent treats some people compared to others. Unlike other prima facie morally wrong acts, such as lying, hurting, or manipulating, one cannot discriminate against someone unless there are others who receive (or who would receive) better treatment at one's hands. I cannot dispute the correctness of an accusation of lying to someone by saying that I did the same thing to this person as I would have done to anyone else in that situation. However, I can rebut an accusation of having discriminated against someone by saying that I would have treated anyone else at least as badly in that situation (relevantly characterized). As we shall see, this has important implications for how we can account for the wrongness of discrimination (chapters 4–6).[7]

The relevant disadvantage is interpersonal, not intrapersonal.[8] Hence, my definition of discrimination implies that an act of discrimination against Y may improve Y's situation relative to the way it was prior to the discriminatory act and relative to how it would have been had the discriminatory act not taken place, and be believed by X to do both. This seems right. If a privately funded

[6] I owe this example to Larry Alexander (1992), "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies," *University of Pennsylvania Law Review* 141.1, pp. 149, 151.

[7] If discrimination were not essentially comparative, we could not say that it is wrong because it incorrectly represents discriminatees as having lower moral status. The denial that discrimination is essentially comparative would therefore strengthen my argument against the disrespect-based account in chapter 5.

[8] Strictly speaking, a discriminatory act need not involve treating different individuals differently. It might involve treating the same individuals differently over time, e.g., withholding favors from someone who has come out openly as gay which one did not withhold prior to his so doing.

aid agency declares that it will devote most, but not all, of its resources to help male, middle-aged, Protestant white people, few would deny that it discriminates offering the fact that it harms no one in either of the intrapersonal senses of harm employed just above as their reason for holding this view.[9]

Some might object to my claim that discrimination is essentially comparative across persons on the following ground: to discriminate against a person it is sufficient to allow one's treatment of that person to be guided by facts about her that are irrelevant to how she should be treated, or to fail to allow one's treatment of her to be guided by facts about her that are so relevant.[10] Thus, a person who complains that she has been a victim of discrimination when her application for a certain job is turned down on grounds of her race need not retract her complaint when she learns that members of all races are equally likely to have their applications turned down on racial grounds because racists of different colors sit occasionally, and with the same frequency, on the hiring committees.

In response, I say that while she can complain that, on this particular occasion, *the decision* not to hire her was discriminatory, she was being treated worse than other individuals because of her race. However, she cannot complain that the *procedure* was discriminatory, since a procedure cannot discriminate against everyone and the procedure involves the same risk for everyone of being rejected on the basis of his or her race.[11] In this sense the example is misleading because its target is unclear. Hence, the example fails to show that discrimination is not essentially comparative across persons.

Another objection to my claim that discrimination is necessarily comparative across persons (in that it requires that the agent treats the discriminatee worse than others) might point out that in some cases treating individuals in the same way can be discriminatory. Suppose, for instance, that a sexist employer pays a highly productive woman the same wage as an unproductive male employee although, except for the present case, the wage the employer pays is determined by the employee's level of productivity. This may well constitute discrimination, although it can be described as treating the two employees equally, i.e., as paying them the same wage.

I retort that this does not constitute a counterexample to my definition of generic discrimination. Any discriminatory act can be described in many

---

[9] On the interesting discrimination case of a French charity serving pig's soup—so-called identity broth—to homeless people to avoid serving soup to Muslims and Jews, see http://www.nytimes.com/2006/02/28/international/europe/28soup.html (accessed October 17, 2012).

[10] Cf. Harry Frankfurt (1999), *Necessity, Volition, and Love* (Cambridge, UK: Cambridge University Press), 146–154, p. 150.

[11] A procedure may be discriminatory in the derivative sense that it allows discriminatory decisions to influence outcomes occasionally where these decisions involve treating some differently from others. However, to say that a procedure is discriminatory in this sense, even if the procedure as such is not biased against anyone, is not to deny that discrimination is essentially comparative.

different ways and the mere fact that an act involves treating people in the same way under one description does not contradict my claim that discrimination is essentially comparative across persons.[12] The case at hand can also be described as a case of treating the two employees unequally, since they are being rewarded differently for a given level of productivity.[13] More generally, any putative case of discriminatory equal treatment can also be described as involving discriminatory unequal treatment and, thus, I see no reason to revise the definition of generic discrimination.

The second feature of my definition is that discrimination involves *disadvantageous treatment*. Disadvantageous treatment should be distinguished from treatment that causes harm. I can treat someone disadvantageously without thereby harming him, e.g., I refuse to hire an applicant who is better qualified and then, as a result of my rejection, lands herself a better job. Also, I may treat some favorably compared to others and yet harm them, e.g., I hire someone despite her worse qualifications and as a result the better-qualified and rejected candidate lands herself a much better job. If actual effects do not (completely) determine whether a certain treatment is disadvantageous, what then does? One way to go here would be to say that whether a treatment is disadvantageous is determined by what the agent believes the likely overall effects of his actions to be. While there might be some truth in this, this does not seem quite right either. Suppose that, out of racial hatred, a Nazi official turns down a Jewish person's application for admission to university in Vienna in 1938 believing that rejection will lead this person to immigrate to the United States and end up better off than had she stayed. This amounts to disadvantageous treatment in relation to the particular goods the distribution of which the official is in charge of—admission to university—even if the official believes (and does not intend) that, overall, the applicant will benefit from being rejected.

The third feature is that there is *no such thing as agent- or objectless discrimination*. We cannot say things like: "There was a lot of discrimination going on in this society, but fortunately no one perpetrated or was being subjected to discrimination."[14] Some will object that there can be discrimination without

[12] *Pace* Hugh Collins (2003), "Discrimination, Equality, and Social Inclusion," *Modern Law Review* 66.1: 16–43, pp. 16–17, it is, thus, not true that in cases where courts rule against indirect discrimination they deviate from the principle of equal treatment. If a certain rule or practice affects members of a certain group much worse than members of others groups, they are being (unintentionally) treated, e.g., affected, unequally even though there also other senses in which they are being treated equally, e.g., impartially assessed on the basis of the same standards.

[13] I am not implying here that, Kibbutz-style, paying all workers the same wage irrespective of effort and productivity is necessarily discriminatory. I am simply responding to a specific objection to my claim that discrimination is essentially comparative.

[14] The present conceptual claim is compatible with the epistemic claim that we can be justified in believing that a certain act of discrimination has taken place even if we cannot identify the individuals who were subjected to the relevant disadvantageous treatment, cf. Sandra Fredman (2011), *Discrimination Law*, 2nd ed. (Oxford: Oxford University Press), p. 286.

subjects of discrimination, i.e., structural discrimination, where it is not individual agents who discriminate. To this it might sensibly be replied that social structures cannot be separated from what individuals do.[15] Rather, they are constituted by and sustained through certain regularities in individuals' acts and states of mind (more on this in appendix 2 in chapter 2). Alternatively, we might say that in these cases the subject of discrimination is the relevant social structures. For present purposes, I can remain agnostic on this issue and would allow that the present definition only requires a very thin, grammatical notion of a subject.

While discrimination cannot be objectless, many different objects will suffice for discrimination to occur. For instance, "Y" and "Z" may refer to possible people. Hence, X may discriminate even if he treats all actual people alike, say, X may believe that there is a certain superior race of human beings and treat all actual human beings worse than he would have treated members of the superior race, e.g., he benefits them less than he would have done had he thought they were members of this superior race and does so in order to save resources for the occasion when he encounters such people.[16] Surely, this amounts to generic discrimination, whether this unusual form of discrimination constitutes a serious moral problem or not. Also, superindividual entities such as private companies, social structures, and states can be subjects (and objects) of discrimination.

Many claim that, unlike persons, animals cannot be the objects of discrimination as opposed to, say, maltreatment. Certainly, there are forms of discrimination and discrimination-related harms from which they cannot suffer, e.g., those involving lowered self-esteem. However, there are also forms of treatment from which persons and animals alike can suffer and which in the case of persons would be classified as discrimination. For instance, persons and human beings who are not persons and might not even have the potential to become ones, e.g., babies with severe forms of congenital disorder, are discriminated against when their treatment is to the detriment of their interests because the discriminator underestimates their moral status. Clearly, this can happen to animals too, even if it may be strained to say that they are being discriminated against. Those of us who are persuaded by Peter Singer's critique of speciesism—the view that the interests of members of the human species counts for more, morally speaking, than the interests of members of other species—will judge that a notion of discrimination that can take any human being, but no

[15] E.g., G. A. Cohen (1997), "Where the Action Is: On the Site of Distributive Justice," *Philosophy and Public Affairs* 26.1, 3–30.

[16] Some might want to describe this as a case of unsuccessful discrimination, i.e., the agent intends to treat members of normal and superior races differently. That presupposes that there actually are members of the superior race and since this presupposition is unsatisfied, he does not actually succeed in treating people differentially. Accordingly, he does not discriminate either.

animal, as an object embodies a distinction that is morally arbitrary.[17] It is a bit like having a term for harm to women and another for harm to men.[18] Indeed, the person who coined the term "speciesism"—the British psychologist Richard D. Ryder—wrote in 1975 that he employs the term "to describe the widespread discrimination that is practised by man against other species . . . Speciesism is discrimination, and like all discrimination it overlooks or underestimates the similarities between the discriminator and those discriminated against."[19]

A fourth feature of discrimination is that it is *actual-properties independent* in that there is no necessary overlap between the properties in virtue of which an individual is subjected to disadvantageous, differential treatment and the properties this individual actually has. So if the discriminator mistakes the sex of a person—say, a sexist employer thinks that the name of an immigrant applicant is a woman's name and thus does not hire him—a man may be subjected to sex discrimination against women. We often forget this distinction. So we tend to think that when members of a certain group are being subjected to discrimination, they are being discriminated against *qua* members of this group. This need not be the case. We might truly say that immigrants suffer from discrimination, even when the reason that is true is simply that all immigrants are women and the country to which they have immigrated is deeply sexist though welcoming toward people *qua* being immigrants.[20] More generally, there is a distinction between members of a certain group being discriminated against in a purely extensional sense and their being discriminated against in an intensional sense, i.e., under something like the description "members of this group."

A fifth feature of discrimination is that it is *treatment specific*. Typically, when people talk about discrimination they have in mind cases where one group of individuals is treated badly compared to other groups across the board, i.e., in all respects. It might well be the case that when someone gets discriminated against in one respect there are other respects as well in which he or she is discriminated against.[21] However, discrimination is not an "all-or-no-dimensions"

---

[17] Peter Singer (1993), *Practical Ethics* (Cambridge, UK: Cambridge University Press), pp. 55–82. In Singer's view, some nonhuman animals *are* persons, e.g., adult chimpanzees.

[18] See Oscar Horta (2010), "Discrimination in Terms of Moral Exclusion," *Theoria* 76.4, 314–332.

[19] Richard D. Ryder (1975), *Victims of Science: The Use of Animals in Research* (London: Davis-Poynter), p. 16.

[20] Suppose a sexist employer hires a male applicant rather than a more qualified female one. We might say that the employer discriminates against the individual female applicant or we might say that he discriminates against women. (We might say both, of course.) In my account, the latter is not true in a strict sense—the employer does not treat each and every woman worse—but it may express certain related truths, e.g., that for any woman he would have treated her worse, or that he treated the actual applicant worse *because* she is female.

[21] Any act of discrimination involves a particular dimension in which the discrimination takes place. Often the dimensions in which discrimination takes place vary across groups. See Alexander (1992), "What Makes?" p. 158.

affair. My definition allows that discrimination could be more complicated than that. For instance, it could be the case that, all things considered, two groups are being treated equally well and yet in all spheres of life either of the two groups suffers from discrimination. To some extent this is the case in relation to sex discrimination. Men who openly display emotions experience censure in a way that women do not. Similarly, women who are assertive experience censure in a way that men do not. No doubt, it is logically possible for members of any given society to fall into two mutually exclusive, jointly exhaustive subgroups: discriminators and discriminatees. However, as a matter of fact often individuals are discriminators in some respects and discriminatees in others. Obviously, this is not to deny that some people—e.g., blacks, lesbians, women, and, in particular, women who are black and lesbian—are subject to much more and much more harmful discrimination than others—e.g., white, heterosexual males.[22]

In some views, discrimination is essentially asymmetric, e.g., because members of dominated groups can suffer from discrimination but not members of dominating groups. Discrimination in the generic sense has no such asymmetry condition built into it. We might, of course, build it into a definition of a more specific kind of discrimination, but I do not think that, as a matter of fact, discrimination in the relevant more specific sense includes some such requirement. Hence, it is possible—on my definition of generic definition as well as the definition of prima facie morally wrong discrimination that I propose below—for someone to discriminate against others who in turn discriminate against the first person in dimensions that either differ from, or are the same as, those in which he discriminates against the latter. It is a common response to being treated on the basis of a stereotype to stereotype one's stereotypers.[23]

Finally, generic discrimination can be *reflexive* in a sense to be explained shortly. Typically, discriminators differ from discriminatees by virtue of their not having the property on the basis of which they discriminate. Yet, it is not rare for some to discriminate against others because of their having some property, which they themselves have. For instance, female employers might favor male over female applicants because of their sex.[24] It might even be that a discriminator discriminates against herself in virtue of some property she has, e.g., a women recommends that one of her male peers be promoted because of his gender despite her being the best qualified. We might label this form of discrimination, where the discriminator and the discriminatee are one and the same

[22] Cf. Fredman (2011), *Discrimination Law*, pp. 139–143, 167.

[23] Cf. Lawrence Blum (2002), "*I'm Not A Racist, but . . .* " (Ithaca, NY: Cornell University Press), esp. pp. 36–39, on the view that "racism = prejudice + power" and, thus, that African Americans cannot be racists.

[24] To say that female employers may discriminate against women is not to say that they may not at the same time themselves be victims of discrimination against women.

individual, "reflexive." Some of the most damaging forms of discrimination are those in which discrimination is, so to speak, internalized by the victims of discrimination and involve some sort of abhorrence of what one is.

Summing up this section, generic discrimination against someone: (1) is essentially comparative across individuals; (2) involves disadvantageous treatment; (3) is agent- and object requiring; (4) is actual-properties independent; (5) is treatment specific; and (6) is potentially reflexive. As we shall see later, some of these features of generic discrimination place significant constraints on an account of the wrongness of discrimination.

3. Irrelevance Discrimination

All forms of discrimination qualify as discrimination in the generic sense. However, as I have indicated, generic discrimination is not what people complain about when they complain about discrimination. The world would not be a better place if discrimination in the generic sense never took place—that is, if the innocent were treated no differently from the guilty, the needy taxed no differently from the wealthy, and so on. So which additional conditions does a case of generic discrimination have to satisfy to qualify as discrimination in the sense that interests us?

There are at least three ways to supplement the definition of discrimination in the generic sense so as to specify the relevant form of discrimination.[25] This and the following two sections will each explore one of these three specific kinds of discrimination. Here is an example that serves to illustrate the first one: Suppose that Anne and Burt are both present at the scene of the crime. Witnesses all point to Burt as the perpetrator. Burt has a knife, the victim's wallet in his hand, and his clothes are blood-stained. None of this is true of Anne. Because of that the police officer arrests Burt, but not Anne. This amounts to discrimination in the generic sense: the officer treat Burt disadvantageously compared to Anne because, unlike Anne, he has the properties mentioned. Yet, no one in her right mind would think that the officer acts wrongfully and, thus, discriminatorily in the way that we are interested in. One suggestion, as to why this is so, is that the police officer acts in the way he does for reasons that should guide his actions under the circumstances. Generally speaking, decisions about whom to arrest ought to be made, in part, on the

[25] A fourth way is "legal discrimination." Roughly, legal discrimination is generic discrimination that satisfies the following additional condition: "P" is a member of a certain legally protected group and there are no legally relevant concerns that legally justify the relevant disadvantageous treatment of Y. If something qualifies as discrimination in this sense, it is legally forbidden. I mention discrimination in the legal sense only to set it aside. My primary interest lies not in the illegality of discrimination as such, but in its immorality, whether this is reflected in legal regulations or not (see chapter 10).

basis of forensic evidence pointing to the perpetrator of the crime. Had he instead arrested Burt, because of his race, the police officer would have acted on irrelevant grounds and the case would be one of discrimination in the specific sense we are concerned with. This suggests that we add a fourth condition to the right-hand side of our basic biconditional of generic discrimination:

(iv) The fact that (X believes that) Y has P and (X believes that) Z does not is irrelevant to whether X ought to Φ i.e., it is neither a justifying reason for, nor a justifying reason against, X's Φ-ing.[26]

Call the specific form of discrimination that satisfies (i–iv) *irrelevance discrimination*. While in many contexts we do think of discrimination as differential treatment on irrelevant grounds where relevant grounds are narrowly circumscribed, e.g., to those properties that affect the strictly economic aspects of a worker's ability to satisfy the job requirements. For instance, this account underpins the view that it is discriminatory to pay men a higher wage than women for the same job, because the only fact that is a suitable reason to differentiate wages is differences in performance, or to hire people with a certain religious or race to perform jobs where religion and race are irrelevant to the job function. Anthony Flew expresses a view quite close to the one embodied in irrelevance discrimination when he writes about racism that it is unjust because it treats people on the basis of traits that "are strictly superficial and properly irrelevant to all, or almost all [remember Iago], questions of social status and employability."[27]

Despite the advantages of the notion of irrelevance discrimination, it can be questioned whether irrelevance discrimination is the notion we are aiming to contrast with generic discrimination here. First, it is not clear that (iv) and (i–iii) give a *sufficient* condition for discrimination in the relevant sense. Suppose that a university hires a less-qualified applicant because the dean's son is married to the applicant and the dean favors members of his family. Most would say that the dean made his hiring decision on irrelevant grounds, but not describe the case as one of discrimination as opposed to, say, one of nepotism.[28]

---

[26] Arguably, what matters here is that "the fact that Y has P and Z does not is irrelevant to whether X ought to φ," not that X believes that Y has P and Z does not have P. Someone who believes, say, that being male is a reason why one should get promoted and, accordingly, treats female employees disadvantageously discriminates, even though his belief in some sense forms a reason for treating men and women differentially, e.g., the employer will be less inclined to rely on his female employees. It is an interesting question, though, whether cases where gender is irrelevant and where the discriminator believes it to be so and nevertheless discriminates against women on grounds of gender is more wrong than those where this fact is irrelevant, but X has no belief to this effect—perhaps despite having good reasons for such a belief.

[27] Anthony Flew (1990), "Three Concepts of Racism," *Encounter* 75, 63–66, pp. 63–64.

[28] See, for instance, Brian Barry (2005), *Why Social Justice Matters* (Cambridge, UK: Polity Press), p. 22; Thomas Scanlon (2008), *Moral Dimensions* (Cambridge, MA: Harvard University Press), p. 70. Both Barry and Scanlon contrast discrimination with nepotism.

Second, treating someone on the basis of irrelevant grounds is not *necessary* for discrimination. Consider the case of "reaction qualifications," i.e., qualifications an employee acquires as a result of the way in which others—customers, co-workers, and others—respond to, and in particular assess, him.[29] (I return to reaction qualifications in chapter 9.) Suppose that a shopkeeper defends his decision not to hire any minority members on the ground that while he personally would not mind employing them, he believes—correctly, let us suppose—that many of his customers prefer being served by a nonminority shop assistant. He is in business to earn a profit, so from his point of view majority membership *is* a relevant ground for differential treatment. We might accept this claim about relevance from a strictly commercial point of view and yet regard his conduct as discrimination, perhaps even wrongful discrimination. It could be replied that a ground is relevant if, and only if, it is one that the relevant kind of decision ought, morally (as opposed to commercially) speaking, to be based upon. But in that case the notion of irrelevance discrimination really rests on a prior account of which aims it is morally unobjectionable to pursue, in which case clarity is better served by formulating the necessary extra clause in terms of what is morally objectionable.

4. The Moralized Concept of Discrimination

In the light of the concluding sentence in the previous section, we might propose the following moralized concept of discrimination:

(iv′) It is morally objectionable that X treats Y worse than Z by Φ-ing because (X believes that) Y has P and (X believes that) Z has not.[30]

Call the specific form of discrimination that (i–iv′) captures *moralized discrimination*. By including (iv′) we build into our notion of discrimination the claim that it is morally objectionable differential treatment that is at issue. Often when the term "discrimination" is used, some such implication is indeed involved. For instance, Lena Halldenius writes that, as she uses the

[29] Alan Wertheimer (1983), "Jobs, Qualifications, and Preferences," *Ethics* 94.1: 99–112; Peter Singer (1983), "Is Racial Discrimination Arbitrary?" in Jan Narveson (ed.), *Moral Matters* (New York: Oxford University Press), pp. 308–324.

[30] Cf. Thomas E. Hill, Jr. (1995), "The Message of Affirmative Action," in Steven M. Cahn (ed.), *The Affirmative Action Debate* (New York: Routledge), 169–191, p, 174. As Singer points out, certain blatant forms of racial discrimination involve treating individuals differently for no good moral reason, i.e., because they involve treating some people's interests as morally more important than other people's. However, he does not want to define discrimination on the basis of considerations about arbitrariness. For simplicity, I ignore a notion of moralized discrimination that includes cases where X *believes* (falsely perhaps) that treating Y and Z differently because Y has P and Z not is morally wrong and yet treats Y and Z differently on that basis.

term, "when an action has been correctly described as an instance of discrimination, it has at the same time been correctly described as unfair."[31] Similarly, when David Wasserman observes that "to call discrimination 'wrongful' is merely to add emphasis to a morally-laden term," he uses "discrimination" in its moralized sense.[32] Finally, few would say that child molesters are discriminated against by the legal system, when, ceteris paribus, they are being treated disadvantageously compared to non-child-molesters. We can explain why, if we assume that people use the moralized conception of discrimination and that people think of this differential treatment as being morally justified.

If we use the moralized notion of discrimination, we cannot meaningfully ask: "Are all kinds of discrimination unjust and morally unjustified?" Yet many scholars have asked this question, and this strongly suggests that in asking this question they were not using the term "discrimination" in its moralized sense.[33] Also, we can dispute whether a case of differential treatment constitutes discrimination in this sense by disputing that it is morally objectionable. For instance, we can deny that elderly patients are subjected to age discrimination when their medical needs are assigned a lower priority than those of younger ones by repudiating the claim that this is morally objectionable. But we cannot say informatively that what makes a particular case of differential treatment morally objectionable is that it constitutes discrimination in this sense. To say that something is discrimination in the present sense simply *is*, in part, to say that it is morally objectionable.

Moralized discrimination is also much too broad a notion. Like irrelevance discrimination, it implies that nepotistic decisions are discriminatory. Moreover, it implies that wrongful, idiosyncratic differential treatment constitutes discrimination. Yvonne and Zack both apply for admission to a graduate program at an Ivy League university. Xavier, the admissions officer, has an idiosyncratic liking for people from Omaha, especially those who did well in sports at high school and who live in a county the name of which contains at least one "s." Yvonne, but not Zack, happens to be distinguished in all these regards and, accordingly, Xavier unjustly admits her rather than Zack. Many would deny that this case of morally objectionable treatment is a case of discrimination. Of course, it would have been a case of discrimination if Yvonne had been white, or heterosexual, or Christian, and Zack had been black, or homosexual,

---

[31] Lena Halldenius (2005), "Dissecting 'Discrimination,'" *Cambridge Quarterly of Healthcare Ethics* 14.4, 455–463, p. 456.

[32] David Wasserman (1998), "Discrimination, Concept of," in R. Chadwick (ed.), *Encyclopedia of Ethics* (San Diego, CA: Academic Press), 805–814, p. 805.

[33] See Matt Cavanagh (2002), *Against Equality of Opportunity* (Oxford: Clarendon Press); Narveson (2002), *Respecting Persons*; and Singer (1983), "Is Racial Discrimination Arbitrary?" 308–324.

or Muslim, and if Xavier admitted Yvonne because he prefers whites, or heterosexuals, or Christians.[34]

5. Group Discrimination

These examples suggest that an alternative condition should be added to (i–iii):

(iv″) P is the property of being a member of a certain socially salient group (to which Z does not belong).[35]

While we often have generic discrimination that satisfies (iv″) in mind when we talk about discrimination—I will say more about the attraction of focusing on generic discrimination that satisfies that (iv″) as well as other conditions shortly—(iv″) needs to be supplemented. Consider the case of love. In the account given so far, it is conceivable, though unlikely, that a heterosexual would discriminate against a person of the same sex seeking to become his lover by rejecting him.[36] Now, this is not as absurd as it may sound and it could happen—indeed, it does happen, I suppose. Imagine a man who, despite falling in love with another man who reciprocates his love, rejects the latter's advances *because* he is a man and would not do the same were he to fall in love with a woman who reciprocated his love. We might, on reflection, call this discrimination

[34] One reason why this case might be different is that such dislikes are, or, at least, have been widespread, and thus would serve to undermine equality of opportunity. However, inequality of opportunity is not necessary for differential treatment constituting discrimination. In a society where race, gender, and religion form the basis of socially salient groups, we are likely to construe differential treatment along these lines as discrimination, even when it is clear that it does not undermine equality of opportunity. (I return to this issue in chapter 5.)

[35] It might be suggested that only historically disadvantaged socially salient groups qualify as objects of discrimination. While, morally speaking, it will often matter whether the object of discrimination is historically disadvantaged, many people are disposed to refer to discrimination against whites or men, when whites or men are treated disadvantageously in ways that clearly would have been discriminatory had the victims been black or women, even when whites (*qua* whites rather than blacks) and men (*qua* men as opposed to women) are not historically disadvantaged. (The parentheses are needed because many white men, e.g., white homosexual men, have been historically disadvantaged.) Yet, it will often be a contested matter whether such cases fall under the label of "discrimination." I return to this issue in Chapter 5. Another interesting suggestion is that a nonsalient group of people, whose members are likely to be mistaken for being members of a certain socially salient group and on that basis treated disadvantageously, may be subjected to discrimination even if, by stipulation, they do not form such a group.

[36] In the context of discussing what makes discrimination wrong, Matt Cavanagh submits that a heterosexual is not being unfair for rejecting sexual relations with a person of his or her own sex, Cavanagh (2002), *Against Equality of Opportunity*, p. 164. See also Narveson (2002), *Respecting Persons*, p. 209. It is to accommodate this sort of objection that I need to add a further condition to generic discrimination and (iv″).

against the potential male partner.[37] In an important respect this is no different from the white person who falls in love with a black person, who reciprocates his love, and nevertheless rejects this person because he is black, which surely is a case of racial discrimination.[38]

However, setting aside such special cases if it just so happens that a person is not sexually attracted to people of the same sex, it seems very strained to say that he discriminates against such a person whose advances he rejects and whom he thereby, treats disadvantageously.[39] To avoid this implication, we could add a clause to our definition, saying that for an act of disadvantageous differential treatment of a person because of his having P to constitute discrimination against him, it must be the case that the relevant act, suitably described, is one of many instances of similarly disadvantageous differential treatment of P-people or some subgroup thereof, and that these many acts together make these people worse off relative to others. In the case at hand, homosexuals would not be better off relative to others if heterosexual men were neither to reject advances from people in whom they are not sexually interested, nor more likely to fall in love with people of the opposite sex.[40] This suggestion will not quite do, however.

First, if just one Catholic employer in a Protestant, fiercely anti-Catholic society refuses to hire Protestant applicants simply as a result of animosity toward Protestants—and not, say, in part or wholly as some kind of response to anti-Catholicism—many would call this discrimination even if the act is the only one of its kind in society.[41] Second, suppose, implausibly, that while it is harmful to a single Protestant worker not to be hired because of a Catholic employer's animosity toward Protestants, it would not harm Protestants, generally speaking, if there were many Catholic employers and if, acting on anti-Protestant sentiments, none of them hired Protestants. Presumably, we would still consider the refusal of the single Catholic employer to hire a Protestant applicant religious discrimination. The same point would apply if the single Catholic employer rejected the application of the Protestant because he

[37] No doubt this is a case where one ought, morally speaking, to have a legal right to discriminate even if one may act wrongfully in exercising that right discriminatorily, cf. chapter 10.

[38] Incidentally, this case is one that persuades most people that discrimination can take place not just in the public sphere, but also in the most intimate areas of private life.

[39] Arguably, an account that ties discrimination to differential treatment that undermines equality of opportunity for welfare can explain this case less well. However, this account cannot accommodate the special case that I addressed in the previous paragraph.

[40] Some might object, not unreasonably, that homosexuals are harmed by the fact that most people are not sexually attracted to people of their own sex because they are thereby denied the benefits of normality, statistically speaking. If we find this objection persuasive, we might be less inclined to consider prevalent patterns of heterosexuality nondiscriminatory.

[41] Note that this claim is consistent with the ruling in *Steel Workers v. Weber*, 443 U.S. 193 (1979). That ruling stressed that the affirmative action plan favoring black employees were adopted "to eliminate traditional patterns of racial segregation."

thought, without any sense of animosity, that Protestants are inferior or that Catholics and Protestants, while being the equals of one another, ought not to intermingle.[42] This suggests we need to add:

(v) Φ is a relevant type of act, policy, or practice, and there are many acts etc. of this type, and this fact makes people with P (or some subgroup of these people) worse off relative to others, or Φ is a relevant type of act etc., and many acts etc. of this type would make people with P worse off relative to others, or X's Φ-ing is motivated by animosity towards or dislike of individuals with P or by the belief that individuals who have P are inferior or ought not to intermingle with others.[43]

Given the assumption that neither men nor women would be better off if people did not treat potential lovers differently depending on whether they have the same gender as them, and the further assumption that rejecting unwanted sexual advances need not involve animosity or any sense of superiority, (v) allows us to deny that conduct informed by heterosexual, or for that matter homosexual, preferences is discriminatory.[44] The same thinking applies to the issue of whether equally resourceful churches discriminate when, for reasons other than religious isolationism or animosity toward people with a different or no faith, they refuse to hire people as priests from other denominations or religions (see chapter 10).[45] Moreover, (v) allows us to say that racially biased, differential treatment in our love lives constitutes discrimination, if, as seems likely, it involves humiliating messages of inferiority. Finally, (v) explains the case of the single discriminatory act against a Protestant applicant, since it is true not only that such acts, if sufficiently common, would be harmful to Protestants, but also that the act in our example is motivated by animosity, a belief that Protestants are inferior, or a belief that Catholics and Protestants should not intermingle.

Call discrimination satisfying (i–iii), (iv″) and (v) *group discrimination*. Unlike immoral discrimination, group discrimination is not wrong by definition.

---

[42] See Blum (2002), "*I'm Not a Racist,*" pp. 8–11.

[43] I need the qualification "relevant" to accommodate the already noted fact that any act instantiates an infinite number of act types, and the fact, presumably, that it is true of some of these types that it is harmful that they are often instantiated, or that it would be so if they were. If, for instance, a white applicant for the role of Martin Luther King in a movie is rejected because he is white, the rejection is an instance of rejecting white job applicants on grounds of their race *and* an instance of rejecting an applicant who is, in a reasonably uncontroversial sense, not suitable for the job (but see chapter 9 on reaction qualifications). If there were many instances of the former, but not the latter, type of act, this would harm whites. Presumably, here we think it is the latter type of act, which is relevant. Obviously more needs to be said about what determines the relevance of act-types.

[44] Compare Cavanagh (2002), *Against Equality of Opportunity*, p. 164.

[45] Compare Narveson (1993), *Moral Matters*, p. 251.

So, to use the terminology of Richard Hare, "discrimination" in the present sense is not a negative, secondarily evaluative term.[46] To say that something is "discriminatory" does not logically commit one to any particular moral evaluation of that thing. Of course, this is consistent with group discrimination always being immoral and with group discrimination always being prima facie morally wrong given the social world in which we live. An important virtue of the notion of group discrimination is that, in the ordinary course of events, almost always, when something is criticized as being discriminatory, the charge concerns group discrimination. This is one reason why I want to stick to group discrimination and will concentrate on this notion of discrimination in the rest of this book. Another reason is that discrimination so construed is a good, but fallible, indicator of moral wrongness. Assuming this is the case, there is a reason why, thinking from a moral point of view, we should be concerned with this notion of discrimination (as opposed to discrimination in the generic sense).

It might be suggested that I should focus on an even more specific notion of group discrimination, i.e., immoral group discrimination or the conjunction of (i)–(iii), (iv′), (iv″), and (v). By definition this kind of group discrimination is morally wrong, so offhand it might seem that this notion better captures what enemies of discrimination are opposed to.

I want to resist this move. The phrase "morally objectionable" in (iv′) either means morally wrong *all things considered*, i.e., morally impermissible, or pro tanto morally wrong. If we were to disambiguate in the former way, we can easily think up thought experiments, e.g.,, involving a Cartesian demon credibly warning us of dire consequences if we do not treat different groups

[46] See Richard Hare (1952), *The Language of Morality* (Oxford: Clarendon Press), pp. 121–126; Richard Hare (1963), *Freedom and Reason* (Oxford: Clarendon Press), 22–25; Richard Hare (1981), *Moral Thinking: Its Levels, Methods, and Point* (Oxford: Clarendon Press), p. 17; Richard Hare (1989), *Essays in Ethical Theory* (Oxford: Clarendon Press), pp. 122–129. By saying that terms like "brave," "racism," "child abuse," and "emancipation" are secondarily evaluative, unlike a term like "good," which is primarily evaluative, Hare indicates that the term has a more specific descriptive meaning than primarily evaluative terms. You could not call anything you think is bad "discriminatory"—for instance, differential treatment is, as I have argued, part of the descriptive meaning of the term—but you could call anything you think is bad "bad." No one would claim that "discriminatory" is a primarily evaluative term, but some would say that it is a secondarily evaluative term (viz. Halldenius). As my section on the moralized concept of discrimination shows, I accept that the term is sometimes used in a way where this commitment is entailed. For instance, that is so in cases of moral reform, as Hare would call them (Hare [1952], *Language of Morality*, pp. 119–120). Thus often, when a new law affects a certain group negatively, people who disapprove of it call it "discriminatory," while people who approve of it do not. Differing assessments of the factual implications of the law need not underlie this difference; there may simply be a disagreement as to whether the law is a good thing. This fact has to be considered along with the facts that "discrimination" is sometimes used in a way that involves no evaluative commitments and sometimes is used to express a positive evaluation as in "reverse discrimination."

unequally, where we would consider such differential treatment of groups discrimination and yet morally justified all things considered. If we disambiguate in the latter way, (iv′) will not do the job it is supposed to do. For instance, if we are pursuing a notion of discrimination implying that affirmative action is not discrimination, (iv′) will not do the trick if affirmative action is morally objectionable in that its absence in a just society (i.e., one that is truly sex- and color-blind) would be morally better.[47]

6. Social Salience

Group discrimination needs clarification in three respects: (1) What is *social salience*?; (2) What is it to *treat* someone in the relevant respect?; and (3) What is it to treat someone *because* this individual has a certain feature that others do not have? The present and the following two sections will address each one of these questions in the indicated order. I propose the following definition of social salience:

> A group is socially salient if perceived membership of it is important to the structure of social interactions across a wide range of social contexts.[48]

Whether one is perceived to belong to the group of people from Omaha who did well at sports in high school and live in a county the name of which contains at least one "s" is irrelevant to almost any kind of social interaction. Hence, this group of people is not socially salient. Similarly, being skilled "at making widgets" is irrelevant to almost any kind of social interaction. Hence, a manufacturer who produces widgets and prefers not to appoint applicants lacking skills in widget making does not, in the present sense, discriminate against such people. More generally, the proposed focus on social salience implies that a meritocratic hiring policy does not amount to discrimination against the less well qualified. This is good, because many would reject the idea that one can be subjected to discrimination by not being hired on account of one's being less well qualified. Having green eyes is also irrelevant in almost any social

[47] Similar objections would apply to a moralized version of irrelevance discrimination.

[48] For a slightly different, subjective account of social salience, see Harriet E. Baber (2001), "Gender Conscious," *Journal of Applied Philosophy* 18.1, 53. Baber writes: "A property is socially salient . . . to the extent that . . . [people] take it to predict and explain beliefs, character traits, tastes or other socially significant psychological characteristics." While the property of being a member of a certain group may in fact structure social interactions even if people do not ascribe explanatory relevance to it (and vice versa), one might suppose that there is a tendency for people to ascribe explanatory significance to being a member of a certain group if in fact perceived membership of it structures a wide range of social contexts (and vice versa).

context. By contrast, an individual's perceived sex, race, or religion will affect social interactions across a wide range of social contexts.[49]

Clearly, social salience is a matter of more or less, and it is impossible to say exactly when a group is so salient that differential treatment on the basis of membership in it constitutes group discrimination. Social salience involves two scalar dimensions: perceived membership in a group may be anything from slightly important to wholly dominant in the structure of social interactions, and it may be important in more or fewer social contexts. In principle, group X may be just as salient, all things considered, as group Y, even though perceived membership in X structures social interactions in fewer social contexts than membership in Y, as long as X carries greater importance in the contexts it actually structures.[50]

To be socially salient, a group need not have a large number of members. It is true, of course, that most salient groups—discrimination against which has been an issue—have had many members. However, this may be so because the mechanisms through which perceived membership of a social group structures social contexts tend to exist only in connection with groups with many members, e.g., the more members a certain group has the more likely it is, other things being equal, that there are stereotypes connected with this group.

A group need not be socially salient independently of the discriminatory treatment. Some of the most infamous forms of discrimination consist in rendering membership of a certain group easily perceived, thereby increasing its importance to social interactions, e.g., the way people classified as Jews were forced by the Nazis to wear the Star of David. Presumably, being constantly reminded of membership in a certain group increases the significance to social interactions of perceived membership of this group.

Having green eyes is special in the sense that, apart from being insignificant in almost any social context, it is very easily perceived in face-to-face encounters. Membership in many other groups will not be immediately visible in this way. A group is socially salient even if membership of it is not easily established as long as perceived membership of it has an important influence on social interactions in a wide range of contexts. Perceptions of such membership may, however, be very important to the structure of social interactions. Being gay in a homophobic society presumably tends not to structure social interactions in very many contexts because gays often have a strong incentive to conceal their

---

[49] Wasserstrom draws a similar contrast; see Richard Wasserstrom (1977), "Racism, Sexism, and Preferential Treatment: An Approach to the Topics," *University of California Law Review* 24, 581–615, pp. 586, 590. A slightly more unusual group in this context is non-human animals (see section 2, this chapter).

[50] To elaborate the notion of social salience further, one would have to say something about the individuation of social contexts and about the sorts of structuring of these that are relevant for present purposes. While it is somewhat unclear when a group is socially salient, this is not a flaw with the formula. The contours of our concept of discrimination are somewhat fuzzy.

sexuality; yet being perceived to be gay in the same society will have an important influence on social interactions. Suppose that in a certain homophobic society nongays will learn that someone is gay only if this person himself reveals it. Suppose also that gays never reveal their group membership to nonmembers, correctly anticipating the hostile reaction this would induce.[51] According to my definition, gays suffer group discrimination even though, there is *a* way in which, gays interact no differently with others than nongays do for if they were seen, which by stipulation they are not, as gays that would affect a wide range of social interactions. But even setting this point aside, it is also the case that gays' way of being and their ways of relating to others in a homophobic society are deeply structured by the constant need to avoid activating the discriminatory dispositions of others.

If membership in a certain group structures social interactions in a wide range of contexts, it is very likely that such membership is also quite central to the members' sense of who they are.[52] Given the former, membership in the relevant group is likely to be seen either, by members, as quite central to their own sense of who they are or, by others, as quite central to who the members are. In the latter case, this will tend to affect members' own sense of who they are by increasing the centrality of group membership to the members themselves.[53] While importance to the structuring of social interactions and centrality to people's self-conception will, therefore, tend to go hand in hand, analytically these are separate issues. Hence, there could be cases of people who are treated disadvantageously on grounds of membership in a nonsalient group, membership of which is crucial to their sense of who they are; and there could be cases of people who are treated disadvantageously on the basis of membership in a salient group, membership of which, however, is insignificant to their sense of who they are.

An example of nonsalient but crucial membership would be nationalists singled out for unfavorable treatment by colonial masters, where being "Indonesian" is very important to the nationalists' self-conception, but, initially,

---

[51] Cf. Erving Goffman (1963), *Stigma: Notes on the Management of Spoiled Identity* (London: Penguin Books), pp. 57–128, on the information management of discreditable, nondiscredited individuals.

[52] Cf. Iris Marion Young (1990), *Justice and the Politics of Difference* (Princeton, NJ: Princeton University Press), p. 186 on social groups; David Edmonds (2006), *Caste Wars: A Philosophy of Discrimination* (London: Routledge), p. 124, on castes. Roughly, a caste in Edmonds's sense is a socially salient group in my sense, where membership in this group is important to its members' self-conception.

[53] Anti-Semitism, for instance, tends to make being Jewish play a more prominent role in the self-conception of those people who are subjected to it. G. A. Cohen (2002), *If You're an Egalitarian, How Come You're So Rich?* (Cambridge, MA: Harvard University Press), p. 34, comments on Jean-Paul Sartre's extravagant claim that "it is the anti-Semite who creates the Jew," that although it is an exaggeration, it contains some truth.

something not recognized by other "Indonesians" as important.[54] An example of salient but insignificant membership would be people in South Africa under apartheid who were brought up to think of themselves as whites but who, due to surprising facts about their ancestry, later in life suddenly came to be thought of by others, including the authorities, as colored people. Given the above definition of socially salient groups, only the latter kind of case might involve discrimination.

There are at least two reasons to tie the definition of discrimination to socially salient groups. First, almost all groups on whose behalf the charge of being discriminated against is voiced are socially salient groups. Women, elderly people, disabled people, gays and lesbians, and ethnic and racial minorities are all socially salient groups in this sense. In fact, it appears that only one kind of discrimination does not involve a socially salient group—the use of genetic information by insurance companies to "discriminate" against people with certain kinds of bad luck in the genetic lottery. This practice does not qualify as discrimination on my account because these people do not constitute a socially salient group (at least not presently). Nothing is lost—in particular, we may still think the practice is usually morally wrong—and perhaps some clarity gained, if we discuss cases of "genetic discrimination" as cases of unjust, albeit (in the less than all-inclusive sense mentioned above) nondiscriminatory, uses of genetic information (see the appendix to chapter 3).[55]

There are cases involving disadvantageous differential treatment that we would probably consider discriminatory if the relevant groups were socially salient, but in which the term "discrimination" is not entirely apt precisely because the relevant group is not socially salient. An employer might be more inclined to hire applicants with green, rather than brown or blue, eyes. This idiosyncrasy might not amount to discrimination in the sense that interests us here, even though, obviously, the employer differentiates between different applicants in a way that is disadvantageous to some of them.[56] This is not to deny that such idiosyncrasies can be as wrong as, and reflect as corrupted a character as, genuinely discriminatory acts. However, they will not seriously harm the disadvantaged party in the great majority of cases, precisely because of their idiosyncratic nature. People with green eyes may find other employers who are indifferent or, in rare cases, idiosyncratically attracted to their eye colors and will, if informed of the reason they did not get the job, infer that the

---

[54] See Benedict Anderson (1983), *Imagined Communities: Reflections on the Origin and Spread of Nationalism* (London: Verso), pp. 110–112.

[55] Because I argue that differentiating insurance premiums on the basis of genetic information is not morally wrong per se (though it often is), and because I claim that discrimination, as we normally construe it, is not wrong per se, there is a way in which I make things harder for myself by denying that differentiating insurance premiums on the basis of genetic information amounts to discrimination in this sense.

[56] Wasserman (1998), "Discrimination," pp. 805, 807.

employer was an oddball; diminished self-respect will not be an automatic effect of the information they receive.[57]

Generally speaking, we are more likely to view differential treatment as discrimination to the extent that the following three conditions are met: (1) membership is evidenced by a dichotomous distribution of individuals in the relevant and contrasting groups; (2) all individuals are a member of only one group; and (3) it is evident whether or not someone is member of a certain group. "People of a similar outlook," for instance, satisfies none of these conditions.[58] For very many sets of three persons, it is very hard to tell whether X or Y is the person with an outlook that is most similar to Z's. Moreover, for any individual there is a group consisting of individuals with a similar outlook and very few, if any, are not members of more than one such group. Things are different with the groups that we most readily think of as being subjected to discrimination. You can be a man or a women; almost everyone is either a man or a woman, but not both. "Physical attraction" is somewhere between people with similar outlook and sex.[59] Although in principle there is wide spectrum of degrees of physical attractiveness, most people use a coarse-grained system of classification as evidenced by the fact that our ordinary language contains rather few predicates relating to people's looks such that these predicates together form an ordinal ranking from best to worst. Moreover, while people vary somewhat with regard to how they classify different people, their judgments tend not to differ very much when they differ and tend to agree at the extremes, which, on my definition, would support the view that there is such a thing as discrimination in favor of the beautiful and against the physically unattractive.

The restriction of the scope of discrimination to differential treatment of socially salient groups explains why we do not talk about discrimination against non–family members or against unqualified applicants. While it is true of each of us that we are not members of most families, there is no group of people made up of non-family members as such. Similarly, everyone is unqualified for some jobs, although, arguably, the category of people who are unqualified for *any* job has no members. In a society divided into a few enlarged families, or clans, where people are treated differently depending on family membership, we

[57] Cf. Alexander (1992), "What Makes?" p. 198: "[T]he social costs of particular kinds of discrimination are a function of how widespread those kinds of discrimination are."

[58] Cf. Alexander (1992), "What Makes?" p. 158. This is not to deny that intolerance toward people with a different outlook might be morally wrong even when it is not discriminatory. Not all morally wrongful acts are discriminatory. Also, in a society where a group of people forms a socially salient group in virtue of their outlook, e.g., communists in McCarthy's USA or Westernized, atheists in Pakistan, my definition allows that they can be discriminated against.

[59] Alexander (1992), "What Makes?" pp. 166–167.

might well approach a situation in which differentiating between people on the basis of family membership involves discrimination to the extent that clan membership is a socially salient group identity. (Of course, it is likely that clan membership would be socially salient in that way.) A similar point may apply to a group of people who are literally unqualified for (virtually) any job where membership of the group is easily detected.

Still, it is not clear that the above account of discrimination in terms of differential treatment of members of socially salient groups is entirely satisfactory. Let Y designate the group of people whom X ought, morally speaking, to treat worse than Z. Assume Y is the group of morally undeserving people, whereas Z is the group of morally deserving people.[60] Suppose X treats Y worse than Z and does so believing that he himself treats Y disadvantageously compared to Z; and suppose that differential group-membership suitably explains why X treats Y worse than Z, e.g., differential group-membership explains why X is motivated by the thought that Y deserves to be treated less well than Y. Suppose, finally, that the morally deserving and the morally undeserving are socially salient groups. This would qualify as discrimination. But X ought morally to treat Y worse than Z; X does what he ought to do by "discriminating against" Y. It follows that, if we do not want to allow ourselves to say that we ought to discriminate against some groups, we will need to revise my account.

In practice, this objection might be ignored, of course. People differ greatly over the question of who is morally undeserving, and even when they do not, they rarely have sufficient information to tell of someone, with whom they interact, whether she is a morally deserving person. Hence, discrimination against the undeserving is unlikely for conceptual as well as empirical reasons. However, if we want our definition of discrimination to apply not just to actual, but also to possible, cases of discrimination, we need to address this challenge. Suppose that, due to an emerging moral consensus concerning standards of moral desert and technological improvements regarding detection of people's level thereof, what people morally deserve is as transparent as their gender. In that situation, it would be in no way odd to say that someone discriminates against the undeserving if he treats them worse than they deserve to be treated, thinks of them in terms of stereotypes, is insensitive to differences between different gradations of undeserving persons, and so on. Of course, if someone treated everyone exactly as he ought to treat them, he would, let us suppose, treat people differently. But barring a cosmic coincidence, his differential treatment would not be suitably explained by the differential group-membership of different individuals, but rather by their individual desert-levels, in which case it would not qualify as discrimination. If it were,

[60] Cf. Alexander (1992), "What Makes?" p. 159.

then the fit between the actual treatment and the deserved treatment could not be perfect.[61]

This completes my account of when a group is socially salient, which—just to recap—it is, if perceived membership of it is important to the structure of social interactions across a wide range of social contexts.

7. Because

The second need for clarification in relation to the concept of group discrimination is this: what does it mean to say that it is *because* (X believes that) Y has the property P and (believes that) Z has not, that X Φ-ies (and thereby treats Y worse than Z)? One suggestion is this:

(a) X treats Y worse than Z by Φ-ing because (X believes that) Y has P and (X believes that) Z does not have P if, and only if, the thought that Y has P and Z does not is part of X's motivating reason for Φ-ing.

Clearly, (a) fits certain paradigm cases of discrimination well—e.g., the hateful racist who refuses to hire someone because he is black. It also fits well the case of the discriminator, who is in no way biased, but nevertheless treats people disadvantageously on the basis of their gender, e.g., an employer bent on maximizing profit is less inclined to hire disabled people as salespersons because, on average, they are less effective in this capacity than nondisabled persons. But if we understand (iii) in the light of (a), we confine discrimination to what is normally called "direct discrimination." Yet it is now common to object to *indirect discrimination*, i.e., rules, institutions, and practices that have different impacts on different groups not only because, or perhaps not at all because, people following these rules, manning these institutions, or engaging in these practices are biased against members of the adversely affected groups or otherwise treat people from different socially salient groups differently, e.g., as in the case of differential treatment on the basis of reaction qualifications. (I return to indirect discrimination in chapter 2 and to reaction qualifications in chapter 9.)

To see what is at issue here, suppose that the members of an employers' association are accused of sex discrimination because a very small percentage of

[61] Suppose Y is very deserving and Z is very undeserving and yet X treats them in a way that is equally beneficial to them, because Z, unlike Y, is a member of a socially salient group with which X identifies. Surely, this constitutes discrimination against Y and yet my definition seems to suggest that it is not, because X acts in a way that is equally beneficial to them and, thus, do not treat Y differently from Z. However, as indicated above, X *does* treat Y and Z differently in one crucial sense, since he is less inclined to reward Z's deserts than Y's and this is suitably explained by their differential membership of socially salient groups. More generally, any putative counterexample of equal treatment constituting discrimination will also involve unequal treatment.

their CEOs are women. On behalf of its members the association responds that its members never refrain from hiring an applicant because this person is a woman, so for no member is it ever a part of their motivating state that they refrain from hiring an applicant because the applicant is a woman. However, the members have rules in place that prevent them from hiring applicants who are considered likely to request legally guaranteed leave in connection with childbirth, and it so happens that the persons who fall into the latter category are all women. The idea is, then, that while the employers' association members might or might not discriminate against people who need leave in connection with childbirth, they do not discriminate against women.

This defense might be dismissed as a piece of sophistry. Employers know that many women will at some point(s) in their life request leave in connection with childbirth, and that few men will, and for that reason they do discriminate against women. However, we have to be careful here. If the employers are defending themselves from the charge of *direct* discrimination against women, their response might be right. There are ways to test this. If, for instance, infertile women are also at a disadvantage relative to male CEO applicants, and if male applicants who are likely to require leave for other reasons, e.g., medical treatment, are not at a disadvantage relative to other male CEO applicants, it looks as though the employers are engaging in direct discrimination against women. However, if neither of these suppositions is correct, the defense against direct discrimination might well be sound. But none of this would show that the hiring procedures do not place women at a disadvantage in a way that qualifies as *indirect* discrimination (see chapter 2). To accommodate this we must broaden our interpretation of "because":

(b) X treats Y worse than Z by Φ-ing because (X believes) Y has P and (X believes) Z does not have P if, and only if, the fact that (X believes that) Y has P and (X believes) Z does not have P causally explains X's Φ-ing.

Since being a female job applicant causally explains whether an employer considers it likely that one will request leave in connection with childbirth, and since that in turn causally explains the employer's decision to hire someone else, the former indirectly, causally explains the latter. Hence (b) implies that, in the case we are imagining, employers do indirectly discriminate against women.

But now a different issue arises. For in our attempt to accommodate the notion of indirect discrimination we may have interpreted "because" too inclusively. Suppose we live in a truly color-blind society with a mixed population of Indians and Chinese, and in which there is equality of resources. Suppose also that there are statistical differences between Indians and Chinese with regard to how well they are represented in different sectors of industry. Suppose next that world markets change. New inventions mean that some sectors in which Indians

are better represented lose out, and as a result of those rules, institutions, and practices that make Indians better represented within the now-ailing segments of industry channel Indians into failing industries with the result that their average earnings fall below those of Chinese people. On the view proposed here, this would, implausibly, mean that Indians are now the victims of indirect discrimination. For the fact that the worse off group are Indians explains why various rules, institutions, and practices (let us call these "institutional arrangements") lead them into certain industries, and that in turn explains why they are worse off than Chinese.[62]

What seems to be lacking here is proper sensitivity to the fact that the existence of the institutional arrangements that work to the disadvantage of Indians is not causally explained, even partially, by direct discrimination against Indians. This explains why we think institutional arrangements in Western societies that presently work to the disadvantage of women and blacks often involve indirect discrimination. We tend to assume that the existence of the relevant arrangements is partly causally explained by direct discrimination against women and blacks. Thus, if we really thought that blacks never endured direct discrimination, would we then think of, say, word-of-mouth hiring practices that disfavor blacks as "discriminatory in operation" as opposed to simply "unjust in operation" because they make blacks worse off through no fault or choice of their own?[63] I doubt it. This suggests we need a third possibility:

(c) X treats Y worse than Z by Φ-ing because (X believes that) Y has P and (X believes that) Z does not have P if, and only if, (i) the thought that Y, and not Z, has P is part of X's direct, motivating reason for Φ-ing, or (ii) the fact that Y, and not Z, has P causally explains X's Φ-ing and this in turn is causally explained by the fact that people with P are often treated worse than those without P in the sense given by (i).

Construed in this way, the group discrimination account is vulnerable to the seemingly powerful challenge of income discrimination. Differential treatment on the basis of income and wealth is ubiquitous, differences in income and wealth have an important influence on social interactions in a wide range of contexts, and yet we rarely hear about income (or wealth) discrimination.

---

[62] I am assuming here that "rules, institutions and practices" are possible subjects of discrimination. Of course, one could deny that the present case involves discrimination because rules, institutions, and practices are not possible subjects of discrimination and, accordingly, there is no proper subject of discrimination and, thus, no discrimination. Doing so, however, would mean giving up on the notion of indirect discrimination and a core ambition of the enquiry here is to uncover what people refer to when they talk about "indirect discrimination."

[63] Compare Jorge L. A. Garcia (2001), "The Heart of Racism," in Bernard Boxill (ed.), *Race and Racism* (Oxford: Oxford University Press), 256–296, p. 284.

The suggestion that this is because one's level of income or wealth, unlike one's religion, race, or sex, is irrelevant to one's identity strikes me as a nonstarter.[64] Witness, for instance, the size of the market for status symbols, expressions such as "nouveaux riches," and the sense of humiliation and exclusion experienced by parents who are unable to pay for their children's participation in out-of-school activities painlessly funded by other children's wealthier parents.

Instead it might be suggested that when, say, the yearly membership fee at a fashionable golf club is $50 000, poor people are not discriminated against for the following reason: they are not being excluded from the club *because* they belong to the group of poor people, but because they do not pay the fee. There is a difference between a golf club that would admit a poor person were he so lucky as to find someone who would sponsor his club membership and the golf club that would not admit such a person despite his ability to pay the fee through someone else's generosity. The former does not discriminate against poor people; the latter does. On reflection, we would consider the latter kind of differential treatment to be income discrimination.

While this reply might work well for direct discrimination, it works much less well in relation to indirect discrimination. For suppose a company defends itself against the charge that its hiring requirements exclude disabled people by applying recruitment criteria that a person in a wheelchair would be unable to satisfy by pointing out that, if an exceptional disabled person were to satisfy the criteria, the company might well hire him. This reply does not seem to differ, in any relevant respect, from the one we imagined the golf club offering to defend itself against the charge of indirect discrimination against poor people—a reply tantamount to saying that the membership fee rule is neutral in intent even if it is discriminatory in operation. Since those who are concerned with indirect discrimination are unlikely to be persuaded by this kind of response in the case of disabled persons, it is unclear why they should be persuaded by a comparable reply to the charge of indirect income discrimination. It appears, then, that we must either adopt a much more restrictive view of indirect discrimination than is normally adopted, or we must put indirect, income discrimination on the list of important forms of indirect discrimination.

I myself am happy with the latter option. There are several reasons why we almost never talk about income discrimination, and these are consistent with

[64] Cf. ". . . the matrix of [the disadvantage of poverty] lies in the very fact of being poor, not in the fact of being perceived as belonging to the group of 'the poor,'" Daniel Sabbagh (2007) *Equality and Transparency: A Strategic Perspective on Affirmative Action in American Law* (New York: Palgrave Macmillan), p. 82. I might not disagree much with this claim if it is comparative, i.e., if it involves a certain claim about the relative size of the disadvantage of poverty as such compared to the size of the disadvantage of being perceived as belonging to the group of poor people, and not noncomparative, i.e., if it involves the noncomparative claim that there are no significant disadvantages from being *perceived* as belonging to the group of poor people.

its being pervasive. Three contenders here are: its elimination would require extremely drastic social changes; people overestimate the extent to which income differs from race, sex, and so on in being a mutable characteristic; and in many contexts the badness of having a low income does not lie primarily in the humiliation and stigma involved in being unable to afford certain goods—proper housing, food, and medical treatment—but simply in not being able to afford these goods. Hence, I stick with (c) as an account of what it is to treat someone disadvantageously because of some property this person has that, on the one hand, does not require that the ascription to this property to the discriminatee is part of the motivating reason for the differential treatment, but, on the other hand, requires differential treatment so motivated in the causal background to the relevant differential treatment.

8. Treatment

I now turn to the last of the three bits of my definition of group discrimination that is in need of elucidation. My definition implies that the discriminator treats the discriminatee in a certain way. But what is it for a discriminator to "treat" a discriminatee?[65]

There are many different discrimination-relevant ways of treating a discriminatee. Initially, one might distinguish between direct and indirect discrimination, where this distinction is understood to be exhaustive of forms of discrimination. Direct discrimination occurs whenever some of the phenomena described below take place. What is common to them is that they all involve representational items—e.g., desires, beliefs, statements, laws—that refer to, or otherwise distinguish between, those who are discriminated against and those who are not in the relevant discriminatory respect. Hence, to directly discriminate against, say, women you need to represent those you are discriminating against *as* women. Roughly, indirect discrimination occurs whenever an individual, institution, or practice acts (or is) in such a way that the interests of some individuals are systematically favored, and yet this does not involve direct discrimination. For instance, many observers would probably say that much of the employment discrimination against women that takes place nowadays is indirect—that women are discriminated against not primarily because men prefer to hire men, but because of the way the labor market is set up, and the connected difficulty of reconciling the demands of one's job with being

[65] There are some affinities between the discussion below and the related one about the "location problem" of racism. For the contours of this debate and a defense of the ecumenical view that racism can be located not only in behavior but also in cognitive as well as noncognitive states, see Joshua Glasgow (2009)," Racism as Disrespect," *Ethics* 120.1, 64–93, pp. 65–71.

pregnant and giving birth to children.[66] The distinction between direct and indirect discrimination is the main topic of chapter 2, so I shall not go deeper into it here.

Discrimination may or may not involve bias. A shopkeeper, say, who refuses to hire women might in no way be biased against women. This might be so if three conditions are satisfied: first, he believes that he maximizes profits by not hiring women; second, he would hold the same belief regarding women if evidence of a strength similar to what is available to him that suggest that he maximizes profits by not hiring women suggested that not hiring men maximizes profits; third, did he hold the belief referred to in the previous condition, he would refuse to hire men. Still, he does discriminate against women by treating them disadvantageously relative to men.

Direct as well as indirect discrimination may involve structural elements, where by "social structure" I mean a regular pattern of social behavior that supervenes on the behavior or dispositions of the members of the relevant social unit. An informal social norm to the effect that girls should not strive to attain higher education or a law to the effect that women are not allowed to drive cars would be examples of direct discrimination involving structural elements. A law to the effect that employees receive 80% of the wage up until $25,000 per year when on parental leave may illustrate indirect structural discrimination, in view of the fact that men are more likely than women to earn more than $25,000 and, accordingly, the law will in effect make it more costly for men than for women to take parental leave.

Direct discrimination involving biases divides into cognitive and noncognitive discrimination.[67] Direct cognitive discrimination occurs when a subject's beliefs are formed in a way that is biased toward certain people, e.g., when comparable evidence for men's and women's ability to become good managers will lead the subject to believe that men are capable of becoming good managers

[66] Suppose that, from the standpoint of the individual employer, it is a matter of "business necessity" that nearly only men are hired to operate the relevant machines, but that the manufacturer who produced the machines could easily and costlessly have designed them otherwise and would have done so in the absence of male dominance. Cf. *Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971), quoted in Gertrude Ezorsky (1991), *Racism and Justice: The Case for Affirmative Action* (Ithaca, NY: Cornell University Press), pp. 39–40. It would seem arbitrary to consider a situation morally objectionable in regard to discrimination if people are hired on the basis of requirements that, although nondiscriminatory in intent, are "discriminatory in operation" and that do not reflect "business necessity," but not consider the state of affairs just described morally objectionable from the point of view of discrimination.

[67] This distinction applies directly only to cases of discrimination in which the relevant representation is a mental state. However, even nonmental representations—e.g., laws—involve mental states in an indirect way, and for that reason one can also, in those cases, distinguish between cognitive and noncognitive discrimination. Thus a discriminatory law that does not recognize the possibility of rape in marriage involves mental states in the sense that it expresses or evinces the mental states of those who proposed and successfully voted in favor of the law.

but will not persuade the subject that women are similarly capable, or when this bias manifests itself in behavior.[68] Subjects are rarely aware that they are cognitive discriminators—it is hard to believe that *p* if you are aware that your assessment of the evidence supporting *p* (and not-*p*) is biased toward believing that *p*—although this is possible.[69] Cognitive discrimination is differential treatment that involves being less inclined to form beliefs somehow (un)favorable to Y than to form beliefs (un)favorable to Z given the same evidence for these beliefs. A job applicant may be a cognitive discriminator if, given the same evidence of excellence, he is more inclined to believe that one company, managed by men, will be successful than he is to believe that another, managed by women, will be.[70] The same applicant is a noncognitive discriminator if he is less inclined to apply for a job at a company with a female executive manager because of a brute desire not to have a female boss.

Often cognitive and noncognitive discrimination go hand in hand, but conceptually speaking they can easily be pried apart. In principle, one can engage in cognitive discrimination in favor of a group while at the same time engaging in noncognitive discrimination against that group, and vice versa. Consider someone who always overgeneralizes and thinks of people in terms of stereotypes. Being brought up in a culture of guilt, he is more inclined to do so when it comes to people of his own kind. So while he makes ungrounded and unfavorable generalizations about people of, say, other religions, he makes even more ungrounded and unfavorable generalizations about people of his own religion. This person is involved in cognitive discrimination *in favor of* people with a religion different from his own. Offhand, this may seem odd. Normally, we assume that a person who makes ungrounded and unfavorable generalizations about people of religious denominations different from his own is involved in *negative*, cognitive discrimination. This, however, is because we assume that, normally, people are not inclined to make ungrounded and unfavorable generalizations about people of their own kind.

---

[68] I say "direct" to allow for indirect cognitive discrimination. Such discrimination may occur even in the absence of direct cognitive discrimination, e.g., when individuals apply norms of enquiry in an unbiased way but where the norms themselves are biased. For an example of structural, cognitive discrimination based on the differential cognitive framework of white and black members of the jury, see Wasserstrom (1977), "Racism," p. 598.

[69] Reverse cognitive discrimination can occur when the agent's desire to see two groups as identical in certain factual respects skews his evaluation of the evidence in this direction. Such a case need not involve noncognitive discrimination.

[70] Elizabeth Anderson (2010), *The Imperative of Integration* (Princeton, NJ: Princeton University Press), p. 59, believes that we tend to "overlook the roles of implicit and automatic cognition, which cause discriminatory treatment even in the absence of discriminatory beliefs or a conscious intention to discriminate." One interesting finding supporting this view is that research subjects are more likely to remember shortcomings of out-group individuals than of in-group individuals despite the fact that the initial assignment of research subjects to groups was randomized; see Sabbagh (2007), *Equality and Transparency*, p. 106.

Noncognitive discrimination obtains when a subject's desires and values are biased, or when this kind of bias manifests itself in behavior.[71] The bias is transparent when the content of the subject's desires and values refers to the group that he discriminates against, e.g., when the subject prefers the company of members of his own race to members of other races or the "subjection to androcentric norms in relation to which women appear lesser or deviant and that work to disadvantage them, even in the absence of any intention to discriminate."[72] Noncognitive bias is nontransparent when this is not the case, e.g., when the subject finds certain ways of avoiding taxation worse than others and the explanation of this is that the "worse" activity is mostly engaged in by members of a certain immigrant community, whereas comparable activities that he judges "less bad" are mostly engaged in by members of the nonimmigrant majority to which he belongs.

Direct discrimination need not involve negative attitudes toward the discriminatees, as is shown by cases involving reaction qualifications. A shopkeeper may discriminate against black applicants, not because he dislikes them, but simply because he wants to earn a profit and, justifiably or not, thinks that, given the racist bias of many of his customers, he is more likely to do so by hiring white assistants. Here his desire not to hire blacks does not reflect an intrinsically racist attitude. Similarly, according to my definition, affirmative action involves discrimination against majority members even though it is unlikely to involve any negative attitudes to, say, majority members.[73]

A distinction cutting across the distinction between cognitive and noncognitive differential treatment is that between valuation-based and non-valuation-based discrimination. Discrimination is valuation-based if, and only if, it satisfies either of the following two conditions: either it reflects the subject's view that X's interests count for more, morally speaking, than Y's; or it reflects the subject's view that, while their interests count equally, X and Y should not relate to each other in the way they should relate to members of their own group, e.g., as in the case of those who subscribe to traditional family values and, accordingly, hold that while the interests of men and women count equally, there is a gender-determined division of labor. I shall refer to the first

---

[71] By "desires" and "values" in the present context I have intrinsic as well as extrinsic desires and values in mind.

[72] Nancy Fraser (1997), *Justice Interruptus* (New York: Routledge), p. 20.

[73] Hence, there is a sense in which, say, standard racial discrimination reflecting racial hostility is discrimination on a different basis than racial discrimination reflecting the aim of reducing unjust racial disadvantage, see Anderson (2010), *The Imperative*, pp. 158–160; see also Sabbagh (2007), *Equality and Transparency*, pp. 13–15. However, logically speaking, the basis could be the same in the two cases. Imagine a strange neo-Nietzschean who is hostile toward members of a certain racialized group for their having suffered unjust racial disadvantages as a result of, and including, being falsely represented as having certain negative traits traceable to causes intrinsic to the group. In one sense, this appalling person discriminates *against* these people on the same basis as the proponents of affirmative action, who discriminates *in favor* of them.

variety of valuation-based discrimination as *hierarchical* and the second as *nonhierarchical*.

Discrimination is non-valuation-based when the subject (i) judges that the interests of the group against which he discriminates are no less important than the interests of the group against which he does not discriminate (almost always the group to which he himself belongs), (ii) does not judge that members of the relevant groups should not relate to one another in the way that they relate to members of their own group, and yet (iii) prefers to marry, employ, accompany etc. people with whom he shares (or does not share) his own group identity. In this case the discriminator has brute discriminating desires. A brute discriminating desire is hierarchical, when it is a desire that members of some group will become, or remain, better off or otherwise superior than others.

Valuation-based and nonvaluation-based discrimination are often confused. Yet, it is important to see that they differ. The way in which they do so depends on the account of valuation we accept. On an influential account, valuation can be explained in terms of second-order desires, i.e., desires with the individual's own desires as their objects.[74] Appealing to this account, we have a case of a valuational nonracial discriminator but nonvaluational racial discriminator in someone who desires to consort with persons of his or her own race, and yet has a second-order desire not to act on this desire.

A final distinction, which I would like to introduce, is that between first- and second-order discrimination. Second-order discrimination is discrimination where the second-order discriminator treats the discrimination performed by different discriminators differently, whereas first-order discrimination is discrimination where the first-order discriminator treats the discriminatee differently where this is not related to discriminatory action on part of the latter. An illustration of second-order discrimination would be one in which people respond differently to gender discrimination depending on which group the discriminator is a member of, e.g., they are indignant about gender discrimination by Muslim immigrants but complacent about gender discrimination by non-Muslim, European men, or indignant about discrimination against Muslims but complacent about discrimination against non-Muslims. As the examples indicate second-order discrimination is not uncommon.[75]

---

[74] David Lewis (1989), "Dispositional Theories of Value," *Proceedings of the Aristotelian Society* suppl. 63, pp. 113–137; Harry G. Frankfurt (1988), *The Importance of What We Care About* (Cambridge: Cambridge University Press), 11–25.

[75] Adrian Piper (2003) uses the term "second-order discrimination" to refer to the phenomenon whereby an agent's disposition to discriminate against people because of their having a certain property, P_1, causes the agent to be disposed to discriminate against people because of their having another property, P_2, the having of which is somehow associated with the having of P_1, in his "Two Kinds of Discrimination" in Bernard Boxill (ed.), *Race and Racism* (Oxford: Oxford University Press), 193–237, p. 215.

The taxonomy that has been introduced in this section is summed up, in part, in the figure below. As indicated by " . . . . " the figure is incomplete from level four and onward. So in my view there are 64 branches at level seven, even if I have only drawn two of them:

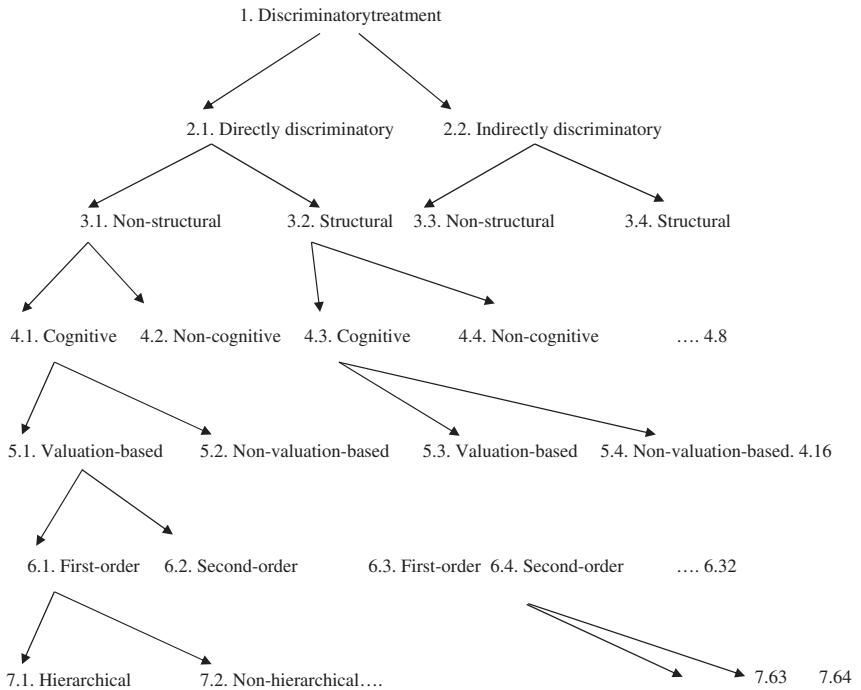

9. Summary

In this chapter I have identified the notion of generic discrimination. "Discrimination" is sometimes used to refer to generic discrimination, but it is important to note that when people complain about discrimination they have in mind a more specific form of discrimination, one that is at least prima facie wrong. I went over three such candidates: irrelevance discrimination, the moralized concept of discrimination, and group discrimination. I defined group discrimination as follows:

Group discrimination
X discriminates against Y in relation to Z by Φ-ing if, and only if, (i) there is a property, P, such that (X believes that) Y has P and (X believes that) Z does not have P, (ii) X treats Y worse than Z by Φ-ing, (iii) it is because (X believes that) Y has P and (X believes that) Z does

> not have P that X treats Y worse than Z by Φ-ing, (iv″) P is the property of being member of a certain socially salient group (to which Z does not belong), and (v) Φ-ing is a relevant type of act etc., and there are many acts etc. of this type, and this fact makes people with P (or some subgroup of these people) worse off relative to others, or Φ-ing is a relevant type of act etc., and many acts etc. of this type would make people with P worse off relative to others, or X's Φ-ing is motivated by animosity towards individuals with P or by the belief that individuals who have P are inferior or ought not to intermingle with others.

One important feature of this concept of discrimination is that it is not moralized, i.e., it does not follow by virtue of the meaning of "group discrimination" that if something is group discrimination, then it is morally wrong or unjust. Having defined group discrimination I clarified (i) what it means for a group to be socially salient; (ii) what it means to treat someone disadvantageously because of her membership of a certain group: (iii) and distinguished between some of the main forms of disadvantageous treatment.

In what follows I shall focus on the concept of group discrimination. In so doing, I am not claiming that this is the only concept of discrimination we have—as just noted, in this chapter I have also introduced the concepts of generic, irrelevance, legal, and moralized discrimination—or that it is, in some sense, the correct concept of discrimination. However, I do think it captures a good deal of what people have in mind when they talk about discrimination—most actual complaints of discrimination are made on behalf of groups that are socially salient—and that it is useful for analytical purposes. Also, I have argued that the concept of group discrimination allows us to make sense of many features of our thinking about discrimination. I will return to how this is the case in chapter 6. Finally, I have argued that discrimination may take very many different forms, e.g., it may or may not be evaluation based. Those who reject my views about the usefulness and importance of the concept of group discrimination must provide some alternative account of what distinguishes discrimination in the relevant sense from mere differential treatment. The challenge of income discrimination, which I mentioned in section 7, is a specific and undertheorized version of this more general challenge that, like the definition challenge, any comprehensive account of discrimination must face. If I am right, discrimination against poor people is an important form of group discrimination, which has undeservedly received much less attention than, say, racial and sex discrimination. It is interesting, at least in part, because many people are strongly opposed to race and sex discrimination, but quite tolerant of inequalities of income. The present chapter motivates the question of whether this stance is coherent.

Appendix 1: Methodology

Ideally, we want to account for the descriptive meaning of "discrimination" in terms of the necessary and sufficient conditions of something being discriminatory. The endeavor to do so faces several challenges. First, as indicated, we use the term "discrimination" to refer to different concepts. This complicates matters, but represents no obstacle in principle for giving the necessary and sufficient conditions for something being discrimination. It simply requires, as has been done throughout this chapter, clarifying which sense of "discrimination" one has in mind. We might even introduce new concepts of "discrimination" and some of them might serve certain theoretical purposes better than existing ones, e.g., they might capture better a specific kind of moral wrong.[76]

Second, even when we discipline our definitional exercise in this way, we might face the problem that no such conditions exist. It might be that, to some extent, *anyone's* use of "discrimination" is guided by fruitful analogies and interesting similarities such that the descriptive meaning of "discrimination" is open-textured, flexible, and unstable over time.[77]

In the light of these observations, I believe that accounts of discrimination such as the one I have offered in this chapter should satisfy the following four desiderata. First, while no account can be expected to fit everyone's classificatory dispositions perfectly, a definition can fit people's use of the term better or worse and the better fit with normal use of the term, the better. As I have indicated, it speaks in favor of group discrimination that most complaints about discrimination are complaints about the way in which members of certain socially salient groups have been treated.

Second, if, as I suspect, there is a common structural core to the descriptive meaning of "discrimination," then it will be desirable for an account of discrimination to capture this core. I take it, for instance, that the claim that discrimination amounts to differential treatment is something that is invariant across pretty much anyone's use of the term, whichever concept of discrimination one has in mind.

Third, apart from capturing structural features, accounts of discrimination should be descriptively enlightening and, for that reason and the one stated in the next paragraph, the definition I am after is not purely descriptive. For instance, if X constitutes discrimination, and if Y by way of the very same social mechanisms renders some people less well off in the very same respects—see chapter 6.6 on stigma and the cumulative nature of discriminatory harms—I want, ceteris paribus, my account to imply that Y constitutes discrimination as well, regardless

[76] The fact that I focus on group discrimination in this book, hence, does not reflect that I am not a conceptual pluralist about discrimination, i.e., that I deny that there is a plurality of useful and equally legitimate concepts of discrimination.

[77] Recall the reference to Wittgenstein's concept of family resemblance in note 1, this chapter.

of whether X and Y are very different morally speaking. This desideratum reflects the autonomy that the descriptive meaning of "discrimination" has in certain contexts and speaks against building immorality or injustice into the very definition of discrimination (because it may have different effects, morally speaking, that some people are made less well-off through a certain mechanism).[78]

Finally, an account of discrimination should be morally enlightening. It should allow us to see, e.g., why different forms of discrimination differ morally (if such is the case). It should help us to see what makes discrimination morally wrong, and why certain practices that are not usually seen as morally analogous to discrimination really should be seen as such and should for that reason be seen as discriminatory—recall the discussion of income discrimination in chapter 1.7. For reasons already indicated, to satisfy this desideratum, a sense of discrimination need not logically entail any evaluative commitment. It suffices if the relevant sense is such that to say that something is discriminatory indicates that it is prima facie wrong, because there is a significant empirical correlation between something being discriminatory and something being wrong. Specifically, the notion of social salience has particular relevance here, because it ties discrimination to a particular kind of harm—stigmatic harms—that are significantly boosted when the discriminatees form a socially salient group. So while my account is neutral about which account of the wrongness of discrimination is the correct one, it highlights a particular feature of discrimination, which is very relevant to the harm-based account of the wrongness of discrimination (see this chapter and chapter 6).[79]

These desiderata are quite vague and, though consistent, will quite possibly point in different directions, when applied to particular cases.[80] However, by

[78] It might be objected that we should not focus so much on the use of the term "discrimination," which is likely to be muddy and confused. Rather, we should focus on which role the concept to which the term refers plays in our scheme of thought, and that role, so the challenge continues, is to identify an injustice of a certain kind. In response, I concede that if we could have a clear grasp of the relevant conceptual role of discrimination prior to and independently of our use of the term "discrimination," then this might be the right way to go. However, a significant part of whatever arguments one can give for the identification of such a role will appeal to facts about word usage. Hence, appeals to conceptual roles are not a way of bypassing the muddiness and confusion of ordinary usage. That being said, it is not as if my account has not in any way been guided by an understanding of what role the relevant concept of discrimination plays in our thinking. Indeed, I said that the concept of discrimination we are interested in is more specific than generic discrimination, because the relevant concept is one, which is used to articulate a moral complaint, and generic discrimination will not do for this purpose.

[79] I admit that the case of animals represents a complication here.

[80] One desideratum that an account of discrimination should not answer to is that the individuals which it implies can be targets of discrimination overlap with those individuals that ought, morally speaking, to be legally protected against disadvantageous differential treatment. My desert accommodating prioritarianism in chapter 6 is the guiding principle here (see the application of it to the issue of the morally desirable legal status of discrimination in chapter 10). Moreover, on the account offered above which groups actually enjoy legal protection against discrimination plays a role in determining which account of discrimination to accept through the ordinary usage desideratum.

listing them, my aim has been to indicate what general considerations my account is informed by and which criticisms I commit myself in advance to consider damaging to my account. Specifically, I hope to have made it clear that my account is not intended as a purely descriptive one that merely sums up ordinary language use of the term "discrimination," and which, thus, will be refuted by any instance of a use of the term that does not conform to my account (see the appendix on genetic discrimination in chapter 3). While it is indeed intended, in part, to capture certain central uses of the term, it is also intended to serve certain theoretical purposes as well and, thus, to serve an explicative function as well.[81]

Appendix 2: Discrimination Skeptics: Oppression and Dominance

According to some theorists, the kinds of morally problematic social relations that are often thought of as employing the concept of discrimination are better thought of as employing alternative concepts, e.g., oppression or dominance. According to them, while the concept of discrimination may be well defined and instantiated, it is nevertheless not very relevant. Prominent theorists posing this skeptical challenge to the concept of discrimination include Iris Marion Young and Catharine A. Mackinnon. In part, to see whether the present book-length enquiry is justified, I want to assess the strengths and weaknesses of their critiques. My general response is that they are based on a *very* narrow conception of discrimination, e.g., as direct, evaluation-based, fault-ridden differential treatment. But once we have in mind, say, the wide extension of the group-based concept of discrimination identified above, these irrelevance-based criticisms can be seen to be compatible with the project of this book. To some extent the apparent conflict between the approach taken here and that of Young and Mackinnon may simply reflect terminological variation. Hence, there is no reason to discard a suitably broad concept of discrimination as being irrelevant to social or normative analysis.

Let us first consider Iris Marion Young's view. She thinks that group injustice should primarily be thought of in terms of oppression and domination rather than in terms of discrimination. Commenting on a so-called paradox suggested by Robert Fullinwider—"If we do not use preferential hiring, we permit discrimination to exist. But preferential hiring is also discrimination. Thus, if we use preferential hiring, we also permit discrimination to exist. The dilemma is that whatever we do, we permit discrimination"—she resists the dilemma, in part because it involves—so she thinks—equivocation on the term "discrimination," in part because it disappears "if proponents of affirmative

[81] Glasgow (2009) makes a similar point in his "Racism as Disrespect," pp. 64–65.

action abandon the assumption that nondiscrimination is a paramount principle of justice, and stop assuming that racial and sexual injustice must come under the conception of discrimination."[82]

In her view, "[s]upporters of affirmative action policies would be less on the defensive . . . if they positively acknowledged that these policies discriminate" and then in turn denied that "discrimination is the only or primary wrong that groups suffer.[83] Oppression, not discrimination, is the primary concept for naming group-related injustice."[84] One problem with discrimination is that it is "primarily an agent-oriented, fault-oriented concept. Thus it tends to focus attention on the perpetrator and a particular action or policy, rather than on victims and their situation. . . . Identifying group-based injustice with discrimination tends to put the onus on the victims to prove a harm is done, case by case. As a concept of fault, moreover, discrimination tends to present the injustice groups suffer as aberrant, the exception rather than the rule. In its focus on individual agents, the concept of discrimination obscures and even tends to deny the structural and institutional framework of oppression."[85]

Young concedes that there is "an expanded concept of discrimination"—presumably she has in mind one that covers direct as well as indirect discrimination—that involves a "results test rather than an intent test of discrimination," which may avoid some of the problems that the narrower concept of discrimination involves. However, she believes that this broader concept has failed to win many friends. Accordingly, it is a "better strategy for addressing the injustice suffered by disadvantaged groups . . . to restrict the concept of discrimination to intentional and explicitly formulated policies of exclusion or preference, and to argue that discrimination is not the only or necessarily the primary wrong that women and people of color suffer."[86] That primary wrong is the wrong of oppression, which she defines as "systematic institutional processes which prevent some people from learning and using satisfying and expansive skills in socially recognized settings, or institutional social processes which inhibit people's ability to play and communicate with others and to express their feelings and perspective on social life where others can listen."[87]

---

[82] E.g., Young (1990), *Justice*, p. 195. The dilemma, if indeed there is one in the first place, does not go away if we follow Young's prescription—i.e., it will still be the case that we permit discrimination whatever we do—but the dilemma may disappear in the practical sense that we will have less reason to worry about it.

[83] Young thinks that discrimination in the sense that affirmative action is discrimination—"conscious preferential practices that favor group members on ground of the group membership," Young (1990), *Justice*, p. 195—"may or may not be wrong, depending on its purpose," Young (1990), *Justice*, p. 197.

[84] Young (1990), *Justice*, p. 195.

[85] Young (1990), *Justice*, pp. 195–196.

[86] Young (1990), *Justice*, p. 196.

[87] Young (1990), *Justice*, p. 38.

The first thing to note about Young's critique of the concept of discrimination is that it seems motivated by the concern to address injustice suffered by certain groups and the conjecture that the concept of oppression would be better to employ in instrumental terms. This line of critique is compatible with the view that, as a matter of fact, all group injustices fall under the concept of discrimination. It is just that the instrumentally best way of remedying or eliminating these injustices involves employing concepts other than the concept of discrimination (see my remarks in the introduction on the relation between political philosophy and political activism). In any case, it is a very complex issue which concept is instrumentally best in terms of addressing group injustice, and Young does very little to substantiate her claims. Offhand, it seems that discrimination is a concept that has won *more* favor and accordingly that it (or its use) has greater efficiency in addressing injustices—in ordinary as well as legal discourse—than the (use of the) concept of oppression when it comes to describing the plight of women and, in Young's and others' awkward expression, people of color.

Second, her claims about the concept of discrimination are false. None of the concepts of discrimination that I have presented in this chapter focuses on faults. Indeed, some of them, including group discrimination, do not even imply that discrimination is wrong. Also, while the concept of discrimination is agent-requiring as I have put it, nothing in the concept of discrimination prevents the agent of discrimination from being an institution or a social structure, e.g., there is nothing infelicitous about the expression "a discriminating social norm." If, indeed, this is the basis for saying that the concept of discrimination is agent-oriented, it should be noted that the concept is also object-(or victim-) requiring, so, by parity of reasoning, one could just as well argue that the concept of discrimination is victim-oriented. Finally, the concept of discrimination in itself is entirely neutral on whether discrimination is the "exception rather than the rule" or whether it is systematic. In my definition, one would not have to retract a claim about group discrimination upon learning that the relevant differential treatment is systematic and widespread. Whether it is, is an empirical question to which the concept of discrimination simply does not speak.

Finally, most processes that would qualify as oppressive in Young's account involve group discrimination. Because oppression is rarely universal, presumably, the relevant oppressive institutional processes will involve differential treatment of different people on the basis of different group membership and, thus, qualify as discriminatory relationships on my account. In sum, Young's critique of the view that discrimination has an important role to play fails. However, this leaves it open for me to agree with her on many substantive points, since what I have taken issue with is her very narrow construal of the concept of discrimination.

Another critic of the concept of discrimination is Catherine Mackinnon. She contrasts the "difference approach" to sexual equality with the "dominance

approach." The difference approach focuses on sex discrimination and contends that sexual equality obtains when no sexual discrimination exists, i.e., when men and women can compete for the same jobs on an equal footing in the sense that considerations about the applicants' gender never in themselves play any role in employers' hiring decisions. Or to put this in another way: when employers simply base hiring decisions on qualifications. The limitation of the difference approach, according to Mackinnon, lies in the fact that men have defined jobs. So even if all arbitrary sex discrimination is eliminated such that men and women are judged by the same criteria on the job market, this does not give us true equality because these criteria are designed by men for the benefit of men. As MacKinnon puts it: "Men's physiology defines most sports, their needs define auto and health insurance coverage, their socially-designed biographies define workplace expectations and successful career patterns, their perspectives and concerns define quality in scholarship, their experiences and obsessions define merit . . . For each of their differences from women, what amounts to an affirmative action plan is in effect, otherwise known as the structure and values of American society."[88] The dominance approach is alert to this situation and to the obstacles to genuine gender equality that it poses.

Again, I do not want to take issue with the substantive points that Mackinnon makes in relation to her favored dominance approach—indeed, I share her view that the mere fact that, in some sense, men and women are judged by the same criteria does not guarantee the absence of unjust inequality (see chapter 9). Rather, I want to point out that the concept of discrimination she has in mind is very narrow and certainly much narrower than group discrimination. Accordingly, the barriers to genuine equality that she identifies would qualify as group discrimination as I have defined it. Take, for instance, the design of attractive jobs in such a way that few people with primary caretaker responsibilities are able to fill them. In my account this may well qualify as indirect group discrimination. Hence, there is no reason why one could not express Mackinnon's complaint in terms of discrimination and, thus, no reason to think that the concept of discrimination is less relevant from the point of view of gender equality.[89]

The basic point is that discrimination can be defined quite narrowly. If it is so defined, there are likely to be analytical aims that cannot be pursued

[88] Catherine Mackinnon (1987), *Feminism Unmodified* (Cambridge, MA: Harvard University Press), p. 36; Fredman (2011) *Discrimination Law*, pp. 11, 30.

[89] A similar point applies to moralized discrimination, assuming that it is morally wrong to design roles and positions in society for the benefit of members of one sex only. The notion of relevance discrimination seems most vulnerable to Mackinnon's critique, but even here one can describe the designing of positions in a way that uniquely suits men as discrimination provided that there are no relevant reasons favoring this one-sided design over one that also accommodates women.

and moral complaints that cannot be expressed in terms of discrimination. However, critics of discrimination who employ a narrow concept of discrimination may not have any substantive disagreements with noncritics who employ a broader concept. Both Young and Mackinnon—two influential critics of discrimination—employ quite narrow concepts of discrimination. They might not—should not, I think—disagree with or find irrelevant a project, such as the present one, that employs a much broader concept of discrimination.

{ 2 }

Indirect Discrimination

1. The Distinction between Direct and Indirect Discrimination

Suppose that an act, policy, or practice in no way reflects any bias on part of the agent.[1] Could it nevertheless involve discrimination? According to a common view, often reflected in legal regulations, it could, because discrimination can be indirect.[2] In the United States, courts have acknowledged indirect discrimination, since the famous 1971 *Griggs v. Duke Power* case in which the US Supreme Court ruled that employment policies of private businesses are indirectly discriminatory against African Americans and thus prohibited, "[i]f an employment practice which operates to exclude Negroes cannot be shown to be related to job performance."[3] This ruling has had a very significant effect on affirmative action policies in the United States even if, in effect, later Supreme Court decisions have ruled against some of these policies. In Europe, in a significant recent ruling the European Court of Human Rights (ECHR) held that "[w]hen a general policy or measure has disproportionately prejudicial effects on a particular group, it is not excluded that this may be considered as discriminatory notwithstanding that it is not specifically aimed or directed at that group."[4]

---

[1] In the interest of brevity I shall often simply refer to indirectly discriminatory acts, rather than to indirectly discriminatory acts, policies, or practices, in what follows.

[2] E.g., Gertrude Ezorsky (1991), *Racism and Justice: The Case for Affirmative Action* (Ithaca, NY: Cornell University Press), pp. 9–27.

[3] *Griggs v. Duke Power* 1971: 431.

[4] *Shanaghan v. U.K.* 2001: para. 129; cf. Sandra Fredman (2011), *Discrimination Law* (Oxford: Oxford University Press), p. 154: "broadly speaking" the concept of indirect discrimination "has three elements: equal treatment; a disproportionately exclusionary impact on those sharing a protected characteristic; and the absence of an acceptable justification." In my view, it is not conceptually impossible to indirectly discriminate against members of a certain group whom one, in one sense, treats more advantageously than others. Also, I am unsure if Fredman's last two elements are logically independent: Could an exclusionary impact be disproportionate in the presence of an acceptable justification of it?

One central task of an account of discrimination is to define indirect discrimination. Call this the *indirect discrimination definition challenge*. A rough answer says that indirect discrimination occurs whenever an act imposes a disproportionate burden on a group of people, where this does not involve *any* bias or discriminatory intention on part of the agent, e.g., when the agent is simply unaware of and indifferent to the effects of her act. Acknowledging that discrimination may exist even in the absence of biases and discriminatory intentions is hugely significant in a legal and, more broadly, political context, because biases and intentions are often hard to prove, whereas it is easy to show that a certain group is "excluded" or suffers from "disproportionately prejudicial effects," when what this requires is showing that the group putatively subjected to indirect discrimination is disproportionately represented, statistically speaking. (I return to the relationship between discrimination, on the one hand, and statistical under- and overrepresentation, on the other hand, in chapters 7 and 8.)

Indirect discrimination is very important today, probably more so than it was in the past. In cultures where discrimination has been on the political agenda for almost 50 years, most find many forms of discrimination to be unjust and would consider it a shameful moral flaw if it turned out or became apparent that they themselves engaged in direct discrimination on the basis of race, gender, or religion. However, many argue that this does not mean that discrimination has been eradicated. It simply means that discrimination now often takes a form where it exists independently of discriminatory intentions, i.e., as indirect dicrimination.[5] In order to assess this claim, we need a firmer grip on the notion of indirect discrimination than the one we presently have.

Terminology can be misleading here. It is natural to see direct discrimination as differential treatment on the basis of a distinction the discriminator is concerned with (e.g., "No women need apply"), and indirect discrimination as differential treatment on the basis of a proxy for the distinction that the discriminator is concerned with, e.g., a sexist might use height as a proxy for gender to discriminate against women, while not caring one iota about height in itself.[6] Yet, intentions—or more broadly, biases—play a central role in the direct-indirect distinction as I expound it below and, accordingly, I would classify both as cases of direct discrimination.

In this chapter I explore the concept of indirect discrimination and how it differs from direct discrimination. I take my point of departure in a definition of indirect discrimination proposed by Andrew Altman in an entry on discrimination in the *Stanford Encyclopedia of Philosophy* (see section 2) and use

[5] This is not to imply some sort of functional explanation of indirect discrimination.

[6] Ben Eidelson (2011), *What is Discrimination and When is it Wrong?*, D.Phil. thesis, Oxford University, pp. 45–47.

it as a springboard to set out some of the complicated issues and challenges that pertain to indirect discrimination (sections 3 to 6). On this basis, I propose a significantly revised definition of indirect discrimination (section 6), which I then relate to my definition of group discrimination in chapter 1 (section 7). One significant claim of this chapter is that Altman, like many others, employs a heterogeneous distinction between direct and indirect discrimination, i.e., the two forms of discrimination are defined along different dimensions. Another significant claim is that we should probably define indirect discrimination in a way that omits any reference to the relevant disparate impact involved in indirect discrimination being, in a sense to be expounded below, "disproportionate".

2. Altman's Definition

Altman distinguishes between direct and indirect discrimination. Direct discrimination occurs in cases where disadvantageous "acts or policies" are aimed at certain socially salient groups. Indirect discrimination occurs in cases where acts or policies are "not aimed—explicitly or surreptitiously, consciously or unconsciously—at persons for being members of a certain social[ly salient] group," but where these acts or policies "have the effect of disproportionately disadvantaging the members of a particular [socially salient] group."[7] Possibly strengthening some of Altman's formulations we can express this slightly more precisely as follows:[8]

A policy or act is indirectly discriminatory against a certain group, G, if, and only if:

(i) It neither explicitly targets nor is intended to disadvantage members of G *qua* being members of G (*the no-intention condition*).[9]
(ii) It disadvantages members of G relative to members of other relevant, socially salient groups (*the disadvantage condition*).
(iii) The relevant disadvantages are disproportionate (*the disproportionateness condition*).
(iv) G is a socially salient group (*the social salience condition*).

---

[7] Andrew Altman (2011), "Discrimination," in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, http://plato.stanford.edu/archives/spr2011/entries/discrimination/, p. 8.

[8] For instance, Altman's formulation may simply offer sufficient conditions for indirect discrimination.

[9] I take it that the bit about "disproportionately disadvantaging" in Altman's definition contrasts with the "not aimed at" bit, thus, offering information about how to interpret what Altman means by not aiming at. Also, Altman (2011), "Discrimination," p. 7, writes: "Without the intent to disadvantage persons based on their race, sex, religion, and so on, there is no direct discrimination; with such an intent to disadvantage, there is direct discrimination . . ."

A company that makes promotion conditional on having a university degree may engage in indirect discrimination against a certain group if: (i) this requirement neither explicitly targets, nor is intended to disadvantage members of this group; (ii) the requirement disadvantages them relative to nonmembers, because fewer members than nonmembers have a university degree; (iii) the requirement is disproportionate in the sense that the achievement of the goals promoted by such a requirement is not sufficiently important to justify the bad effects, comparatively speaking, on members of this group; (iv) and the relevant group is a socially salient one.

Not only must the discriminatee be (a member of) a socially salient group. Those groups relative to which the discriminatee is disadvantaged must themselves be socially salient groups and not, say, the artificial group of people who are better off than the discriminatees as a result of the putatively indirectly discriminating policy. They also need to be a relevantly contrasting socially salient group (men-women, black-white, gay-straight etc.), e.g., one does not establish that women are subjected to indirect discrimination by showing that they are disadvantaged compared to immigrants even if all immigrants are men.[10] The group of individuals who indirectly discriminate need not form a socially salient group: e.g., if it is common that women are subjected to disadvantageous treatment, but that the people who discriminate against women do not form a socially salient group, e.g., if it is not men, but people with even passport numbers, who treat them disadvantageously, women might still be victims of indirect discrimination.

Altman's definition captures central elements in many of the other definitions and characterizations of indirect discrimination that can be found in the literature and has the virtue of being clear.[11] In sections 3 to 5, I comment in numerical order on each of its first three conditions. My aim is to work through various problems with Altman's definition to reach a better one. This will help us later in the book—see chapter 6—to see which moral significance the charge of indirect discrimination has.

As the discussion in chapter 1 shows, I have no quarrel with the social salience condition as such. However, it should be noted that identifying the relevant object of indirect discrimination—call this *the discriminatee identification challenge*—is even less straightforward in the case of indirect discrimination. In the case of direct discrimination the mental states of the discriminator may

[10] Hence, by definition, if a social group has no relevantly contrasting social group, it cannot be subjected to indirect discrimination. As it happens, all socially salient groups I can think of do have relevantly contrasting, socially salient groups. However, it might be suggested that the disadvantaged-relative-to-a-relevantly-contrasting-socially-salient-group-condition is too strong, because it suffices if the socially salient group being subjected to indirect discrimination is worse off than, say, the population average as a result of indirect discrimination.

[11] I compare Altman's definition to some other definitions and characterizations in Appendix 1.

serve as a basis for determining which groups suffer discrimination—say the discriminator is a misogynist and accordingly the group of people whom he discriminates against is women.[12] No such mental states need be involved in indirect discrimination, e.g., by definition someone who indirectly discriminates against women does not, qua indirectly discriminating against them, act on an intention to exclude women.[13] Hence, the object of indirect discrimination must be determined on other grounds.

One obvious suggestion is the view that the relevant group of individuals that is harmed is the group of those individuals who are harmed. (This may sound trivial but is not.) Suppose that for reasons wholly independent of sex discrimination primary schools operate in a way that make boys do worse on average than girls. Suppose also that there is a significant subgroup of boys that benefit in that they would do worse under any scheme that would make boys and girls do equally well. Suppose finally that there is a subgroup of girls who would do better on the gender-equality scheme.[14] The question then is: should we say that the scheme indirectly discriminates against boys, or should we say that it indirectly discriminates against the large subgroup of boys and the small subgroup of girls who are harmed by the present policy? If we do the former, obviously, we cannot say that groups of victims of indirect discrimination are identified on the basis of their members suffering from harm from such discrimination. We would then have to say something else about the basis for the identification of the relevant groups and, as a matter of fact, which groups are identified as victims of indirect discrimination is determined by whether they form a socially salient group. So boys can be the object of indirect discrimination in my example because gender is the basis for a socially salient group, but the majority of boys and the minority of girls who would do better under a different school system do not form a socially salient group.[15]

---

[12] This may be more problematic than indicated here for reasons relating to the distinction between discrimination in the extensional and in the intensional sense (chapter 1.2; see also chapter 9's ironic case of discrimination). Also, people's discriminatory responses might respond to more fine-grained distinctions than the ones we usually have in mind in relation to direct discrimination, say, men and women, e.g., women are not intentionally treated disadvantageously compared to gays.

[13] In my account of indirect discrimination it is true by definition that no such mental states are involved.

[14] In almost all cases at least, where we say that a certain practice indirectly discriminates against a certain group, some—sometimes perhaps only a few—of its members actually benefit from the scheme.

[15] Admittedly, differential treatment of socially salient groups is not always considered discrimination. One case that springs to mind here is disadvantageous treatment of low income groups. However, as noted in Chapter 1.7 we can explain why people tend not to talk about income discrimination in a way that is compatible with its being widespread.

3. The No-Intention Condition

The no-intention condition should be revised in three ways. First, it needs to be revised to accommodate a case where, say, someone hires on the basis of a test that does not explicitly target minority members, nor has any intention of disadvantaging minority members, but has the nonexplicit, i.e., nonpublic, intention of excluding minority employees and, thus, harbors a bias against them. He intends to ensure that all his employees are majority members, but he would welcome that minority members are not disadvantaged as a result of his exclusion of them. Intuitively, this qualifies as a case of direct discrimination, and we should revise Altman's definition accordingly.

Second, the no-intention condition distinguishes between policies that are aimed explicitly and policies that are aimed surreptitiously at certain social groups. A policy may be one without being the other. If a job announcement states that "No women can apply," then that announcement explicitly targets women. This might be so even if the relevant employer intends to hire women only and if she knows that some women—women of the sort she would like to hire—will apply anyway despite the wording of the announcement. Conversely, a policy may be formulated in wholly neutral terms, serve a legitimate aim (in the sense that the policy is causally conducive to the achievement of an aim the achievement of which justifies the policy), and yet be adopted in order to disadvantage women, in which case it is surreptitiously aimed at them. If so, it will count as a case of direct discrimination on Altman's definition. This is unproblematic. However, the former case suggests that what really matter are intentions (or, as I will discuss shortly, biased mental states in general), not the formulations used although, obviously, they matter indirectly and importantly as indicators of underlying intentions.

The employer, who uses the "No woman can apply" formulation as a means of achieving her end of hiring the right sort of women, does not in my view engage in direct discrimination against women. Nor does she engage in indirect discrimination according to Altman's definition. On the assumption made in chapter 1 that any instance of discrimination is either an instance of direct or indirect discrimination, we must draw the distinction between direct and indirect discrimination differently from how Altman draws it. I suggest that we disregard formulations such as those used in job advertisements and only focus on the intentions of the discriminator.

Third, I want to broaden the range of mental states to take into account such that not only the intentions of the discriminator bear on whether she is engaged in direct or indirect discrimination (see also chapter 4). I want to say that any bias on part of the discriminator makes her actions directly discriminatory.[16]

[16] "Mental states" is used here to refer not only to states of which the agent is aware, but also to mental dispositions.

So, for instance, an employer who hires on the basis of a certain test that disadvantages minority applicants and who fails to see this because of her dislike of minority applicants—she would have noticed this effect of the test had those being disadvantaged by it been majority applicants—engages in direct discrimination even if she does not intend to disadvantage women but, say, just welcomes the fact that they are disadvantaged and this bias of hers plays a causal role in explaining why she treats male and female applicants differently. Michael Dummett writes that "indirect discrimination may occur when some rule unintentionally disadvantages members of some racial group; plainly, when this has happened by accident, it is due to thoughtlessness rather than racism."[17] The contrast drawn presupposes that racism is necessarily tied to intentions, but thoughtlessness on part of a person may reflect racist biases such that Dummett's understanding of direct, racist discrimination is much too narrow.

Admittedly, direct discrimination is normally defined in terms of the discriminator's intention to treat members of some groups disadvantageously relative to others, and indirect discrimination is then defined as cases of differential treatment where such intentions are absent.[18] One motivation for this way of making the cut between direct and indirect discrimination is to tie the wrongness of direct discrimination, and direct discrimination only, to the discriminator's objectionable, because disrespectful, intentions (see the discussion in chapter 4).[19] However, to the extent this is the underlying motivation the distinction between direct and indirect discrimination should be drawn differently. This is so because an agent's action may reflect objectionable, because disrespectful, mental states other than intentions.[20]

Take, for instance, Frances Kamm's distinction between doing something *in order to* make something the case (e.g., throwing a party in order have fun with one's friends) and doing something only *because of* the fact that something else will happen (e.g., throwing a party because of the fact that one's friends will help clean up the mess afterwards—one does not throw a party in order to make it the case that one's friends will help clean up the mess and yet one might not have done it otherwise).[21] A discriminator who treats

[17] Michael Dummett (2004), "The Nature of Racism," in Levine and Pataki (eds.), *Racism in Mind*, (Ithaca, NY: Cornell University Press), 27–34, p. 30.

[18] Recall that Altman's definition quoted at the beginning of section 1 refers to what the discriminator aims at. See also Appendix 1, this chapter.

[19] E.g., Larry Alexander (1992), "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies," *University of Pennsylvania Law Review* 141.1: 158–163.

[20] Alternatively, we might retain an intention-focused account of direct discrimination and then say that indirect discrimination can be no less disrespectful than otherwise comparable forms of direct discrimination because of the mental states that they involve. We would then have to distinguish between those forms of indirect discrimination that involve disrespectful mental states and those that do not.

[21] Frances Kamm (2007), *Intricate Ethics* (Oxford: Oxford University Press), pp. 92, 94–95.

Catholics differently not in order to disadvantage them but to achieve some other aim, but does what she does only because of its effect of disadvantaging Catholics, acts from an objectionable, because disrespectful, mental state. That mental state, however, is not an intention and, accordingly, from the perspective of a mental-state-based account of the wrongness of direct discrimination, the narrow focus on intentions is unwarranted (cf. the discussion in chapter 4).[22]

On the basis of the three criticisms above, I would substitute the following statement for Altman's formulation of the no-intention condition.

(i*) It reflects no bias on part of the discriminator against members of G on account of their being members of G (*the revised no-intention condition*)

My anti-Catholic, and my employer who excludes minority applicants without intending to disadvantage them, both act in ways that reflect their biases against those whom they treat disadvantageously and thus satisfy the revised no-intention condition. My feminist employer who intends to hire a certain kind of woman by using the "Women need not apply" formulation is not biased against women and, thus, may be involved in indirect discrimination against women.[23]

4. The Disadvantage Condition

I now turn to the disadvantage condition and how it should be revised. Before proceeding I should note that the challenge of specifying the relevant sense of disadvantage involved in indirect discrimination—call it *the disadvantage specification challenge*—is one that any satisfactory account of indirect discrimination must face. This challenge has several subchallenges. First, it is unclear whether the disadvantage condition concerns local or global disadvantage, i.e., whether it suffices that the relevant act or policy is disadvantageous in some respect, or whether it needs to be disadvantageous all things considered

[22] Incidentally, this explanans also motivates omitting the "explicitly targets" bit in the no-intention condition. A feminist shopkeeper who out of sheer incompetence writes in his job advert that "Women need not apply" has no objectionable, because disrespectful, mental states regarding women. Yet, he is involved in direct discrimination against women based on Altman's definition.

[23] For (i*) to be satisfactory, "bias" would have to be understood in such a way that an employer who intends to exclude members of a certain group out of a "separate-but-equal" attitude, even though he bears no hostility towards them and has an unbiased assessment of their qualifications, is biased in the relevant sense. Otherwise, he might be engaged in indirect discrimination. Alternatively, (i*) should be revised so that it says: "It reflects no bias or 'separate-but-equal' attitude on part of the discriminator . . . "

(cf. chapter 6). Suppose that the admission rules at a prestigious humanities faculty require that applicants speak three European languages. This disadvantages immigrants, who are, let us suppose, mostly non-European, in that they will find it harder than nonimmigrants to be admitted, because as a matter of fact, a smaller proportion of them speak three *European* languages. Suppose that, setting aside their effects on groups, these rules are not unjust to anyone, e.g., speaking three European languages is a relevant requirement for the courses offered. Suppose, finally, that this rule leads immigrants to apply for education in engineering science, law, and medicine, where people end up with better jobs and lead lives that are better in terms of welfare than people with different university degrees. So while the rule disadvantages immigrants locally in that it makes it harder for them than for nonimmigrants to study at the faculty of humanities, it makes them better off, globally, relative to nonimmigrants.

There are three possible views here: only local disadvantages matter; only global ones do; and, finally, local as well global disadvantages matter.[24] I find the local-only view quite implausible. If, say, African Americans were made significantly worse off, globally, by policies that benefited them locally, we might well classify these as involving indirect discrimination. Similarly, we would not call a ban on foxhunting indirect discrimination against the British upper class even if such a ban involved a local harm for these people.[25]

The "global-disadvantages-only" view is also somewhat unattractive. Suppose that unlike all other sectors of society the educational system has a history of direct discrimination against women. While such direct discrimination has been eradicated, admission rules are still in place that perpetuate women's disadvantageous situation within the educational system, e.g., they place great weight on test results in areas where women tend to score less well than men. Such rules might well be indirectly discriminatory even if they do not result in women being disadvantaged, globally (whatever the relevant metric of global advantage, e.g., welfare or capabilities, consists in), relative to men. (Suppose the labor market is such that people actually have a lower lifetime income, the better educated they are and that more education significantly increases the risks of suicide, divorce, alcohol and drug abuse, and depression etc. resulting in women being better off, globally speaking, than men as a result of their being indirectly discriminated against in the educational system.) More generally, the "global-disadvantages-only" view implies that there could be no

[24] The same three options are relevant to how we understand (ii) in chapter 1's definition of group discrimination, i.e., that X treats Y worse than Z.

[25] Another reason why this is so might be that the locally harmful effects of upper-class people would not be disproportionate in the light of what is at stake for foxes. Ultimately, however, I think the disproportionateness condition should be omitted from the definition of indirect discrimination.

indirect discrimination in a society where everyone ends up equally well-off, globally speaking, and this implication is implausible.[26] For these reasons, to capture what we mean by "indirect discrimination," the most reasonable way to specify the disadvantage condition is to say that local as well as global disadvantages count for whether indirect discrimination occurs.

Second, however we settle the local versus global disadvantage issue, we also need to address the question of what determines whether a particular group is disadvantaged by a certain act or policy. A natural view here—*the average view*—is to say that one compares (a) the average level of benefits that accrue to members of this group relative to members of other groups under the relevant policy with (b) the average level that would accrue to them relative to others under some alternative policy. If, and only if, the gap between the average level of benefits accruing to the discriminated group and the average level of benefits accruing to other groups under the relevant policy is higher than the gap under all other relevant alternative policies, this policy may be indirectly discriminating against the worse-off group.

While the average view is a natural view, it is not uniquely attractive. To see this compare a situation where a certain group's average level of benefits is slightly higher than that of another group. Under the alternative provision, its average level of benefits would be slightly lower than those of the other group, but the variation in the level of benefits within the former group would also be much smaller, say, because while a few, very well off people within the group are much worse off every other member of the group are slightly better off. In this case, we might want to say that the group suffers indirect discrimination under the policy that makes it better off on average than the other group on the ground—*the numbers view*—that the majority of all members of it are harmed by the relevant provision.[27]

The average and the numbers view will converge in terms of which groups they identify as being disadvantaged, because acts that make most members of a group worse off will also tend to be acts that lower its average level of benefits. Still, there is no necessity about this, and in those cases where the two measures of disadvantage diverge we will need to determine which is relevant for the purpose of determining whether indirect discrimination takes place. However, I suspect that our concept of indirect discrimination is indeterminate at this point and that any attempt to specify it will amount to a sheer stipulation. Still, to have a precise concept of indirect discrimination we will need to stipulate one way or another.

---

[26] I am indebted to Lea Ypi at this point.

[27] Cf. Oran Doyle (2007), "Direct Discrimination, Indirect Discrimination and Autonomy," *Oxford Journal of Legal Studies* 27.3, 537–553, pp. 540–541. There are many more views here than the average and the numbers views. For instance, a third view—call it *the weighted view*—gives extra weight to how policies affect the worse off within groups.

On the basis of the two points made above, I would substitute:

(ii*) It relevantly disadvantages, i.e., on average or in the case of most, members of G, locally or globally, relative to other relevant, socially salient groups (*the revised disadvantage condition*)

for Altman's formulation of the disadvantage condition as it brings to the fore two important issues of specification in connection with indirect discrimination. These will become important when we turn to the question of whether indirect discrimination identifies an additional wrong.

There is a third consideration to be taken into account when determining whether a certain group is being relevantly disadvantaged: namely, what is the relevant baseline to compare with? I do not address this question in any comprehensive way here, because, while relevant to how (ii*) should be understood, it is irrelevant to how Altman's definition should revised. I return to the issue of the baseline in Chapter 6. Suffice to mention that there are least three possible baselines here.[28] First, in assessing whether new policies are indirectly discriminatory we might say that disadvantage is determined on the basis of how well-off the relevant group was prior to the introduction of the relevant policy (*the status quo baseline*). On this account, a company that changes its paid maternal leave policy for new employees from six months to three months disadvantages women applicants in a way that may amount to indirect discrimination. To see how the status quo baseline can be problematic, compare the company just mentioned to one that changes its parental leave policy from no months to three. Now they both have the same policy, but one is indirect discrimination against women and the second one is not, simply due to the nonchange baseline. This seems problematic. A different problem can be seen, if we consider a case involving normal market transactions, e.g., someone sets up a booze store and members of some particular socially salient group are overrepresented among his customers and end up worse off than they would have, had he not opened the store. Few, I suspect, would allow that members of the relevant group has been disadvantaged in a way that bears on indirect discrimination.

Second, another suggestion is that disadvantage is measured by comparing the actual situation to one where benefits and burdens are distributed as they ideally should be distributed (*the moralized baseline*). This suggestion avoids the two problems just mentioned, i.e., status quo plays no role on this account and if normal market transactions are morally unproblematic the group that is overrepresented among booze store customers is not harmed relative to the relevant normative baseline. However, the obvious challenge here is to provide a well-grounded normative account of this sort. Another problem is to deal adequately with cases involving partial compliance situations, i.e., situations in

[28] I am indebted to Mathew Coakley at this point.

which a certain policy is harmful relative to one, in which everyone acts as they should do on the relative normative account, but where, as a matter of fact, many do not.

Third, the relevant baseline might be those policies a well-informed unbiased agent could not reasonably reject (*well-informed, unbiased agent baseline*). Like the previous baseline this is a normative one, but it might account better for situations, which involve partial compliance because a well-informed unbiased agent would reject or accept policies partly in the light of his beliefs about the compliance or noncompliance of others. However, like in the case of the moralized baseline there is a problem in explaining exactly, which is the relevant baseline here, i.e., which policies that a well-informed unbiased agent could and could not reasonably reject.

5. The Disproportionateness Condition

Finally, I address the disproportionateness condition, i.e., the last of the three conditions in Altman's definition that I critically discuss in this chapter with the aim of proposing a better definition. First, ideally, one would like: (1) the distinction between direct and indirect discrimination to be clearly drawn along one dimension; (2) direct and indirect discrimination to be mutually exclusive—such that a case cannot amount to direct as well as indirect discrimination against the very same group, at least not when one's act as well as the discriminatee are identified under the same descriptions;[29] and (3) that the distinction is exhaustive—such that there are no cases of discrimination that are neither direct, nor indirect, discrimination. Altman's disproportionateness condition means that the distinction between direct and indirect discrimination is drawn along different dimensions. That is, putative direct discrimination still counts as direct discrimination even if it does not disproportionately disadvantage anyone. Such cases are possible, e.g., a nonsecular state refuses to hire a female applicant for the job as a bishop to accommodate ultra-conservative religious fanatics and thereby avoid civil war, which will otherwise eventuate. This implies that there are possible cases that qualify neither as direct, nor as indirect discrimination, i.e., cases that satisfy the no-intention and the

[29] Cf. Fredman (2011) *Discrimination Law*, p. 190. To see the need for the "at least" qualification, consider the following case. I intend to exclude female applicants and intend to do so by hiring on the basis of a very stringent abilities test: say, I believe falsely that strenuous tests will scare off female applicants. Unbeknownst to me, the test, while apparently neutral, disadvantages female applicants, because of its focus on knowledge about "male" sports. Here my policy of hiring on the basis of this test is both directly and indirectly discriminating against women. My act is directly discriminating when identified under the description "hiring on the basis of a strenuous test," and indirectly discriminating under the description "hiring on the basis of a test focusing on 'male' sports."

disadvantage conditions (and, for that matter, the revised versions of these conditions too), but where the disadvantages to the discriminatee are proportionate. Hence, omitting the disproportionateness condition would give us a cleaner distinction between direct and indirect discrimination.

Second, what are the relata of the disproportionateness, i.e., the bad and good aspects of the relevant policy that are to be disproportionate to one another for the latter to amount to indirect discrimination, and on what basis are the relata to be compared? Clearly, a disproportionateness condition often appears in legal definitions of indirect discrimination, where the relata are assessed from a legal point of view. However, here we are interested in the moral qualities of indirect discrimination and, accordingly, we should assess how the relata compare from a moral, not a legal, point of view.[30] For a similar reason the relata are not restricted to legally recognized concerns. So let us turn to the relata.

It seems obvious that one relata is the degree to which the victim of indirect discrimination is disadvantaged relative to other groups. Presumably, the other relata are net benefits to all those benefited by the relevant policy, including perhaps and as discussed above some members of the group suffering indirect discrimination itself.[31] One problem here is that, according to some moral views, the first relata—i.e., the gap between relative positions of discriminatees and others—is not a morally relevant concern at all. That is, only for strict egalitarians is this consideration as such a morally relevant one. Accordingly, for nonegalitarians, no form of indirect discrimination will be disproportionate insofar as it involves net benefits for individuals and, thus, by Altman's definition, nonegalitarians should deny that indirect discrimination exists! This suggests that the disproportionateness condition should be eliminated.

Third, to avoid the implication just stated—it should be avoided, because nonegalitarians will want to talk about indirect discrimination—one might suggest that the disadvantages that the disproportionateness condition concerns are different from the ones the disadvantage condition concerns. That is, the disadvantages relevant for the former condition are not comparative, but absolute. In this view, a policy that, say, makes men better off than women, but where any alternative policy will leave women worse off than under the former policy, does not satisfy the disproportionateness condition. In this way, egalitarians as well as nonegalitarians may acknowledge the existence of indirect discrimination. However, so revised, indirect discrimination has been turned into a very different notion from what most people have in mind. That is, indirect discrimination is then in principle compatible with very significant

[30] Morally relevant considerations may not be legally recognized just as morally irrelevant considerations might be legally relevant considerations that bear on whether disadvantages are disproportionate, legally speaking.

[31] Admittedly, among the benefits that count one might also find gains in impersonal values.

inequalities between groups, provided these inequalities benefit the disadvantaged groups. Moreover, even so amended, the absolute disadvantage-focused, disproportionateness condition seems problematic. Based on the present reading of it, as well as on Altman's, groups can be proportionately as well as disproportionately affected by a certain policy. Imagine that a particular unbiased employer hires on the basis of a test that for some reason only 10% of female and male, immigrant applicants pass, whereas 20% of male, nonimmigrant applicants pass. Suppose, moreover, that only women are *disproportionately* disadvantaged, because while many other employers use different tests on which both women and male immigrants have better scores, these employers engage in direct discrimination against women. Hence, women find it difficult to get a job within the relevant industry, unlike male immigrants, who simply find it difficult to get a job with this particular employer. Yet, how can the fact that other employers engage in direct discrimination against women imply that this particular employer does not indirectly discriminate against them?

Note that the present problem arises whether we say that disproportionateness is determined by all relevant moral concerns, or only on a subset of them. For instance, one might suggest that the relata are benefits and harms to, on the one hand, the discriminatee and, on the other hand, the discriminator. Presumably, these concerns are morally relevant, but they do not exhaust the set of concerns that are relevant for assessing the wrongness of the relevant policy etc.[32] One reason why this is so is that because harms and benefits to parties other than the discriminatee and the discriminator are morally relevant as well.

Finally, one striking implication of building a morally disproportionateness requirement into indirect discrimination, when there is no comparable condition for something being direct discrimination, is that, at least in one respect, indirect discrimination is by definition *morally worse than* direct discrimination. The reason is that, by definition, there is no goal the fulfillment of which is morally proportionate to the disadvantage that indirect discrimination results in. Direct discrimination, on the other hand, may ensure the achievement of goals that morally justify the relevant goals promoted through direct discrimination. Suppose a shop owner will go bankrupt if she hires a minority applicant, because her customers are prejudiced. Suppose that the costs for her and others of her going bankrupt are so great that she is justified—not just excused—in intentionally refraining from hiring any minority applicant (in which case she

[32] This has the advantage of not making the concept of indirect discrimination a moralized one, i.e., one need not know if something is morally unjustified to know if it is indirect discrimination. This is an advantage, because, presumably, often it is controversial whether something is morally justified and people may agree that something amounts to indirect discrimination even if they disagree about its moral credentials (cf. the discussion of moralized discrimination in chapter 1).

engages in justified, direct discrimination).[33] Omitting the disproportionateness condition will avoid this implication and, thus, accommodate most people's belief that direct discrimination is, if anything other than equally morally bad, morally worse than indirect discrimination.

On the basis of the problems discussed in this section. I think the disproportionateness condition should be given up. This leaves us with the following revised version of Altman's definition of indirect discrimination:

A policy or act is indirectly discriminatory against a certain group, G, if, and only if:

(i*) It reflects no bias on part of the discriminator against members of G on account of their being members of G (*the revised no-intention condition*).

(ii*) It relevantly disadvantages, i.e., on average or in the case of most members of the relevant socially salient group, locally or globally, relative to other relevant, socially salient groups (*the revised disadvantage condition*)

(iii) G is a socially salient group.

This definition is neither weaker, nor stronger, logically speaking, than Altman's. It drops the disproportionateness condition and eliminates the "explicitly targeting" condition making it easier to satisfy—i.e., making it logically weaker—but it also broadens the scope from intentions to mental states in general, thus making it harder to satisfy—i.e., making it logically stronger.[34]

6. Sufficient for Indirect Discrimination?

I now want to ask whether satisfying conditions (i*), (ii*), and (iii) is sufficient for indirect discrimination. I shall argue that it is not and will suggest two additional revisions such that I end up with a definition of indirect discrimination

[33] In this example, one might appeal to something like an agent-centered prerogative, see Samuel Scheffler (1982), *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions* (Oxford: Clarendon Press). It is because persons are morally justified in giving special weight to their own interests that the harm imposed on minority members is not disproportionate. Alternatively, one might think that disproportionateness is determined from a wholly agent-neutral perspective in which case significant economic costs to the shop owner may not be enough to outweigh small harms to many minority members from a discriminatory hiring policy. More generally, the disproportionateness condition can be filled out in various ways, e.g., (i) in terms of the relevant perspective (agent-relative or agent-neutral), and (ii) in terms of which benefits and harms count. If moral permissibility is determined by weighing harms and benefits and if all harms and benefits (perhaps weighted according to their moral significance) together determine proportionality, then, by definition, there can be no such thing as morally permissible indirect discrimination.

[34] I am indebted to Christian List at this point.

that states the necessary and conjointly sufficient conditions of indirect discrimination.[35]

First, suppose that a certain act, e.g., implementing some kind of affirmative action program, is intended to advantage an otherwise disadvantaged socially salient group, but happens to disadvantage the intended beneficiaries. According to both the revised definition and Altman's initial one, this group may be the victim of indirect discrimination.[36] I suspect that we would not describe this as a case where the relevant group is subjected to indirect discrimination and accordingly propose that the revised no-intention condition be amended as follows:

(i**) It reflects no bias on part of the discriminator against members of G on account of their being members of G, and it is not the case that the relevant policy or act reflects a bias, or unbiased attitude, of the discriminator in favor of members of G qua members of G (*the expanded no-intention condition*).[37]

Again, I suspect our concept of indirect discrimination is somewhat indeterminate here and, accordingly, I concede that the expanded no-intention condition may to some extent reflect stipulation on my part.

Second, (i**), (ii*), and (iii) may obtain, and yet there might be no indirect discrimination. Suppose we have a society in which both direct and indirect forms of discrimination exist. Its government hires a dream team of superlawyers, social engineers, etc. and assigns them the task of eliminating all forms of indirect discrimination. They are told to disregard direct discrimination. Conceptually speaking, there is nothing that would prevent this team from being successful. If so, we can have direct discrimination without indirect discrimination.[38] Given a certain background assumption, I am skeptical of the reverse situation being possible, however.

Suppose that in a society where no direct discrimination exists and has never existed in the past, there are some inequalities between groups. Since the fact that some groups are disadvantaged relative to others is not miraculous,

---

[35] Obviously, the problem of insufficiency is deepened because the revised definition eliminates the disproportionateness condition. However, while this condition narrows down the range of cases that amount to indirect discrimination, it does so in the wrong way for the reasons suggested.

[36] Some affirmative action programs have "targeted" minority applicants in that they have, say, given extra points to such applicants. However, since this is supposed to benefit them it does not involve targeting them in the relevant sense.

[37] Note that if the discriminator intends to benefit the discriminatee through the affirmative action program, but has a biased, e.g., condescending, view of what benefits it, because of (i**), my amendment does not rule out that this qualifies as a case of direct discrimination against the intended beneficiary.

[38] A related question is whether a group that has always been directly discriminated in favor of, and only that, can be indirectly discriminated against.

presumably, whenever this is the case there will be some—let us suppose, apparently neutral—sets of events or even some sets of acts that produce these inequalities. The problem facing us then is what we might call *the limitation challenge*: among the groups that are worse off, how do we identify those that are being subjected to indirect discrimination and those that are worse off than other groups for some other reason? Ex hypothesi, we cannot meet the limitation challenge on the basis of which groups have been subjected to direct discrimination in the past, or whether the relevant acts perpetuate the harmful effects of prior direct discrimination.

One might suggest that any inequality between socially salient groups reflects indirect discrimination, whereas inequalities between all other groups are independent of any kind of discrimination. This suggestion, however, will not do (see chapter 7).[39] Suppose that in a society untainted by sex discrimination people are given numbers randomly and, despite the improbability of this happening, almost all women get even numbers. If the state imposes a rule making greater demands on people who have been given even numbers (not knowing that most of these people are women), presumably, this would not qualify as a case of indirect discrimination.[40]

It is equally clear, however, that the existence of direct discrimination against some group or other is not sufficient for there being indirect discrimination against a particular group. Rather, it must be the case that (i) there has been, or presently exists, direct discrimination against the group being subjected to indirect discrimination and (ii) that the indirect discrimination is suitably related to these instances of direct discrimination. David Wasserman has suggested something along these lines.[41] In his account of indirect discrimination, he describes

---

[39] This is essentially Matt Cavanagh's point when he writes that people "who are concerned primarily with how things like race and sex show up in the overall distribution [of jobs] have no business saying that their position has anything to do with discrimination. It is not discrimination they object to, but its effects; and these effects can equally be brought about through other causes" (Cavanagh [2002], *Against Equality of Opportunity* (Oxford: Clarendon Press), p. 199. However, some people do not just object to the effects, when they object to indirect discrimination. If these effects came about in a random way, they might think that the situation was less objectionable than if the same effects came about as a result of indirect discrimination. For instance, Pogge would seem to adopt a position like this in relation to distributive justice. For a competing view, see Andrew Morris (1995), "On the Normative Foundations of Indirect Discrimination Law: Understanding the Competing Models of Discrimination Law as Aristotelian Forms of Justice," *Oxford Journal of Legal Studies* 15.2, 199–228.

[40] If it does, then surely this is a case of indirect discrimination where the fact that it amounts to indirect discrimination adds nothing to its being unjust (relative to the distribution of burdens being unjust on general grounds of distributive justice).

[41] "[A] regime that did not express hatred or contempt toward disadvantaged groups might be unfair, but it would arguably not be discriminatory"; DavidWasserman (1998), "Discrimination: Concept Of," in R. Chadwick (ed.), *Encyclopedia of Ethics* (San Diego, CA: Academic Press), pp. 805–814, p. 809. Doyle's definition of indirect discrimination on the disparate impact model has a similar implication; see Doyle (2007), "Direct Discrimination," p. 538.

indirect discrimination as follows: "Allocative actions and policies that are neutral in appearance and design, but that have the effect of perpetuating the harmful effects of deliberate exclusion."[42] A natural reading of this passage suggests that for Wasserman not just any disadvantage for one group relative to others qualifies for the purpose of identifying indirect discrimination. Only disadvantages that perpetuate disadvantages resulting from past or present direct discrimination (i.e., deliberate exclusion) do so.[43] I believe this is on the right track and, at any rate, I suspect that few complaints have ever been raised on behalf of groups that have never been subjected to direct discrimination.[44] Accordingly, I suggest a fifth condition:

(iv) The disadvantages referred to in (ii*) would not have occurred in the absence of past or present direct discrimination by the agent of the relevant act or policy against G (*the causal condition*).

Note that this requirement implies that it is not enough for an agent's indirectly discriminating against G that some agent or other engaged in direct discrimination against G and that disadvantages for G of the relevant agent's act or policy would not have occurred in the absence of some other agent's direct discrimination against G. Suppose that Finnish immigrants in the United States experience certain kinds of disadvantages despite the fact that there never has been direct discrimination against Finns in the United States. Suppose also that Finnish immigrants in Sweden have been subjected to various kinds of direct discrimination, but that they are irrelevant to the forms of disadvantages Finnish immigrants experience in the United States. In this case, Finnish immigrants could not be subjected to indirect discrimination in the United States, even though they could elsewhere.[45] Note that in saying this I am not suggesting that, as a matter of fact, less often than people normally think immigrants are subjected to indirect discrimination. Most immigrants are as a matter of fact subjected to various forms of direct discrimination and,

---

[42] Wasserman (1998), "Discrimination: Concept Of," p. 811.

[43] Alternatively, one might require that the disadvantages may fall on members of the group being subjected to indirect discrimination in some sense *because* they belong to this group.

[44] European Americans have complained about being subjected to direct discrimination, but not indirect discrimination.

[45] Consider a variation of the *Griggs v. Duke Power* verdict where the relevant, European American–run company had a history of discriminating in favor of African Americans. Whether the company's promotion rules would be indirectly discriminatory would depend on whom the agent of direct discrimination is: the group of European Americans in North Carolina, to which those people who run the company belong, or the company CEOs, whom we imagine have never discriminated against African Americans. A trickier question here is what to say about a case where the direct discrimination of Finns in Sweden by Swedes for some weird reason is a causal prerequisite for the disadvantages that Finnish immigrants in the United States experience. Here the harmful effect of past discrimination by some people is perpetuated by the policies of others, not just other individuals but also other collectivities.

accordingly, pointing to the plight of actual immigrants involves no objection to my definition.[46]

We have now reached a revised version of Altman's definition of indirect discrimination:

> A policy or act is indirectly discriminatory against (members) a certain group, G, if, and only if:
>
> (i**) It reflects no bias on part of the discriminator against members of G on account of their being members of G, and it is not the case that the relevant policy or act reflects a bias, or unbiased attitude, of the discriminator in favor of members of G qua members of G.
>
> (ii*) It relevantly disadvantages, i.e., on average or in the case of most members of the relevant socially salient group, locally or globally, relative to other relevant, socially salient groups.[47]
>
> (iii) G is a socially salient group.
>
> (iv) The disadvantages referred to in (ii*) would not have occurred in the absence of past or present direct discrimination by the agent of the policy or act in question against G.[48]

Each of these conditions must be satisfied for indirect discrimination to obtain and together they are sufficient.[49]

7. Direct vs. Indirect Discrimination

Having proposed this definition of indirect discrimination I should relate it to my definition of group discrimination proposed in chapter 1. I want the distinction between direct and indirect discrimination to be exhaustive, direct and indirect discrimination to be mutually exclusive, and my definition of indirect discrimination to capture a form of discrimination that falls under the definition of discrimination proposed in chapter 1. To accommodate the mutual exclusion desideratum I propose:

---

[46] Another implication of my definition is that there might be such a thing as indirect, income discrimination.

[47] Suppose European Americans pursue a policy that benefits African Americans more than it benefits European Americans, but benefits African Americans less than it benefits Asian Americans. This policy might be indirectly discriminatory against African Americans even though they are benefited relative to the group whose members indirectly discriminate against it.

[48] I said that ideally the distinction between direct and indirect discrimination should be drawn along the same dimension(s). Like Altman's, my revised definition fails to satisfy this desideratum, e.g., direct discrimination does not require something like (v).

[49] This definition is stronger than the one sketched in the previous section. However, like that one it is still neither weaker, nor logically stronger than Altman's definition.

> A policy or act is directly discriminatory against (members of) a certain group, G, if, and only if (i) it amounts to discrimination according to my definition of group discrimination proposed in chapter 1 and (ii) it does not amount to indirect discrimination according to my definition proposed in section 6.

To accommodate the remaining two desiderata I propose adding the following condition to my definition of indirect discrimination:

(v) It satisfies my definition of group discrimination.

Given this condition and the above definition of direct discrimination, my distinction between direct and indirect discrimination is exhaustive, because any form of group discrimination that is not indirect is direct. Because of (v) any act that qualifies as indirect discrimination necessarily qualifies as group discrimination. Admittedly, this leaves us with two oddities. First, victims of indirect discrimination may form a certain socially salient group, G, whose disadvantageous treatment would not have occurred had they not been discriminated against in the past, albeit they were not discriminated against qua members of G. (Recall the distinction between discrimination in the extensional and in the intensional sense.) Second, because my definition of indirect discrimination does not require that for an act to be indirectly discriminatory it must be of a type that is often instantiated and has harmful consequences and because indirect discrimination need not involve any particular kind of motivating state, in principle there could, in the absence of (v), be a case of indirect discrimination that does not qualify as group discrimination. However, because of (v) such an act would not qualify as indirect discrimination on my account and, arguably, this seems to be pure stipulation if not counterintuitive. Because these two complications are very rarely instantiated—so I conjecture—I am content to set aside for future work, how the relevant definitions should be tightened.

8. Conclusion

Indirect discrimination has a solid foothold in our thinking about discrimination. This might well be a good thing. However, it is unclear what exactly indirect discrimination amounts to, and in this chapter I have named four challenges that any account of indirect discrimination must face: (1) How should indirect discrimination be defined?; (2) Which particular worse-off group is the discriminatee in a case of indirect discrimination?; (3) What is the relevant sense of disadvantage involved in indirect discrimination?; and (4) By which criterion do we determine which groups, among those groups that are worse off, are worse off as a result of indirect discrimination and those that are worse off for other reasons? In response to these challenges among others, I have proposed

and motivated a revised version of a definition offered by Altman. My definition is meant to imply that, with the exception mentioned in the previous section, any form of group discrimination is either direct or indirect, and that indirect discrimination against a certain group is parasitic on the fact that this group has been subjected to prior direct discrimination. The latter implication reflects the need to avoid that any group disadvantage qualifies as reflecting indirect discrimination or, to it differently, to provide a satisfactory response to the limitation challenge, i.e., (4) above. Unlike Altman's definition and that of many others, my definition omits any disproportionateness condition such that a group of people may be subjected to indirect discrimination even though the comparative disadvantages this involves for them are quite small relative to the benefits this treatment involves for others. This should make my definition congenial to those who are not telic egalitarians in that how people fare relative to one another has no noninstrumental moral relevance and who nevertheless think that indirect discrimination takes place. Also, I have clarified the nature of the relevant disadvantages involved in indirect discrimination and tied it to absence of bias of individual agents rather to their intentions, unlike what most accounts of indirect discrimination do. My later discussion of the morality of indirect discrimination—see chapter 6—is based on the definition proposed in this chapter. Other theorists might reject my account of indirect discrimination, but their alternative accounts must struggle with the same four challenges that I have taken up in this chapter.

Appendix 1: Some Other Definitions of Indirect Discrimination

In this appendix, I compare Altman's definition of indirect discrimination to some other descriptions or characterizations present in the literature. My aim is to show that it captures some features commonly present in these and, thereby, to motivate my use of Altman's definition as a springboard for my own attempt to delineate indirect discrimination.

Consider first a definition of indirect discrimination from the European Commission: "Indirect discrimination occurs when an apparently neutral provision, criterion or practice would disadvantage people on the grounds of racial or ethnic origin, religion or belief, disability, age, gender or sexual orientation unless the practice can be objectively justified by a legitimate aim. An example of indirect discrimination is requiring all people who apply for a certain job to sit a test in a particular language, even though that language is not necessary for the job. The test might exclude more people who have a different mother tongue."[50]

[50] http://ec.europa.eu/employment_social/fdad/cms/stopdiscrimination/resources/glossary/index.html?langid=en#I (accessed May 21, 2011); cf. Article 2(2) of Council Directive 2000/43 EC quoted in Doyle (2007), "Direct," p. 539.

While this definition simply states a sufficient condition of indirect discrimination, something similar to all four of Altman's conditions are found in this definition. The no-intention and the disadvantage conditions seem clearly present in it, albeit it does not mention intentions explicitly but simply refers to an "apparently neutral provision" that, somewhat problematically in my view, introduces an epistemic element in the definition, i.e., apparent to whom? The disproportionateness condition is found in a more specific version in that the definition requires "objective justification" by a legitimate aim. Presumably, "objective justification" means that it is not enough that the aim is legitimate, but the relevant provision objectively speaking will promote the achievement of this aim. The social salience condition appears as a list of protected groups. My contention is that any account of this definition, which is satisfactory from a philosophical point of view, will have to explain why the groups mentioned are on the list and social salience is the best such account.[51] Note, finally, that the European Commission's definition is unclear in that what comes after "would disadvantage . . . " may either be read as a specification of what it means for a provision not to be neutral or be read as stating an extra condition in addition to the requirement of a provision being neutral. On the former conception, if a rule to the effect "No women need apply" for some absurd reason results in no disadvantage to any gender, then that rule may be neutral even though there is a sense in which it is not "apparently neutral," i.e., it appears nonneutral.

Second, the ECHR holds that "[w]hen a general policy or measure has disproportionately prejudicial effects on a particular group, it is not excluded that this may be considered as discriminatory notwithstanding that it is not specifically aimed or directed at that group."[52] Unlike the previous definition this one explicitly mentions intentions, i.e., the no-intention condition. Also, the disadvantage condition is clearly present, while the disproportionateness condition is hinted at in that the definition uses the term "may" reflecting that there may be a suitable justification for the relevant differential effects.

Third, in the *Griggs v. Duke Power* verdict, which has been very important in the American context, the US Supreme Court held that "[i]f an employment practice which operates to exclude Negroes cannot be shown to be related to job performance, the practice is prohibited" under the heading of indirect discrimination.[53] Interesting here is the requirement of job performance relevance, which can be seen as a very weak version of the disproportionateness requirement. Suppose an employment practice, which excludes a certain group, is

[51] From a legal point of view, it may suffice to have a precise and finite list of clearly defined protected groups. This is so, because two very important (and, in part, related) *desiderata* legal regulations must meet are transparency in terms of what they forbid and permit, and simplicity. A list of the sort offered by European Commission might meet this criterion better than a general and abstract criterion that would need to be interpreted, whenever it is applied to concrete cases.

[52] *Shanaghan v. UK* 2001: para. 129.

[53] *Griggs v. Duke Power* 1971: 431.

related to job performance, but results in very significant disadvantages for a particular ethnic group. On a moderately weak version of the disproportionateness condition this may qualify as indirect discrimination, but it may not on the condition invoked by the Supreme Court. However, one could also imagine a case where an employment practice satisfies the disproportionateness condition but fails to satisfy the job performance relatedness condition. So suppose that a certain employment practice is unrelated to job performance and results in the slightly disadvantageous exclusion of a certain ethnic group. Suppose also that the only alternative and job performance related employment practice will result in massive disadvantages for some other ethnic group. In this case one might think that the employment practice satisfies the disproportionateness condition—since the job performance irrelevant practice is only slightly harmful to the excluded group and the job performance related practice is very disadvantageous to another ethnic group—even though it fails to satisfy the Supreme Court's job performance related condition.

Fourth, John Gardner characterizes indirect discrimination as follows: "Indirect discrimination occurs when someone applies unjustifiable standards which people of one sex or race find it harder to comply with than people of another, to the detriment of someone who is a member of the former sex or race."[54] Again, one would want to know why sex and race are mentioned but other group markers are not, and I suggest that something like the social salience condition offers the best answer. Admittedly, the best answer here is not a satisfactory answer, because socially salient groups other than those defined by sex or race are subjected to indirect discrimination.

Fifth, Gertrude Ezorsky writes that "institutional discrimination" "is exemplified when an organization (e.g., a business firm) uses an intrinsically bias-free selection procedure that has disproportional adverse impact on minorities or women."[55] As I see it, Ezorsky defines *institutional*, indirect discrimination and not indirect discrimination per se, since individuals too may engage in indirect discrimination. Unlike the previous definitions, Ezorsky's definition contains no disproportionateness condition, although her statement of exemplification contains the other three elements in Altman's definition.

Finally, Lena Halldenius defines indirect discrimination as involving "a selection criterion that looks neutral" but "systematically leads to disadvantage for a particular group" where this criterion is not "independently motivated."[56] Again, this definition has an epistemic focus and does not mention intentions, so it only partly captures Altman's no-intention condition. Also, Halldenius's

[54] John Gardner (1996), "Discrimination as Injustice," *Oxford Journal of Legal Studies* 16.3.1, 353–367, p. 355.

[55] Ezorsky (1992), *Racism*, p. 264.

[56] Lena Halldenius (2005), "Dissecting 'Discrimination,' " *Cambridge Quarterly of Healthcare Ethics* 14.4, 455–463, p. 459.

definition does not require that the relevant disadvantaged social group be socially salient, even though it will often be the case that for a group to be systematically disadvantaged it will have to be socially salient. Still, the "independently motivated" requirement is of the same sort as the disproportionateness condition, and the disadvantage condition is found in a strong form—disadvantages have to be systematic—in Halldenius's definition as well. Finally, one wonders what the qualification "systematically" refers to. Suppose that the application of a certain selection criterion for idiosyncratic reasons leads to disadvantages for women in some years, but not in others, and that there is no discernible pattern over time regarding which years involve disadvantageous outcomes. Would this not qualify as indirect discrimination because of the unsystematic nature of the production of the relevant disadvantage?[57]

Appendix 2: Institutional and Structural Discrimination

It is common to talk about institutional and structural discrimination. What are these forms of discrimination? How do they relate to one another and to the distinction between direct and indirect discrimination? Take first institutional discrimination. Assuming that an institution, e.g., a state, a company, the Catholic Church, or a university, is a collective agent organized along certain more or less formal lines, institutional discrimination is discrimination performed by a collective agent. This makes it different from structural discrimination on the assumption that structures do not constitute agents even if they have causal powers. Moreover, on the assumption that we can ascribe (biased) mental states such as intentions, beliefs, and desires to collective agents, institutional discrimination can be direct as well as indirect.[58] Indeed, we can even imagine that an institution directly discriminates against some people even if a majority or perhaps all of its members are in no way biased against these people, e.g., because of the discriminatory rules and stated purposes of the institution. So, for instance, the Ku Klux Klan would still be an institution engaging in racial discrimination, even if all its members over time happened to lose all forms of racial biases and—in a way that would still be objectionable despite the absence of racist biases—retained their membership simply to remain part of the relevant social network of members.

Structural discrimination obtains where, and only where, the social structures are such that certain socially salient groups are disadvantaged relative to others and where at least part of the explanation why these structures are in

[57] Cf. section 7 above on putative, one-off cases of indirect discrimination.

[58] For a defense of such ascriptions to collective agents, see Christian List and Philip Pettit (2011), *Group Agency: The Possibility, Design, and Status of Corporate Agents* (Oxford: Oxford University Press).

place appeals to the fact that these groups are subjected to or have been subjected to various forms of direct discrimination. Different social theorists will give different accounts of what a social structure is and for present purposes I need not address the huge question of what exactly a social structure is. A very abstract account says that social structures are enduring orderly and patterned relationships between the elements of a society. A less, but still very, abstract account says that social structures are rules that constitute and regulate the major sectors of life such as family relations, property ownership and exchange, and political powers and responsibilities.[59] Since neither relationships, nor rules, are agents in any substantive sense, by its very nature structural discrimination is indirect discrimination. Also, on the definition I have offered the present structure of society might discriminate against a certain group, even if no present member of society engages in direct discrimination against members of this group. Finally, while institutions may involve and be part of a social structure, institutions are not social structures and vice versa. Accordingly, we can, in principle, have structural discrimination without institutional discrimination—e.g., because institutions are short-lived and ever changing and, thus, do not involve sufficiently orderly and patterned relationships to involve social structures—and institutional discrimination without structural discrimination—e.g., because the discriminating institution does not constitute or regulate major sectors of life and, thus, involves no social structures.

[59] Thomas Pogge (2008), *World Poverty and Human Rights*, 2nd ed. (Malden, MA: Polity Press), p. 37. Pogge uses the term "social institutions" but he also makes clear that he does not have in mind organized collective agents.

{ 3 }

Statistical Discrimination

1. Introduction

In the previous chapter I distinguished between direct and indirect discrimination. Now I will expound upon a different distinction, i.e., the distinction between statistical and nonstatistical discrimination (section 2). I shall relate it to the distinction between direct and indirect discrimination (section 3) and identify a number of features that many actual cases of statistical discrimination have, but which are not features of statistical discrimination per se (section 4).

One very controversial form of statistical profiling, to which I will return in chapter 11, is racial profiling by the police. In some countries it is common to use statistics regarding particular crime rates within different racial groups as a basis for deciding how to use scarce policing resources, e.g., to do more stop-and-searches, relatively speaking, of people belonging to racial groups in which, say, illegal drug dealing or certain other kinds of criminal activity is (believed to be on the basis of crime statistics or simply on the basis of general impression) more common. For instance, research from the Equality and Human Rights Commission (EHRC) on the use of police stop powers where officers do not require concrete suspicion of involvement in crime—so-called section 60 stops—shows that in the UK "black people were 37 times more likely to be stopped and searched under section 60 than white people in 2010–11."[1] Such statistical differences seem appalling to many. As David Wasserman puts it: "[W]e may feel outraged by being held liable merely because our association

[1] http://www.guardian.co.uk/uk/2012/jun/12/police-stop-and-search-black-people. The EHRC's report also concludes that the arrest rates in connection with section 60 stops are similar for blacks and whites, but says that: "The lack of a significant difference [in arrest rates] does not prove that black people are not inappropriately targeted." See also Sandra Fredman (2011), *Discrimination Law* (Oxford: Oxford University Press), pp. 60–61.

with a group, or our past conduct, makes us especially likely to have engaged in criminal or tortuous acts."[2]

Racial profiling, however, is just one form of statistical discrimination, and many other forms are neither controversial nor even commonly thought of as discriminatory.[3] For instance, for obvious reasons riot police pay more attention to young males at Premier League matches than to elderly females at botanical garden displays even if the latter gather in large, unruly crowds. This is profiling too and, ultimately, relies on statistical information about the likelihood of engaging in criminal activities such as hooliganism provided one is a young male or an elderly lady, respectively. Yet, no one finds this particular form of statistical discrimination problematic.

More broadly, all of us engage in statistical discrimination in that we treat people differently on the basis of explicit or implicit statistical generalizations pertaining to the group to which they belong; native speakers speak more slowly when talking to nonnative speakers (which, generally speaking, is quite nice and facilitates understanding); women walking home at night respond differently to an approaching lone stranger if this person is male than if she is a female; racial minority members are more alert to signs of racial bias when speaking to a majority member than when speaking to another minority member.[4] Indeed, acting in a social world without relying on statistical information about socially salient groups seems impossible. As Frederick Schauer puts it: "What some people pejoratively call stereotyping is an essential part of our cognitive and decision-making apparatus. It is simply how we think."[5] It may be that our use of statistical generalizations is often biased and the generalizations unsound—and in some cases in morally discreditable ways—but the fact that we all constantly rely on them is undeniable. The omnipresence of actions based on statistical beliefs makes it important to clarify what statistical discrimination is in the first place. Call this the *statistical discrimination definition challenge*. Any account of why it is wrong when it is, obviously must be based on a prior account of what it is.

---

[2] David Wasserman (1991), "Morality of Statistical Proof and the Risk of Mistaken Liability," *Cardozo Law Review* 13, 935–976, pp. 940–941. I am less certain about the claim Wasserman makes regarding inferences from our past criminal conduct; see David Edmonds (2006), *Caste Wars: A Philosophy of Discrimination* (London: Routledge), pp. 11–33.

[3] Cf. F. Schauer (2003), *Profiles, Probabilities, and Stereotypes* (Cambridge, MA: Harvard University Press), p. 5: "In the vast majority of instances the actuarial behavior of the insurance industry is accepted simply as a fact of life." See Schauer (2003), *Profiles*, p. 154 for more examples.

[4] Affirmative action motivated by a concern for corrective justice similarly amounts to statistical discrimination, when the state indulges in the generalization involved in "[e]quating blackness with victimhood"; Daniel Sabbagh (2007), *Equality and Transparency: A Strategic Perspective on Affirmative Action in American Law* (New York, NY: Palgrave Macmillan), p. 17.

[5] Schauer (2003), *Profiles*, p. 75.

2. Statistical Discrimination vs. Nonstatistical Discrimination

Obviously, statistical discrimination is a form of discrimination.[6] Hence, it will have to satisfy my general definition of discrimination offered in chapter 1. So, for instance, like discrimination in general, statistical discrimination is essentially comparative with respect to groups. It is conceptually impossible to engage in statistical discrimination against members of one group of people, if there is no other group of people who are treated better in the relevant respect. However, the fact that a certain group is being statistically discriminated against and thus treated worse does not mean that it is being treated badly or that it is harmed relative to how well off it was prior to the introduction of the statistically discriminatory practice. In principle, a practice of distributing benefits where some groups are given a proportionately larger share of the benefits may be said to be statistically discriminatory against those groups that are given a proportionately smaller share even if every group benefits relative to the initial situation and all end up quite well off. The issue here is that of the baseline, which I have touched upon in chapter 2.4. Its form in relation to statistical discrimination is no different from its form in relation to other forms of discrimination.

Similarly, to engage in statistical discrimination those against whom one discriminates must form a socially salient group, i.e., one's differential treatment must satisfy the social salience condition. The satisfaction or nonsatisfaction of this condition is what distinguishes statistically based differential treatment as such from discriminatory, statistically based differential treatment. Almost all forms of differential treatment that have been labeled "statistical discrimination" satisfy the social salience condition. There is, however, a notable exception, namely genetic discrimination, i.e., cases where insurance companies differentiate premiums on the basis of information about insurers' genes and the probabilities of their developing various illnesses later in life. I return to this exception in the appendix to this chapter. With this word of caution I propose the following:

> X engages in statistical discrimination against Y in relation to Z by Φ-ing if, and only if:
>
> (i) X group-discriminates against Y in relation to Z by Φ-ing (i.e., satisfies the conditions as stated in chapter 1) and
> (ii) It is because (X believes that) Y has P and (X believes that) Z has not, and because (X believes that) P is statistically relevant, that X treats Y worse than Z by Φ-ing.

[6] Like with discrimination in general, one can statistically discriminate in favor of, against, or between. If an aid agency is to distribute clothes to a group of men and women and acts on the basis of statistical evidence showing that men, on average, need larger sizes than women, it discriminates between men and women, but not against either. What matters here is that the relevant policy neither favors men over women so affected, nor vice versa and yet treats them differentially.

This means that nonstatistical discrimination can be defined as group discrimination where (ii) is false. So while, for instance, a misogynist who treats women worse than men is likely also to engage in statistical discrimination against women—hating women, the misogynist is likely to accept all sorts of (biased) statistical views about how women differ from men and treat women worse on the basis thereof—we could imagine a "pure" misogynist who only engages in nonstatistical discrimination against women, e.g., he treats them worse than men only because he hates women and not because he thinks gender difference is a statistical proxy for any other properties he thinks form the basis of how people should be treated.

By way of further illustration: an employer who refrains from hiring a better-qualified female applicant on grounds of this applicant having a higher probability of wanting to take leave in connection with childbirth engages in statistical sex discrimination, whereas an employer who refrains from hiring a better-qualified female applicant on the ground that this applicant has informed him that she intends to have a child and take parental leave engages in nonstatistical sex discrimination. As this example brings out it is not as if nonstatistical sex-discrimination relies on no statistical information—the employer will not employ the female applicant only because he relies on his statistical knowledge of the likelihood of people actually having children and taking parental leave if, at a job interview, they inform their possible future employer of their intention of so acting.[7] Or, alternatively, the employer might rely on his statistical knowledge of the likelihood of *women* actually having children and taking parental leave if, at a job interview, they inform their possible future employer of their intention of so acting.

Even if the motivating reason is not a statistical belief, there will always be statistical beliefs in the background without which the relevant motivating belief would not function as such. For instance, if we had no idea of the likelihood that people would act as they declare they will, the fact that they state at a job interview that they intend to apply for parental leave would give an employer no reason to act any differently from the way they would act had this not been stated. Observations such as these has led Frederick Schauer to question, in effect though not in these words, the coherence of the distinction between statistical and nonstatistical discrimination: "[T]he distinction between the use of the profile and the use of so-called direct evidence is far more illusory than real. Inferences drawn from observations or from physical evidence are themselves based on probabilistic generalization, and the cumulative set of inferences that produces a purportedly 'direct' conclusion or observation is nothing

[7] In this case, the employer might engage in indirect, statistical discrimination against women. The discrimination is indirect because the relevant statistical information does not concern women, but applicants in general, and because acting on this information disadvantages women relative to men.

more than a collection of inferences drawn from generalizations known to be reliable. Just like a profile."[8] Hence, we face a challenge here. On the one hand, we might experience some pressure to agree that the employer, who is motivated by the information supplied by the applicant, because of his statistical knowledge of women actually applying for parental leave if at a job interview they state their intention of doing so, is engaged in some form of statistical discrimination. On the other hand, if we admit that the relevance of this statistical background belief implies that the case is one of statistical discrimination, probably too many cases will count as cases of statistical discrimination. In effect, we here face a challenge pertaining to statistical discrimination that is similar to *the limitation challenge* pertaining to indirect discrimination, i.e., that almost any kind of discrimination amounts to statistical discrimination.

One possible exception to the claim that all forms of discrimination involve statistical background beliefs might be cases of differential treatment where the agent experiences animosity against members of the group against whom he discriminates. However, even in such cases, typically there are statistical beliefs in the background such that if the agent did not have these beliefs, he would not be motivated to act in a discriminatory way, nor experience animosity.

Another possible exception can be brought out through imagining a slightly more complicated example. Suppose an employer has interviewed 100 applicants for a position. Each and every one of them truly and trustworthily stated whether he or she intended to apply for parental leave at some later stage. The employer now has to make a decision about whom to hire. She made a note of which applicants intend to apply for parental leave, but unfortunately she has lost this note. All she remembers is that another member on the interviewing board declared that 45 out of 50 female applicants intended to apply for parental leave, while only 5 out of 50 male applicants intended to do so. She remembers she was baffled by this observation and did the calculation to see if it was true, which it was. If she now treats being a female applicant as a fallible but quite reliable proxy for intending to apply for parental leave, is she then engaged in statistical discrimination against women? Clearly, her reasoning does not fit Schauer's description of "generalization": "On the basis of a characteristic of some members of a class, we reach conclusions or make decisions about the entire class."[9] The employer knows by enumeration the distribution of intentions to apply for parental leave across men and women, and does not make any inductive inference about this, or about the likelihood of any individual applicant's intentions, from knowledge of the distribution within some smaller subset of the entire class of applicants. Also, the employer does not rely

[8] Schauer (2003), *Profiles*, pp. 171–172. Cf. Schauer on individualized testing, pp. 65–67; Edmonds, *Caste Wars*, pp. 11–33.

[9] Schauer (2003), *Profiles*, p. 4.

on any statistical facts about women (and men) as such—rather, she bases her reasoning on knowledge that she acquired through a deductive inference from the conjunction of facts about the self-professed intentions of each and every applicant. Intuitively, this case is not a case of statistical discrimination. For one thing, the relevant statistical information does not pertain to a socially salient group. Yet, to the extent that we object to cases of statistical discrimination because such cases do not treat each individual on her own merits, we should find the employer's conduct in this case objectionable for the very same reason.[10] The same complaint would apply if all applicants had been divided, on the basis of coin-flipping, into two equally large groups both of which have equally many male and female applicants and it just so happens that many more members of the first group than members of the second declare that they intend to apply for parental leave, and the employer, having lost information about each individual applicant's stated intention, acts on the basis of her knowledge of which of the two groups to which the applicant belongs.

To meet the definition challenge, one has to offer an account of the difference between cases involving foregrounded and backgrounded statistical beliefs. One suggestion here is that we can distinguish between cases where the relevant belief is formed through a process where the relevant fact that the belief concerns has actually played a causal role in the formation of the belief and cases where the relevant fact is causally inert.[11] So, for instance, if an applicant informs the employer at a job interview that she intends to have two children and to apply for parental leave, the fact that she so intends will then be part of the cause of why the employer forms the belief that this applicant intends to apply for parental leave in the future. If instead the employer forms the very same belief on the basis of and, hence, foregrounded statistical evidence suggesting that the majority of female applicants intend to apply for parental

[10] Suppose all 50 female applicants had stated that they intended to apply for parental leave and no male applicant had done so. Since the employer has forgotten what each applicant said—setting aside that she is not relying on statistical information about men and women as such—does she engage in statistical discrimination when she prefers male to female applicants? Can a rejected female applicant complain that the employer ought to have asked her about her intentions (again) even though the employer could logically infer from what she knew to be true of the group of female applicants what the rejected female applicant's intentions are? If not, why could a rejected female applicant so complain if 49 out of 50 applicants stated that her intention was to apply for parental leave such that the employer could be 98% certain that the rejected applicant intended to apply for parental leave? After all, even the former employer has no "individualized" evidence about the female applicants that she rejects. All she knows is that they are members of a group all of whose members intend to apply for parental leave, so she is relying on a generalization, albeit a universal one.

[11] Schauer (2003), *Profiles*, p. 246 on how to distinguish between two kinds of presumed offenses. See also Judith Jarvis Thomson (1986), *Rights, Restitution, and Risk* (Cambridge, MA: Harvard University Press), pp. 192–250; Joel Feinberg (1984), *Harm to Others* (Oxford: Oxford University Press), p. 201.

leave once hired, the fact that this particular applicant intends to do so has played no causal role in the employer's formation of this belief. Note, however, that in the former case it is not as if the employer's belief formation takes place independently of statistical beliefs—it is merely backgrounded as it were. If, for instance, she thought that people are very likely to state such intentions and not stick to them, she would give little weight to such statements. And if she thought that people who did not state such intentions were in fact even more likely to have an intention to apply for parental leave in the future, she would be more, rather than less, inclined to hire an applicant who avowed the relevant intention. The general point here is that even if we can draw the distinction between statistically based and nonstatistically based beliefs, the labeling of this distinction is misleading because virtually all beliefs are to some extent statistically based. As Schauer puts it: "[E]ven the processes that initially appear to us to be 'direct,' 'actual,' or individualized turn out to rely far more on generalizations from past experience than is often appreciated."[12] Note also that the distinction between statistical and nonstatistical beliefs strictly speaking fails to apply to beliefs about the future, assuming that facts about the future are causally inert with regard to beliefs that we hold now. The employer is concerned with whether the applicant as a matter of fact will apply for leave in the future—she is indifferent about whether the applicant now has an intention of doing so in the future per se. Yet, acting on the basis of information obtained through the agent's avowed intentions about future actions would not be thought of as cases involving acting on statistical beliefs.

Even supposing that such an account could be provided, and I think it could, it would then be a further question whether this distinction is morally relevant. One challenge that must be addressed here is that while most of us intuitively believe that beliefs formed through processes where statistical beliefs are backgrounded are more reliable than beliefs formed through processes where statistical beliefs are foregrounded ("naked statistical evidence"), this is not so. Observational evidence is notoriously fallible—people often make mistakes about what they see or hear—and in some cases their observations may not be very credible.

Another problem is that people ignore the problem of baseline probabilities. Drawing on work by Tversky and Kahnemann, Schauer illustrates this by imagining an example where 85% of the taxis in a city are from the Green Cab Company and 15% are from the Blue Cab Company. There is a traffic accident and a witness is certain that the offending driver drove a taxi and is also confident that it was blue. The fact that he is confident but not certain reflects that there is a 20% chance that his observational report is mistaken. If people are asked from which company they believe the taxi was from, most will respond

[12] Schauer (2003), *Profiles*, p. 105.

that it was from the Blue Cab Company. After all, the witness is quite certain that it was blue. However, this is an error, which is due to the fact that, when offered testimony based on observation, people ignore the much greater baseline probability that the taxi was green—after all, almost all taxis in the city are green. Once this is taken into account, we see that it is most likely that it was green.[13]

What does it mean for membership of a socially salient group to be statistically relevant? All it means is that the probability of having some other feature—say, applying for parental leave, possessing illegal drugs, or applying for admission at a medical school—varies on the basis of which particular socially salient groups one is a member of. This means that membership can be used as a proxy or, to put the same point differently, a statistical indicator for other features. That is, a discriminator may be wholly indifferent to membership in the relevant socially salient group per se but still treat people on the basis thereof because membership is tied to the probability of having other features about which the discriminator cares. So in this sense, statistical discrimination may be said to be indirect, even if it is a different sense of "indirect" from the sense employed in the previous chapter.[14] For P, i.e., membership in a socially salient group, to be statistically relevant, it is necessary and sufficient that the probability of having a certain interesting feature given that one is a member of this group is higher (or lower) than the probability of having this feature given that one is not a member of this group.[15] An implication of this is that P may be statistically relevant even if it is only a small minority of the members within a certain socially salient group that has the relevant property provided that the corresponding minorities in other socially salient groups constitutes a smaller or larger proportion of these groups.

Whether a given policy constitutes statistical discrimination depends crucially on how it came into existence. Suppose that statistical evidence available to employers suggests that Indian applicants are less well qualified for jobs within a certain industry and suppose that they are underrepresented within that category of jobs. Suppose that even if the evidence had suggested that they are better qualified than average, they would still have been underrepresented. They are underrepresented because most employers within that industry are

[13] There is a .12 probability that the cab is blue and the witness will report it as such (.15 × .8) and a .17 probability that the cab is green and the witness will misreport it as blue (.85 × .2). Hence, there is a .29 probability that the witness will report it as being blue, but note that the probability of a misreport (.17) is higher than the probability of a correct report (.12), when the witness reports the taxi to be blue. Accordingly, the probability that the taxi is blue given that the witness reports it to be blue is the true positives (.12) divided with the total positives (.29).

[14] Suppose I know that all Danes are lazy, do not care about being Danish as such, but dislike lazy people, whom however I find hard to identify. If I treat Danes worse for this reason, this would qualify as a limiting case of statistical discrimination, since here my statistical indicator of laziness is perfect assuming I really *know* that all Danes are lazy.

[15] To put it formally: p (H/P) > p (H/not-P).

non-Indians who are prejudiced against Indians. In the present account, Indians do not suffer from statistical discrimination, since the reason they are being treated less well has nothing to do with the relevant statistical evidence. No doubt, they might well suffer from other kinds of discrimination in my example, e.g., racial discrimination, and in all likelihood employers might attempt to misrepresent direct racial discrimination as statistical discrimination.

I am not suggesting that statistical discrimination cannot also be a case of racial discrimination. Suppose the relevant statistical evidence explains why Indians are badly represented within the industry, but that employers would not respond to comparable statistical evidence about non-Indians by not hiring non-Indians. In that case, the underrepresentation of Indians would reflect racial, statistical discrimination. Nor am I suggesting that racial discrimination may not be causally rooted in statistical discrimination.[16] Conceivably, the employers' hostility against Indians may be causally rooted in differential treatment motivated by statistical evidence suggesting that Indians differ in various ways from non-Indians. However, it is not clear that it suffices for the policy amounting to statistical discrimination.

3. Direct vs. Indirect, Statistical Discrimination

In chapter 2, I distinguished between direct and indirect discrimination. How does this distinction relate to the distinction between statistical and nonstatistical discrimination? Taxonomically, it is crucial to know whether all forms of statistical discrimination constitute direct discrimination. Morally, it is important too. If, for instance, direct and indirect discrimination are morally wrong for different reasons and statistical discrimination can be direct as well as indirect, presumably, different kinds of statistical discrimination will be wrong for different reasons too. If, alternatively, all forms of statistical discrimination are direct and some forms of statistical discrimination are not morally wrong, it follows that some forms of direct discrimination are not morally wrong. This last implication is very striking, because, as a matter of fact, many people believe:

1. Some forms of statistical discrimination are not morally wrong.
2. Statistical discrimination is a form of direct discrimination.
3. All forms of direct discrimination are wrong.

This set of beliefs form an inconsistent triad, so at least one of them will have to go. Determining which one(s) that is, I shall call the *trilemma challenge*.

[16] Cases in which the causal relation has the opposite direction are discussed in the literature under the label "displaced biases," see Larry Alexander (1992), "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies," *University of Pennsylvania Law Review* 141.1," 170.

Being clear on taxonomy is one way of bringing such challenges up to the surface.

Obviously there can be direct, statistical discrimination. In fact, I believe this is the most common form of statistical discrimination. So, in many cases, individuals treat people differently in part because they believe that their membership of a certain socially salient group is a good statistical indicator of other things in which they are interested, e.g., insurance companies charge young men higher premiums for car insurance because of the (true) belief that young men are more likely to be involved in traffic accidents than are other people.

However, indirect statistical discrimination is possible too, although perhaps not very common. Suppose the CEO board of a major city decides to employ racial profiling not knowing which racial group will be subjected to increased police scrutiny as a result. Suppose also that the relevant members of the board have taken part in direct racial discrimination against the racial group that, as a result of their decision, becomes the object of increased police focus—something which perpetuates the disadvantages they suffered as a result of past direct discrimination against them. Suppose, finally, that members of the board are no longer racially biased. These members made a decision that disadvantages members of the relevant racial group without this decision in any way reflecting any bias on part of these members or that they knowingly treat members of different racial groups differently, albeit in an unbiased way. True, they do not know that they disadvantage members of the relevant racial minority, but indirect discrimination does not require that the agent is aware of the fact that she disadvantages those against whom she indirectly discriminates.

By way of further support for the possibility of indirect statistical discrimination consider a case where statistical evidence suggests that Y-people differ from non-Y-people statistically speaking, Y-people are treated worse for that reason, and it so happens that X-people are more inclined to be Y-people than non-X-people.[17] So, by way of illustration, consider a case

[17] This may not be sufficient. Take the case of hiring according to height. Presumably, there are an infinite number of groups other than women who, on average, have a lower height, e.g., very old men, immigrants from Vietnam, people who weigh less than 55 kilograms, and people who had an inadequate diet as children. It is not clear that the mere fact that they are less likely to be tall is enough to say that they are subjected to indirect, statistical discrimination. One suggestion is that the additional necessary condition of being identified as a victim of indirect, statistical discrimination (needed to form a conjunction of conditions sufficient for indirect, statistical discrimination) is that one is a member of a socially salient group, members of which are harmed by the relevant, statistically based differential treatment. This problem is akin to what I called the *discriminatee identification challenge* in relation to indirect discrimination (see chapter 2).

where employers base hiring decisions on a certain general ability job test where people who pass the test are more likely to succeed in the job than those who do not and whites are more likely to pass the test than blacks. In some such cases, the use of the test might be considered indirect statistical discrimination against blacks (and not direct statistical discrimination against people, whites as well as blacks, who fail the test), say when the employer is unaware of the fact that the test results are correlated with race and where direct racial discrimination plays a role in the explanation of why employers use the relevant test.[18] Hence, I conclude that the distinctions between direct and indirect and between statistical and non-statistical discrimination cut across one another:

Discrimination	Statistical	Nonstatistical
Direct	The insurer who charges young males higher car insurance premiums	The misogynist who discriminates against women because he dislikes them
Indirect	The CEO board deciding to use racial profiling not knowing which racial groups will be targeted as a result thereof	Promoting on the basis of seniority in a company that excluded minorities in the past and not because seniority is used as a statistical indicator of something else

In short, any form of statistical discrimination is either direct or indirect discrimination and some forms of statistical discrimination are direct, others indirect.

4. What Statistical Discrimination Is Not

Actual cases of statistical discrimination may often involve other kinds of discrimination or, for that matter, other kinds of morally problematic features that are not essentially tied to statistical discrimination. To have a clear picture of what exactly statistical discrimination per se is, these other features should be set aside. They are useful to have in one's mind when one morally evaluates a concrete case involving statistical discrimination, but confusing to have in one's mind when one assesses morally statistical discrimination *as such*.

[18] See Mark Kelman (1991), "Concepts of Discrimination in 'General Ability' in Job Testing," *Harvard Law Review* 104, 1157–1247.

First, it is often claimed that statistical discrimination is based on statistical evidence that is false or, whether false or not, insufficient.[19] So, for instance, if a certain general ability test is used in relation to hiring applicants and it turns out that the members of certain socially salient groups pass the test in much smaller numbers, it is often submitted that the test is unreliable. It may be that the test is skewed in a way that reflects cognitive discrimination on part of the designers of the test (see chapter 1) or, more simply, the test may just be a bad one that, or more precisely, the use of which, accidentally happens to disadvantage members of some socially salient groups. No doubt, generally speaking policies ought not to be based on false or insufficient statistical evidence. However, sometimes the relevant statistical evidence suggesting that X-people differ from non-X-people in dimension—D—is overwhelming and correct—e.g., that children are less able to understand complex political issue than people above the age of, say, 15 or 18, or that people above the age of 60 on average have slower reaction times, worse hearing, worse vision, and a greater "risk of sudden and unexpected incapacitation" than those between, say, the age of 30 and 60[20]—it is not a necessary feature of statistical discrimination that it is based on insufficient statistical evidence. Obviously, from a moral point of view, cases of statistical discrimination where the relevant statistical evidence is sufficient and corresponds to the facts are most interesting.[21]

Second, sometimes relevant statistical evidence is used selectively (or, alternatively, some statistical evidence is used disproportionately) and, thus, some groups, but not others, are subjected (disproportionately) to statistical

---

[19] E.g., J. Angelo Corlett (1993), "Racism and Affirmative Action," *Journal of Social Philosophy* 24.1, 163–75, p. 164; David Wasserman (1996), "Racial Generalizations and Police Discretion" in JohnKleinig (ed.), *Handled with Discretion: Ethical Issues in Police Making Decisions* (Lanham: Rowman & Littlefield), 115–130, pp. 119, 123. See also Schauer (2003), *Profiles*, pp. 35–36. Schauer believes that, ultimately, the morality of acting on "a statistically sound but nonuniversal generalization" boils down to whether it is "morally worrisome in those cases in which the generalization does not hold," Schauer (2003), *Profiles*, p. 40. One possibility here is to say that in such cases the agent treats the patient unjustly, but that, given the costs to all of us that would result from refraining from acting on generalizations, it is morally permissible, all things considered, or legitimate in many cases to act on statistically sound but nonuniversal generalizations; see G. A. Cohen (2011), *On the Currency of Egalitarian Justice and Others Essays in Political Philosophy* (Princeton, NJ: Princeton University Press), pp. 124–143, 225–235. The fact that Schauer admits that in acting on such generalizations one will inevitably make "mistakes" (Schauer [2003], *Profiles*, p. 54) suggests that he should accept some such view. For what do the "mistakes" Schauer refers to consist in, if not in treating the patient unjustly?

[20] Schauer (2003), *Profiles*, pp. 108–130.

[21] Schauer (2003), *Profiles*, p. 18, takes the use of "statistically sound but nonuniversal applications" to be the "primary concern." He does so, in part at least, because he wants to contest the idea that "All human beings . . . deserve to be treated as individuals and not simply as members of a group," p. 19, and to defend the view that the use of statistically sound generalizations are often, but not always, morally permissible.

discrimination, e.g., when it is known that members of a minority group are more inclined to commit a particular kind of crime, police officers dealing with that sort of crime devote much more attention to this group, but when it is known that members of a majority group are much more often engaged in a comparable kind of crime no use is made of that—let us suppose—equally useful information.[22] While use of statistical information may often be selective—specifically, it may often be used to target minorities—the use of statistical information per se need not be selective or objectionably selective.[23]

Third, sometimes the available statistical evidence may be correct and used in a non-biased way. However, it might be clear that the available statistical evidence is not the best available and that given reasonable efforts much more evidence could be obtained. Suppose that the London Metropolitan police in response to the 2005 bombings had started giving special attention to anyone with an appearance different from the typical English one, e.g., by giving as much attention to Buddhists, Japanese tourists, Inuits, and elderly Pakistani women as to tense young Asian males carrying rucksacks. Given the sparse statistical evidence available—four suicide bombings all in all and some background information—it was reasonable to assume that a suicide bomber would not have a typical English appearance. However, it would also have seemed just as reasonable to assume that there was more fine-grained statistical evidence on which to base police tactics at the time. For instance, Buddhists, Japanese tourists, Inuits, and elderly Pakistani women were, I take it, as unlikely to be suicide bombers as people of typical English appearance. Accordingly, having a biased and much too broad focus in the case might easily be seen as sending a message of hostility and dislike. The failure to obtain the best available statistical basis given reasonable resource constraints may be morally objectionable for a number of reasons.[24] But, again, it is not true of the use of statistical information

[22] Schauer (2003), *Profiles*, pp. 186–8, notes that race tends to be frequently "overused" to the exclusion of other sound statistical indicators, so that even if law enforcement officers have a "profile that includes race" instead they simply apply a "racial profile," Schauer (2003), *Profiles*, p. 195. See also Iris Marion Young (1990), *Justice and the Politics of Difference* (Princeton, NJ: Princeton University Press), p. 59.

[23] In any case, the selective use objection never supports a conclusion of the form "This kind of statistical discrimination is morally wrong." At most, it supports a conclusion of the form "It is morally wrong to engage in that kind of statistical discrimination while abstaining from this kind: either you do both, or you do none" in which case the conclusion concerns not statistical discrimination as such, but a particular (perhaps common) form of statistical discrimination. Note also that, arguably, if one employs the notion of irrelevance discrimination (see chapter 1) such statistically informed differential treatment might not amount to statistical *discrimination*.

[24] See Laurence Thomas (1992) on "indiscriminate generalizations" in his "Statistical Badness," *Journal of Social Philosophy* 23.1, 30–41.

about socially salient groups as such that it involves not using the best information available.[25]

Fourth, a common objection to policies of statistical discrimination is to concede that the aims served by the relevant policies are worthy aims, but assert that these aims are better served by other means.[26] So, for instance, it might be conceded that while it is a worthy aim to prevent terrorist acts this is better done by building relations of trust between the police and communities from which terrorists are typically recruited rather than by singling out these communities for surveillance and preventive searches. A related objection says that while certain aims may be better achieved through statistical discrimination, the statistical facts are such that the relevant aims are only marginally better achieved and that this will be more than outweighed by the harms involved in setting aside and targeting a whole group of people. For instance, the statistical differences between groups may be small, in which case the use of statistical information may involve little benefits even if the properties for which one has a statistical indicator are very common. Suppose that 788 out of 1,000 white motorists drive a car with illegal brakes whereas 789 out of 1,000 black motorists do so. In that case checking only black motorists for illegal brakes might well (be perceived to) send an unpleasant message that is not worth a marginal improvement in the detection rate for motorists using cars with illegal brakes. Similarly, the relevant behavior that one wants to prevent through the use of statistical discrimination may be very rare, in which case the use of statistical information may involve little benefits even if the difference in how common the relevant property is in different groups is very large. Suppose that 1 out of 10 million Saudi Arabian air travelers flying to a US destination is a hijacker, whereas 1 out of every 100 million non-Saudi-Arabian air travelers flying to a US destination is a hijacker. Suppose, unrealistically, that these ratios are unaffected by a practice of screening Saudi Arabian passengers only and that there is no falling marginal benefits of the use of screening resources. Even then screening Saudi Arabian passengers only, but more thoroughly than done presently, would at best mean a reduction in an already tiny risk of boarding a plane about to be hijacked, a reduction that would have to be weighed against the costs involved in singling out Saudi Arabians for screening, e.g., the sense of humiliation and stigmatization experienced by people singled out for scrutiny as well as the costs involved in any harmful response on their part fuelled by a sense of resentment

[25] Which statistical evidence is the best available depends not only upon how reliable a predictor a certain feature is, but also upon how costly it is to establish whether a certain individual possesses the relevant predictor-feature. A very reliable predictor is of little use if it is very costly to establish whether a person possesses the predictor.

[26] E.g., Chana Berniker Cox (1993), "On Levin's 'Responses to Race Differences in Crime,'" *Journal of Social Philosophy* 24.1, 155–162, pp. 158–159.

over such differential treatment.[27] Obviously, whenever such objections are correct, there is no case for statistical discrimination and presumably a strong one against. But this leaves us with the question of what to say about cases where abstaining from statistical discrimination really does involve serious costs, say, in terms of the amount of serious crime prevented, taxes collected etc. and cases where being from the relevant socially salient group is just one of many features included in a profile. Often it may be controversial whether a particular case is of the sort where employing statistical discrimination is not self-defeating, but I assume that there are some such cases. At least, it is not true of statistical discrimination per se that it is self-defeating, e.g., because people will adjust to expectations in a way that statistical discrimination generates the very same statistical pattern that are then used to justify discrimination, or will in other ways fail to achieve its purposes.[28]

Fifth, in some cases the uses of statistical discrimination cluster around marginalized or impoverished minorities in such a way that either it is a sign of the fact that we do not relate to one another as equals or it is in itself something that causally contributes to making this the case.[29] In those cases, statistical discrimination may be said to involve subjecting certain socially salient groups to exclusion. No doubt, statistical discrimination clustered around badly off minorities will often involve such bad effects. When, for instance, African-Americans see themselves singled out by the police, unsurprisingly, they often see this as carrying the message that they are second-class citizens.[30] However, none of these features are necessary features of statistical discrimination. Suppose that all socially salient groups commit different kinds of crime and accordingly all groups are targeted by statistical discrimination to an equal extent. Under such circumstances being targeted by the police may not be seen

[27] To the extent that some of these costs are under the control of people subjected to statistical discrimination they should not simply be taken as given, for the purpose of moral evaluation, but themselves be subjected to moral evaluation. To see this in a less controversial and not wholly analogous case suppose that I am a young, male fan of my local football club known for its hooliganism. Suppose I attend matches wearing the relevant tribal dress carrying objects that potentially can be used for throwing at other people. If I feel humiliated and stigmatized when subjected to heightened attention by the riot police, this is a cost others should not simply take as a fact when addressing the justifiability of the relevant police tactics: they can ask me to reflect on the matter from the perspective of those caught up in mass fights etc. and to resist my resentment, which I could if I would.

[28] There might also be statistical generalizations where their use is self-fulfilling. This involves moral problems of their own, but lack of accuracy is not among them.

[29] For some recent formulations of this ideal, see Samuel Scheffler (2003), "What is Egalitarianism?" *Philosophy and Public Affairs* 31.1, 5–39; Samuel Scheffler (2005), "Choice, Circumstance, and the Value of Equality," *Politics, Philosophy, and Economics* 4.4, 5–28; Elizabeth Anderson (1999), "What is the Point of Equality?" *Ethics* 109.2, 287–337; David Miller, "Equality and Justice" in Andrew Mason (ed.), *Ideals of Equality* (Oxford: Blackwell, 1998), pp. 21–36.

[30] See Wasserman, "Racial Generalizations," pp. 116–117.

as a symbol of inferiority: it is something everyone is subjected to in some contexts.[31] Similarly, if neither the groups targeted nor those who carry out the discriminatory policy, are aware of the fact that statistical discrimination takes place, the relevant discriminatory policies may not involve any problems of stigmatization and messages of inferiority undermining our standing as equals. Second, not all kinds of statistical discrimination threaten our social standing as equals. For instance, statistical discrimination by tax units targeting tax avoidance by the rich or statistical discrimination involved in preventive police work focusing on violent crime by young males are hardly wrong for this reason. On this account, statistical discrimination is essentially connected neither with hostility against, nor with beliefs about, the inferior status of those who are subjected to discrimination.[32] In principle, those who implement a policy of statistical discrimination may be *more* favorably inclined towards those whom they subject to statistical discrimination than towards those whom they do not so treat and think of the former as having a superior status.

Finally, it might be thought that in relying on statistical information agents are more likely to treat patients in way that they should not be treated than they would be were they to use individualized information. While this may often be so, it is a contingent feature of discrimination. Where individuals are prone to make many errors when they form an individualized judgment they might be less likely to get things wrong than if they simply relied on broad generalizations and crude guidelines based thereupon.[33]

In sum: while statistical discrimination often appears together with other kinds of discrimination and, for that matter, features that are objectionable in various other ways, these features do not pertain to statistical discrimination as such. Accordingly, in so far as statistical discrimination is morally wrong per se, it cannot be wrong due to its having these features with which, at most, statistical discrimination contingently appears together with.

5. Conclusion

In this chapter I have defined statistical discrimination; distinguished it from nonstatistical discrimination; and presented a number of challenges to be addressed by any satisfactory account of statistical discrimination, to wit, (1) the

[31] Even when statistical discrimination is morally bad for these reasons, one will also have to take into account the possible benefits forgone by abstaining from statistical discrimination. If, for instance, the aim of preventing violent crime really could be achieved much better through various kinds of statistical discrimination, it would then be a further question whether reducing serious crime might not be worth some measure of stigmatization, i.e., it is not clear that the goal of ensuring a social environment in which we relate to one another as equals should take lexical priority over other aims such as reducing violent crime.

[32] Samuel R. Bagenstos (2003), "'Rational Discrimination,' Accommodation, and the Politics of (Disability) Civil Rights," *Virginia Law Review* 89.5, 825–923, p. 851.

[33] Schauer (2003), *Profile*, p. 260.

statistical discrimination definition, (2) the discriminatee identification, (3) the trilemma, and (4) the limitation challenges.[34] While I have argued—drawing on a distinction between foregrounded and backgrounded statistical beliefs—that it is possible to draw a distinction between statistical and nonstatistical discrimination and, thus, answer the limitation challenge, I have also argued that most acts rely on statistical information of some kind, and that statistical discrimination is widespread, if not practically unavoidable (which is neither to concede that any particular form of statistical discrimination is unavoidable, nor to concede that it is not unjust). I then distinguished between direct and indirect statistical discrimination arguing that the latter is a real possibility. Finally, I noted that statistical discrimination often appears together with other morally objectionable features, e.g., the selective use of relevant statistical information, but in section 4 I argued that these are at best contingently associated with statistical discrimination. Accordingly, insofar as statistical discrimination is wrong per se, it cannot be wrong in virtue of *these* features. In the next three chapters, I shall look at three general accounts of the wrongness of discrimination.

Appendix: Genetic Discrimination and Social Salience

As mentioned in chapter 1 there is one form of statistics-based differential treatment that is normally referred to as statistical *discrimination* even if the relevant group of discriminatees clearly does not form a socially salient group. This is the case of genetic discrimination, i.e., the case where insurers charge insurance takers different premiums on the basis of genetic information about the different health risks they face. Undoubtedly, the term "genetic discrimination" has won favor so in not allowing it as a form of discrimination, my account deviates from part of ordinary language, or at least from the relevant, specialized parts of ordinary language.

As noted in appendix 1 to chapter 1, reflecting ordinary language is not the only desideratum that a suitable definition of discrimination should meet. In this particular case, it would simply be theoretically unfruitful to admit that differentiation of insurance premiums on the basis of genetic information as such is a form of discrimination. There are two reasons why this is so: first, descriptively speaking, the social mechanism involved in this form of differential treatment is very different from standard cases of discrimination, and second, morally speaking, if such differentiation is wrong it is wrong for reasons different from the reasons why discrimination in general is wrong. To admit differentiation of insurance premiums on the basis of genetic information

[34] (1), (2), and (4) are similar to the relevant challenges pertaining to indirect discrimination.

as a case of discrimination would be a bit like noting that, apparently, some describe a ban on pit bulls as "canine racism," because it was motivated by the statistical generalization that this breed of dogs is more dangerous than other breeds, and then allow that people who favor such a ban manifest a particular form of racism and, accordingly, that an account of wrongness of racism must also accommodate this particular species of racism.[35] It seems theoretically much more fruitful to note that "canine racism" is here being used simply because of its rhetorical force and then seek to explore the wrongness of standard forms of racism independently of the one putatively at stake in the dispute about a ban on pit bulls.[36]

To back up this claim about theoretical unfruitfulness consider the following argument:

1. Genetic discrimination is best seen as a form of discrimination only if, by its very nature, it involves the same kind of injustice as ordinary forms of discrimination do by their very nature.
2. The injustice that, by its very nature, genetic discrimination involves is the injustice of some people being disadvantaged relative to others as a result of bad brute luck.
3. It is not the case that, by their very nature, ordinary forms of discrimination involves the injustice of some people being disadvantaged relative to others as a result of bad brute luck.
4. Hence, it is not the case that genetic discrimination is best seen as a form of discrimination.

The first premise rests on the assumption that discrimination is a concept that plays a particular normative role and that acts etc. can be wrong for reasons other than their being discriminatory. The second premise is true, because it is bad brute luck if one turns out to have bad genes and accordingly one will be disadvantaged, insurance-premium-wise, relative to others if insurers can differentiate insurance premiums on the basis of genetic information in a way that involves the injustice of bad brute luck disadvantages. There is no reason to think that, by its very nature, genetic discrimination will involve any other kind of injustice. For instance, differences in overall genetic health risks are not significantly correlated with membership of socially salient groups (even if particular health risks, e.g., the risk of developing sickle-cell anemia, are) such that it may be stigmatizing etc. The third premise is true because at least some

[35] Schauer (2003), *Profile*, p. 56.

[36] I would say the same about the term "geographical discrimination among citizens," i.e., letting people who live in areas that are prone to natural disasters like earthquake carry the burdens of so doing themselves; see Elizabeth S. Anderson (1999), "What Is the Point of Equality?" *Ethics* 109.2, 296. (By "discrimination" Anderson probably has in mind what I have referred to as "generic discrimination.")

ordinary forms of discrimination are unjust even if they do not involve the injustice of bad brute luck based injustice. For instance, it would be unjust to engage in religious discrimination even if one's religion is not a matter of bad brute luck.

It might be replied that the first premise of the argument presented is too strong and that it suffices for genetic discrimination to be usefully seen as a form of discrimination that, typically but not necessarily always, is wrong for the same reason as ordinary forms of discrimination. I am sympathetic to the view that this premise is too strong. However, genetic discrimination is not even typically wrong for the same reason as standard forms of discrimination. Overall genetic risks are more or less evenly distributed across socially salient groups and accordingly it is hard to see how genetic information-based differentiation of insurance premiums could result in stigmatization etc. Indeed, it might even result in more people realizing that there are no significant overall genetic differences across different racial groups. Accordingly, even a suitably weakened version of premise one is false.

By way of further support for my claim that genetic discrimination is not best seen as a species of discrimination in any morally relevant sense, suppose that I am wrong and that it should be seen as such. In that case, one might ask why differentiating insurance premiums on the basis of ordinary medical tests or the insurance taker's medical history should not also be seen as a case of discrimination. Genetic tests and ordinary medical tests etc. may differ in terms of the *content* of the information *typically* acquired through the relevant tests. Genetic tests rarely show that the person being tested has a present medical problem, only that the person has a certain risk of developing a certain health problem later in life, whereas medical tests sometimes (unfortunately) reveal just this. However, this is not a necessary difference between genetic tests and ordinary medical tests. The information made available through nongenetic, medical tests or medical records, e.g., that the person is a smoker, obese, or has an alcohol problem, bears on the likelihood that the person will develop certain health problems later on in life. Hence, the only necessary difference between genetic information and nongenetic medical information concerns the *source* of the relevant information. But it is puzzling that whether the use of a certain piece of information by an insurer is discriminatory depends not on what this information is information about—e.g., that there is a 20% probability that the person will suffer a stroke within the next 20 years—but depends on the source of this information—e.g., whether the information about this risk is obtained through genetic testing or through ordinary medical tests or ordinary records pertaining to the insurance taker's medical history. No one would think that differentiating insurance premiums on the latter basis is a form of discrimination and, as indicated, it is difficult to see why one should think differently about differentiating insurance premiums on the basis of genetic information.

Finally, my point here is not that differentiating insurance premiums on the basis of genetic information is just or morally justified, all things considered—it is typically neither precisely because it is a matter of making some worse off than others as a result of bad brute luck—but simply that, generally, it is wrong for reasons other than typical forms of discrimination.[37] Not all agree with this claim and, for instance, Deborah Hellman has argued otherwise.[38] On her view discrimination is wrong whenever it demeans the discriminatee and genetic discrimination demeans those who are charged higher premiums. Specifically she thinks that information obtained through nongenetic testing is different from information so obtained on the ground that, given our past history of eugenics and Nazism, the use of genetic information expresses a certain wrong-making objectionable meaning that the use of nongenetic information does not express: "genetic discrimination does indeed express disrespect for people with genetic mutations that especially dispose them to disease or disability."[39] I return to Hellman's account of the wrongness of discrimination in chapter 5.2–3 and as will become apparent there I do not accept it. However, even if her account of the wrongness of discrimination were correct, genetic discrimination does not necessarily demean the discriminate thereby undermining the second premise in my argument above.

Hellman submits that genetic discrimination expresses the meaning that the lives of [the] "flawed" individuals with genes that dispose them to develop certain diseases later in life are "less important and less valuable."[40] However, here it is crucial to distinguish between the claim that an individual is less important or valuable in the sense that she has a lower moral standing than others have, on the one hand, and the claim that her life has less value for her than the life of others have for them, on the other hand. It might be argued that genetic discrimination and, indeed, taking out insurance against having a certain genetic disposition for developing a certain illness reflects the view that one's life has less value for one with the illness than without it, without in any way implying that one has a lower moral status than others have. Hellman does not show why it is not simply the latter, and presumably unobjectionable, meaning that genetic discrimination expresses.

---

[37] Typically, but not always, differentiation of insurance premiums on the basis of genetic information is unjust because it involves not compensating for bad brute luck. People who face higher health insurance premiums because genetic tests suggest that they will live longer and, hence, all other things being equal, will need more medical assistance paid for by the supplier of the insurance than people who live shorter lives, are charged a higher premium in a way that reflects their having *good* brute luck.

[38] Deborah Hellman (2003) "What Makes Genetic Discrimination Exceptional?" *American Journal of Law and Medicine* 29.1, 77–116.

[39] Hellman (2003), "What Makes Genetic Discrimination," p. 106.

[40] Hellman (2003), "What Makes Genetic Discrimination," p. 108.

Another challenge to Hellman's view becomes apparent if we assume that everyone is genetically disposed to an equal degree to develop illnesses later in life though it will differ across persons *which* diseases they are likely to develop. Accordingly, everyone will pay the same if they want to insure against all genetically determined illnesses although what they will pay for insuring against a particular genetically determined disease might vary a lot. In this scenario it would seem very unclear how genetic discrimination could express the meaning that people have unequal moral value. Admittedly, the assumption made—that everyone is genetically disposed to an equal degree to develop illnesses later in life—is false, but it may well contain a significant grain of truth. Add to this fact the fact that many persons who are demeaned in other respects, e.g., in virtue of their race, are likely to be charged very low health insurance premiums—a fact, which on Hellman's reasoning, would seem to carry the deracializing message that these people are almost "flawless" individuals unlike those "flawed" individuals belonging to an otherwise nonstigmatized racial group who will have to pay high health insurance premiums—and it becomes unclear why basing insurance premiums of genetic information would have the social meaning ascribed to it by Hellman. For these reasons, I do not think that Hellman's argument undermines my case for not regarding genetic discrimination as a form of discrimination in the sense that interests us here.

This page intentionally left blank

{ PART II }

The Wrongness of Discrimination

This page intentionally left blank

{ 4 }

Mental-State-Based Accounts

1. Introduction

The first part of this book focused on the concept of discrimination and distinguished between various kinds of discrimination. This part explores the moral qualities of discrimination. Clearly, many acts of group discrimination are seriously wrong. However, it is unclear what makes group discrimination morally wrong, when it is. The book's second part neither surveys all accounts of the wrongness of discrimination that could be offered, nor all accounts that actually have been offered.[1] But it does review most of the more influential, recent accounts on offer.

By claiming that something makes an act of discrimination wrong, I do not mean that it makes the action impermissible.[2] I mean that it is a feature of the action that counts as a reason in favor of its being impermissible to perform. However, there might also be reasons for the action being morally permissible. So as I am using the term, an act of discrimination might be wrong in virtue of having a certain feature, even if it is, all things considered, morally permissible to perform it. If such an act is performed it will, of course, in one respect be morally problematic, because in virtue of a particular feature it is a morally wrongful act.

Chapter 4 primarily assesses accounts of the moral wrongness of discrimination that focus on the nature of the mental state reflected in discriminatory acts. Having briefly commented on the general issue of how mental states relate to wrongness (section 3) and having provided an overview of different mental-state accounts of the wrongness of discrimination (section 4), sections 5 to 7 take a closer look at one such very influential account—the one expounded by Larry Alexander in a recent classic. However, before mental-state accounts enter center stage, I shall reflect briefly on the nature of the question about the

[1] Some of the accounts of the wrongness of discrimination discussed below have not been offered as accounts of the wrongness of group discrimination in my specific sense. However, the notions of discrimination employed are sufficiently close to group discrimination for my discussion in this chapter and the next one not to be off target. (The harm-based and the desert prioritarian accounts discussed in chapter 6 are offered as accounts of the wrongness of group discrimination.)

[2] Note 8 in the introduction.

wrongness of discrimination (this section) and quickly run through some of the more often canvassed, but sometimes less well articulated, accounts thereof (section 2).

Discriminatory acts can be morally objectionable for reasons other than the fact that they are discriminatory. In some cases, the reason why a particular instance of discrimination is wrong is clearly not primarily that it is an instance of discrimination even though it clearly *is* a case of wrongful discrimination. The Nazis discriminated against Jewish (and other groups of) people by singling them out for the death camps. However, the wrongness of Holocaust is due at least partly (if not mostly) to the wrongness of killing millions of innocent people. Indeed it would seem belittling the immorality of these events to focus on discrimination alone. Where discrimination is wrong as such, it is wrong in virtue of properties that are true of discriminatory acts by definition; and were we to conclude that discrimination is not wrong as such, we would still be able to agree that many discriminatory acts are seriously morally wrong—not because they are discriminatory, but because of other features they have, like being such as to humiliate the victims or to deny them access to important goods.

When we ask what makes discrimination wrong, what we are really asking is what makes discrimination morally worse than nondiscrimination. Hence, the question involves an implicit comparison. But what discrimination is to be compared with is unclear. If we focus on (iv″) in my definition of group discrimination, i.e., that P is the property of being a member of a certain socially salient group to which Z does not belong (see chapter 1), one suggestion is that the relevant comparison is with acts where people are not treated differently, period. Another suggestion is that the comparison should be with acts where people are treated differently on the basis of differential group membership but where the relevant groups are not socially salient. In the latter view, to ask what makes discrimination wrong is to ask what reasons, if any, speak against differential treatment on the basis of *socially salient* group membership, but do not speak against differential treatment on the basis of *nonsocially salient* group membership. If this is the relevant comparison, it seems there is nothing wrong with discrimination *as such*. While there are, as I will explain in chapter 6, important instrumental reasons why differential treatment tracking membership of socially salient groups tends to be morally worse, because more harmful, than differential treatment tracking membership of nonsocially salient groups, the mere difference in whether others respond differently to a particular case of group membership cannot make an intrinsic moral difference. Hence, I shall assume that the relevant comparison is with equal treatment whatever kind of group membership is at issue.

By way of a further preliminary remark, I note that the next three chapters focus on the wrongness of discrimination as such and, according, largely, though not completely, ignore moral differences between different forms of

group discrimination. For instance, I do not at length address the fact that one may to some extent control whether one is obese and that this might make discrimination against obese people less wrong than sex discrimination, ceteris paribus. Similarly, by and large I ignore that discrimination against disabled people differs from discrimination against the elderly in that it increases lifetime inequalities. For present purposes, I can remain neutral on whether such morally relevant differences exists—I think they do—but I shall assume, for the purpose of enquiry, that there is at least one feature common to all forms of wrongful discrimination by virtue of which they are morally wrong. This is the feature I want to identify. Admittedly, the assumption that such a feature exists could be denied. However, an analogous one may seem plausible in relevantly similar cases, e.g., terrorism. In any case, the assumption is merely one that structures my enquiry, not one that predetermines its results. Specifically, it does not guarantee that a "wrong-making feature" common to all cases of discrimination can be identified. As a matter of fact, one significant conclusion that will emerge from part 2 of the book is that there is no such feature in virtue of which group discrimination is morally wrong per se. This conclusion may be unsurprising when it comes to the desert prioritarian account discussed in chapter 6—it implies that acts that do not maximize the sum of moral value are the only acts that are morally wrong per se—but it is more surprising in relation to the views discussed in chapters 4 and 5 and to certain versions of the harm based account discussed in chapter 6, since these are, or are often seen as, deontological views.

2. Some Common Accounts

I start by taking a look at some popular accounts of the moral wrongness of group discrimination as such. The first, and probably best-known, account says that discrimination is wrong because it involves treating people arbitrarily or irrationally.[3] Frederick Schauer, for instance, seems to endorse this account when he writes that it would be wrong to use gender as a proxy for upper-body strength and then not hire women for jobs as computer programmers, when "it is fairly obvious that upper-body strength has little to do with the skills we would expect in our computer programmers."[4] Call it *the irrelevance account*.

[3] If one uses the notion of irrelevance discrimination defined in chapter 1 it is true by definition that discrimination involves treating individuals on arbitrary grounds.

[4] F. Schauer (2003), *Profiles, Probabilities, and Stereotypes* (Cambridge, MA: Harvard University Press), p. 142; cf. *Ghaidan v. Godin-Mendoza* quoted in Sandra Fredman (2011), *Discrimination Law*, 2nd ed. (Oxford: Oxford University Press), p. 131.

Unfortunately, as pointed out by Matt Cavanagh, this account is problematic for a number of reasons.[5] First, as I have already argued, in some cases one's membership of a socially salient group is (in a commercial or some other respect) relevant to how one should be treated. This seems, for example, to be true in cases of discrimination involving reaction qualifications. But if we consider differential treatment wrongful discrimination in such cases, it follows that discrimination cannot be wrong as such simply because it involves treating people differently on irrelevant grounds (see chapter 9).

Second, we are not in *general* required, morally, to treat people on the basis of relevant grounds: my acquiring a friend on a whim is not morally objectionable. Perhaps there are special contexts in which we have a moral obligation to attend to relevant factors in making decisions, but if there are morally objectionable cases of discrimination outside these areas—e.g., in decisions about which neighborhood to live in—these cannot be objectionable because they involve treating people on the basis of irrelevant desiderata.

Third, to treat someone badly on idiosyncratic grounds—say, because they have a family name with six letters—is to be guided by arbitrary considerations. Hence, if discrimination is morally objectionable, because it involves treating people differently on irrelevant grounds, these two cases are morally on a par. Yet many would unhesitatingly regard them as morally different.

Note, finally, that the present criticisms do not suggest that there are not very many contexts in which treating people differentially on the basis of their membership of socially salient groups is not treating them on irrelevant grounds or arbitrarily in some sense. It is just that treating them on this basis is not what *makes* these actions morally wrong.[6]

Another influential account says that it is morally wrong to discriminate against people on the basis of preferences regarding salient group membership such as race or sex, because people do not choose which salient groups to be members of.[7] Call it *the choice account*. On a weaker and a more plausible view, discrimination on the basis of race or sex is partly wrong, because the discriminatees did not choose their race or sex. So while this is not the only

[5] Cavanagh (2003), *Against Equality of Opportunity* (Oxford: Clarendon Press), pp. 99–101, 156–157. Cavanagh writes: "If we actually came across someone who discriminated against people on the basis of the number of vowels in their name, I suspect we would think them mad, rather than bad," p. 156. Setting aside that people with the same number of vowels in their names do not form socially salient groups, how we would assess the discriminator's character is irrelevant to the wrongness of an act of discrimination. It is possible to engage in wrongful discrimination even if one is not a bad person (and possible for a mad person to be bad).

[6] Obviously, this is consistent with irrationality being an indicator of some other wrong-making feature, e.g., if employers do not employ wheelchair users even if being a wheelchair user is obviously irrelevant to the job at hand, then, assuming that people generally act rationally, the non-employment of wheelchair users might be an indicator of discomfort-based disadvantageous treatment of disabled people and this in turn might be morally wrong.

[7] See Cavanagh (2003) *Against*, pp. 161–166, 186–187, 192–193.

reason discrimination on such grounds is wrong, discrimination on the basis of unchosen or uncontrolled properties is morally worse than discrimination on the basis of chosen or controlled properties.[8] Peter Singer writes: "The person who is denied advantages because of his race is totally unable to alter this particular circumstance of his existence and so may feel with added sharpness that his life is clouded, not merely because he is not being judged as an individual, but because of something over which he has no control at all. This makes racial discrimination peculiarly invidious."[9]

Clearly, there is something to this view, albeit it may not provide us with a full and satisfactory account of the wrongness of discrimination. Generally speaking, treating people badly on the basis of properties the possession of which they themselves do not control, e.g., whether they are African American, harm people more than treating them badly on the basis of properties the possession of which they do control, e.g., whether they smoke. The obvious reason is that one's treatment of them gives them a reason to change and they may well succeed in doing so if they exercise some control over whether they have the property on the basis of which people are treated badly. One can opt out of (putative) discrimination against smokers by not smoking, but cannot opt out of racial discrimination by changing race.[10] So, as a matter of fact, discrimination that targets unchosen characteristics is more harmful than discrimination that targets chosen characteristics (or, for that matter, characteristics that are not chosen but which change randomly across time).[11]

Still, this account of the wrongness of discrimination is vulnerable to three objections. First, treating people on the basis of properties that they have not themselves chosen (or, whether they have themselves chosen it or not, do not control whether they will continue to have) is not in general wrong—it is not wrong for doctors to treat patients differently on the basis of their unchosen illnesses—and if even if it were, some kinds of wrongful discrimination targets

[8] Something may result from my past choice and yet not now be subject to choice on my part.

[9] Peter Singer (1983), "Is Racial Discrimination Arbitrary?" in Jan Narveson (ed.), *Moral Matters* (New York: Oxford University Press), p. 317. I take it that Singer thinks that racial discrimination is "peculiarly invidious" because the feeling, which Singer refers to, manifests a strong preference against not being judged as an individual. At least, given Singer's acceptance of preference utilitarianism this seems to be the best way to make sense of his remark.

[10] This might not be entirely true due to the possibility of immigration.

[11] Refining the objection one might distinguish between characteristics that result from choices in the past (or, more precisely, *relevantly informed* past choices) and characteristics the persistence of which are affected by present and future choices (whether or not they result from past choices). Arguably, discrimination on the basis of characteristics acquired in the past in a choice-independent way may be less wrong than otherwise, if present or future choices affect the persistence of the relevant characteristic. Also, discriminating against people on the basis of properties, which they do not control whether they retain, but which they control whether they acquire in the first place, may be less wrong, because people can avoid becoming discriminatees by not acquiring the relevant property in the first place.

people on the basis of characteristics that in the sense which is relevant here *are* chosen—e.g., one's religious beliefs.[12]

Second, suppose that in the future scientists discover a pill that causes one to be a heterosexual and another pill that causes one to be a homosexual. The pill is cheap and widely accessible. People can then indirectly choose their sexuality. Surely, this would not render discrimination against homosexuals morally permissible. If so, the wrongness of discrimination lies not simply in the fact that the discriminator treats people differently on the basis of unchosen characteristics, so this defeats the strong version of the choice account.

Finally, it might be argued that much (but not all, viz., the previous objection) of the appeal of the view that there is something particularly morally wrongful about treating people disadvantageously on the basis of characteristics that they do not control whether they have, may derive from the natural, but false, assumption that when they are so treated, they are harmed in a way that they do not control. However, this does not follow.[13] Suppose that one does not control one's sexuality, but one has full control over whether one reveals it, and, unlike in the actual case of, among others, gays and lesbians, one is not harmed in any way by having not to reveal it in order to avoid discrimination. I conjecture that in this hypothetical case, we would not consider discrimination on the basis of sexuality as wrongful as we do under the actual circumstances.

A third account of the moral wrongness of discrimination in certain specific contexts commonly offered claims that discrimination violates a meritocratic norm according to which the best qualified ought to acquire whatever benefit is at issue, e.g., the job in question. Call it *the meritocratic account*. A prominent proponent of this view is David Miller: "Justice demands that the job be offered to the best-qualified applicant. We express this by saying that the best-qualified applicant deserves the job or, in a slightly different formulation, that the principle involved is one of merit. This is the principle that condemns discrimination on grounds of sex, race or religion when hiring employees."[14]

At best, this account is partially correct. First, it has a rather narrow scope, unless we adopt unusually wide notions of qualifications and jobs, or more broadly (see chapter 10), positions. Discourse about qualifications is at home in competitive contexts such as the labor market or university admissions, but many forms of discrimination take place outside these contexts. Consider discrimination in relation to house purchases—where, say, buyers are reluctant to move into a neighborhood with a significant proportion of people of a race

[12] Indirect control suffices for choice in the relevant sense.

[13] For a similar point made in relation to luck egalitarianism, see Richard J. Arneson (1990), "Liberalism, Distributive Subjectivism, and Equal Opportunity for Welfare," *Philosophy & Public Affairs* 19.2, 158–194, p. 186; G. A. Cohen (2004), "Expensive Taste Rides Again," in Justine Burley (ed.), *Dworkin and His Critics* (Oxford: Blackwell), 3–29, p. 7.

[14] David Miller (2001), *Principles of Social Justice* (Cambridge, MA: Harvard University Press), p. 156.

other than their own. On the assumption that acting on such inclinations may be morally objectionable, the meritocratic norm provides at best a partial explanation of the wrongness of discrimination.

Second, violations of the meritocratic norm seem, in general, morally different from discriminatory decisions about whom to hire. Suppose a company decides to hire a less well-qualified, immigrant applicant because it wants to do its share to improve the conditions of immigrants. This case seems relevantly different from an otherwise comparable case in which a company decides to hire a less well qualified male applicant because they do not want to hire women. Since the cases involve equally grave violations of the meritocratic norm, our differential moral evaluation of the two suggests that something else is at stake as well.

Finally, the meritocratic norm may in itself be problematic and certainly more problematic than the claim that discrimination is wrong. Even to the extent that the meritocratic norm is not problematic, its underlying rationale might be one of sheer efficiency. From within a Rawlsian perspective, for instance, it has been argued that it is unclear why individuals behind the veil of ignorance would not prefer an improvement of the prospects of the worse off to a violation of the democratic principle of fair opportunity.[15] Similarly, then, if the meritocratic norm is ultimately grounded in considerations of efficiency, it seems that there could be no objection to forms of discrimination that increase efficiency and yet many would think that such discrimination can be wrongful.[16] (I return to the relation between meritocracy and discrimination in chapter 9.)

A fourth account of why discrimination is wrong has it that discrimination increases inequality of opportunity and that this is morally objectionable. Call it *the equality of opportunity account*. In a recent article, Shlomi Segall offers a sophisticated defense of this view claiming that "[d]iscrimination is bad as such . . . because and only because it undermines equality of opportunity."[17]

Measures to reduce discrimination may well increase equality of opportunity across different socially salient groups; and the claim that this account explains the wrongness of a considerably wider range of cases than the meritocratic norm does is credible. However, it is not difficult to conceive of cases in which discrimination reduces inequality of opportunity. Consider discrimination

---

[15] Richard J. Arneson (1999), "Against Rawlsian Equality of Opportunity," *Philosophical Studies*, 93.1, 77–93.

[16] It is a common mistake to believe that discrimination is inefficient and that firms which practice it will not survive in a free market, e.g., Narveson, *Moral Matters*, p. 250. For a critique of the underlying assumption here, see Cass Sunstein (1991), "Why Markets Won't Stop Discrimination," *Social Philosophy and Policy* 21.2, 22–37; Richard A. Epstein (2007), *Forbidden Grounds: The Case against Employment Discrimination Laws* (Cambridge, MA: Harvard University Press), esp. pp. 59–78. While Epstein does not think that no forms of discrimination will exist on a free market in the long run, he also thinks that state intervention to eliminate or neutralize the effects of discrimination is counterproductive.

[17] Shlomi Segall (2012), "What's so Bad about Discrimination?" *Utilitas* 24.1, 82–100, p. 82.

against an otherwise well-off group that serves to offset other advantages enjoyed by this group and, thus, to reduce inequality of opportunity—e.g., racially motivated discrimination against better-off Chinese in favor of the generally worse-off Malayans in Malaysia. If we assume that racial discrimination is morally objectionable, because it violates equality of opportunity, we may have no grounds on which to object to this variety of discrimination.[18] A similar point applies to a case where everyone discriminates against everyone as a result of which everyone is harmed, but harmed in a way that leaves intact equality of opportunity (see the example in chapter 5.5, second last paragraph). Obviously, friends of the equality of opportunity account might bite the bullet and say that this case involves no wrongful discrimination.[19] Indeed, discussing a case in some ways analogous to the last one, Segall concedes that "perhaps" there might be something bad about it, "but if so, it does not seem to me to be for reasons of discrimination."[20] I am worried about this move. In the light of cases such as the ones presented, many will say that this case shows that what is common to all incidents of wrongful discrimination is not inequality of opportunity, but something else. In saying that this instance may be bad but not for reasons of discrimination, Segall does not, as far as I can see, give those who do not already accept his equality of opportunity-based account any reason to revise their views.

A fifth suggestion is that discrimination is wrong, because it involves social exclusion of discriminatees, where social exclusion involves being prevented from "participating in the benefits of citizenship or membership of society" as a result of which discriminatees end up disadvantaged in the noncomparative sense that they do not enjoy the "essential elements of 'well-being'" on some perfectionistic account thereof.[21] Call this the *social exclusion account*.

---

[18] On the harm-based account and, more specifically, on the desert-prioritarian view we do have reason to think that this kind of action is wrong, since, surely, it is harmful to some and it reduces the desert levels of others (see chapter 6).

[19] Related problems arise in connection with discrimination against disabled people provided that the disabilities in question are ones that everyone faces an equal risk of acquiring, e.g., as a result of a car accident or disease, and with age discrimination provided that everyone has the same age profile.

[20] Segall (2012), "What's so Bad?" p. 95. Segall would also be willing to bite the bullet in relation to the first case, if racial discrimination really did reduce overall inequality of opportunity (personal communication). He thinks, however, that intuitions about this case might reflect that such discrimination would exacerbate inequality of opportunity between the Chinese and non-Chinese well-off segments of Malaysia: see Shlomi Segall (forthcoming), *Equality and Opportunity* (Oxford University Press).

[21] Hugh Collins (2003), "Discrimination, Equality, and Social Inclusion," *Modern Law Review* 66.1: 16–43, pp. 22–24. While the specifics of Collins's account are not common, it is much more common to think that the wrongness of discrimination derives from the social exclusion it involves. Also, Collins seeks primarily to account for the "aim" of antidiscrimination legislation, not the wrongness of discrimination, and, obviously, avoiding social exclusion might in large part be what motivated legislators to enact anti-discrimination law even if social exclusion does not offer the best account of the wrongness of discrimination.

One problem with the social exclusion account is that discrimination can be wrongful even when directed against people who are not socially excluded. For instance, racial discrimination against members of a dominant racial group is wrong, even though the discriminatees in no way suffer social exclusion.[22] This problem is particularly pressing if one construes social exclusion along the non-comparative sufficientarian, perfectionist lines. In this construal, the social exclusion account is compatible with *any* form of discrimination not being morally wrong as long as everyone enjoys the essential elements of well-being.

Another problem is that the social exclusion account is unable to explain the wrongness of discrimination against socially excluded discriminatees, where, for some odd reason, discrimination does not strengthen, or perhaps even weakens, the social exclusion of these people.

A final problem is that one can be socially excluded in ways that do not involve discrimination, e.g., because of one's being convicted of child abuse or because of one's holding extremist, morally offensive political views. The social exclusion account is unable to explain why the social exclusion of such groups of people are not wrongful in exactly the same way as the social exclusion involved in paradigmatic cases of discrimination.

Having surveyed five common, usually cursorily mentioned, but problematic, accounts of the wrongness of discrimination, the rest of this part of the book will explore some accounts that have been developed in greater detail. These accounts need not be rivals in the sense that discrimination, or least many instances of discrimination, cannot have several wrong-making features, including the objectionable mental state of the discriminating agent, the objective meaning of the discriminatory act as well as their harmfulness.[23] Often, however, theorists focus on versions that are rivals rather than accounts that are combinable into some hybrid account of the wrongness of discrimination, e.g., because it claims that it gives the one and only reason why discrimination as such is wrong, or because it explains what makes discrimination distinctively wrong. In the rest of this chapter I will focus on mental-state accounts, specifically one proposed by Larry Alexander.

3. Mental States and Permissibility

Mental-state accounts contend that discrimination is wrong, when it is, because it reflects a certain morally objectionable mental state on behalf of the discriminator. The idea that mental states determine the wrongness of an action is well

[22] Collins (2003), "Discrimination," p. 42, acknowledges this problem.

[23] Acts of discrimination could also have more wrong-making features than the three features mentioned here, e.g., the features mentioned above as well as others.

known outside the context of discrimination. For instance, the doctrine of double effect holds that actions that involve some bad effect might be permissible when this effect is merely foreseen, but not permissible when it is intended. The doctrine is often invoked to explain why terrorism differs morally from conventional warfare involving a similar number of innocent casualties.[24] Similarly, the Kantian principle that one must never treat someone merely as a means seems to involve a link between permissibility and mental states. More generally, Anderson and Pildes have defended the view that "what makes an action morally right depends on whether it expresses the appropriate valuations of . . . persons" where " 'expression' refers to the ways that an action or a statement (or any other vehicle of expression) manifests a state of mind."[25]

However, setting aside certain complications other moral theorists argue that mental states, or at least intentions, never affect the moral permissibility of actions at a suitably fundamental level. For instance, Thomas Scanlon and Judith Thomson defend this claim in relation to intentions on the ground that it would be "odd" or "weird" to offer the following advice to an agent who is deliberating about whether it would be morally permissible for her to perform a certain action: "That depends on the intention with which you will perform the action, if you do so. If you will intend to bring about the harmful effects involved in the action, it is impermissible. If you will not intend them, but merely foresee them, then it is permissible."[26] While this argument—the deliberative perspective objection—seems to apply not just to intentions but to mental states in general and, thus, seems to support the conclusion I will defend below, I find the argument unconvincing and accordingly will not draw support from it.[27] In any case, in the present context addressing the relation between permissibility and mental states in its full generality would take us too far away from the specific issue of wrongness of discrimination.[28]

---

[24] Jeff McMahan (2009), "Intention, Permissibility, Terrorism, and War," *Philosophical Perspectives* 23.1, 345–72, p. 346.

[25] E. S. Anderson and R. H. Pildes (2000), "Expressive Theories of Law: A General Restatement," *University of Pennsylvania Law Review* 148.5, 1503–75, pp. 1504, 1506. Because what matters from Anderson's and Pildes's point of view is the mental state of the agent and not the action's objective meaning, I take their account to be a mental-state-based account and not an objective meaning account (see chapter 5).

[26] Thomas Scanlon (2008), *Moral Dimensions: Permissibility, Meaning, and Blame* (Cambridge, MA: Belknap Press of Harvard University Press), pp. 19–20; Judith Jarvis Thomson (1991), "Self-Defense," *Philosophy & Public Affairs* 20.4, 283–310, esp. pp. 292–96; and Judith Jarvis Thomson (1999), "Physician-Assisted Suicide: Two Moral Arguments," *Ethics* 109.3, 497–518, pp. 515–16; Hellman (2008), *When Is Discrimination Wrong?* (Cambridge, MA: Harvard University Press), p. 153.

[27] See Kasper Lippert-Rasmussen (2010), "Scanlon on the Doctrine of the Double Effect," *Social Theory and Practice* 36.4, 541–564.

[28] For a recent defense of the view that motives can affect the permissibility of an action, see Steven Sverdlik (2011), *Motive and Rightness* (Oxford: Oxford University Press).

4. Different Mental-State Accounts

Assuming that mental states affect the moral permissibility of actions, there is a wide range of objectionable mental states, which an account of the wrongness of discrimination may invoke. For instance, mental-state accounts may ground the wrongness of discrimination in noncognitive states such as hostility or in cognitive states such as the belief that members of a certain group are inferior. More generally, mental-state accounts vary in terms of which propositional attitudes—e.g., believing, expecting, hoping, doubting—and in terms of which contents—i.e., object—say, "I will win the lottery"—believed, expected, hoped for, or doubted, etc.—they contend are the *loci* of the wrongness of discrimination. Also, they may vary in terms of how they flesh out the relation of reflection that obtains between the mental state and the relevant discriminatory action.

Not all mental-state accounts are concerned with equality. For instance, Richard Arneson proposes that "discrimination that is intrinsically morally wrong occurs when an agent treats a person identified as being of a certain type differently than she otherwise would have done because of unwarranted animus or prejudice against persons of that type."[29] A person who treats Copts differently than Muslims because of unwarranted animus or prejudice against Copts responds differently to members of those two groups. But she might have no mental states to the effect that Copts have lesser worth than Muslims, e.g., because she is an unreflective person who does not bother to rationalize her aversions. Similarly, Matt Cavanagh believes that discrimination, which involves treating people "with unwarranted contempt," is wrong.[30] Being contemptuous need not clash with equality of moral status. I can feel contempt for someone, say, whom I consider a superior person, on account of how she has squandered her talents. Similarly, I might think of a nonhuman animal as having a lower moral status and yet harbor no contempt for it. With this broader picture as background, I now turn to Alexander's account.

5. Alexander on Disrespect and Discrimination: The Falsehood Account

In an article that deservedly has become a standard reference in the last two decades' writings on discrimination, Alexander observes: "Discrimination is not one thing, but many."[31] Most forms of discrimination are wrong, when they

[29] Richard Arneson (2006) "What Is Wrongful Discrimination?" *San Diego Law Review* 43.4, 775–807, p. 779.

[30] Cavanagh (2003) *Against*, p. 166.

[31] Larry Alexander (1992), "What Makes Wrongful Discrimination Wrong?" University of Pennsylvania Law Review 14.1, 149–219, p. 153.

are, for contingent, consequentialist reasons. However, in Alexander's view, one species of discrimination is intrinsically wrong for a different and Kantian reason, namely discrimination "premised on the belief that some types of people are morally worthier than others," where for some to have greater moral worth than others is for them to merit "greater moral concern" than others.[32] Here, Alexander submits, the consequentialist account must be supplemented by a respect-based account. Presumably, by "intrinsically wrong" Alexander does not rule out the possibility that, all things considered, some such actions may be morally permissible, e.g., because they minimize the overall number of such intrinsically wrongful acts. Hence, I conjecture that his belief that consequentialist accounts of the wrongness of discrimination must be supplemented reflects that respect-based accounts are often embedded in a deontological moral framework, which involves an agent-centered duty to avoid disrespecting other persons. In principle however, consequentialist theories can incorporate concern for disrespect by treating it as non-instrumentally valuable that actions that express respect are performed.[33] For simplicity, I shall set aside this possibility in what follows.

While Alexander focuses on harmful and disrespectful cases of discrimination, his account would seem to allow that acts of discrimination can be morally wrong, because disrespectful, even if they harm no one in any morally relevant respect. This seems particularly plausible, as far as an account of respect goes, if the relevant sense of harm is loss of well-being.

The ambition behind Alexander's account is to explain the moral wrongness of a rather limited part of the full spectrum of discrimination cases. This seems well motivated. Many forms of discrimination involve no mental disrespectful states at all. This is true by definition in the case of indirect discrimination (see chapter 2) and it is plausible to say that the same point applies to direct statistical discrimination per se (see chapter 3). Hence, at best mental-state-based accounts of the wrongness of discrimination can at most cover only a limited range of cases.[34]

---

[32] Alexander (1992), "What Makes?" pp. 160–161. Alexander's notion of moral worth is a specific kind of what may be labeled "moral status."

[33] David Cummisky (1990), "Kantian Consequentialism," *Ethics* 100.3, 586–516.

[34] Friends of a mental-state-based account might add that while statistical and indirect discrimination are often wrong, they are not wrong in the way that discrimination, or at least the set of "core" cases of discrimination, is distinctively wrong. One problem with this move, however, is that it is unclear on what ground one can declare a way of being wrong as one that is not the way in which discrimination is distinctively wrong other than one's prior commitment to the mental-state-based account. Why, for instance, assume that there *is* a way of being wrong that is distinctive (in a sufficiently robust sense) of some subset of acts of discrimination? (A similar move applies to a similar move by friends of the objective meaning account—see chapter 5.) Another problem specifically tied to Alexander's account is that acts other than discriminatory acts, e.g., idiosyncratic disadvantageous treatment, may involve the sort of disrespect, which he thinks makes a certain class of discriminatory acts wrong for, as he puts it, non-consequentialist reasons.

Also, it seems possible to be a cognitive discriminator even if you have no beliefs about the moral status of those you are discriminating against, and it is unclear how the mere biased formation of a belief presupposes anything about the moral status of the individuals represented in one's beliefs. In fact, it is even possible, albeit psychologically extraordinary, to be firmly convinced that individuals belonging to a group have a *higher* moral status than in fact they have and yet, in some contexts, be involved in cognitive discrimination against them.

Finally, nonvaluational, nonstatistical discrimination simply involves brute desires and hence need not rest on any value judgments about moral status. To steer around these problems one should focus on the case in which the respect-based approach to the wrongness of discrimination is most plausible: valuational, noncognitive discrimination (see chapter 1). This is the kind of discrimination where Alexander's account seems most plausible.

Alexander's account is a member of the family of respect-based accounts of the moral wrongness of discrimination. On these views, discrimination is morally wrong, because it involves a failure to respect discriminatees.[35] On Alexander's particular account, an act or practice fails to respect an individual if, and only if, it involves an underestimation of that individual's moral status relative to that of others.[36]

Obviously, there are other accounts of what respect and disrespect involves. The respect-based account is really a family of accounts. Its adherents may differ, for instance, over the properties by virtue of which we merit respect and over when an act or a practice presupposes that someone has a lower moral status. On some accounts only those properties by virtue of which one qualifies as a person are properties by virtue of which one merits respect. On this view, discrimination against human beings not qualifying as persons (perhaps newborn infants or people in a late stage of Alzheimer's) is not morally wrong for the same reason as discrimination against human beings who are persons and thus merit respect.

---

[35] Absence of respect does not imply disrespect. For example, in certain contexts, such as public transportation, not showing concern for one's co-passengers is neither respectful nor disrespectful.

[36] Commenting briefly on an earlier version of my argument below, Glasgow (2009), "Racism Is Disrespect," *Ethics* 120.1: 64–93, p. 83*n*44, ascribes to me the view that "respect-based accounts of discrimination . . . entail that discrimination can be bad only when someone is 'actually being disrespectful' and that the only way out of this problematic entailment is to have the discriminator falsely represent the discriminate as having a lower moral status." If I understand Glasgow correctly this is surely false. I specifically note that respect-based accounts are "really a family of accounts" and that, in part, I *stipulatively* define disrespectful acts or practices as acts or practices that involve false representations of the discriminatee as having a lower moral status, see Kasper Lippert Rasmussen (2006), "The Badness of Discrimination," *Ethical Theory and Moral Practice* 9.2, 167–185, p. 178. When I write "in part," the reason is that I discuss Alexander's account, which *does* tie disrespect to false representations of moral status.

In a different account of respect proposed by Harry Frankfurt, "[t]reating a person with respect means . . . dealing with him exclusively on the basis of those aspects of his particular character or circumstances that are actually relevant to the issue at hand."[37] In this account, respect is essentially nonrelational. Hence, if discrimination is wrong because it involves a failure to respect in this sense, discriminatory acts may be wrong for the same reason as nondiscriminatory acts that involve treating people equally, but on the basis of equally irrelevant properties. In the account explored here, if failing to respect an individual is to treat him as if he has a lower moral status than he in fact has, then the concern for respect is also nonrelational. The basic reason why it is morally wrong to treat one person as if he has a lower moral status than another, when in fact they have the same, is that one treats at least one of them as if he has a moral status different from the one he in fact has. It is only derivatively morally wrong that one treats them differently. Hence, in Frankfurt's account of respect, it is not clear that, say, people who have been discriminated against in relation to punishment—see chapter 8—have been treated wrongly provided that they got the punishment they deserve when others got off more lightly, e.g., in virtue of their sex or race, and provided that the "aspects of [an individual's] particular character or circumstances" does not include facts about other defendants who are guilty of the same crime are punished.[38]

Joshua Glasgow has proposed a respect-based account of racism, which might be thought, mutatis mutandis, to speak to the wrongness of discrimination in general (assuming that racism is or includes a particular form of discrimination). In Glasgow's view any form of racism is morally objectionable because disrespectful, where disrespect is "something like a failure to adequately recognize autonomous, independent, sensitive, morally significant creatures."[39] Such inadequate recognition may be located in the racist's attitude, doxastic states, or behavior, and the racist need not be a person, but might be a collective entity like a company or a state. Assuming that disrespect is *pro tanto* wrong and assuming that behavior can be disrespectful independently of their manifesting disrespectful mental states, it follows that respect-based accounts overlap only partially with mental-state-based accounts. Alexander's account is both respect- and mental-state-based, but as Glasgow's account reminds us, some might think that acts can be disrespectful independently of the mental states they involve. Moreover, in principle at least, some might think that an act

[37] Harry G. Frankfurt (1997), *Necessity*, *Volition, and Love* (Cambridge, UK: Cambridge University Press), p. 150.

[38] If such facts were relevant from the point of view of respect, Frankfurt's account of respect would seem to incorporate the value of equality and despite, what he says (e.g., "It is easy to confuse being treated with the sort of respect in question with being treated equally," Frankfurt (1999), *Necessity*, p. 150), treating with respect would logically require equal treatment in a range of cases.

[39] Glasgow (2009), "Racism," p. 85.

may be wrongful because of the mental states they involve, even if these states involve no disrespect. While, as I noted, Glasgow's account is offered as an account of racism, one could build on his rough account of disrespect to provide an account of the wrongness of discrimination. However, the prospects for such an account are not good. Obviously, we need to know more about what "adequate recognition" involves. If it involves recognizing that the discriminatee is an "autonomous, independent, sensitive, morally significant creature," it seems that very many kinds of discrimination are not disrespectful, e.g., a statistical discriminator might certainly be acutely aware that those against whom she discriminates are "autonomous, independent, sensitive, morally significant" creatures and express this belief in many ways through her actions. If however it means recognizing and treating the discriminatee in accordance with the moral requirements that one is subjected to, the suspicion is that respect turns out to be parasitic on a prior account of what these moral requirements are and that, in any case, some forms of discrimination—e.g., age discrimination in relation to health-care priorities—are not disrespectful to the discriminatee, because the discriminator respects her moral claims.[40]

Returning to Alexander's account, one attractive feature of it is that its account of (dis)respect is that it is more elaborate than the one hinted at by Glasgow, who explicitly makes no pretention of providing an account of what respect involves. Another attractive feature is that some paradigm forms of racist discrimination have involved the relevant disrespectful belief, e.g., the belief held by the Nazis that Jews are subhuman. It is unclear, however, that all paradigm forms of discrimination involve beliefs about unequal moral status. Paternalistic sexism holds that men and women have different functions, but *in principle* this is compatible with the view that they have equal moral worth.[41] Similarly, discrimination against disabled persons often involves various forms of condescending treatment, but those who engage in it hardly believe that disabled people merit less moral concern than the able-bodied.

Alexander suggests three reasons why this kind of discrimination is intrinsically morally wrong. First, he submits that the underlying biases "are intrinsically morally wrong because they reflect incorrect moral judgments."[42] Similarly, he writes: "[B]iases premised on the belief that some types of people are morally worthier than others are intrinsically morally wrong because they

[40] For a different account of respect, see Stephen Darwall (1977), "Two Kinds of Respect," *Ethics* 88.1, 36–49. For the view that expressions of disrespect are epiphenomena of actions being unjustified, see Shelly Kagan (1989), *The Limits of Morality* (Oxford: Clarendon Press), p. 175.

[41] Indeed, paternalistic treatment of women may even be motivated by the view that women have a higher status than men. Orthodox Judaism often couches the maxim that women stay at home and not work in phrases that suggest that they are closer to God than men are, and hence of higher moral status than men.

[42] Alexander (1992), "What Makes?" p. 161.

reflect incorrect moral judgments."[43] On the view suggested by such passages, the fact that the bias reflects a false moral judgment explains why it is morally wrong. Call this *the falsehood account*.

According to the falsehood account, valuational discrimination is wrong, because it is based on a false moral judgment. This has several noteworthy implications. First, in this account, discrimination against and discrimination in favor of are morally symmetrical. Favorable treatment based on a false belief that someone has a higher moral worth than he has is no more, and no less, based on a false moral judgment than unfavorable treatment based on a false belief that someone has a lower moral judgment than he has.

Second, if X has a positive bias in favor of Y and this is based on his incorrect judgment that Y is more worthy than he is, then this is no less wrong than if X thought himself to be worthier than Y, since both involve a false judgment about moral worthiness.[44] So, by way of illustration, it is no more wrong for a male sexist to be biased in favor of men, himself included, due to his false belief that men are more worthy than women than it is for an oppressed woman to be biased in favor of men due to her false beliefs that men are more worthy.

Third, incorrect moral judgments about a person's moral worth can be absolutely or relatively incorrect (or both). This makes no moral difference on the falsehood account. An individual, X, may incorrectly consider himself morally slightly more worthy than Y and at the same time consider himself *and* Y to be morally much more worthy than they in fact are. The falsehood account does not have the resources to explain why discrimination based on this false belief is morally less problematic than discrimination based on a slight absolute error and a great comparative error.

However, it is not in general true that because one acts on the basis of false beliefs, one's action is intrinsically morally wrong. If I hold the false but justified belief that people like myself are less inclined to act altruistically, and treat people who are different from me better than myself for this reason, my act is ill informed but it is not intrinsically morally wrong. This suggests that insofar as the falsehood account is true, it reflects a specific moral truth about discrimination and not a general moral truth. To assess the falsehood account so construed we must clarify two issues. First, we need to clarify what it is for treatment of someone to be based on a false judgment about moral worth. One suggestion is that an act that is disadvantageous to the target, or believed to be so by the agent, can be said to be based on the target's moral status if, and only if, that moral status is the lowest possible compatible with this treatment's being

[43] Alexander (1992), "What Makes?" p. 161. Alexander thinks that such disrespectful acts of discrimination are wrong regardless of the gravity of its effects, see Alexander (1992), "What Makes?" p. 159.

[44] This, as well as the previous, observation is true of the two other accounts discussed in section 6 and 7.

morally permissible. Hence, if I treat someone in a way that would be morally permissible only if this individual was a nonhuman animal, my act is based on the belief that his moral status is lower than it is in fact. Thus, I fail to respect him.[45]

This suggestion should be rejected. First, it implies that one does not treat another person disrespectfully when one acts toward this individual in a morally permissible way. Yet it would seem possible for me to treat someone in a way that is morally permissible, but for reasons that make my treatment disrespectful. For example, as a borrower I might punctually and despite my dire situation repay someone the tiny sum I owe him, thereby signaling my contempt for his, let us suppose, unreasonable insistence on repayment whatever the circumstances. Second, because of false information I may act impermissibly in treating an individual in a certain way and yet have the deepest respect for this individual. In such cases, saying that I fail to treat this individual with respect seems to add nothing to the claim that I act impermissibly. Finally, the account of presupposing something about someone's moral status implies that whenever an act is morally impermissible, it is based on a falsehood about the discriminatee's moral status. However, it also implies that the act is not morally wrong in virtue of this believed falsehood. For this to be the case, being based on a falsehood about the person's moral status would have to be something other than being morally impermissible. Yet friends of the respect-based account want to say that being based on a falsehood about moral worth is different from, and hence something by virtue of which, acts are morally impermissible. Accordingly, an act can be based on an assumption about the moral worth of the affected individual if, and only if, this act is somehow motivated by the actor's judgment of the individual's moral worth.[46] In this account, a sufficiently unreflective wrongdoer may make no assumptions about the moral status of his victims.

Second, we need to clarify the concept of moral worth and who can possess it. One possibility is that some sentient beings have moral worth and the rest have none. When Kant famously contrasted persons that have worth with mere things that only have a price, he had this possibility in mind.[47] Kant thought that all rational beings have worth by virtue of their rationality and that their worth does not vary, say, with varying cognitive abilities. Variations above (and

[45] On another suggestion, an individual's moral status is roughly determined by the moral weight given to his interests. Hence, having a false view about an individual's moral status is to have a false view of the moral weight that her interests should be given.

[46] One can make a false moral judgment concerning someone's moral status, treat this person as if this person had a lower moral status, and yet not treat this person as such *because* one holds this false moral judgment (it may be causally inert). We can leave as undecided whether such an act is based on the relevant incorrect judgment about moral worth.

[47] Immanuel Kant (2002[1785]) *Groundwork for the Metaphysics of Morals* (New Haven, CT: Yale University Press), p. 53.

below) the threshold of rationality are irrelevant to personhood and do not affect moral status. Although almost all adult human beings are persons, not all human beings are persons in the relevant moral sense, i.e., rational beings that have the capacity for self-consciousness (fetuses, newborn infants, and irreversibly comatose people are exceptions; of course, fetuses and newborn infants have the potential to become persons).[48] Hence, if we tie the wrongness of discrimination to a failure to respect, and if this in turn is tied to the moral worth of persons, we cannot explain the wrongness of certain kinds of discrimination. Consider, for instance, someone who does not disrespect persons whatever their sex, but discriminates against fetuses and newborn infants (potential persons) on the basis of sex. Imagine that this person will not perform an abortion when the fetus is male but is happy to do so when it is female. Defenders of a respect-based account are, of course, free to say that this kind of discrimination is wrong, not because of any failure to respect the worth of the fetuses, but because it is hurtful and thus harmful for some existing persons, in this case women. This, however, would not explain why it would be wrong to screen fetuses to prevent homosexuals from coming into existence assuming that so far no homosexual has existed.

Another way to think about moral worth is to hold that all sentient beings have moral worth, but in varying degrees. This is likely the sense of worth Alexander has in mind. He explicitly allows for, but does not commit himself to, the view that virtuous and deserving persons have a greater moral worth than vicious and undeserving persons.[49] According to most versions of this view, the function that takes us from a given degree of virtuousness to a resulting degree of moral worth is continuous, so different individuals have different levels of moral worth. Moreover, there is a huge difference between the most deserving (say, Gandhi) and the least deserving (say, Hitler).

With these clarifications in mind, the question becomes whether the mere falsity of the judgment of moral worth on which one's act is based renders that act morally wrong. Consider first a case in which X treats Y worse than Z in a way that is impermissible because, through no fault of his own, X holds the false belief that Y is very undeserving and Z very deserving. Consider next a case in which X treats Y worse than Z in a way that is impermissible because, through no fault of his own, he does not falsely believe that Y is very undeserving and Z very deserving. The question then is: do these cases differ, morally speaking, and if so, which is worst?

In answering this question we should ignore two confusing factors. First, harmful behavior that reflects a judgment of lower moral status is often more harmful, other things being equal, than behavior that reflects no such judgment,

[48] Jeff McMahan (2002), *The Ethics of Killing: Problems at the Margins of Life* (Oxford: Oxford University Press), p. 6.

[49] Alexander (1992), "What Makes?," p. 159.

and for that reason it is often morally worse. This is merely a contingent truth and as such does not support Alexander's account. Second, when we think of cases in which someone misjudges the moral worth of others, we tend to have in mind cases involving epistemic vice—e.g., cases in which the discriminator applies epistemic double-standards in assessing evidence of moral worth. Such vices, however, should be disregarded; they need to be considered instead in relation to the irrational comparative falsehood account.

Returning to our question, one view is that the two cases do not differ morally speaking, since while X's beliefs about Y's and Z's levels of desert are relevant to our evaluation of his moral character, they are irrelevant to our moral evaluation of his act. This view seems promising. However, for present purposes we can set this aside and note that if X's beliefs about Y's and Z's levels of desert are relevant to evaluating his act, it would seem that what X does in the second case is worse than what he does in the first. Although X's actions in the first case were wrong, he at least believed (or had beliefs to the effect) that the person he treated badly was thoroughly undeserving. In the second case not only were X's actions wrong, but he also lacked the belief that the person he treated disadvantageously was thoroughly undeserving. If this makes *any* moral difference, it suggests that the latter act is worse. But if one's discriminatory activities are in some cases less bad when accompanied by a false judgment about the moral worth of the person affected by those activities, it follows that discrimination cannot be bad simply because it reflects an incorrect judgment of moral worth. Hence, we should reject this account.

6. Alexander on Disrespect and Discrimination: The Comparative Falsehood Account

One might emphasize a different aspect of the previous citation from Alexander's article, namely that the relevant falsehood concerns relative moral worth. Such false beliefs might plausibly be held to constitute "a failure to show the moral respect due to the recipient."[50] The underlying thought seems to be the Kantian one that persons have a duty to respect one another's moral worth and that in holding a person to have less moral worth than others, one violates this duty. Call this *the comparative falsehood account*.

But again: acting on the basis of false comparative beliefs is not wrong in general. Suppose I believe that I deserve less than others. Suppose, moreover, that this belief is false and yet I act on it. Setting aside controversial moral duties to self, here my act is based on a false comparative belief without being morally wrong. One might suspect that there is some aspect of the mind of the

[50] Alexander (1992), "What Makes?," p. 159.

mental states of the Nazi bias holder that explains why she fails to grasp the truth of equal moral worth, but then it is this mental state that renders her discriminatory act morally wrong, not her failing to believe a clearly true proposition.

Moreover, if the argument above shows that we should reject the falsehood account, a similar conclusion is soundly drawn here. Comparing a case in which X treats Z better than Y and this treatment is based on the false belief that Z is more worthy than Y, with a case of differential treatment in which X has no such false belief suggests the false belief is either irrelevant to our moral evaluation of the acts or reduces the moral wrongness of the act.

7. Alexander on Disrespect and Discrimination: The Irrational Comparative Falsehood Account

Commenting on Nazi biases, Alexander writes: "Their biases were intrinsically morally wrong because Jews are clearly not of lesser moral worth than Aryans."[51] One might read this passage as specifying something that was implicit in the passages on the wrongness of certain kinds of discrimination by Alexander that I quoted above. It is not the mere fact that a bias reflects a false comparative moral judgment that renders it intrinsically wrong, but rather the fact that it is a false, comparative judgment that is *clearly* false.[52] This suggests an explanation that is more specific than the previous one. Discrimination is not intrinsically morally wrong simply because it involves a bias based on a false belief, but because it involves a bias that is based on a comparative belief about status that is *clearly* (presumably, even to the discriminator) false. The Nazis had available to them abundant reasons for thinking that people of different races (however defined) have identical moral worth, yet they managed to form spectacularly false beliefs about the matter. Call this *the irrational comparative falsehood account*.

This is the most promising interpretation of Alexander's view and one I explore below. Specifically, I focus on three important questions that it raises. First, why is the content of the belief that renders discrimination intrinsically morally wrong one that concerns unequal moral worth and that only? Second, why is the only relevant propositional attitude to this content the attitude of

[51] Alexander (1992), "What Makes?," pp. 158–159.

[52] This passage could be read in other ways as well. First, it might be suggested that what is morally relevant is that the bias reflects a false moral judgment to the effect that some have a lower moral worth than others. Someone who fails to believe what his beliefs entail may falsely judge some to have a higher moral status without judging that others have a lower moral status. This complication is ignored below. Second, it might be suggested that what is morally relevant is that the discriminator's moral judgment favors people of his "own kind" and that the individuals judged to have different moral worth in fact have the same moral worth.

believing? Third, how must the relevant belief relate to the relevant discriminatory act for it to be intrinsically wrong? I address these questions in the order just listed.

First, I address the question relating to the content of the relevant belief. Compare a case in which X treats Z better than Y and this treatment is based on the false *and irrational* belief that Z is more worthy than Y, with a similar case of differential treatment in which X has no such belief. Is the former morally more wrong than the latter? One possibility that needs to be explored is whether the falsity of the judgment of Y and Z's comparative worth drops out of the picture as irrelevant to moral evaluation of differential treatment. To see what motivates such a line of enquiry, consider a case in which the discriminator correctly judges the comparative moral worth of Y and Z, based on an irrational, false moral judgment about the relevance to one's moral worth of membership in the socially salient groups to which Y and Z belong. So suppose a discriminating, sexist agent falsely believes that Y is a particularly morally virtuous person—indeed so morally virtuous that the discriminator irrationally believes Y, despite being a woman, is more worthy than Z, who is a man. Suppose additionally that Y is in fact morally worthier than Z. Here, we might assume, X makes an irrational, but correct judgment about the comparative moral worth of Y and Z. From the point of view of respect, this case seems no different from an otherwise comparable case in which X does not falsely believe that Y is a particularly virtuous person and thus makes a false judgment about Y and Z's comparative moral worth. This suggests that the falsity of the judgment about the comparative worth of Y and Z is in itself irrelevant from the point of view of respect. What matters is the fact that the comparative judgment is held irrationally.

However, as noted above sheer irrationality seems intrinsically morally irrelevant and, thus, cannot explain the moral wrongness of discrimination. However, it may be different with irrationality tied to moral flaws.[53] Assume X irrationally believes that Z is more worthy than Y because X has an unjustified hatred of Y because of Y's membership in a certain socially salient group. It is no longer immediately clear what we should say about the moral credentials of our two cases of differential treatment. One possibility is to say that the fact that a given differential treatment is based on irrational beliefs about differential moral worth, where the irrationality involved reflects a moral flaw, renders it intrinsically morally wrong. Another possibility is to say that the fact that a given differential treatment is so based is irrelevant to the moral evaluation of the *activity* of differential treatment but is highly relevant to the moral (and epistemic) evaluation of the *character* of the agent involved.

[53] Irrationality may reflect moral virtues. For example, if you continue believing that your friend is telling you the truth, despite evidence to the contrary, you may be a better friend than if your view of the matter simply tracked the evidence.

The second of these possibilities allows us to say, in some cases, that an agent did what he ought to do, morally speaking, despite his vicious, irrational false beliefs about comparative moral worth. Suppose a sexist male professor denies a female colleague sabbatical on account of his irrational, false beliefs about the differential moral worth of men and women. Suppose, moreover, that unbeknownst to the professor, his colleague is a thoroughly undeserving person with a long history of misconduct and ought, morally speaking, to be denied the sabbatical. Here we might like to say that the professor's action was morally right; it just happened to be based on an irrational and false belief about the differential moral status of men and women which, in the circumstances, led to a correct judgment about the moral status of the colleague.[54] Of course, the professor's motivation displays a flaw of character.[55]

In response to this example it might be said that, on a central construal, moral worth is a very Kantian notion and that the example in the previous paragraph does not involve moral worth in this important sense. Roughly, the idea is that in virtue of being a rational, self-conscious, self-legislating individual one should be treated as an end in oneself and not merely as a means.[56] Non-personal animals like mice are not rational beings and, thus, do not have moral worth. Accordingly, they can be treated very differently from persons. Rational beings may vary in terms of how rational they are, but variations do not affect their moral worth. Nor does what they do affect their level of moral worth.[57]

Judging someone to have moral worth is one species of assessment of someone's moral status. Alternatively, one might assess people's Aristotelian virtuousness or deem persons morally more or less deserving. One has moral worth regardless of what one does, but this is true of neither virtue nor moral deservingness. Suppose a discriminator rejects the notion of moral worth, but assesses people in terms of moral desert and discriminates against Swedes on the basis of the clearly false belief that Swedes are less morally deserving than Danes are. In Alexander's account, unlike discrimination based on a clearly false judgment about moral worth, this does not amount to an intrinsically wrongful type of discrimination. Yet, this difference needs to be explained. Given that

---

[54] Hence, here we have a case of a correct judgment of differential worth that is based on an incorrect judgment of differential worth. Alexander's account implies that a bias is wrong, not just if it is immediately grounded in a false judgment of differential worth, but also if any link in the chain of grounding beliefs of a bias involves an incorrect judgment of differential moral worth. See Alexander (1992), "What Makes?" pp. 217–18.

[55] Admittedly, this move is controversial and, in a different context, I myself have argued that friends of the doctrine of double effect do not confuse assessment of wrongness of actions and assessment of the agent's character or deliberations leading up to the action, see Lippert-Rasmussen (2010), "Scanlon."

[56] Immanuel Kant (2006 [1798]), *Metaphysics of Morals,* in P. Kleingeld, *Toward Perpetual Peace and Other Writings Politics, Peace, and History* (New Haven, CT: Yale University Press), p. 91.

[57] But see Alexander (1992), "What Makes?," p. 200.

discrimination manifesting biases based on clearly false beliefs about unequal *moral worth* renders discrimination intrinsically wrong, why is discrimination that manifests biases based on clearly false beliefs about unequal *moral status in general* not intrinsically morally wrong?

Alexander does not address this question. He mentions the case of a bias in favor of the "morally virtuous" over the "morally vicious" and suggests that this bias may be "morally required."[58] If so, discrimination manifesting such a justified bias (or better perhaps because "justified bias" seems like a contradiction in terms: a justified disposition to give more concern to the morally virtuous relative those who are not) could hardly be intrinsically wrong. But this does not speak to the question of discrimination manifesting unjustified biases rooted in clearly false assessments of the relative moral virtue of different socially salient groups.

Second, turning to the issue of propositional attitudes, suppose that Alexander's discriminator does not *believe* but *assumes* that Jews have less moral worth than non-Jews and acts on this basis. Provided that it can be disrespectful to make such an assumption, as indeed it can, it is unclear why this agent's discriminatory act may not be intrinsically wrong in the same way as that of the agent who actually holds the relevant belief about unequal moral worth. Moreover, suppose the agent believes that Jews have the same worth as non-Jews, but wishes that this will be different in the future, say, through a reduction of the cognitive powers of all Jews to a level below the threshold necessary for personhood. If Alexander is right, it is difficult to see why actions motivated by such desires are not disrespectful too, although, in a perverse way, they manifest recognition that the discriminatees *have* equal moral worth. Is it not disrespectful to want to reduce a rational being to a nonrational being? Indeed, if, as some claim, acts can be disrespectful irrespective of their being manifestations of disrespectful mental states, Alexander's account should be developed in a way that makes it cross the boundaries of mental-state accounts to allow for disrespectful, intrinsically wrongful acts of discrimination that reflect no disrespectful mental state, e.g., carelessly using a gesture that, by convention, signals disrespect.[59]

Third, let me finally address the issue of how the relevant false belief must be related to discrimination for it to make the latter intrinsically morally wrong. My critique here will be different from the two previous ones. Their general thrust was that Alexander's account should be broadened to include other disrespectful contents and propositional attitudes. Now my critique will question the focus on mental states in the first place.

---

[58] Alexander (1992), "What Makes?," p. 159.

[59] See Glasgow (2009), "Racism is Disrespect," *Ethics* 120.1, 64–93, pp. 83–84.

One can hold a clearly false belief about unequal moral status and yet disregard this belief in one's deliberations about how to act. A neo-Nazi bureaucrat may treat a Jewish citizen no differently from an impartial, non-prejudiced bureaucrat. Indeed she may treat her better than non-Jewish citizens for fear of complaints of discrimination and consequent dismissal. Yet, in both cases the agent acts with no less disrespect than in the case where she actually discriminates against those whom she believes to have lesser moral worth.[60] Hence, if what renders discriminatory acts intrinsically morally wrong is the disrespect these involve, there could be nondiscriminatory acts, or even acts of discrimination *in favor of*, that are just as intrinsically wrong as discriminatory acts.

Moreover, there could be discriminatory acts that are not based on the relevant belief about unequal moral worth. Suppose the neo-Nazi bureaucrat manages to set aside her belief about unequal status when dealing with citizens, but that her conduct is influenced by an unconscious bias triggered by some property other than being Jewish where, however, this property is one many Jews have and which is causally unconnected to the bureaucrat's belief about unequal moral worth. The question is: is the bureaucrat's discriminatory act intrinsically wrong in the way it would have been had it been based on her belief about unequal moral worth? It is hard to believe that the mere fact that an agent believes that a discriminatee has a lower moral status should render her discriminatory action more wrongful compared to that of an agent who performs exactly the same action and for the same reason, but has no inert belief of the relevant sort. In support of this claim, note that to hold a certain belief it is not required that it is present to the agent's mind at the time of acting. Everything that the former agent thought of, responded to, perceived, and felt at the moment of performing the discriminatory action may be identical to what the latter agent thought of, responded to, perceived, and felt. It is just that if you had asked the former agent if she thought that Jews have a lesser moral worth and she had answered honestly, she would have responded affirmatively. This comparison suggests that if a judgment of unequal worth has no causal influence on the agent's action, it does not render it wrong.

So suppose this is correct. We should now compare two agents, both of whom have the relevant beliefs about unequal moral status—we can even imagine that the belief comes to the mind of both agents at the time of acting; one acts partly on the basis of this belief whereas the other one, like the neo-Nazi bureaucrat a few paragraphs ago, does not. Apart from that their actions are identical in all other respects. Clearly, we would want to hold both agents accountable for their disrespectful beliefs of unequal worth and for whatever

[60] Generally, acting with disrespect depends crucially on which attitude is manifested through or lies behind our act, not on whether the act happens to have beneficial consequences for the patient or whether the agent intends the action to have beneficial consequences for the patient.

flaws in their deliberation these beliefs result in, but it is unclear that we might ever want to say that one of them acted wrongly, whereas the other did not.[61] If so, beliefs about unequal moral worth are irrelevant to permissibility (which is not to say that it is not relevant to other types of assessment or other objects of assessment, e.g., the discriminatee's moral character or the quality of the agent's deliberations).

As further support for this conclusion, consider a Nietzschean case of discrimination against the strong by the weak and resentful. Suppose a Norwegian believes idiosyncratically that Danes merit greater moral concern than Norwegians and resents them for it, which manifests itself in her treating Danes disadvantageously compared to Norwegians. Compare this to a case of a Norwegian treating Danes in exactly the same adverse way, because she believes Danes have a lower moral status. To prevent perceived victimhood from influencing our intuitions here, suppose both Norwegian discriminators live equally bad lives on account of their ethnic prejudice. To eliminate any pollution of our intuitions about how the discriminating act might carry different messages depending on the underlying thought, suppose that the official anonymously makes a decision about whether to admit the relevant Danish applicant to a university and identical letters of rejection are sent out in both cases. If Alexander's account is correct, these two cases should differ in terms of wrongness. But it seems they do not. More generally, arguments analogous to the one offered here can be offered against a wider range of mental-state accounts.

8. Conclusion

My discussion of Alexander's position can be summed up in two conclusions. First, insofar as we accept that discrimination based on beliefs about unequal moral worth is intrinsically wrong, we should allow that discrimination based on other inequality involving mental states, e.g., clearly false beliefs about unequal moral deservingness or desires about unequal moral worth, is intrinsically wrong. Second, it is unclear, however, that the fact that discrimination is rooted in beliefs about unequal moral worth renders it intrinsically wrong. This is neither to deny that it is unjust to discriminate, nor to claim that there is nothing wrong with the character, motivation, or deliberations of an agent who discriminates on the basis of such a belief. Such claims are open to friends of the objective-meaning (chapter 5) and the harm-based account (chapter 6) and, thus, are not a matter of contention between those who accept and those who reject mental state based accounts and, thus, do not support mental-state-based accounts.

[61] Cf. Scanlon (2008), *Moral Dimensions*, p. 22.

As I have noted, Alexander's account is not the only possible equality-focused mental state account of the wrongness of discrimination, let alone the only possible mental-state-based account. Yet, the problems presented above raise serious doubts about the viability of the mental-state-based account of the wrongness of discrimination.

Note finally that perceived lack of respect often affects the degree of harm involved in differential treatment, the desert level of individuals, and thus the moral qualities of discriminatory acts.[62] Hence, friends of the harm-based account, and in particular the desert prioritarian version of the harm-based account, may agree that acts that are perceived to be disrespectful tend to be harmful and, thus, wrong. Presumably, a similar claim is open to friends of objective meaning accounts, who might conjecture that such acts tend to be objectively demeaning and, thus, wrongful.

[62] Accordingly, harm-based accounts are compatible with cases where "doing X" has a different moral status than "doing X with the intention of Y," "doing X out of motive Z," and "doing X thereby expressing an attitude of W." The latter cases may have different consequences than the former. See Arneson's discussion of thin versus thick act descriptions in Arneson (2006), "What Is Wrongful Discrimination?" pp. 780–783.

{ 5 }

Objective-Meaning Accounts

1. Introduction

Objective-meaning accounts claim that the wrongness of discrimination results from the objective meaning carried by it. As such they differ from subjective-meaning accounts, i.e., accounts according to which the wrongness of discriminatory acts derives from what the discriminator means by her discriminatory act, and from harm-based accounts, which I will focus on in chapter 6. Subjective-meaning accounts can be seen as a subspecies of mental-state accounts in that any account of the subjective meaning of an action must ascribe subjective meaning to acts on the basis of a subset of the agent's mental states, e.g., the agent's communicative intentions.[1] Not all mental states, however, are relevant to the subjective meaning of an act, e.g., explaining everything very carefully and twice to elderly people might be disrespectful even if the agent is motivated solely by a desire to make sure that the receivers of her message understand it, which presumably is irrelevant to its disrespectfulness. To see how the subjective meaning and the objective meaning of an act can come apart, consider a case in which an agent acts within a cultural environment with which he is unfamiliar, e.g., in an antisexist society a man holds the door for a women, the man coming from a culture where this signals respect.

Objective-meaning accounts can concede that acts of discrimination can be wrong for reasons other than that they involve an objectionable objective meaning, e.g., that they harm the discriminatees, or that the act expresses a disrespectful attitude. However, it is crucial to objective-meaning accounts that they hold that acts of discrimination can be wrong, even when they neither

[1] Admittedly, there is conceptual space for subjective accounts that locate the wrongness of acts of discrimination in the subjective meaning ascribed to it by the discriminatee (or for that matter, some third party). However, such subjective meanings will similarly depend on the mental states of the discriminatee.

harm anyone, nor involve any objectionable mental states.[2] Friends of mental-state accounts and harm-based accounts are likely to deny that such acts are wrong. Supporting this claim, they might say that acts having objectionable, objective meanings almost always reflect objectionable mental states on the part of the agent and will result in harm to the discriminatee. Once we set the wrong-making features aside, they will continue, it is implausible that an act having an objectionable, objective meaning is wrong simply on that account. So, for instance, imagine someone who addresses male professors by formal title and female professors informally, where this person's differential treatment is caused entirely by the fact that this person's mother was very informal in personal encounters, whereas his or her father was much more formal. Because everyone knows this, no one is offended or otherwise harmed by this. The claim that this person acts wrongly now looks weak. Friends of mental-state accounts may say that normally people will know, or ought to know, when their acts have an objectionable meaning and, accordingly, that when they perform such acts despite this fact, this will normally reflect an objectionable mental state of theirs, e.g., negligence or disregard. Similarly, friends of the harm-based account will say that normally, when an act has an objectively objectionable meaning, it is harmful in certain ways and the agent will typically know this. Accordingly, much of the initial attraction of objective meaning accounts may really stem from features other than those they identify as wrong-making.

Just as the idea that mental states affect wrongfulness is not limited in its scope to acts of discrimination, the idea that the objective meaning of an act determines wrongfulness is not so restricted. Jean Hampton defends the view that to wrong someone in general is to treat her in a way "that is objectively demeaning."[3]

Objective-meaning accounts form a large family whose members differ in terms of which objectionable meaning the wrongfulness of discrimination derives from. One crucial distinction here is whether the relevant objectionable objective meaning is comparative across (groups of) individuals or not, e.g., whether it involves the message that a certain group of individuals is prone to crime or the message that one group is more prone to crime than members of other groups. Of course, often when the objective meaning of discriminatory acts is an objectionable noncomparative one, this implicitly involves the message that members of the relevant group are worse in the relevant dimension than members of other groups, i.e., it is not just that they are lazy, they are lazier than members of other groups. Still, an objective meaning can be objectionable

[2] In principle one could imagine an objective meaning account according to which acts can be wrong only when they involve harm or objectionable mental states, but where the fact that an act has an objectionable meaning "boosts" the wrongness of the discriminatory act (see Section 5 in this chapter).

[3] Jean Hampton (1988), "Forgiveness, Resentment and Hatred," in J. G. Murphy and J. Hampton (eds.), *Forgiveness and Mercy* (New York: Cambridge University Press), p. 52.

even without any implicit comparison. In this chapter I will focus on two accounts—one by Deborah Hellman (sections 2 to 3) and another one by Thomas Scanlon (sections 4 to 7)—according to which the wrongfulness is tied to the relevant discriminatory action's objective meaning that some kind of inequality between the discriminatee and others obtains, e.g., that latter has a higher status or is more socially acceptable than the former.

Unlike mental-states-based accounts, it is not as if the impermissibility of certain kinds of statistical discrimination, e.g., indirect discrimination, cannot be accounted for in terms of objectionable objective meanings. Given a suitable context, in principle any discriminatory act can have an objectionable objective meaning. So in this sense, the objective-meaning account is broader than the mental-state-based account. However, it is also clear that, as a matter of fact, many discriminatory acts or outcomes are not perceived to involve any objectionable objective meaning. To the extent that this is the case, it is not clear what the objective-meaning account implies about such cases. Suppose, for instance, that employers pay men more than women, but that this is not known to anyone. If it were known, people would rightly see it as sending out an objectionable message about the unequal worth of men and women. But since this state of affairs is unknown, no one ascribes any objectionable meaning to this wage differential. At this point objective-meaning accounts divide into two groups: those that require that the objectionable meaning of a discriminatory act is known or believed to be conveyed by someone for it to render discrimination impermissible and those that impose no such condition, e.g., it is enough that such a meaning could reasonably be ascribed to its bearer. Either position comes with a cost. The former position implies that by making sure that no one perceives the objective meaning of a discriminatory act, it becomes permissible (at least as far as the objective meaning account goes). The latter position implies that an act may be impermissible in virtue of something—a free-floating, uninterpreted cultural meaning—that does not really affect people's lives.

2. Hellman's Account: Demeaning Others

Probably the most detailed objective-meaning account of the wrongness of discrimination to date has been offered by Deborah Hellman. In this section, I will present her account and then in section 3 offer some challenges to it. According to Hellman, the mental state of the discriminator is irrelevant to the distinctive wrongfulness of discrimination, as is the kind of harm to the discriminatee that it may involve.[4] On her account, wrongful discrimination is differential

[4] Recall my reservation regarding claims about the distinctive wrongness of discrimination expressed in chapter 4.1.

treatment that is demeaning to the discriminatee, where an act or omission demeans someone if, and only if, the following two conditions are satisfied. First, the relevant individual is treated as "not fully human" or "not of equal moral worth," debased or degraded, and she is not merely insulted but also "put down," diminished and denigrated.[5] To demean is to "treat another as a lesser."[6] A demeaning act "expresses that the other is less worthy of concern or respect."[7] Second, the agent who does the demeaning must be in a position of power or have a superior status.[8] A boss might demean an employee, but an employee cannot demean her boss, unless she has some employment-unrelated power or status.

Whether these two conditions are satisfied depends on the context: "Whether classification demeans depends on the social or conventional meaning of drawing a particular distinction in a particular context."[9] Having a sign on restroom doors saying "Men Only" is not demeaning, whereas a sign saying "White Men Only" would be. This need not be so in all contexts, but in a cultural context where the memory of Jim Crow lingers on, the latter is demeaning even if there is another equally suitable facility with a "Black Men Only" sign on the door. Intentions of the discriminator and the perceptions of the discriminatee are not decisive, if relevant at all, for whether discriminatory acts demean. Similarly, whether the second condition is satisfied depends on "history and the current social status" of the relevant groups of people.[10]

In Hellman's view, discrimination that satisfies these two conditions is demeaning and for that reason wrong, because it clashes with the bedrock principle of equal moral worth. One implication is that it does not even appear to be part of the definition of discrimination in that it involves disadvantageous treatment besides carrying a demeaning message. One might treat a discriminatee advantageously—e.g., use more resources in a way that signals one's belief that they have more difficulties learning than others—all things considered and in all respects other than the message that one's treatment of them carries.[11]

Another implication of Hellman's view is this: because the distinctive wrongness of discrimination consists in its objective meaning, in principle, whenever we are faced with wrongful discrimination we may either eliminate

---

[5] Deborah Hellman (2008), *When Is Discrimination Wrong*? (Cambridge, MA: Harvard University Press), pp. 35, 29.

[6] Hellman (2008), *When Is Discrimination*? p. 29.

[7] Hellman (2008), *When Is Discrimination*? p. 35.

[8] Hellman (2008), *When Is Discrimination*? pp. 35, 57.

[9] Hellman (2008), *When Is Discrimination*? pp. 29, 7.

[10] Hellman (2008), *When Is Discrimination*? p. 14.

[11] Hellman (2008), *When Is Discrimination*? pp. 27–28. When imprisoned on Robin Island, along with other black prisoners, Nelson Mandela was required to wear shorts, while white and "colored" prisoners could wear long trousers. This differential treatment was demeaning, because infantilizing, to black prisoners even if wearing shorts on Robin Island probably was more pleasant.

the differential treatment or change its objective meaning. Changing social meanings may be unfeasible, e.g., because social meaning is determined by history and even if historical consciousness can change, history cannot. Suppose, however, that we change the social meanings of a society blighted by racist and sexist discrimination such that the relevant differential treatment persist unchanged to the highest degree compatible with rendering it nondemeaning. In Hellman's account, such a society would have eliminated wrongful discrimination.

A third implication of Hellman's account is that demeaning acts of discrimination against people who no longer exist—e.g., members of a certain extinct religious group—may be wrongful in the distinctive way that demeaning acts of discrimination against people who do exist are. A shopkeeper who falsely believes that there are still people around who belong to a persecuted and suppressed religious minority and puts up a sign saying that these people are not welcome in her shop demeans these (now deceased) people. Arguably, no one is harmed (or perhaps even offended), but since (neither) harm (nor offense) plays any role in Hellman's account, she cannot deny that this case involves wrongful discrimination on this ground. Also, it is surely possible to demean people after they have ceased to exist, e.g., Ward Churchill demeaned the victims of 9/11 working in the twin towers by referring to them as "little Eichmanns."[12]

Hellman might, of course, avoid this implication simply by adding to her account a requirement to the effect that acts of demeaning are wrongful only when there are presently existing people who are demeaned by the relevant sort of differential treatment. However, she would also need to explain an additional requirement to this effect to avoid ad hocness and in so doing she could hardly appeal to the factor that would seem to most readily be at help here: namely, the fact that no one will be harmed by being demeaned. After all, she is trying to offer an account of the wrongness of discrimination, which is independent of considerations about harm.

Before proceeding to my criticisms of Hellman's account, I want to make a friendly amendment. As noted above, she offers many, nonequivalent phrasings of the first condition. For instance, to debase or degrade someone is different from not treating them as fully human. To say of someone that he is not very intelligent may degrade him, but it is not to suggest that he is not fully human. Also, by treating someone as not fully human a deep ecologist may not debase or degrade. This is possible because in the view of someone who subscribes to deep ecology there is nothing debased or degraded about not being (treated as) fully human.

Some of Hellman's phrasings imply that discrimination is not tied to unequal treatment. For instance, she says that demeaning discrimination involves

[12] For Ward Churchill's essay, "'Some People Push Back': On the Justice of Roosting Chicken", see http://www.kersplebedeb.com/mystuff/s11/churchill.html (accessed June 5, 2013).

the relevant individual being treated as "not fully human." While this may be demeaning, it need not be discriminatory as is shown by the possibility that one treats all human beings—including oneself—as if they are not fully human. Yet, she writes that "demeaning is an inherently comparative concept."[13] Also, it is unclear what it means to treat someone as not "fully human." Does it mean not treating them as if they belong to the human species? Or does it mean treating them as if they are not persons? Probably only the most extreme forms of discrimination involve treating discriminatees as if either is the case. Accordingly, I shall understand "treating someone as not fully human" in Hellman's sense—and this is the friendly amendment—as treating someone as if she is not of equal moral worth, i.e., X treats Y as not fully human because, given how X treats Z, treating Y in the way X treats Y would be permissible only if Y had lesser equal moral worth than Z. Pedantically, I would add that by the complaint about treating someone as if she is "not of equal [my emphasis] worth," strictly speaking, Hellman means the complaint of treating someone as if she is of "lesser worth than other people" (and not "higher"). So construed, and if we set aside the hierarchy condition, we have an account that is related to Larry Alexander's in its focus, i.e., it focuses on attributions of a lesser moral status to some people. The difference lies in the fact that Alexander concentrates on the beliefs of the discriminator, whereas Hellman focuses on the objective meaning of the differential treatment. As noted in the introduction, these two can come apart. An agent can perform an act without any demeaning objective meaning despite her objectionable mental state, and an agent can perform an act with an objectionable objective meaning despite her innocent underlying mental state.

3. Some Challenges to Hellman's Account

Is the objective-meaning account with this friendly amendment satisfactory? No doubt, it is relevant to the wrongfulness of discrimination that it is objectively demeaning. This often results in stigmatic harms, and for that reason friends of a harm-based account might accept many of the insights in Hellman's account.[14] However, the issue is whether discriminatory acts are wrongful *simply* because they are demeaning, i.e., independently of any stigmatic harm. I am skeptical of this view and offer five challenges below.

The first is a challenge to the rationale Hellman provides for her account. As noted, Hellman offers various nonequivalent phrasings of the first condition of demeaning discrimination. This reflects a dilemma for Hellman's account. She wants to tie the wrongfulness of demeaning someone to the "bedrock principle

[13] Hellman (2008), *When Is Discrimination*? p. 33.

[14] Paul Brest (1976), "In Defense of the Antidiscrimination Principle," *Harvard Law Review* 90.1, 1–54.

of equal moral worth of persons."[15] The advantage is that if the claim that discrimination that demeans is wrong per se follows from a bedrock principle, presumably we can be very confident about the implied principle as well. But if, on the one hand, we interpret this principle in such a way that it has a plausible claim to be a bedrock principle, not many wrongful demeaning acts of discrimination will violate it. If, on the other hand, we interpret it broadly such that any demeaning, discriminatory wrongful acts will violate it, it becomes a very implausible candidate for a bedrock principle.

To illustrate this dilemma, consider the notion of demeaning acts. Acts can be more or less demeaning, when this term is used in its ordinary sense. For instance, there is a difference between simply ignoring what adolescents say in discussions about politics at family gatherings and paying slightly less attention to it than to what adults say. Both dispositions are demeaning to adolescents and arguably wrongful, but clearly the latter disposition is less demeaning and it would seem shrill to complain that in paying slightly less (or even no) attention to the views of teenagers, you fail to treat them as having the same worth as other persons. In responding to such a complaint, you might reasonably say that you hold no such view and that many of the other features normally associated with the principle of equal moral worth of persons are features that you think apply to them. So, for instance, you would not consider it less wrong to kill an adolescent than an adult and, in general, you are not less inclined to give deliberative weight to the interests of adolescents than to the interests of adults. While this might show that you do not violate the principle of equal moral worth of persons, it does not get you off the hook with regard to the charge of demeaning adolescents. Hence, you can demean someone in the ordinary sense of the term without denying that they possess equal moral worth. It follows that if discrimination is differential treatment that is wrong because it demeans in this sense, wrongful discrimination need not involve denying the equal moral worth of discriminatees. But this flatly contradicts Hellman's account.

To avoid this challenge we might expand our notion of equal moral worth such that paying slightly less attention to the expressed political views of teenagers does involve a failure to respect the principle of equal moral worth. This habit is morally wrong, but I doubt that any moral principle, which implies this, can have the status of a "bedrock principle."

Further support for this skeptical claim about the status of Hellman's basic principle comes from the fact that there are many ways, other than Hellman's, in which one can cash out the equal moral worth of persons. For instance, it might be held to imply that everyone's interests count equally or that everyone should end up equally well off unless they are responsible for ending up worse off. If, in a way that is counterfactual and unrealistical but conceptually

[15] Hellman (2008), *When Is Discrimination*? p. 30.

possible, we were to imagine a scenario in which no one is demeaned but many are much worse off than others in nondemeaning ways through no responsibility of their own, it is hard to see that this conforms better to the principle of equal moral worth of persons than one in which many are objectively demeaned, but in ways they simply do not care about and, let us suppose, are not harmful to them, and no one is worse off than others through no responsibility of their own.[16] And if Hellman makes this claim, the equal moral worth of a person cannot be the relatively uncontroversial "bedrock" moral principle that it is often thought to be.

One might respond that Hellman leaves open the possibility that there are other ways of offending against "the equal moral worth of persons" than through demeaning others.[17] Hence, she might agree that the former scenario clashes with the principle of equal moral worth. She might even concede that it does so to a higher degree than the latter scenario. But if so, she cannot claim to identify wrongful discrimination on the basis of which kinds of differential treatment are demeaning, because she would then have conceded that differential treatment might clash with the principle of equal moral worth, even if it is not demeaning.[18]

To sum up: Hellman's account faces a dilemma. Either we understand the principle of equal moral worth in a rather narrow way that makes it plausible to assert it as a basic moral principle, or we understand it in a very broad way. If we do the former, to the extent possible, it becomes doubtful that it justifies a prohibition on nonharmful acts that, conventionally speaking, are demeaning. Many acts that are demeaning in the ordinary sense do not involve denying the equal moral worth of the persons one demeans. If we do the latter, such that demeaning acts need not involve denying the equal moral worth of persons, the principle is very controversial and any account of the wrongfulness of discrimination that builds upon it will be correspondingly controversial.

Second, independently of the problem of whether the equal moral worth principle explains the prohibition on demeaning discrimination, there is the problem of epiphenomenality, i.e., while any wrongful act of discrimination is demeaning, some underlying factor explains both the fact that they are wrongful and that they are demeaning—the fact that discriminatory acts are demeaning does no explanatory work in relation to the wrongfulness of these acts.[19] In effect, Hellman's account says that it is impermissible to discriminate because

---

[16] Cf. Hellman (2008), *When Is Discrimination*? p. 49; Mathew Adler (2000), "Expressive Theories of Law: A Skeptical Overview," *University of Pennsylvania Law Review* 148, 1363–1501, pp. 1462–1493; Frej Klem Thomsen (unpublished manuscript on file with author), "Discrimination and Disrespect," http://www.google.dk/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&ved=0CDEQFjAA&url=http%3A%2F%2Fwww.ps.au.dk%2Ffileadmin%2Fsite_files%2Ffiler_statskundskab%2Fsubsites%2Fnnpe%2Fdokumenter%2FThomsen_-_Disrespect__intuitions_and_discrimination.pdf&ei=pxjtUarIDci1PM2jgFg&usg=AFQjCNEJDTpNzOxGe0VVnaH-FpuLbGnYZw&bvm=bv.49478099,d.ZWU (accessed July 22, 2013).

[17] Hellman (2008), *When Is Discrimination*? p. 31.

[18] Hellman (2008), *When Is Discrimination*? p. 29.

[19] Sophia Moreau (2010), "What Is Discrimination?" *Philosophy & Public Affairs* 38.2, 143–179, pp. 177–178.

discrimination involves demeaning the discriminatee where this amounts to treating the discriminatee in a way that would be permissible only if she had a lower moral worth than other people, which she has not. But then, presumably, there are facts about her in virtue of which the discriminatory action is impermissible and these facts are those in virtue of which the discriminatory action is impermissible. So, for instance, when South Africa under apartheid treated blacks impermissibly by denying them access to higher education, it treated them in a way that would be permissible only if they had a lower moral status, i.e., they were not persons. Yet, this is not what makes the discriminatory policy wrong. What makes it wrong is, say, that they are persons like other South Africans and that if the latter have access to higher education then so should the former.

Third, if discriminatory actions are wrong in virtue of being demeaning, the question arises whether all wrongful forms of discrimination carry a demeaning message. If not, at best Hellman's account provides a partial account of the wrongness of discrimination. In fact, many forms of discrimination do not per se seem to involve any objective meaning of lesser moral worth. Forms of indirect discrimination involving job-relevant language tests used for nonsuspect reasons, for instance, might disadvantage immigrants in unjust, but nondemeaning, ways. Here Hellman might concede that this involves wrongful differential treatment because "(s)ometimes institutions or actors have special obligations that derive from their mandates or their roles that require them to employ particular criteria in differentiating among people. Failing to comply with these norms is wrong, but is not the wrong of wrongful discrimination."[20] In Hellman's view, the wrong of wrongful discrimination is that it demeans the discriminatees.

In part, I agree with Hellman's view here, but it involves the concession on her part that different cases of wrongful discrimination are wrongful for different reasons, and some cases of wrongful discrimination are not wrongful because they demean. Hence, her account should not be seen as a general account of the wrongfulness of wrongful discrimination, but as a partial account that explains the wrongfulness of an important subclass of cases of wrongful discrimination. The aim to provide such an account is perfectly legitimate and perhaps this is all Hellman intends her account to offer, but it is less ambitious than that of explaining the wrongfulness of discrimination in general. More importantly for our purposes, it implies that not all forms of wrongful discrimination clash with the equal moral worth of persons.

Fourth, it is problematic to tie discrimination to a failure to "treat those affected as persons of equal moral worth," if we want to allow for wrongful discrimination of nonpersons.[21] Assuming that small children are not yet persons, treating boys and girls differently cannot be accounted for in terms of

[20] Hellman (2008), *When Is Discrimination?* p. 137.
[21] Hellman (2008), *When Is Discrimination?* p. 7.

a failure to respect equal worth deriving from personhood. We might, of course, say that sex discrimination of babies, unlike sex discrimination of persons, is wrong only for indirect reasons, i.e., it amounts to a failure to respect the "inherent dignity and worth" of, say, female persons.[22] But, first, sex discrimination of babies seems wrong independently of a moral concern for existing persons. Second, it is unclear that discrimination against female babies need involve disrespect for female persons. It could be performed by persons who think that women have greater moral worth than men have.

Finally, Hellman's hierarchy requirement is problematic. Take racist hate crimes, where the victim belongs to a group of people, which historically has had and presently has a higher status than the group to which the perpetrators belong. One would suppose that if anyone is ever demeaned, a person who is beaten up or killed simply because her assailants believe members of her race have an inferior moral status is demeaned. Hellman seems compelled to say either that this is not so, because it all depends on which race the victim has, or that what matter here is the local hierarchy, which was established, when the assailants overpowered their victim.[23] The former option is problematic on the assumption that it is not morally less wrong to beat up or kill whites simply because of their race than to kill blacks when all other things are equal except for the general facts about history and status. Suppose you hear about a racially motivated crime in some faraway country. Would you inquire whether the victim belonged to a dominant group before you formed a judgment about the wrongfulness of the act?[24] The latter option is problematic but it implies that to demean is something that automatically follows from successfully performing certain discriminatory acts. In any case, it seems that this is not the option Hellman prefers, since she often stresses whether a group is subordinated "in our culture," which suggests that she disregards the local context.[25] The basic problem here is that, offhand, there is no requirement attached to showing "lack of respect for another's equal moral worth" that the one who expresses this lack of respect is in a position of power or higher status. When Hellman writes that "[t]o demean, rather than merely to insult, requires a certain degree of power," she is employing a technical distinction between demeaning and insulting.[26] Moreover, it is unclear that the distinction is morally relevant.

---

[22] Hellman (2008), *When Is Discrimination?* p. 6.

[23] Cf. Lawrence Blum (2002), "*I'm Not a Racist, But . . .* " (Ithaca, NY: Cornell University Press), pp. 33–52.

[24] This rhetorical question has no bearing on whether lethal hate crimes are worse than killings otherwise motivated, e.g., because of the objectionable motivation involved or because of the additional harms hate crimes involve compared to otherwise comparable crimes. Note, specifically, that you may be a member of an oppressed minority and kill a member of the oppressing majority out of hatred, not caring whether this person is in any way responsible for the oppression in question or the like.

[25] Hellman (2008), *When Is Discrimination?* p. 37.

[26] Hellman (2008), *When Is Discrimination?* p. 36.

To sum up: Hellman's objective-meaning account of the wrongfulness of discrimination may not be the best one available, e.g., because its conditions of hierarchy and history may be unjustifiable; that the mere fact that an act of discrimination is demeaning is not what explains its wrongness; that the wrongness of demeaning discrimination cannot be based on the bedrock principle of equal worth; and that not all wrongful discriminatory acts are demeaning.[27] I agree that in many contexts discrimination *is* demeaning and for that reason wrongful, because demeaning is often harmful. Indeed, systematically demeaning certain groups will often involve great harm to very many people. But these points are not matters of contention between harm-based accounts and Hellman's and, accordingly, not points that support her account.[28]

4. Scanlon on Racial Discrimination and the Meaning of Actions

In a number of recent works Thomas Scanlon has argued that while intentions are relevant to the assessment of an agent's deliberations, at a fundamental level, they are irrelevant to the permissibility of actions. Yet, Scanlon is not entirely happy with the implications of this view when it comes to racial discrimination on the job market: "there does seem to be something particularly objectionable about discrimination on racial grounds" and by "particularly objectionable" he has in mind permissibility (also), not (merely) the moral qualities of the discriminating agent's deliberative processes.[29] This present section sets out Scanlon's account of the wrongness of racial discrimination, while section 5 exposes an important ambiguity in the account pertaining to what role objectionable meanings play in relation to the harmful effects of discrimination. Based on this exposition, sections 6 and 7 identify a number of problematic features of Scanlon's account.

To see Scanlon's worry regarding his general stance regarding the irrelevance of intentions to permissibility when it is applied to racial discrimination, compare two nearly identical cases of hiring. In both cases an employer needs to fill a vacant position and in both cases there are several applicants who are no less qualified than all other applicants. In the first case the employer picks someone

---

[27] Of course, Hellman agrees that not all wrong discriminatory acts are demeaning. I am claiming that some wrongful acts of discrimination that we want our account of the wrongness of discrimination to explain are not explained by her account (and—irrelevant to the present point but worth mentioning for the sake of avoiding misunderstanding nevertheless—that some acts of discrimination which on her account are wrong in virtue of being demeaning are either not wrong, or wrong for a different reason).

[28] Cf. Matthew Adler (2000) "Expressive Theories of Law: A Skeptical Overview", *University of Pennsylvania Law Review* 148.5, 1363–1501, pp. 1375, 1377.

[29] Scanlon (2008), *Moral Dimensions: Permissibility, Meaning, and Blame* (Cambridge, MA: Belknap Press of Harvard University Press), p. 72.

among this group for nondiscriminatory reasons—no other applicant is better qualified, someone must be selected and the one selected has a charming personality; in the second case the employer picks the identical counterpart applicant but for discriminatory reasons, e.g., that the other best qualified applicants have a disfavored racial, ethnic, gender, or religious identity or a disfavored sexual orientation (henceforth I will use "a minority applicant" as a placeholder for an applicant who has any of these identities or some other identity that in a given context tends to elicit discrimination), and the employer considers such people "inferior, and not to be associated with."[30] While Scanlon thinks the latter employer may well act impermissibly even if the former employer does not, he rejects the view that this is to be accounted for simply in terms of their different intentions. Hence, he must provide an alternative account of the moral distinction between the two cases.

One suggestion that might seem obvious is that the discriminatory employer violates a requirement of fairness when he hires on the basis of racial prejudice because fairness requires that all the best qualified applicants get an equal chance of being hired—so ideally the employer ought to decide on the basis of a lottery—and in refusing to hire minority applicants he in effect deprives these people of a fair chance of being hired.

On reflection, however, it is not clear that an appeal to fairness offers the right kind of explanation here. Suppose that almost all best-qualified applicants are nonminority applicants and the employer decides to select among them on the basis of a fair lottery, because he knows that the risk—as he sees it—of hiring a racial minority applicant is minimal and because by this hiring procedure he will avoid any legal suits on grounds of discrimination. This employer acts with a discriminatory intention and, arguably, this means that the employer acts impermissibly if there is a difference in permissibility between the initial cases of the two employers.[31] But if that is the case, we cannot explain the differential permissibility on the basis of a requirement to treat all best qualified applicants fairly, since both employers do so by employing a lottery. Moreover, Scanlon suggests—and many will agree—that the violations of a requirement of fairness involved in racial discrimination seem particularly objectionable and, obviously, cannot be accounted for simply by an appeal to the fairness requirement itself.[32]

So instead of appealing directly to differences in intentions or fairness, Scanlon proposes that the specific objectionableness of racial discrimination can be accounted for by the fact that intentions may bear on the permissibility

[30] Scanlon (2008), *Moral Dimensions*, p. 73.

[31] But see Alec Walen (2006), "The Doctrine of Illicit Intentions," *Philosophy & Public Affairs* 34.1, 39–67.

[32] Scanlon (2008), *Moral Dimensions*, p. 72.

of actions albeit indirectly. First, intentions have what he calls predictive significance. We can, for instance, expect an employer with discriminatory intentions to violate the meritocratic injunction to hire the best qualified much more often than an employer without such intentions, all other things being equal. Accordingly, the intentions of the discriminatory employer tell us something about what "effects of the agent's action on the world around her" one can reasonably expect.[33] While the predictive significance of intentions is no doubt important, it is unlikely that this explains the asymmetric moral evaluation of the two employers. For instance, if the employer harboring discriminatory intentions is so bad at executing his intentions that, invariably, he will end up employing the best qualified applicant irrespective of race, we will not for that reason think that there is nothing particularly objectionable about this employer's discriminatory decisions.

Second, in certain cases intentions are indirectly relevant because they affect the meaning of an action, and the meaning of an action, i.e., "the significance of this action for the agent and others . . . ," affects the permissibility of it. It does so either by strengthening "the agent's case for being permitted to so act" or by boosting "the objections that others have to such an action."[34] To see what Scanlon has in mind by the meaning of an action, consider the following case: "My call to a sick relative has a different meaning for the recipient (and hence also for me), depending on the reason for which I call. It is one thing if I call because I am genuinely concerned about his welfare, another if I am indifferent to his welfare but want to please my wealthy grandfather by appearing to be concerned about my relatives, and yet another if I hate this relative and am calling to have the pleasure of hearing how weak he sounds."[35] More generally, "the meaning of an action for a person . . . [is] the significance that person has reason to assign to it, given the reasons for which it was performed and the person's relation to the agent."[36]

Applying this line of thought to the initial pair of employers, we might say that rejected applicants who are members of the disfavored (racial) minority have reason to assign different significance to their not being hired in the two cases. In the case of the discriminatory employer, they have reason to see their rejection as an insulting act carrying the message that, say, they are inferior. In the case of the nondiscriminatory employer who rejects them, but for some other, nondiscriminatory reason their rejection has no such meaning.[37]

---

[33] Scanlon (2008), *Moral Dimensions*, p. 13.

[34] Scanlon (2008), *Moral Dimensions*, p. 52, p. 62.

[35] Scanlon (2008), *Moral Dimensions*, p. 52.

[36] Scanlon (2008), *Moral Dimensions*, p. 54.

[37] Of course, it might involve the insulting message that their personality is less charming comparatively speaking. I return to this issue below.

Discriminatory decisions "involve a kind of insult—an expression of the view that certain people are inferior or socially unacceptable."[38]

To the extent that employers are morally required not to insult applicants by signaling that they are inferior, we can explain the distinction in permissibility between the two cases in a way that involves no direct appeal to intentions. Admittedly, intentions play an indirect role in that they determine the differential meaning of the rejections in the two cases, i.e., whether they convey an insulting message of inferiority and that in turn determines permissibility. However, because the intentions only matter through what significance recipients of actions have reason to ascribe to actions performed for the relevant reasons—as opposed to, say, what significance the recipients of the actions as a matter of fact ascribe to the action—Scanlon's account is best understood as an objective meaning account of the wrongness of discrimination.[39]

To explore the nature of Scanlon's account further consider the following challenge: normally we do not consider people with, comparatively speaking, less charming personalities inferior: it is just that we have brute—sometimes strong—desires in favor of persons with charm. But now consider an unusual employer idiosyncratically obsessed with charm. Let us stipulate that she considers charming people superior to not so charming people in the same way the racially discriminatory employer considers white people superior to racial minority people. Now, both cases would seem to involve the same kind of insult. Accordingly, if we explain why racial discrimination in hiring is particularly objectionable because of the insulting message that it tends to involve, we seem to be forced to conclude that the idiosyncratic employer obsessed with charm acts in a way, which is particularly objectionable in the same way that racial discrimination is particularly objectionable.[40] Yet, this is not so.

Responding to this challenge, Scanlon allows that a prejudice in favor of Ivy League graduates may express the insulting view that "certain people are inferior or socially unacceptable" and yet the somewhat idiosyncratic discrimination against non–Ivy League graduates would not seem to involve the same

[38] Scanlon (2008), *Moral Dimensions*, p. 72. Presumably, Scanlon has a specific kind of insult in mind—unwarranted insults, I conjecture—since, arguably, there is a sense in which a person who is turned down for a job on the ground that she is incompetent, because stupid, may find this decision insulting. It might even be that, on a reasonably meritocratic job market, such decisions are common and have serious consequences for the group of people deemed incompetent, because stupid. Yet, we evaluate this situation very differently from an otherwise comparable situation involving racial discrimination. I owe this point to Philippe van Parijs.

[39] Depending on how one cashes out the notion of what significance one has reason to ascribe to an action, the relevant account might be labeled "idealized subjectivist" instead. My point here is that even such an account incorporates some degree of objectivity such that my label—"objective meaning accounts"—is not misleading.

[40] Scanlon (2008), *Moral Dimensions*, p. 13.

kind of distinctive wrongness as racial discrimination.[41] Hence, the insult involved in racial discrimination cannot be the whole explanation. Scanlon writes:

> One thing that seems crucial to racial discrimination in particular is that the prejudicial judgments it involves are not just the idiosyncratic attitudes of a particular agent but are widely shared in the society in question and commonly expressed and acted on in ways that have serious consequences . . . The basis of the wrong of discrimination lies in the moral objection to this kind of harm. No one can be asked to accept a society that marks them out as inferior in this way and denies them its principal benefits . . . [Individual acts of discrimination] are thus wrong because of their consequences—the exclusion of some people from important opportunities—and because of their meaning—the judgment of inferiority they express and thereby help to maintain.[42]

Incidentally, the claim just quoted also suggests how Scanlon would account for the wrongness of forms of discrimination that do not involve any insulting messages of inferiority, since, provided these forms of discrimination are systematic, surely they will even in the absence of any insulting messages deny people subjected to discrimination access to important opportunities. This kind of discrimination would be wrong, although it would not be wrong in the same "particularly objectionable" way racial discrimination is. So, by way of illustration, consider the wrongness of discrimination on the basis of differential reaction qualifications—i.e., "abilities or characteristics which contribute to job effectiveness by causing or serving as the basis of the appropriate reaction in the recipients"[43]—rooted in customers being racist (see chapter 9). Suppose an employer harbors no racist sentiments, but her customers do and they will respond negatively to her hiring a racial minority applicant and sales will drop. Suppose the employer refrains from hiring a minority applicant for this reason alone. Here the rejected applicant will not have reason to think of her

---

[41] Scanlon writes that "(m)any of the likes and dislikes of our fellow citizens may be things that we just have to live with," Scanlon (2008), *Moral Dimensions*, p. 73. Racist judgments to the effect that certain groups of human beings are inferior are different, partly because they are widespread, partly because they are not just brute likes or dislikes but involve hierarchical value judgments.

[42] Scanlon (2008), *Moral Dimensions*, p. 73. By "serious consequences," presumably, Scanlon means serious consequences for those who are subjected to prejudicial judgments and not for those who hold the judgments. Also, that a judgment of inferiority is maintained or helped maintained by an individual act of discrimination is a consequence of this action. Yet, from the context it appears that Scanlon has a particular subset of consequences in mind that does not include the social reproduction of the view that the members of certain groups are inferior (which arguably is just one more bad consequence along with others).

[43] Alan Wertheimer (1983), "Jobs, Qualifications, and Preferences," *Ethics* 94.1, 99–112, p. 100.

rejection—as opposed to the background customer responses motivating the employer's decision—as involving an insulting message of inferiority: that clearly is not warranted by the reasons for the employer's actions. While such a rejection might have other meanings that might be relevant for permissibility—e.g., the employer's unwillingness to bear economic costs to counter unjust responses of the employer's customers might constitute a hurtful, even insulting, message of indifference to the plight of minority people—arguably it is more objectionable to send a message of inferiority than a message of indifference.

5. An Important Ambiguity

While Scanlon's account of the wrongfulness of racial discrimination on the job market is sophisticated, there is a crucial ambiguity hidden in it. More specifically, we can construe it in two different ways. First, it might be read as identifying two factors that each contribute to the permissibility of discriminatory action independently of one another. In this view, insulting discrimination in the job market would be wrong even if it denied no groups "important goods and opportunities" and job discrimination would be wrong when it denied certain groups "important goods and opportunities" even if this sort of differential treatment involved no insults to anyone.

Second, while the denial of "important goods and opportunities" in itself renders discrimination wrong, the fact that it is insulting might not, when this fact is considered in isolation from the harmful effects of discrimination, make it wrong. However, if there is another reason why discrimination is wrong, i.e., that it denies certain groups "important goods and opportunities," the fact that discrimination is insulting affects wrongness by boosting the independent "objections that others" already "have to such an action."[44] Hence, in this view, while an insulting act of discrimination would render the agent's deliberations objectionable, the act would not be impermissible by virtue of its being insulting if such discriminatory acts had no serious consequences for anyone and, thus, no one had any objection to these acts independently of the insult.[45]

To see the difference between these two accounts, consider a scenario in which a society consists of 1,000 different, equal-sized, equally wealthy, and equally powerful racial groups. The first group is racially prejudiced against the second racial group and that group only. The second group is racially prejudiced against the third, and so on and so forth. Accordingly, members of all

[44] Scanlon (2008), *Moral Dimensions*, p. 62.

[45] Hence, on this interpretation the factor of insult interacts with other morally relevant factors in a way that differs from the additive model of how factors bear on moral permissibility; see Shelly Kagan (1988), "The Additive Fallacy," *Ethics* 99.1 (1988), 5–31.

groups are occasionally subjected to racial discrimination—1 out of 1,000 times they apply for a job—but no members of any group are "denied access to important goods and opportunities" and of no group is it true that "society marks them out as inferior," since 999 out of the 1,000 racial groups of which the society consists are not racially biased against it.[46] Based on the second interpretation, Scanlon's account does not allow us to say that this society is more wrongful than a society identical to this one in all other respects, except that no group has racial prejudice against any other group. The first interpretation implies that the former society is wrong in a way that the latter is not, since in that society people are sometimes subjected to insulting messages of inferiority.

I am unable to tell from Scanlon's account which of these two views, if any, he accepts. (He might simply be undecided.) On the one hand, he says that idiosyncratic "likes and dislikes of our fellow citizens may be things that we just have to live with," which suggests (not: entails) that insulting, differential treatment not tied to the denial of important goods and opportunities to any groups does not make this kind of treatment wrong, i.e., the employer who idiosyncratically hires an applicant who is best qualified and whom the employer considers superior by virtue of his or her charm does not act impermissibly.[47,48] Note also that on the present understanding, Scanlon does not really object to the initial case of the discriminatory employer presented in section 4, since it seems that applicants not hired for discriminatory reasons have no independent claim to be hired instead of the applicant actually hired: the rejected applicants are no better qualified than the applicant actually hired. Hence, it is not as if the discriminatory intentions of the employer boost the strength of an objection they have against the successful applicant being hired where this complaint exists independently of the employer's discriminatory intentions. On the other hand, many do seem to consider it wrong to insult others, so, presumably, many will consider it wrong to visit an ailing elderly relative for the sort of strategic reasons Scanlon mentions, even if there are no independent reasons for others to object to this visit, e.g., that the ailing relative will find a visit so motivated hurtful. Since I am unable to tell from Scanlon's work which of these two views he subscribes to; since which we accept, if any, makes a crucial difference to the moral qualities of some kinds of discrimination; and since each appears worth discussing in itself, I shall consider the significance of the distinction between these two views below.

---

[46] Cf. Scanlon (2008), *Moral Dimensions*, p. 73.

[47] Cf. Scanlon (2008), *Moral Dimensions*, p. 73. However, Scanlon might not be suggesting that these acts are not wrong, only that there are wrong actions that we cannot eradicate.

[48] Perhaps Scanlon thinks that what one cannot reasonably accept is being designated as inferior by the society that denies one important goods and opportunities, and thereby renders one's position inferior in important respects.

6. Some Worries About Scanlon's Account

I now want to express three worries about Scanlon's account of what *makes* racial discrimination in jobs wrong (not: his claim that it *is* wrong). First, Scanlon's second conjunct in his account of the wrongness of racial discrimination says that individual acts of discrimination are wrong "because of their meaning—the judgment of inferiority they express and thereby help to maintain."[49] It is not clear from this passage whether the second conjunct of Scanlon's *explanans* is essential to the explanation that he offers, i.e., if the reproduction of the judgment of inferiority is necessary to it. Imagine a case where someone, by uttering a judgment of inferiority, undermines that judgment, e.g., because he is widely considered an outcast and people want to distance themselves from him and for that reason will revise their views to minimize agreement with outcasts such as this person. If it is necessary, this person's discriminatory acts would not qualify as "particularly objectionable" in Scanlon's sense.

Yet, when people find discriminatory acts particularly objectionable, they often do not have a view on whether these will undermine the relevant judgments of inferiority. Alternatively, they will even find acts of discrimination particularly objectionable in cases where its practitioners hold views that are so obviously extreme and absurd that their engaging in these acts of discrimination will undermine, rather than maintain, the relevant judgments of inferiority. Hence, if intended to capture common-sense views, Scanlon's account should probably be revised to allow that racial discrimination on the job market that reflects judgments of inferiority can be "particularly objectionable" may make a discriminatory act wrong even though it does not "help maintain," perhaps even undermines, such judgments.[50] Having actual cases of discrimination in mind, one can see why Scanlon's account invokes the reproduction of judgments of inferiority, but it is also easy to see that it is not a necessary component in what renders discrimination wrong and his account is easily reformulated to accommodate this.

My second worry about Scanlon's account is more substantial. It is not clear how exactly expressing a judgment of inferiority relates to not hiring someone with the intention to avoid hiring members of a minority deemed inferior. Surely, one can express a judgment of inferiority even in the absence of such an intention, e.g., if I hire a person from that group simply because he is a member of this group and I believe that they are inferior and that a case for employing any of them could not rest on purely job-related reasons. Also, it seems that not all kinds of intentional discriminatory acts involve a judgment of inferiority, e.g., a patriarch who avoids hiring a certain applicant simply because she is a

[49] Scanlon (2008), *Moral Dimensions*, p. 73.
[50] Scanlon (2008), *Moral Dimensions*, p. 73.

woman and thereby intends to avoid hiring women, not because he thinks, as most patriarchs do, women are inferior, but because he thinks that there is a clear division of labor between the sexes and women's place is in the home.[51] A further complication derives from the fact that an employer may refrain from "giving a person a certain benefit," i.e., hiring him, because he thinks that the job is inferior and that it ought only to be performed by inferior persons. Should one say here that it is permissible not to hire the applicant—say, a Brahmin—considered superior by the employer, but not permissible not to hire the Brahmin with the intention to avoid hiring a superior person for an inferior job—say, a job the employer deems suitable only for Dalits? Clearly, in a sense the unsuccessful Brahmin applicant is discriminated against. However, he is not being deemed inferior and, thus, so I take it, not insulted, albeit not hiring him on the indicated grounds involves an insult against those who are deemed inferior, i.e. Dalits, and some of whom are hired for this very reason.

This possibility represents a problem for Scanlon's account. For in principle, we could imagine a society in which all employers offering less attractive jobs insist on hiring only people whom they deem suitably inferior for the jobs offered and that all employers offering attractive jobs have a nonsuperiority-based preference for applicants deemed inferior by other employers, e.g., as a way of compensating them. Suppose this has been going on for a considerable amount of time. Accordingly, we have a situation in which members of one group of people are deemed superior by many employers and yet are denied "access to important goods and opportunities," i.e., money and jobs, while those who are deemed inferior by many employers have better "access to important goods and opportunities." Members of the former find it harder to get a job regardless of how attractive that job is. This would be a situation in which a certain group is "marked out as superior by society" and yet is "denied access to goods and opportunities" as a result. The flip side of the coin is that those members of society who are marked out by society as inferior enjoy better access to goods and opportunities than do others. Accordingly, this, admittedly unrealistic, case puts pressure on Scanlon to identify more clearly how the factor of being marked as inferior interacts with the factor of disadvantage. One way to go would be to say that even in the present case marking someone out as inferior may be morally wrong even if, as it happens, they benefit from being so marked. Another way to go would be to say that marking someone out as inferior is wrong, only if it is harmful to these people.[52] If so, judgments of

[51] Perhaps Scanlon thinks that there is nothing "particularly objectionable" about this kind of sexist discrimination in the way there is about racial discrimination based on hierarchical value judgments.

[52] It is a further question if it makes a crucial difference whether an individual idiosyncratically marks members of a certain group as inferior or whether society does so. Scanlon seems attracted to the latter view.

inferiority may in effect fall out of the account of why racial discrimination is "particularly objectionable."[53]

I now turn to my third worry about Scanlon's account. Scanlon suggests that his idea of discrimination is "unidirectional. It applies only to actions that disadvantage a group that has been subjected to widespread denigration and exclusion. In the case of race, it applies to actions or policies that disadvantage blacks, but not to all policies that employ race-based criteria. So when discrimination is understood in this way, "reverse discrimination" is an oxymoron."[54]

Now, I agree that negative differential treatment of European Americans, to stay within the US context Scanlon has in mind, is less objectionable than otherwise comparable cases of discrimination against African Americans, precisely because such negative differential treatment is much less widespread and thus does not systematically disadvantage European Americans. Still, if one thinks that racial discrimination is "particularly objectionable" in part because it involves the expression of a judgment of racial inferiority, there is no reason why acts of disadvantageous differential treatment (which Scanlon would not refer to as cases of discrimination) cannot be morally wrong for exactly the same reason as acts of discrimination against African Americans.[55] Also, it is not clear how we should understand Scanlon's "access to important goods and opportunities" requirement. For suppose that some members of a certain minority group actually and perversely benefit from their group being deemed inferior—they end up better off, all things considered, than if their group had not been stigmatized—and some members of the group deemed superior by many are harmed by this fact. If we understand the "access" requirement in such a way that it applies to groups generally, and not to individuals, it implies that whereas privileged members of the former group may suffer discrimination when they are not hired for a job due to the employer's prejudice against them, underprivileged members of the latter group cannot suffer from discrimination when not being hired for a job because the employer is one of the few who is prejudiced against them. That does not seem right. Or, at least, if this is how we identify discrimination, a case of discrimination may be no worse, morally speaking, than an otherwise comparable case of nondiscriminatory differential treatment.

---

[53] I say "may" because Scanlon might think that it is more wrong to mark some group out as inferior, when this harms its members, than to do so when the relevantly similar harms fall on members of some other group, *ceteris paribus*.

[54] Scanlon (2008), *Moral Dimensions*, p. 74.

[55] When Scanlon writes that "not all deviations from relevant criteria are open to the charge of discrimination" (Scanlon (2008), *Moral Dimensions*, p. 74) he seems to suggest that at least some forms of race-based disadvantageous differential treatment of European-Americans cannot constitute discrimination because European-Americans have not "been subject to widespread denigration and exclusion" (Scanlon (2008), *Moral Dimensions*, p. 74) whereas a comparable race-based disadvantageous treatment of African-Americans constitutes discrimination.

To support the conclusion defended in the previous section consider discrimination in love. Suppose whites generally deem blacks inferior or socially unacceptable and for that reason frown upon racially mixed love relationships. Now, on the assumption that potential white and black partners do not differ as such in terms of attractiveness and on the assumption that for any person, black or white, there is a sufficient number of attractive, same-race potential partners such that racial discrimination in love does not deny blacks "access to important goods and opportunities," or, to the extent that it does so, it does the same to whites. That is, persons—black or white—who fall in love with different-race persons are worse off as a result of love discrimination. To the extent that we find the sort of difficulties these persons face objectionable for the same reason that we find racial job discrimination objectionable, this suggests that individual harms, not group harms, matter (insofar as harmless discrimination is not particularly objectionable in the way Scanlon thinks racial discrimination is).

7. The Moral Distinctiveness of Discrimination Based on Judgments of Inferiority

The sort of racial discrimination in hiring that Scanlon considers and considers "particularly objectionable" is not the only possible form.[56] That is, intentional (racial) discrimination need not involve a value judgment of inferiority or any value judgment at all for that matter (see chapter 1.8). Presumably, people may prefer people of a certain race even if they do deem any race superior to any other. It may well be odd for such a preference to exist in the absence of evaluational racism and most forms of racial discrimination on the job market that we know of may well involve such kinds of racism, but surely nonevaluational, racial discrimination is possible.

Given that assumption, the question arises as to whether there may be a difference in terms of permissibility between evaluation-based discrimination involving judgments of inferiority and otherwise comparable nonevaluation-based discrimination. Scanlon's treatment of discrimination in terms of insulting messages of inferiority suggests that he shares this view. It is by no means an uncommon view, and elsewhere in this book we have already seen similar views canvassed by Larry Alexander, Richard Arneson, and Deborah Hellman (see this chapter and chapter 4).[57]

---

[56] Scanlon (2008), *Moral Dimensions*, p. 72.

[57] " . . . Biases premised on the belief that some types of people are morally worthier than others are intrinsically morally wrong because they reflect incorrect moral judgments," Larry Alexander (1992), "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies," *University of Pennsylvania Law Review* 141.1: 149–219, p. 161; "Discrimination that is intrinsically wrong occurs when an agent treats a person identified as being of a certain

To address the question raised in the previous paragraph, let us consider a situation in which people generally do not care much about whether others deem them inferior. They do, however, care very much about whether others have brute desires in their favor. Hence, if someone prefers to be with people like me, but nevertheless considers people like me inferior, this will not damage my self-respect. However, if someone prefers not to be with people like me, even though he or she in no way deems me inferior that will seriously undermine my self-respect. Roughly: I care a lot about whether people desire to be with me, but I am indifferent as to whether they deem me worthy of such a desire. In Scanlon's account, only the latter kind of discrimination will involve an insulting message of inferiority (or social unacceptability) and only the latter kind, thus, involves the sort of distinctive wrongness of racial discrimination. Yet, when we take away all the negative effects that we generally associate with being the object of insulting judgments of inferiority—indeed, here we imagine a reverse situation, since people generally care more about being deemed worthy of a desire for their company than about being the object of a desire for their company—and just consider these judgments in themselves, my confidence in this kind of discrimination being distinctively wrong, in a way that brute preference-based discrimination is not, disappears.

I do not want to overstate the significance of this. After all, others may respond differently to this thought experiment than I do. Still, it shows that at least there is an argumentative burden to lift for those who, like Scanlon, put special emphasis on how (intentions reflecting) judgments of inferiority insult in their account of the wrongness of discrimination—a burden, which has not been lifted by Scanlon.

A different way to test the insult-based account would be to imagine a case where the employer acts with the same intention as the discriminating employer and yet the rejected minority applicant has no reason to think of the rejection as an insult that sends a message of inferiority. Suppose that, strangely enough, minority applicants have good reason to believe they are inferior—say, they have good reason to rely on testimony from people who tell them that they are—and no reason not to think of themselves as inferior—perhaps they really are inferior, not moral-status-wise of course, in a number of factual ways, e.g., in terms of income, achievement, and education, as a result of unjust social conditions, but regrettably they have no knowledge of the social conditions

---

type differently than she otherwise would have done because of unwarranted animus or prejudice against persons of that type," Arneson (2006), "What Is Wrongful Discrimination?" *San Diego Law Review* 43.4: 775–807, p. 779; " . . . it is morally wrong to distinguish among people on the basis of a given attribute when doing so demeans any of the people affected," Hellman (2008), *When Is Discrimination?* p. 7. As we saw above, Hellman thinks that differential treatment may be insulting without being demeaning and that intentions are irrelevant to whether treatments are demeaning and, thus, constitute wrongful discrimination, Hellman (2008), *When Is Discrimination?* pp. 138–168. Accordingly, my challenge below does not apply to her account.

that produce their inferiority. In such a case, rejected minority applicants might have no reason to see their rejection carrying a message of insult: they might even think it is formulated in flatteringly respectful terms, etc.[58] Analogously, a person who is punished for a crime he justifiably, but falsely, believes he has committed has no reason to see his sentence as an insult in light of the fact that the judge released another person charged with a similar crime and whom our convicted person justifiably believes not to have committed the crime for which he stood trial. Hence, if we—as I am inclined to—consider this kind of discrimination no less wrong than an otherwise comparable case of discrimination which does involve an insulting message of inferiority, the insult that racial minorities rightly take racial discrimination to involve is not what explains what makes racial discrimination "particularly objectionable."[59]

8. Conclusion

In this chapter I have taken a critical look at two attempts to account for the wrongness of discrimination in terms of the objective meaning of discriminatory acts. First, I had a look at Hellman's account, according to which discrimination is wrong whenever demeaning. I argued that there is an important plasticity in the notion of demeaning such that this account is either quite controversial or, alternatively, one that applies to very few cases of wrongful discrimination.

Next, I explored Scanlon's account, according to which acts of racial discrimination are particularly objectionable because the significance the discriminatees have reason to ascribe to such acts. I argued that not only is his account ambiguous in a crucial respect, when we abstract from those features that are only contingently tied to adverse differential treatment on the basis of evaluative judgments of inferiority, it is not clear that such discrimination is distinctively wrong in a way that other forms of discrimination are not. While this may constitute a criticism of specific points in Scanlon's account, it should also be noted that it is in line with his overall view of the irrelevance of intentions to permissibility.

One way to tie the wrongness of discrimination to equality is by contending that it derives from treating someone in a way that reflects a lesser moral status, either as a matter of the discriminator's beliefs about the status of the

[58] That is, unless we say that in some objective sense independently of the relevant parties' beliefs—a sense such that one cannot infer "Y's treatment of X is not insulting" from "X does not believe and has no beliefs that would motivate seeing Y's treatment of X as being insulting"—the relevant message is an insult.

[59] Scanlon (2008), *Moral Dimensions*, p. 72. This argument allows that it may be the fact that the content of judgment is the assertion of inferiority, not the insult that the assertion may constitute, that explains the particular wrongfulness of racial discrimination.

discriminatee or as a matter of the objective meaning of the discriminator's act. I have challenged the two most sophisticated versions of these views, arguing that neither provides a satisfactory account. However, there are other ways of tying the wrongness of discrimination to equality, e.g., by accounting for it in terms of the unequal treatment of the claims of different people. Alternatively, we might think that the wrongness is not to be explained in terms of unequal treatment or resulting inequality, but, say, in terms of the harm to discriminatees (or others) that discrimination involves. This is a topic to be explored in the next chapter.

{ 6 }

Harm-Based Accounts

1. Introduction

Having looked at various accounts that explain the moral wrongness of discrimination in terms of the mental states of the discriminator and in terms of the objective meaning of discriminatory acts, I now turn to an account that explains the wrongness of discrimination in terms of how it harms individuals, in particular, victims of discrimination. Initially I unfold a generic harm-based account (sections 2 to 5). I then proceed to propose a particular harm-based account of the wrongness of discrimination, i.e., the desert-prioritarian account (sections 6 and 7). Section 8 uses indirect discrimination as a test case for harm-based, in particular desert-prioritarian, accounts of the wrongness of discrimination.

One initial clarification is in order. Deontology is the view that there are moral "restrictions on action which have the effect of denying that there is any non-agent-relative principle for ranking states of affairs such that it is always permissible to produce the best available state of affairs so construed."[1] An example of such a moral restriction is the moral restriction on killing innocents—it is impermissible to kill an innocent even if by doing so one would prevent more killings of innocents and thereby bring about a less bad state of affairs from an agent-neutral perspective. Deontology is often contrasted with consequentialism, according to which it is always permissible, indeed morally required, to bring about a state of affairs that ranks below no other state of affairs according to an agent-neutral principle for such a ranking. Like with deontology, there are many versions of consequentialism. The most well known version is utilitarianism. Utilitarianism holds that one state of affairs is better than another if, and only if, it contains more welfare and that agents are always

[1] Samuel Scheffler (1982), *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions* (Oxford: Clarendon Press), p. 80.

morally required to bring about as much welfare as possible. Because harm is tied to welfare and because considerations about moral status and respect are normally central for deontologists, it is tempting to think that while the generic harm-based account of wrongness of discrimination is consequentialist, the mental-state and the objective-meaning accounts are deontological, or at least, nonconsequentialist. However, this temptation should be resisted. Consequentialists may endorse core elements from either or both of the accounts proposed in chapters 4 and 5, e.g., they might ascribe negative, noninstrumental value to demeaning actions being performed. Similarly, in order to accept a harm-based account, one need not be a consequentialist, e.g., one might think that there is a moral restriction against harmful discrimination such that one may not harmfully discriminate against someone even if that will prevent more cases of harmful discrimination. Someone who holds this view thinks that discriminatory acts are wrong in part, at least, because of the harm they involve even if she does not assess acts only on the basis of an agent-neutral ranking of states of affairs.

While there are consequentialists as well as nonconsequentialists within the broad family of harm-based accounts, this is not so within the narrower desert-prioritarian account discussed in sections 6 to 7. This account is a consequentialist one and, accordingly, has all of the drawbacks and advantages of consequentialism as such. Obviously, consequentialism, like other normative theories, is controversial and faces forceful objections. It is beyond the scope of this inquiry to deal adequately with those general objections in addition to those that the desert-prioritarian account faces as an account of the wrongness of discrimination. Accordingly, I do not propose the desert-prioritarian account as one that is clearly the right one, but as one that is better as an account of the wrongness of discrimination than the two previous ones. This might be seen as a somewhat disappointingly unambitious conclusion, but the strength of my arguments does not bear one that is any stronger than that. Still, I shall argue that the desert-prioritarian account has a good deal of plausible implications in relation to the morality of discrimination and this is significant, because many think that consequentialism in general is unable to accommodate core moral intuitions in relation to discrimination.

2. The Essentials of the Harm-Based Account

The harm-based account of the wrongness of discrimination says that

> an instance of discrimination is wrong, when it is, because it makes people worse off, i.e., they are worse off given the presence of discrimination

> than they would have been in some suitable alternative situation in which the relevant instance of discrimination had not taken place.[2]

This account applies to all the forms of discrimination that I distinguished between in part 1 of the book. Unlike the mental-state and objective-meaning accounts, it has no problem in explaining the wrongness of indirect and statistical discrimination. Typically, these forms of discrimination involve neither wrongful intentions, nor objectionable objective meanings, but there is little reason to think that these forms of discrimination are less harmful than direct discrimination.

The harm-based account is consistent with discrimination being only contingently wrong. Assuming that discrimination is not necessarily harmful, some instances of discrimination may not be harmful to anyone—and perhaps may even be beneficial to some or even everyone—in which case the harm-based account is compatible with there being no reason why this instance of discrimination is morally wrong.[3] While some may think that this is a problem with the harm-based account, it is important to note that friends of the mental-state and the objective-meaning accounts must—though for different reasons—accept (and, in the case of Larry Alexander and Deborah Hellman, do in fact accept) a similar implication of their views.[4] That being said, it should be noted that very few forms of discrimination involve no harm to anyone.

While the harm-based account of wrongness says that acts of discrimination are wrong because they harm individuals, it does not imply that acts of discrimination are always wrong insofar as they harm certain individuals. First, it might be that the harm some people suffer as a result of a discriminatory act is morally outweighed by the benefits enjoyed by other people such that the act is, all things considered, morally permissible (recall my example of the Cartesian demon issuing a credible warning of disastrous consequences if one disobeys his command to engage in direct discrimination, see chapter 1.5). Second, the account is consistent with—but does not entail—the finding that an act of discrimination may harm some individuals and yet not be wrong in virtue of this

---

[2] Note that the harm-based account is origin-neutral in the sense that the harm in question need not be imposed, in the ordinary sense of this term, by the discriminating agent, e.g., it might be imposed by others who see that the agent is discriminated against and who would otherwise not have harmed the discriminatee. See Nils Holtug (2002), "The Harm Principle," *Ethical Theory and Moral Practice* 5.4, 357–389, p. 360.

[3] Discrimination necessarily involves disadvantageous treatment of some people relative to others, but these people need not be harmed relative to how well off they would otherwise have been, nor need the relevant disadvantage involve a global disadvantage (cf. chapter 2.4), and one might think only such disadvantages (or harms) matter.

[4] Recall that Alexander thinks most forms of discrimination are not disrespectful in his sense and that Hellman similarly thinks that some instances of discrimination do not demean.

fact, e.g., because the person harmed deserves to be harmed because he has often harmed others through discriminatory acts.

The harm-based account says that "an instance of discrimination is wrong, when it is, because it makes people worse off." This formulation rules out that instances of discrimination can be wrong when they do not harm anyone. However, it does not rule out that instances of harmful discrimination may be wrong for reasons other than that they harm people, e.g., that a particular act of discrimination is a violation of a promise not to discriminate. "Because" in the formulation of the harm-based account is not to be read as "only because" and, accordingly, the harm-based account of the wrongness of discrimination does not rule out that one should be a pluralist about the sources of the wrongness of acts of discrimination (or about the sources of wrongness of acts in general for that matter). Admittedly, some will think that the harm-based account is flawed, because it rules out instances of discrimination that are wrong even though they harm no one. To accommodate their view one could further weaken the above formulation, e.g., by inserting a "typically" in front of "an instance of discrimination." The resulting account would be more immune to objections, but it would also be less interesting. Accordingly, I shall stick with the present formulation to see how far it takes us.

The harm-based account is a mere formula and to give it content, one will need to specify its key components. First of all, something needs to be said about what is the relevant counterfactual situation one has to compare the actual situation with in order to determine whether a certain instance of discrimination has been harmful to anyone. This is the baseline issue that I have already addressed in different contexts (see chapters 2.4 and 3.2) and that I will address again in the next section. Second, one will have to specify the relevant notion of harm to be employed. For instance, we will need to know if an instance of discrimination to be wrong must reduce someone's well-being or whether it might also be wrong if it, say, frustrates the non-well-being-related preferences of an individual even if it reduces no one's well-being. I address this issue in section 4. Because the harm-based account needs to be specified in these two ways, what may at first seem to be an objection to the harm-based account per se may turn out to be an objection to a particular version of it, i.e., one where the relevant baseline situation has been incorrectly specified or one where the relevant notion of harm is implausibly narrow.[5] Section 5 responds to some general challenges to harm-based accounts.

[5.] Suppose that in order to explain the wrongness of a certain type of discrimination in accordance with the harm-based account one has to appeal to a particular version of that account; that in order to explain the wrongness of another type of discrimination one has to appeal to another version; and that these two accounts are incompatible. If this were so, then the attempt to account for the wrongness of discrimination in terms of its harmful effects would fail. However, as far as the second of the two issues that I discuss below is concerned, there is no inconsistency

3. The Baseline Issue

Consider first the question of the relevant baseline. One could simply say that an instance of discrimination is bad when it makes the discriminatee worse off than she would have been had she not been subjected to it and we hold everything else other than the discriminatory act and what follows causally downstream from it constant. Call this *the Straightforward Account*.[6]

This account implies that when someone who is subjected to discrimination that perversely makes her better off than she would have been had she not been so subjected, the discrimination is not wrong. So suppose that a certain head of a German university in the 1930s in Nazi Germany, in making a particular decision of promotion, discriminates against a Jewish employee, making him immigrate to the United States, where he ends up much better off than he would have ended up had he stayed at his university. Yet, intuitively, the head of the university does act wrongly in making a discriminatory decision.

If we want to accommodate this objection, we can either try to explain the wrongness of the action through something other than the harm it involved, e.g., we can say that the act was wrong because of the unfairness of the decision made or because of the wrongful mental state of the head of university, or we can understand the notion of harm differently from how I proposed. Focusing on the latter option, there are least several moves one could make here.

First, one could say that the challenge presupposes that we understand harm in terms of how an act affects the individual's overall well-being or something like this. For clearly, there was one respect in which our discriminatee was harmed, i.e., he had a preference for being promoted and by discriminating against him the head of university frustrated this preference. While it is true that as a consequence of the discriminatory act the discrimatee ended up having a higher degree of overall preference satisfaction, this does not change that he was harmed in a certain respect, and if this suffices for the harm-based account to apply, the case is not a counterexample to the account as such. (I return to this issue in section 4.) One problem with construing harm in this way, however, is that supposing that the head of university had strong reason to believe that although his discriminatory act would harm the discriminatee promotion-wise,

---

involved in multidimensional, harm-based accounts that accommodate all the dimensions in which discrimination may be harmful and, thus, wrong. As far as the baseline issue is concerned, it is unlikely that anyone will think that the explanation of what they see as the wrongness of different types of discrimination will require appeals to different baselines.

[6] Some might suggest that there is a baseline, which is even more straightforward, i.e., how well off the discriminatee was prior to the discriminatory act. However, this baseline involves an implausible conception of harm; see Holtug (2002), p. 368. Specifically, it implies that in a case where an applicant does not get the job due to discrimination, but is made no worse off than she was prior to the employer's decision to reject her application the applicant is not harmed by being subjected to job discrimination.

it would further his well-being, all things considered, and this was part of the reason why he made the decision, it is no longer intuitively clear that his act was wrong and yet this is what the harm-based account implies on the present construal of harm.

Second, one could respond to the counterexample by saying that it was entirely coincidental that the Jewish university lecturer ended up better off and that the head of university presumably had good reason to think that the result would have been different. Along these lines, one might suggest that what matters to wrongness is not the actual consequences of the discriminatory act, but the consequences that the discriminating agent believed, or had good reason to believe, would ensue from his act. So construed, the harm-based account implies that the head of university acted wrongly.[7] Now, this shows that the present counterexample connects with a rather fundamental discussion in moral theory about whether actual or believed (or reasonably believed) consequences (or all of them) matter to moral permissibility. This discussion is one that I shall set aside here noting that one ecumenical response is simply to distinguish between different types of wrongness, e.g., subjective and objective wrongness.[8]

Third, we concede that the present example shows that the relevant baseline is not the one identified by the Straightforward Account. Rather than comparing the actual outcome to the one that would have obtained if the discriminatory act had not been performed and we assume everything else to be equal other than what follows causally downstream from the discriminatory act, we should compare the actual outcome to the one that would have obtained if the discriminatory act had not been performed and no else had performed discriminatory acts in the future.[9] Call this *the No-Discrimination Baseline Account*. The motivating idea behind this is that the head of the university was one among very many Germans in the 1930 who were in the grips of a virulent anti-Semitic ideology and as a result they discriminated against Jewish people. Together, these people harmed very many people a lot and it cannot be the case, so one might argue, that the fact that others would have harmed a certain Jewish individual even more had some German people not treated them

[7] As noted in the previous paragraph if he had reason to believe that he acted for the discriminatee's best, it becomes less clear that it is a counterexample to the harm-based account, because it becomes less clear that he actually acted wrongly.

[8] One can substitute " . . . it makes or is reasonably expected to make the discriminatee worse off . . . " for " . . . it makes the discriminatee worse off . . . " in my formula above. This should appeal to those who, unlike Hare, think that wrongness and, thus, badness is tied to the reasonably expected, rather than to the actual, consequences of acts. See the helpful discussion in Brad Hooker (2000), *Ideal Code, Real World* (Oxford: Clarendon Press), pp. 72–75.

[9] An account that moves the relevant baseline even closer to an ideal moral one would be one which said that the discriminatee is not made worse off than she would have been had everyone else acted morally permissibly from the time of the discriminatory act and henceforth. (Call *the Ideal Baseline*.) This view is like the No-Discrimination Baseline View a moralized view.

disadvantageously in certain discriminatory ways makes these acts not harmful and perhaps not wrong. The fact remains that the set of anti-Semitic acts harmed very many Jewish people seriously even if it might be true of some members of this set that the relevant disadvantageous differential treatment ended up making the discriminatee better off in the Straightforward Account. In the No-Discrimination Baseline Account, we should compare how the university lecturer ended up as a matter of fact with how he would have ended up, had he been promoted and no future acts of discrimination, e.g., the actions constituting Holocaust among many others, had taken place.

In this view, it does seem plausible that the university lecturer was harmed by not being promoted and, accordingly, that the harm-based account can handle the challenge. Moreover, the No-Discrimination Baseline Account has several attractive features. First, almost always when someone raises a complaint on behalf of a discriminatee, the complainer alleges that the discriminatee is worse off than he should be. In the Straightforward Account, one could be harmed by a discriminatory act and yet it could be the case that one still has more than one should have, in which case it seems that one has at best a weak complaint about being discriminated against (see the discussion of comparative injustice in chapter 9), e.g., because while one is harmed by the individual act of discrimination this does little to undo the many benefits one unjustly gains from unjust discrimination against others. The No-Discrimination Baseline Account motivates this focus.[10]

Second, the No-Discrimination Baseline Account also explains why there is a moral asymmetry between affirmative action (or, as it was sometimes called in the past, reverse discrimination) and (nonreverse) discrimination.[11] Affirmative action tends to close the gap between how well-off those who benefit unjustly from discrimination are and how well-off they would be if no discrimination took place henceforth. Discrimination (or if you like: other forms of discrimination than affirmative action) tends to widen this gap.[12]

---

[10] In fact, the present consideration motivates going even further, since one has little complaint too if one still is much better off through the unjust, albeit nondiscriminatory, acts of others and would have been much worse off than one actually is if these unjust acts are not performed.

[11] Many describe affirmative action as involving a form of discrimination, e.g., Iris Marion Young (1990), *Justice and the Politics of Difference* (Princeton, NJ: Princeton University Press), p. 195.

[12] It might be said that the No-Discrimination Baseline View only takes into account future discrimination and not past discrimination. Yet, one's complaint about being discriminated against seems as weakened by the fact one still ends up better off than one would have ended up did one not enjoy that fruits of past unjust discrimination as by the fact that one will eventually enjoy the fruits of future unjust discrimination. One response here would be to adopt a variant of The No-Discrimination Baseline Account that compares how well off one ends up to how well off one would have ended up in a world that never did and never will contain discriminatory acts. While this view might accommodate the present worry, I believe it has problems of its own.

Third, the No-Discrimination Baseline Account is compatible with the view that some of the forms of discrimination that I have identified tend to be worse, because more harmful, than others. Such accounts may distinguish between being a discriminatee and being a victim of discrimination, i.e., a discriminatee that is harmed as a result of being discriminated against. Not all discriminatees are victims of discrimination and some victims of discrimination are harmed more than others; and normally, just how much harm it does to a group to be discriminated against depends crucially on its level of income, education, and cultural self-confidence.

Fourth, denying that a certain discriminatory act that involves no harm is wrong implies neither that the agent cannot be criticized for performing it—since he might, e.g., have had reason to believe that the act would be wrong and, thus, may be blameworthy for performing it—nor that it is not bad that the agent has a character that disposed him to perform the discriminatory act.[13] Hence, if a discriminatory act makes someone worse off, but not worse off relative to how well-off she would be if no one else discriminated, we can still criticize the discriminator for his conduct, his moral reasoning, or his character even if we cannot say that the discriminatory act as such is wrong.[14] One can do so whether one subscribes to the Straightforward Account or to the No-Discrimination Baseline Account. My main aim in introducing the No Discrimination Baseline is to show that even if one is persuaded by putative counterexamples to the conjunction of the harm-based account and the Straightforward Account, this may simply show that one should change one's view as to which baseline account is the right one rather than giving up on the harm-based account.

4. The Metric of Harm

I now turn to the second dimension in which harm-based accounts may vary—that is, in respect of how we determine how well-off a person is. Here there are several subdimensions in which answers may vary. First, in some harm-based accounts what matters is whether some people are made worse off in terms of how their lives go as a whole as a result of discrimination (see chapter 2). This may seem a natural view to hold. However, some theorists reject it, or insist that it needs to be hedged or supplemented. Consider ageism. Many people

[13] See Hare's distinction between judgments of rightness of the act, judgments of moral rationality of the act, and judgments of character of the agent in Richard Hare (1989), *Essays in Ethical Theory* (Oxford: Clarendon Press), pp. 212–230.

[14] For a related point about intentions and moral impermissibility, see Judith Jarvis Thomson (1999), "Physician-Assisted Suicide: Two Moral Arguments," *Ethics* 109.3: 497–518, esp. pp. 515–518, and Thomas M. Scanlon (2000), "Intention and Permissibility," *Proceedings of the Aristotelian Society*, Suppl. Vol. 74.1, 301–317.

think this is wrong. However, if we adopt the whole life account, then, on the (optimistic) assumption, that all of us eventually grow old, this form of discrimination is one from which we will all eventually suffer. Hence, in principle we could be ageists and no one need be worse off than others in terms of how his life goes as a whole. To accommodate the view that ageism is wrong, one could adopt a life-segment version of the harm-based account, i.e., the view that discrimination is bad if it renders someone worse off in one or more segments of his life even though there is no one whose lives it makes worse off overall.[15]

Second, in multidimensional harm-based accounts, discrimination might be bad when it does not harm the discriminatee overall, provided that it harms him in some particular dimension, e.g., social status, legal recognition, income, education, or freedom from subordination and oppression; and this may enable moralized baseline versions of the harm-based account to accommodate the following sort of case. Suppose that in a just state there would be no bull fighting, and that in the present unjust state only men are allowed to be bull-fighters. This constitutes discrimination against women, and we may think this is bad even if women are not made worse off overall relative to how well off they would be in a just state (because neither men nor women would be allowed to be bull fighters in a just state) as a result. Adopting a multi-dimensional account, we can say that legal discrimination against women bull fighters is bad, because it subjects women to the harm constituted by lack of legal recognition that they would not suffer in a just state.

Third, some observers think that how well-off people are in the relevant sense depends on their level of (opportunity to use, or access to) resources. Others think it depends on their level of welfare. These views have very different implications. If we think that a just distribution is one in which everyone's life contains the same amount of welfare, a demonstration that, say, men generally command more resources than women will not in itself show that women suffer wrongful discrimination. Things, of course, would be different if we were concerned with resources rather than welfare. (We could, of course, be concerned with both.) Note also that if we hold the view that the form of harm that is relevant to the wrongness of discrimination is harm that consists in reducing someone's welfare, there is a further issue about what determines a person's level of welfare. It is common here to distinguish between mental-state accounts—a person's life goes better the more pleasurable mental states that he has; preference-based accounts—a person's life goes better the more of his preferences are satisfied; and objective-list accounts—a

[15] For the related view that equality requires not just that people should be equally well off in terms of how their lives go as a whole, but that they should also be equally off in different segments—e.g., corresponding or simultaneous segments—of their lives, see Temkin (1993) *Inequality* (Oxford: Clarendon Press), 232–244.

person's life goes better the greater the degree to which his life is characterized by certain objectively valuable features.[16] While it is reasonable to think that often when a discriminatory act makes someone worse off on one of these accounts, it makes her worse off on the others as well, e.g., frustrating someone's desire will often lead this person to have unpleasant mental states. However, the accounts are different and accordingly, there are bound to be cases where discrimination will harm on one account, but not on the others. Note, finally, that most agree that not all of our preferences matter to our degree of welfare. If, for instance, I have a preference for the preservation of the Amazon into the next century, whether it is satisfied will make no difference to my welfare: only preferences that are in a certain (surprisingly elusive) sense concerned with my own life do. If the frustration of such preferences counts as harm for the purpose of applying the harm-based account, it will be able to explain the wrongness of acts of discrimination that reduces no one's welfare. The drawback is that if the frustration of such preferences counts as harm, almost any act will involve harm but it will often be unclear whether these harms are morally relevant.

Fourth, there is the question of whether the units that may be subjected to harm are individuals or groups (or both). Some have argued that while black persons who suffered discrimination in the past and are now dead cannot be compensated, justice requires reverse discrimination because justice is also, at least, a matter of the equality or compensation of groups over time.[17] One challenge that group-based accounts of justice face is to explain how we identify the groups that are relevant from the point of view of justice.[18] Whatever the force of this challenge, there are both individualistic and group-centered versions of the harm-based account of the wrongness of discrimination, so this dispute is internal to the family of harm-based accounts.

To sum up: In this section I have explored four dimensions along which the harm involved in discrimination can vary, and for each dimension I have distinguished between various views one might hold. I have not defended any particular view regarding which is the right way to assess harm along any of these four dimensions. In part this reflects that this is a huge discussion in itself. In part it reflects that my primary aim in this section has been to demonstrate the broadness of the harm based account and, thereby, to show that many objections that might be seen as objections to the harm-based account really are objections to particular versions of the harm-based account.

---

[16] Derek Parfit (1984), *Reasons and Persons* (Oxford: Oxford University Press), pp. 493–502.

[17] For a defense of the latter view, see Paul W. Taylor, "Reverse Discrimination and Compensatory Justice," *Analysis* 33 (1973), 177–182.

[18] See George Sher (1999), "Diversity," *Philosophy and Public Affairs* 28.2, 85–104, pp. 90–93.

5. Some Challenges to the Harm-Based Account

In this section I review four challenges to the harm-based account. One challenge to the harm-based account of the wrongness of discrimination has it that the account is *trivial*. It is about as helpful as an account of the wrongness of pushing buttons that says that an instance of button pushing is wrong, when it is, because it harms someone (as it does, e.g., when the red button is the one said to be placed in the suitcase which the American president always has close by).

In response to this challenge, note, first, that, in contrast with the case of button pushing, many would say that discrimination is noncontingently wrong. Hence, the comparison with button pushing is misleading. Second, many would say of paradigm cases of discrimination that, even if they are not necessarily wrong, they are wrong for reasons independent of any harm caused. Finally, again unlike in the case of button pushing, certain common, salient, noncontrived types of discrimination are wrong.[19] Hence, given the dialectical setting of the debate about the wrongness of discrimination, the comparison with button pushing is misleading.

Second, my response to the previous argument might be met with the counterclaim that it ignores that the fact that an act is discriminatory is an *aggravating* factor. Perhaps only harmful discriminatory acts are wrongful, but the fact that a harmful act is discriminatory makes it *more* wrong than it would have been had it simply been harmful to an equal degree. This distinguishes discrimination from button pushing in that the latter obviously is not an aggravating factor.

In response, I note that if one can either prevent someone from imposing a discriminatory harm on someone else or prevent another from unfairly imposing a nondiscriminatory harm on a fourth person but not both, and one is slightly more likely to be able to be able to prevent the latter imposition, but otherwise all other morally relevant factors are equal, then one should prevent the latter. This suggests that, as far as wrongness of actions is concerned, discrimination is either not an aggravating factor—it might be in relation to the character of the agent—or is only weakly aggravating.[20] Admittedly, this

[19] Consequentialist friends of the harm-based account, who endorse something like Hare's two-levels theory of morality—see Richard Hare (1981), *Moral Thinking: Its Levels, Methods, and Point* (Oxford: Clarendon Press), pp. 44–64–might say that discrimination differs from button-pushing in that, consequentialistically speaking, our intuitive moral thinking ought to include a principle to the effect that one should not discriminate, but no principle to the effect that one should not push buttons. Hence on their view, it is morally justified that we subscribe to intuitive moral principles, according to which discrimination is morally wrong per se, even if, at the fundamental level of morality, it is not.

[20] Admittedly, this leaves open the possibility that unfairness is a significantly aggravating factor. I allow that the case does not suggest that discrimination is not a weakly aggravating factor, because our failing to distinguish any moral difference between the discrimination and the unfairness variants of the case might simply reflect that detecting weakly wrong-making factors is difficult.

argument relies on the assumption that the wrongness of an action is correlated with the strength of the reasons others have to prevent that action, and some deny this assumption, e.g., on the ground that a similar test would implausibly suggest that killing is no more wrong than letting die. However, in the test I imagine the nondiscriminatory act with which I compare the discriminatory act may involve unfairness too, so the test provides at least some evidence for discrimination not being an aggravating factor.

The third challenge is that the harm-based account is so broad that it becomes *all-inclusive*. Take, for instance, the view that discrimination is wrong in virtue of the agent's hostility toward those whom he discriminates against, i.e., a particular version of the mental-state-based account that I discussed in chapter 4. Now suppose someone treats people with a certain religion disadvantageously because of religious hatred, where as a matter of fact this does not result in anyone being harmed. Perhaps the religious discriminator does not even expect his discriminatory act to result in anyone's being harm, e.g., because he expects the relevant discriminatees to find jobs elsewhere that are no worse. Offhand, one would expect this to be a test case that one might use to weigh the *pro*s and *cons* of mental-state-based and harm-based accounts. However, given the inclusiveness of the harm-based account, it is not. For one version of the harm-based account says that people are harmed when their preference for not being treated with hostility is frustrated and, accordingly, some versions of the harm-based account could agree with people subscribing to the mental-state-based account that this case may be wrong, even though no one's welfare is reduced and the agent did not expect his action to reduce anyone's welfare.

In response to this objection, I note first that it is misarticulated when it is framed as an objection to the harm-based accounts. If the objection has any bite, it has bite against mental-state- as well as objective-meaning-based accounts too. These accounts are also wide-ranging in the sense that there are, say, many different accounts of which mental states render actions wrong within the broad family of mental-state-based accounts. Second, the first answer is limited in the sense that there are instead things that mental-state-based account cannot explain, e.g., the wrongness of discriminatory actions that involve no objectionable mental states whatsoever, whereas literally anything can count as a harm. However, even if this is true such that for any set of acts that constitutes wrongful discrimination on a mental-state-based account there is a version of the harm-based account that implies that the very same set of acts are acts of wrongful discrimination, this does not render the two accounts indistinguishable even if they might be co-extensional in the sense explained here. It is the fact that someone is harmed that renders the act wrong on the harm-based account, whereas this effect on the discriminatee, according to the mental-state-based account, is not what explains the wrongness of the act of discrimination.

The fourth challenge says that the harm-based account is *seriously incomplete*. Almost any act, whether discriminatory or not, will harm some and benefit others. Accordingly, insofar as harm bears on the wrongness of actions we will then have to weigh harms and benefits to different individuals against one another, and the harm-based account offers no way of doing this. But, ultimately, this is what we really need to do.

Unlike the three previous challenges, this one is a forceful one. It does not show that the harm-based account is false—incompleteness is not the same as falsehood. Also, it does not show that the harm-based account is worse than the mental-state-based and the objective-meaning accounts. They too are incomplete in that they do not in themselves tell us, e.g., whether it is wrong to engage in a demeaning discriminatory act if this is the only way to prevent many more equally demeaning discriminatory acts. It does show, however, that the harm-based account fails to give us all we want from an account of the wrongness of discrimination. To meet this challenge the next section will sketch a particular version of the harm-based account, which I find promising. It is not satisfactory in all respects—I respond to some objections to it in section 7—but it might be true that it is better than many of its competitors.

6. A Desert-Prioritarian Account

The particular harm-based account that I want to defend is a version of prioritarianism. Standard prioritarianism says that benefits that accrue to individuals with less well-being, not comparatively but absolutely speaking, have greater moral value than benefits to people with a higher level of well-being.[21] This means that smaller benefits to worse-off people might have greater moral value than greater benefits to people who are already better off and this is what gives prioritarianism a claim to be seen as an egalitarian theory, broadly speaking. Prioritarians do not give absolute priority to the worse off, so sometimes greater benefits to better-off people will have greater moral value than smaller benefits to worse-off people. Different prioritarians give different accounts of how much extra moral weight benefits to worse-off people have. No one has offered any other argument for why one weighting is better than another than that it has greater intuitive fit. Prioritarians, however, agree that all benefits have positive moral weight. Hence, unlike telic egalitarians, prioritarians avoid the so-called leveling-down objection: according to prioritarians there is nothing bad about going from an unequal state of affairs to an equal one where some are worse off and no one is better off.[22] Finally, standard

[21] Derek Parfit (1998), "Equality and Priority," in Andrew Mason (ed.), *Ideals of Equality* (Oxford: Blackwell Publishers), pp. 1–20.

[22] See the last section in this chapter.

prioritarianism says that individuals are morally required to bring about as much value as possible.

Standard prioritarianism, however, may be vulnerable to another objection, which I will try to accommodate. Suppose that someone engages in an act of blatant, racial discrimination against another. As a result someone is harmed or made worse off than the other.[23] In one variant of the case it is the discriminator himself who suffers the harm. In the other variant it is the discriminatee who suffers an equally great harm. According to prioritarianism these outcomes contain equal amounts of moral value. Hence, if a third agent can make the harm fall either on the discriminator or on the discriminatee, she has no reason to do one rather than the other as far as prioritarianism goes. Assuming, as we normally do, that people can be more or less deserving, this seems implausible. All other things being equal, it seems the discriminator is less deserving than the discriminatee and, accordingly, that it would be better, morally speaking, if the harm falls on the discriminator rather than on the discriminatee. Saying this is not saying that it is in some way good if the harm falls on the discriminator—it is simply less bad if it falls on him.

To accommodate this objection, I want to explore a version of prioritarianism that incorporates a concern for desert, desert-accommodating prioritarianism. As Richard Arneson expounds this view, in another context, it says the following:

> An act is morally right if, and only if, it maximizes moral value. The moral value involved in an action depends on three factors: (i) the greater the well-being for individuals affected by the act, the greater the moral value, (ii) the lower the level of well-being of those to whom additional units of well-being accrues, the greater the moral value, and (iii) the more deserving those to whom additional units of well-being accrues, the greater the moral value.[24]

In this view a given amount of well-being has greater moral value when it accrues to a badly off, deserving person than it does when it accrues to a well-off, undeserving person. Extra benefits to people always increase moral value. Hence, an act that leaves some worse off and no one better off cannot be morally right (unless it affects people's levels of desert).

---

[23] Note that I can make someone worse off than others without harming this person relative to how well off this person was prior to my act, e.g., I could benefit both persons but one more than the other.

[24] Richard J. Arneson (1999), "Egalitarianism and Responsibility," *Journal of Ethics* 3.3, 225–247, pp. 225, 239–240. Strictly speaking, desert-accommodating prioritarianism is not a particular moral principle, but a family of moral principles whose members differ from each other in terms of how much weight they give to the three factors held to be intrinsically morally significant. I do not attempt to specify these weights. Such specifications can only be done in a rough and intuitive way. See Parfit (1998), "Equality and Priority," p. 12.

Like utilitarianism and, for that matter, other consequentialist theories, desert-accommodating prioritarianism enjoins each agent to bring about as much moral good as possible. However, the two theories are different in a number of respects. First, desert prioritarianism leaves open whether well-being should be construed along perfectionist lines. Second, it ascribes intrinsic, not just instrumental, moral significance to considerations about desert.[25] Finally, it is not indifferent to the distribution of well-being among an act's beneficiaries. All other things being equal, desert-accommodating prioritarianism favors a more equal distribution over a less equal distribution. Equality in itself, however, has no moral significance in this view, not even when the desert levels of different individuals are equal. Hence, desert-accommodating prioritarianism may favor a more unequal distribution, if some people are better converters of resources into well-being, or if equality literally unavoidably erodes incentives (see chapter 11) and thus reduces the pool of resources. This feature of desert-accommodating prioritarianism has a very important implication in relation to discrimination. Because we cannot take for granted that equality is the morally best distribution of welfare, merely pointing out, say, that on average women are worse off than men as a result of discrimination will not establish that any wrongful discrimination has been going on (see also chapter 7). The distribution that maximizes the weighted sum of the distribuendum may be one in which women (men) are worse off than men (women).[26] Of course, there are good reasons for thinking that this is not the case and that inequality between men and women is evidence that the weighted sum of the distribuendum has not been maximized. However, this is different from saying that inequality between men and women is a criterion, or simply a necessary condition, for the existence of wrongful discrimination.

Consider next the baseline and the metric issues that I addressed in the sections above in relation to harm-based accounts. On the baseline issue, desert-prioritarians takes the Straightforward Account of the baseline. This means that they will have to resort to one or more of the responses to the objection regarding the Jewish lecturer that I sketched above, e.g., they would say that the discriminator acts subjectively wrong or they would have to say that desert-prioritarianism applies to sets of acts in addition to individual acts. On the metric issue, desert prioritarianism, on the most straightforward understanding, is concerned with whole lives and not just part of lives; it applies to overall harm to individual, not harm in particular dimension; harm refers to

[25] This raises the worry of a problematic circularity: what X deserves at t_1 depends on whether X performs the right act at t_1; that fact in turn depends partly on what X deserves at t_1. One view that avoids this problem says—as many do—that what one deserves depends wholly on what one did in the past. I thank Larry Alexander for pressing me on this point.

[26] Janet Radcliffe Richards (1980), *The Sceptical Feminist* (Harmondsworth: Penguin), pp. 124–127.

deficiencies in well-being; and desert prioritarianism focuses on individuals, not groups.

Desert-accommodating prioritarianism explains some of the moral distinctions that we want to make between different ways of discriminating. Consider first the issue of "reverse discrimination" or affirmative action. Assuming that such discrimination tends to favor people who are worse off and no less deserving than others, desert-accommodating prioritarianism can explain why it is morally different from the kinds of "nonreverse discrimination" that we know of and that, almost without exceptions, tend to harm those who are worse off.[27]

Second, we tend to consider discrimination worse the more harm it imposes on discriminatees and the worse off they are. The present account explains why discrimination against vulnerable and stigmatized groups tends to be morally worse than idiosyncratic discrimination against privileged and nonstigmatized groups (and against otherwise privileged but stigmatized groups).[28]

Third, suppose, contrary to what I argued in chapter 1, that the group of unqualified applicants is a socially salient group and, thus, that they can be subjected to discrimination, e.g., by not getting hired. In that case, the desert-accommodating prioritarian account does not imply that discrimination against unqualified applicants (based on their lack of qualifications) is necessarily morally wrong. For it is reasonable to assume that, setting aside complications pertaining to reaction qualification among other things (see chapter 9), for reasons of efficiency, maximization of moral value requires there to be a general fit between a person's qualifications and the job that he or she obtains.

The desert-prioritarian view (like other accounts of the wrongness of discrimination that appeal to considerations of harm) provides a satisfying, two-pronged, justification for why we have a concept marking differential treatment based on membership of socially salient groups. It seems harder to do that in the mental-state-based and the objective-meaning accounts. One can idiosyncratically deem members of a nonsalient group to have a lower moral status and one's treatment of them can be demeaning. First, stigma is a major type of harm involved in discrimination—perhaps the most grievous harm distinctive to discrimination. Stigmatic harm renders attributes of the victim "deeply discrediting" thereby reducing the bearer "from a whole and unusual person to a tainted, discounted one."[29] It is unlike pain in that one

[27] This is not to deny that other accounts apart from the desert-accommodating prioritarian account may explain this as well.

[28] Admittedly, there are other related intuitions, which sit less well with desert-accommodating prioritarianism; see my remarks on hate crime in chapter 4. Desert-accommodating prioritarians might say that the desert level of perpetrators of hate crimes is reduced, but this does not seem to be the reason why hate crimes are more wrong than ordinary crimes.

[29] See Erving Goffman (1963), *Stigma: Notes on the Management of Spoiled Identity* (London: Penguin Books), p. 12.

individual's pain will not harm another individual, unless this individual cares about the first individual; for example, because they are related as child to parent.[30] In contrast, stigmatic harms spread across persons belonging to socially salient groups even when they do not harbor external preferences for the well-being of other members of the group. If an African American person is treated as a "tainted, discounted one" by virtue of being African American, then the stigmatization will likely harm other African American persons; they will also be seen by some as "tainted, discounted" individuals.[31] But when a person with green eyes (or another socially nonsalient property) is treated in a likewise manner on account of his green eyes, other people with green eyes are unlikely to suffer.[32]

The second reason we have a concept marking differential treatment of people based on membership of socially salient groups is that, when individuals suffer disadvantageous treatment of this kind, the "individual acts of discrimination" are likely to "combine into a systematic and inequitable frustration of opportunity"[33] By contrast, when individuals suffer disadvantageous treatment based on membership of socially nonsalient groups (or on the basis of most so-called "individual properties") no such cumulative harm is likely to ensue. (The same point goes for most "individual properties.") The harm involved in each act of discrimination may be very close to zero.[34] However, the marginal harm from discriminatory acts is likely to rise steeply, when a certain number of discriminatory acts against a particular individual is reached.[35] Individual acts of discrimination may, were they to take place in an otherwise

[30] Hence, to care about whether others of one's own kind are subjected to stigmatic harm, one need not have a preference that is intrinsically external, that is to say, intrinsically about how others fare.

[31] Goffman (1963), *Stigma*, p. 12.

[32] I am offering an explanation of why we have the concept of discrimination here. Hence, I am not suggesting that discriminating against someone where no one notices and, accordingly, no stigmatic effects on others occur is not morally wrong.

[33] Paul Brest (1976), "In Defense of the Antidiscrimination Principle," *Harvard Law Review* 90.1, p. 10. Stigmatic harm tends to be a species of cumulative harms; the stigmatic harm is not a result of a singular act or a small number of singular acts, but the cumulated result of a large set of acts.

[34] Indeed, each such act may, considered on its own rather than as a member of a large set of discriminatory acts, be beneficial as a result, say, of the draconian social sanctions imposed on people belonging to different, hierarchically ordered groups who relate to one another in a nondiscriminatory way.

[35] What will at first constitute a violation of formal equality of opportunity will turn into a serious violation of substantive equality of opportunity. Unless one is extremely unlucky and differential treatment on whimsical grounds common, differential treatment on whimsical grounds is unlikely to turn into a serious violation of substantive equality of opportunity. See Richard Arneson (2002), "Equality of Opportunity," *Stanford Encyclopedia of Philosophy*, http://plato.stanford.edu/entries/equal-opportunity/; see also Paul Woodruff (1995), "What's Wrong with Discrimination?" in Steven M. Cahn (ed.), *The Affirmative Action Debate* (New York: Routledge), pp. 39, 40–41. Woodruff discusses disrespect-related harms.

discrimination-free social context, harm no one or involve only imperceptible harms.[36] Hence, when applying the desert-prioritarian principle, we should not apply it merely to individual acts of discrimination; we need to apply it to sets or series of acts. A single act that maximizes moral value might nevertheless be a member of a set of acts that is wrong, because there are other sets of acts that bring about more moral value.[37]

Summing up these two observations, the harm involved in disadvantageous differential treatment based on membership of socially salient groups is likely to spread across individual acts and to accumulate across individual acts. The same is not true of disadvantageous differential treatment based on membership of socially nonsalient groups or individual properties.[38]

This section has described some implications of desert-accommodating prioritarianism vis-à-vis the morality of discrimination. These implications are welcome, and support desert-prioritarianism. However, some of the implications combine equally well with moral positions other than desert-prioritarianism,[39] so this discussion of the morality of discrimination does not alone provide sufficient grounds for accepting desert-prioritarianism. In the next section, I shall address six discrimination-related challenges to the view.

7. Some Objections

One challenge to the desert-prioritarian account is that if discrimination is wrong, simply because it fails to maximize the sum of moral value, then there is nothing distinctively wrong with discrimination. It is wrong in exactly the same way that all other sorts of harmful or desert-level reducing acts are wrong. This cannot be right, because we know discrimination to be wrong in distinctive ways. For instance, Matt Cavanagh objects that an account of the wrongness of discrimination that appeals to the ideal of equality "does not seem to

[36] An applicant who is denied a job as result of discrimination in an otherwise discrimination-free social context is, of course, harmed in that he does not get the job. However, the applicant who in fact got the job would otherwise have been harmed in this way. The distinctive kind of harm involved in discrimination tends to arise only in social contexts permeated with discrimination.

[37] Compare Frank Jackson (1987), "Group Morality," in Philip Pettit et al. (eds.), *Metaphysics and Morality* (Oxford: Blackwell Publishers) pp. 91, 98–103; Parfit (1984), *Reasons and Persons*, pp. 70–73. Parfit argues that individual acts which are not morally wrong when considered in isolation can be made wrong by membership in a wrongful set of acts.

[38] Obviously, this claim does not prevent me from claiming that other factors affect the wrongness of discriminatory actions, e.g., it may be more wrong, because more harmful, to discriminate against members of an already badly off group than against members of a better off group.

[39] Whether some of the implications support respect-based accounts as well is discussed later.

capture what is distinctively wrong with discrimination."[40] (Call this objection the *nondistinctiveness objection*.)

This challenge can be answered.[41] There are two ways in which discrimination might involve a distinctive wrong. First, it might violate a basic moral principle that pertains to discrimination in the same way that some deontologists think that basic moral principles pertain to lying or to the killing of innocents.[42] Second, it might often cause or involve kinds of harm or reductions of levels of desert that are associated only with discrimination; for example, a sense of shame regarding one's social identity or a sense of being inferior or, at least, as being regarded as inferior.[43] It is question begging to object that desert-prioritarianism does not show that there is anything distinctively wrong with discrimination in the first sense, because it straightforwardly follows from desert prioritarianism that there is nothing distinctively wrong with discrimination in the present sense and, accordingly, the objection seems convincing only to someone who has already given up on the desert-prioritarian view. To object that it implies that there is nothing distinctively wrong about discrimination in the second sense is plainly wrong. As we have seen above, defenders of the desert-prioritarianism and other harm-based accounts can say that the harm involved in discrimination is distinctive, either in terms of the kinds of harmful effects that it involves or in terms of the kinds of harm that being discriminated against may be thought to constitute.

Second, another challenge says that insofar as discrimination is wrong because someone is harmed, primarily it must be wrong because it harms those who are subjected to discrimination. So any account tying the wrongness of discrimination to harm must make sure that it is somehow tied especially to the harm suffered by those who are subjected to the discrimination.[44] Despite the fact that desert prioritarianism incorporates a concern for desert, it might be

---

[40] Cavanagh (2002), *Against Equality of Opportunity* (Oxford: Clarendon Press), p. 155. Note that the objection is different from the triviality objection to the harm-based account for it does not say that the account of the wrongness of discrimination by desert prioritarians is uninformative as such.

[41] Despite Cavanagh's challenge, it turns out that even he does not think that there is anything distinctively wrong with discrimination. He thinks discrimination is wrong when it involves treating someone with unwarranted contempt, Cavanagh (2002), *Against*, p. 166. But if that is the case, then presumably wrongful discriminatory acts are wrong for the very same reason as other acts that are nondiscriminatory but involve treating someone with unwarranted contempt.

[42] See Charles Fried (1978), *Right and Wrong* (Cambridge, MA: Harvard University Press).

[43] See R. Wasserstrom (1977), "Racism, Sexism, and Preferential Treatment: An Approach to the Topics," *University of California Law Review* 24: 581–615, p. 593; see also Blum (2002), "'I'm Not a Racist, But . . . '" (Ithaca, NY: Cornell University Press), pp. 8–11, on inferiorization caused by racism.

[44] Cavanagh (2002), *Against*, p. 157, makes this objection against meritocratic accounts of the wrongness of discrimination. For an account of the harms caused to the discriminators by apartheid, see Desmond Tutu (1999), *No Future Without Forgiveness* (London: Rider), pp. 154–155.

denied that the desert-prioritarian account does that. (Call this objection the *discriminatee-focused objection*.)

In response to this objection, suppose X discriminates against Y, and third-party Z is neither subjected to discrimination nor discriminates. There are then two cases to consider: whether it is worse that Y is harmed rather than Z, and whether it is worse that Y (or Z) is harmed rather than X. Assuming that one does not necessarily become morally more deserving by being subjected to discrimination, it follows that, all other things being equal, it could be morally irrelevant whether the harm involved in Y's being discriminated against falls on Y rather than on Z.

This implication is not damaging. Suppose an ethnic majority discriminates against an ethnic minority. Suppose we initially regard this as morally wrong, but we then discover that, contrary to what we had supposed, the discrimination does not really harm the ethnic minority. The harm we thought fell on this minority actually falls on another ethnic minority, one that is *not* discriminated against. I doubt that we would then conclude with some relief that the discrimination in question is less bad than we had thought initially. A second point is this: assuming that the actions one performs determine desert, being discriminated against (or being subjected to unjust treatment in general) cannot in itself affect one's desert status. This, of course, is consistent with the expectation that if two people enjoy the same level of benefits and one of them has been subjected to discrimination, this latter person is likely to have a higher level of desert than the other because he had to strive harder to obtain his benefits.

This clarification implies that it is morally irrelevant whether X or Y (or Z) is harmed; one might urge that this implication is damaging. However, in cases in which we think it matters morally whether harms or benefits fall on the discriminator, X, or on others (including the victim of discrimination, Y, or a third party, Z), the discriminator is presumably morally less deserving by virtue of being a discriminator, e.g., one is less deserving when one discriminates in response to racist or sexist value judgments, it is (other things being equal) morally better that the resulting harm falls on the discriminator rather than on the discriminatee.[45] Given this, it follows from the proposed account that it is worse for a given harm to fall on Y or Z rather than X.[46] As I hinted above, to many this will seem to favor desert prioritarianism over straight

[45] Note that if X blamelessly discriminates against Y and thereby violates, let us assume, a right of Y, it might well be the case that if the relevant harm were to fall on Z instead of Y, X will have violated Z's rights as well, in which case even one who subscribes to a right-based view might think that it is no worse if the harm falls on Y rather than Z.

[46] In cases where discrimination is not wrong, being a discriminator presumably does not render one less deserving. Accordingly, it is not clear that one may convincingly object to the proposed view on the grounds that it does not discount harms to the discriminator in such cases.

prioritarianism (and other views according to which desert has no intrinsic moral relevance).[47]

Third, it could be argued that desert prioritarianism has counterintuitive implications in a case where someone is subjected to discrimination and, perversely, this maximizes moral value on the specifications described above; perhaps the discriminatee benefits because adversity strengthens his will to succeed and dramatically increases his level of desert;[48] perhaps the benefit to others more than compensates for, morally speaking, the harm to the discriminatee. In desert prioritarianism, such discrimination is not morally wrong. This implication is bound to be seen as problematic, just as some would say that, *pace* the desert-prioritarian principle, that the mere fact that a given act of lying, deceiving, or killing an innocent maximizes moral value does not render it morally permissible. (Call this objection the *deontological objection*.)

Friends of desert prioritarianism can concede that something may well be morally amiss in cases involving beneficial discriminatory acts, but insist that what is morally amiss is *not* that the act is wrong. To deny that a certain discriminatory act that maximizes moral value is bad is not to imply that the agent cannot be criticized for performing it; the agent might, for example, have had reason to believe that the act would on balance harm the discriminatee, and thus attract blame for performing the act. Nor need it oblige us to deny that it is bad that the agent's character predisposes him to perform the discriminatory act; on most occasions, discriminatory acts far from maximize moral value, so we should cultivate a character that does not lead us to perform them.[49] Hence, for reasons similar to those mentioned in section 3 above if a discriminatory act maximizes moral value, in all likelihood we can still criticize the discriminator for his conduct, his moral reasoning, or his character even if the discriminatory act is not morally impermissible.[50] Admittedly, there could be cases in which we could do neither of these things, but then in such cases—such as a case in which the discriminator intends his act to strengthen the discriminatee's resolve and

---

[47] Some may object that the proposed view does not go far enough. For example, they may retributively think it morally better that the discriminators' well-being is lowered; it is not merely that the moral value of a given level of well-being for them is reduced. Alternatively, they may reject that in cases in which all other things are not equal—e.g., the discriminators are much worse off than those whom they subject to racial discrimination—it may be better that the resulting loss of well-being falls on those who are subjected to discrimination, rather than on those badly off people who discriminate.

[48] See Larry Alexander (1992), "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies," *University of Pennsylvania Law Review* 141.1: 149–219, p. 188.

[49] See Hare (1989), *Essays in Ethical Theory*, pp. 212–230, on the distinction between judgments of rightness of the act, judgments of moral rationality of the act, and judgments of character of the agent.

[50] Cf. Thomson (1999), "Physician-Assisted," esp. pp. 515–518, and Scanlon (2000), "Intention and Permissibility."

thereby benefit him—we are unlikely to have any firm moral intuitions to the effect that there really is something morally amiss with the discriminatory act, and accordingly we cannot base an objection to desert-prioritarianism on consideration of such cases.

Fourth, desert prioritarianism implies that harmful discrimination is morally permissible in cases where it maximizes moral value. So suppose—unrealistically—the only way to maximize moral value is through formal discrimination—for example, legal discrimination forbidding some socially salient groups to do what other socially salient groups are permitted to do—and substantive discrimination—that is, differential treatment that makes members of some socially salient groups worse off than they would be in the absence of such treatment. On the view proposed here, discrimination would be morally required under such circumstances. This is bound to strike some as implausible, because it seems to take a much too tolerant view on discrimination. (Call this the *discrimination-tolerant objection*.)

This objection owes at least part of its force from a misunderstanding, which is best set aside at once. It may be suspected that anyone who considers a certain kind of action (such as lying, killing, or discriminating) morally permissible under circumstances very different from those we live in is somehow committed to the view that this kind of action is not so morally wrong, after all, in our actual circumstances. Although the claim that a certain type of action would be morally wrong under any possible circumstance is often taken to imply that the action is seriously morally wrong, this does not strictly follow.[51] One could, in principle, consider a certain kind of act to be mildly morally wrong whatever the circumstances. In that case, someone who thought the relevant kind of act was wrong only under some circumstances might think that the act is morally more wrong under the circumstances in which they are wrong than the person who finds them mildly morally wrong whatever the circumstances. Therefore, the present view does not involve a morally lax view of discriminatory acts performed in our actual circumstances—circumstances under which neither formal, nor substantive, discrimination maximizes moral value.

To evaluate the allegation of implausibility properly, we need to consider it in more detail. Let us focus on formal discrimination, e.g., the discrimination involved in the very formulation of sexist voting laws, since many people view this as even more wrong than substantive discrimination. Suppose that sexist laws forbidding women to do certain kinds of work are required to maximize value. Several reasons could explain this requirement: (1) because in the absence of these laws people would be less deserving and, hence, benefits to them would count for less morally speaking; or (2) because in the absence of these

[51] This conversational implicature is thought to hold because the most obvious explanation of why something is always morally wrong is that the moral reasons by virtue of which the act is morally wrong are infinitely strong.

laws people would benefit less. Those who would not be willing to harm women to secure nonsexist laws would not be bothered by at least some versions of the first scenario.[52]

Things might be different if men are better off than women and women will benefit only marginally from sexually nondiscriminatory laws, whereas men will lose a lot. If men are better off than women, their losses will have to exceed the gains of women to outweigh the latter on the desert-prioritarian approach. If men are worse off than women, it may not seem so counterintuitive that for some size of benefits to men and for some size of loss to women, formally discriminatory laws might be morally permissible. To insist that this implication is implausible, one would have to subscribe to an absolute deontological constraint against formal discrimination, and most observers rightly consider absolute deontological constraints implausible.[53] In fact, the situation would be akin to our own situation in which many of us favor efficient, sexually discriminatory laws implementing affirmative action in favor of women.

Fifth, desert-prioritarianism does not condemn discrimination as wrong in itself and any plausible account must make some such claim. (Call this the *intrinsic wrongness objection*.)

Obviously, it is indeed an implication of a desert-prioritarian account that discrimination is not in itself wrong. However, this is not a weakness of the account and, in any case, it is a feature of many of accounts of the wrongness of discrimination and, accordingly, not one that one can appeal to as a reason to favor one of these accounts over a desert-prioritarian one.

Consider first consequentialist theories in general. There is a broad range of consequentialist views. One has it that any instance of discrimination is in itself bad. On this view discrimination is always regrettable. However, it need not be impermissible, since a regrettable discriminatory act may prevent more such acts or other acts that are equally bad, or even worse. In any case, the absence of discriminatory acts is not a standard ingredient in axiology. This strikes me as plausible. A world in which French favor French, Turks favor Turks, and Chinese favor Chinese contains no less intrinsic value than one in which no one is favored on grounds of nationality, provided (as is unlikely to be the case) that all other things are equal. Hence, let us instead consider the large majority of consequentialist theories incorporating the notion that what is in itself good is *well-being*. On such views, discriminatory acts are impermissible if, and only if, they fail to maximize the (morally weighted) sum of well-being. Since it is possible that some discriminatory acts maximize the sum of well-being, some discriminatory acts are not morally impermissible according to consequentialism. Perhaps there is a greater sense of community, and thus more well-being, when

[52] Cf. Richard J. Arneson (1999), "Against Rawlsian Equality of Opportunity," *Philosophical Studies* 77: 77–93, p. 93.

[53] See Shelly Kagan (1998), *Normative Ethics* (Boulder, CO: Westview Press), p. 79.

(given a just background) everyone favors their co-nationals rather than treating co-nationals and foreign nationals no differently.

Presumably, many people would say that the welfarist-consequentialist story about what makes discrimination wrong, when it is, is highly implausible. Since I am inclined to accept this kind of account, let me indicate two key attractions it has. First, consequentialist accounts derive support from the fact that, generally speaking, we find discrimination *more* morally objectionable as it harms people more. Second, for reasons similar to the one mentioned in my discussion of the baseline issue discussed in section 3 above the consequentialist can explain why differential treatment on the basis of race, as in the case of affirmative action, is morally different from racist discrimination. Affirmative action tends to favor people who are worse off.[54] Accordingly, given decreasing marginal benefits from additional resources or given that our favored consequentialist theory is of the prioritarian type, consequentialism can explain why it is morally different from the kinds of "nonreverse discrimination" which we know of and which, almost without exceptions, tend to harm those who are worse off. Finally, it can explain what is wrong about indirect discrimination. Mental-state-based and objective-meaning-based accounts seem unable to do so, because by definition indirect discrimination does not involve any objectionable mental states and often it has no objectionable objective meaning. I will return to indirect discrimination and desert prioritarianism in the next section.

Finally, it might be argued that discrimination is wrong, in part at least because it is unjust or unfair in the sense that it involves not taking people's moral claims equally into account.[55] Hence, even in cases where no one is harmed, there is something wrongful about treating disabled people disadvantageously when they apply for jobs even if one's disadvantageous treatment of them ends up not harming anyone. (Call this *the unfairness objection*.)

This objection may well contain a kernel of truth, but it does not force us to give up the desert-prioritarian or, *a fortiori*, the harm-based account. First, arguably the unfairness objection does not apply to all forms of discrimination, so at most it offers a partial explanation of the wrongness of discrimination. For instance, many forms of indirect discrimination do not involve discriminators not taking people's moral claims equally into account. Also, assuming that individuals have no claim not to be treated on the basis of statistical facts about groups of which they are members, harmful statistical discrimination may not raise any fairness issues. Second, even in cases where the unfairness issue does

[54] Admittedly, the mental state and the objective meaning accounts can explain this difference too, since agents of affirmative action do not deem anyone to have a lower moral worth and affirmative action as such does not seem to have any demeaning objective meaning.

[55] John Broome (1990), "Fairness," *Proceedings of the Aristotelian Society* 91: pp. 87–101; John Broome (1994), "Fairness versus Doing the Most Good," *The Hastings Center Report* 24:36–39.

arise, it may be accommodated by desert prioritarianism's desert component, i.e., benefits to people who culpably treat others unfairly has less moral weight than people who do not. Thirdly, even in cases where considerations of unfairness cannot be so accommodated, unfairness does not appear to make a very significant moral difference. If, for instance, I am about to treat a group of people unfairly but in such a way that everyone benefits from the unfair treatment, it is not clear that others should prevent me from acting unfairly or indeed that I myself should refrain from the unfair treatment. Finally, the harm-based account allows that there might be other sources of wrongness than harm and while desert-prioritarianism in the formulation I have given it does not, one could modify this formulation to allow such sources, while retaining one's claim that it captures the most general and most important source of the wrongness of discrimination.[56]

In this section I have considered six objections to a particular version of the harm-based account of the wrongness of discrimination, the desert-prioritarian account. While I think that the objections sketched above can be met, by and large, I do not claim that I have considered all plausible objections to desert-prioritarianism.[57] Still, desert prioritarianism appears to be a strong contender among accounts of the wrongness of discrimination. In the next section, we will see that one virtue of it (and harm-based accounts in general) is that it coheres well with a plausible view about the morality of indirect discrimination.

8. A Test Case: Moral Wrongness of Indirect Discrimination

Often when accounts of the wrongness of discrimination are tested, theorists have in mind direct discrimination. However, it is useful also to consider the plausibility of the harm-based and, more specifically, the desert-prioritarian account of the wrongness of indirect discrimination. To test the moral relevance of indirect discrimination, consider the following thought experiment: there are two scenarios that have the same distributional profile, where, however, one involves indirect discrimination and the other one does not. Suppose

[56] A similar reply could be made to luck egalitarians who would object that desert prioritarianism fails to identify the wrong of some ending up worse off than others through responsibility of their own.

[57] For instance, I have not considered the objection that because desert prioritarianism is aggregative it implies implausibly that for any great harm imposed through discrimination on a small number of people there is a large number such that preventing trivial harms to this number of people produces more moral value than preventing the great harm to the small number of people and, thus, is the harm that one ought to prevent, when one cannot prevent both. This objection applies to aggregative theories in general and is not tied specifically to the morality of discrimination.

that in the first scenario the educational system involves indirect discrimination against women. In order not to pollute the example with concerns about direct discrimination, suppose that while women in the past experienced various forms of direct discrimination such discrimination belongs to the bad old days only. Suppose in the second scenario the educational system involves indirectly disadvantageous differential treatment of people with certain personality types. For some odd reason the types do not form a group of people that has any particular shape, so we do not think of them as a group of individuals in any sense other than that these people have the same scores on personality tests (often to everyone's great surprise). Thus, they do not form a socially salient group and, thus, do not form a possible object of group discrimination. Suppose that the very mechanism, whereby the relevant disadvantageous results are produced, is more or less the same in the two cases and suppose that the harm suffered in the two cases by the groups of disadvantaged people is of the same size. In this comparison it is hard to see that one case differs morally from the other. If so, indirect discrimination is not more wrong than a comparable case of rules with differential effects that, however, do not amount to indirect discrimination, because in this case the victims of the relevant mechanism do not form a socially salient group. This implication is consistent with desert-prioritarianism.

One might resist the view that these two cases do not differ morally and, thus, think that the fact that desert prioritarianism implies that they do not speaks against it. One suggestion as to why they differ morally is that women as a group have a claim to compensation for the direct discrimination that they suffered in the past, and hence the injustice they suffer from by being victims of indirect discrimination is double, so to speak. Whatever one thinks about the group-focused claim to compensation for historic injustices, this is a separate issue from the one I address here. So in comparing the two cases we should assume that women have been compensated for past injustice.

Another suggestion is that the first case is morally wrong because women are an already disadvantaged group, whereas a group of random individuals is not. In response, we might simply stipulate that, for some weird and unexpected reason, it turns out that the group of individuals picked out randomly also forms an already disadvantaged group. I submit that this will not lead us to assess the two cases differently and accordingly that discrimination need not be worse than comparable differential treatment that, however, does not amount to discrimination.

Third, it might be suggested that the indirect discrimination that women suffer is morally worse because it is tainted by its being a causal effect of past direct discrimination. If, for a moment, we set aside whether this makes a moral difference, note that something similar may be true of the group of people with the relevant personality trait: their being disadvantaged now may be the causal result of such people being unjustly disadvantaged in the past. But then the causal taint may apply to both cases and will not constitute a morally relevant

difference between them. So on the present suggestion, at least some cases of indirect discrimination may be no worse than otherwise comparable disadvantageous treatment.

It might be replied that all the present line of argument shows is that the social salience condition is morally irrelevant per se. That is, the fact that disadvantages fall on a socially salient group does not make the act from which these disadvantages result more wrong than an otherwise comparable act that results in disadvantages falling on a nonsalient group, all other things being equal. This, however, does not show that the other conditions stated in the definition of indirect discrimination are not morally relevant per se. However, the expanded no-intention condition (see chapter 2) does not look like a good candidate here. The *absence* of bias hardly makes indirect discrimination involve an additional wrong and, presumably, can also be true of cases that involve no kind of discrimination. This leaves us with the revised disadvantage and the causal condition.

As I have already noted the revised disadvantage condition is unlikely to concern those who are not egalitarians in a strict sense. So, for instance, desert prioritarians may think that practices that disadvantage certain groups are often morally wrong, because there are alternatives that produce a greater sum of weighted benefits, there might well be cases where the highest such sum is produced if some socially salient groups are worse off than others.

So let us consider the revised disadvantage condition from the perspective of an egalitarian in a strict sense. One thing to note here is that many egalitarians are concerned about inequalities between individuals and not inequalities between groups, whether socially salient or not.[58] Hence, from the perspective of these egalitarians the revised disadvantage condition does not capture a morally relevant feature.

Even if we set this aside, there is a further issue. The revised disadvantage condition leaves open whether the relevant disadvantages are local or global.[59] However, many egalitarians would deny that local disadvantages matter per se. I share this intuition and as noted above, desert prioritarianism holds a similar position on this matter. In support of it suppose we have a society with all sorts of rules that, when considered in isolation from the other rules, disproportionately disadvantage some groups and not others, but such that, all in all, no group is better off than any other, i.e., the disadvantages imposed on different groups cancel one another out so to speak.[60] Suppose, moreover, that there is

[58] Temkin (1993), *Inequality*; Anderson (2010) *The Imperative*, p. 67.

[59] Cf. Kagan (1998), *Normative Ethics*, pp. 86–87.

[60] Note that this argument does not presuppose a welfarist metric of interpersonal comparison. We might think that the metric is the range of available options or deliberative freedoms—see Sophia Moreau (2010), "What Is Discrimination?" *Philosophy & Public Affairs* 38.2: 143–179, and the appendix to this chapter—and yet agree that under the circumstances described, there is no reason to eliminate indirect discrimination that is locally disadvantageous.

no direct discrimination in such a society, even if there may have been discrimination in the past. It need not be the case that, *in this society*, we should eliminate all indirect discrimination in light of the fact that indirect discrimination does not in any way affect the relative overall positions of the groups involved. Indeed, on the assumption that eliminating some kinds of locally disadvantageous forms of indirect discrimination, but not all, such that eliminating some forms only would disadvantage some groups globally, then we should refrain from eliminating locally disadvantageous indirect discrimination. In fact, we could even strengthen this assumption and assume that no one's relative positions would change by eliminating rules that, locally speaking, favor some groups over others. Here it would seem pointless to eliminate indirect discrimination.[61] This supports the view that it adds no distinct complaint to that of distributive injustice to say of something that it involves indirect discrimination at least in so far as one has local disadvantages in mind.[62]

Shifting the focus from local to global disadvantage, certainly this is relevant from the perspective of all egalitarians. So this feature of indirect discrimination may be one in which indirect discrimination is unjust per se. However, my claim here is not that it is not. My claim is that two distributions that involve the same distributional profile, and where one distribution but not the other involves indirect discrimination, are equally wrong. Since the revised disadvantage condition does not really say more than that the relevant policy or act disadvantages a socially salient group, and since, presumably, for all disadvantaged socially salient groups there are some acts that disadvantage them, and not all such disadvantages are morally wrong, this feature of indirect discrimination cannot be what makes it *more* wrong than comparable cases of non-indirect-discrimination-related disadvantages.

So let us finally consider the suggestion that the causal condition is morally relevant per se. What then to say about cases of indirect discrimination that are the causal result of past direct discrimination against those globally

[61] This point applies with even greater force to a case where indirect discrimination results in local disadvantages for groups that are better off, globally speaking. In such cases one would hardly think that justice is promoted if these local disadvantages are eliminated as a result of which global inequality is increased. For instance, if one subscribes to a Rawlsian principle of fair equality of opportunity, one would resist this move.

[62] It might be thought that in making this reply I am presupposing the anti-Walzerian view that there is no plurality of spheres of justice, only one overall currency of justice. In fact, I need not rely on this assumption. I could simply restrict my claim to indirect discrimination within a certain sphere of justice. So if higher education is one such sphere, even friends of the view that there are different spheres of justice might concede that there is nothing unjust about indirect discrimination of the sort I describe in my example, i.e., one in which there inequalities in terms of particular sectors of the sphere of education. I thank Nils Holtug for pointing out the need to address this point. Note also that in most cases where someone complains about indirect discrimination involving a local disadvantage they do so in a setting of global disadvantage for the relevant group.

disadvantaged by indirect discrimination. Myself I do not think that this makes a moral difference. However, whether this view of mine is correct is irrelevant. For even if I am wrong, as already mentioned cases of disadvantageous treatment that do not amount to indirect discrimination because the victims do not form a socially salient group can also be the causal result of past direct discrimination. Perhaps this is quite unlikely to be the case but this is neither here nor there. The question that concerns me is whether the features, which distinguish cases of indirect discrimination from otherwise comparable cases of disadvantageous differential treatment, are wrong-making per se.

In response, it might be argued that such forms of disadvantages are morally wrong too. Hence, while indirect discrimination may not be worse than comparable cases of differential disadvantages falling on non-socially salient groups, they are both worse than otherwise comparable cases where the discriminatees are worse off than other socially salient groups and this inequality would not have occurred in the absence of past or present direct discrimination by the discriminating group against the discriminatees.

One thing that might indicate that the former kind of inequality is worse is that, sometimes, it will involve an additional and symbolic harm, because the relevant rules that are indirectly discriminatory symbolizes the direct discrimination, which the disadvantaged group was or is presently subjected to. While this is a valid observation it does not show that indirect discrimination is worse per se, since it is not part of the definition of indirect discrimination that it involves any symbolic harms. Indirect discrimination may take place without anyone noticing, in which case it is unlikely to involve any symbolic harm. Also, desert prioritarianism seems able to deal with symbolic harms. On this view, these matter insofar as they affect anyone's well-being and do not if they do not.

Second, assuming that any act of direct discrimination as such is unjust—something which I grant in this paragraph for the sake of argument—it might be said that the discriminators have a stronger obligation to their victims to eliminate the relevant disadvantages than they would have had, had these disadvantages not resulted from the discriminator's unjust—past or present—acts towards them. If this is the underlying reason why indirect discrimination identifies an additional wrong, then indirect discrimination is not morally worse than disadvantages that result from unjust acts etc. other than direct discrimination. Still, assuming disadvantages can arise in the absence of unjust acts this would be a substantial point. So the question then is: is this a valid point?

Note first that if this suggestion should explain the wrongness of paradigm cases of indirect discrimination, it will have to be group-focused. For instance, rules that disadvantage African Americans may be indirectly discriminatory even if it were the case that no presently existing African American had been subjected to direct discrimination by any presently existing non–African American. In this scenario, no presently existing non–African American would

have a stronger obligation to any African American to eliminate the relevant disadvantages. Hence, to explain the existence of such obligations on the basis of an obligation to remedy the bad effects of one's past wrongdoing, we would have to claim that people would have such a duty, not in virtue of their individual actions, but in virtue of their being members of a certain group whose past members acted unjustly. This is a problem. For justice focuses on individuals, not groups.

Obviously, not everyone agrees that justice is only concerned with individuals, but here is what I think is a pretty decisive counterexample to group focused views. Suppose that in the past Protestants engaged in direct discrimination against Catholics. There are a number of rules in place that involve indirect discrimination against Catholics. Suppose, finally, that since the old days of direct religious discrimination all Protestants have converted to Catholicism and all Catholics have converted to Protestantism. Intuitively, present indirect discrimination against Catholics here does not identify any additional group-related wrong. After all, the individuals who suffer the disadvantages from indirect discrimination are the very same individuals who in the past directly discriminated against other members of the community and the disadvantages that they experience *qua* Catholics would not have occurred had they not done so. However, if groups matter from the point of view of justice, we must—absurdly—say otherwise. On that assumption, the relevant fact is that the group of Catholics is disadvantaged and that it would have been so in the absence of past discrimination against Catholics.

It might be replied that it has some plausibility that an individual who acted unjustly towards another has a stronger duty to eliminate disadvantages that result from this act than other kinds of disadvantages. Hence, one might respond to the present argument simply by revising (v) in the definition of indirect discrimination (see chapter 2):

(v*) The disadvantages for present individuals belonging to G referred to in (ii*) would not have occurred in the absence of past or present direct discrimination by the very same individuals being agents of the policy or act in question against these very same individuals (*the individualized causal condition*).

So revised indirect discrimination may for all I have said involve an additional wrong. However, the revision also comes with a very high cost. For in almost no case of indirect discrimination, do people bother to show—let alone, believe that they ought to bother to show—that those individuals who are disadvantaged would not have been disadvantaged in the absence of direct discrimination against them by those whom they are disadvantaged relative to.

In conclude that indirect discrimination is not unjust in ways that otherwise comparable cases of disadvantageous, differential treatment are not. Moreover, this claim is compatible with harm-based accounts and more specifically so

with the desert-prioritarian account. I have not claimed that it is a reason to favor either of these accounts over mental state- or objective meaning based-accounts as these would appear to have similar implications.

9. Conclusion

Let me briefly summarize the main points made in chapters 4 to 6. First, it is not the case that group discrimination is morally wrong per se. This is so on none of the accounts that I have discussed. As noted in chapter 4.1, it is not surprising that this conclusion follows from desert prioritarianism, because on this view only acts that fail to maximize the morally weighted sum of benefits are morally wrong per se. However, it is significant that the claim that discrimination is not morally wrong per se would appear to be robust—that is, it seems to be consistent with various moral theories and with the different common and, presumably, deontological accounts of what makes discrimination morally wrong. It is significant, in part because one standard objection to consequentialist views is that they imply that, in principle, discrimination might not be morally wrong; in part because it conflicts with the untutored common sense view. Probably, the latter conflict arises even if we bear in mind that I have allowed that we sometimes employ a moralized concept of discrimination (see chapter 1.4) such that if something is discrimination in this sense, it is morally objectionable.

Second, there are several kinds of discrimination, and while some are much more morally problematic than others some seem not to be morally problematic at all. This claim is consistent with strong moral condemnation of standard cases of discrimination. In fact, it is even consistent with the view that we tend to underestimate the seriousness of the moral wrongness involved in some standard cases, such as racial discrimination and discrimination against women.

Third, I have argued that in those cases where discrimination is wrong, mental-state-based and objective-meaning-based accounts offer unsatisfactory accounts thereof.

Finally, I have proposed that the harm-based and, more specifically, the desert-prioritarian accounts offer the best account of the wrongness of discrimination. These accounts may not offer a satisfactory, complete account of the sources of wrongness of discrimination and I have indicated that unfairness may play a role as well. So in this sense I am open to a pluralistic account of the wrongness of discrimination.

It may seem somewhat unsatisfactory to end at a conclusion as open as this. However, while the present topic is the wrongness of discrimination, this topic cannot be addressed in a way that circumvents entirely the more general question of what makes acts morally wrong. This is the core question in normative ethics and a daunting one too. Even those who would have expected a firm,

monistic answer to the question of the wrongness of discrimination might think that the present discussion is helpful in that it identifies a number of crucial issues at stake. I shall end this chapter by pointing to one issue in particular. Consider the following argument, modeled on the leveling-down objection to egalitarianism:[63]

(1) Possibly, there are cases of discrimination that harm no one and benefit some (or even everyone).
(2) Necessarily, if an act or omission is pro tanto unjust or pro tanto morally wrong, there is someone for whom it would have been better in some respect had this act or omission not taken place.
(3) Therefore, possibly, there are cases of discrimination that are neither pro tanto unjust nor pro tanto morally wrong.

The first premise of this argument is surely uncontroversial. The possibility of benevolent discrimination where someone acts supererogatorily to benefit others—e.g., at serious risk to himself he enters a burning building several times to save those trapped and succeeds in so doing, but saves his co-religionists before he saves non-co-religionists as a result of which they suffer burns (it was unavoidable that someone would be burned)—shows that one can discriminate against someone even if one benefits him relative to how well off he would otherwise have been, to how well off he already is, or whatever. More generally, this follows from the fact that discrimination is—like the ideal of equality—essentially comparative.

The second premise is attractive, but controversial. It is the deontic version of Temkin's Slogan, i.e., "One situation cannot be worse (or better) than another if there is no one for whom it is worse (or better)."[64] Its attraction derives partly from the fact that it seems to underlie a wide range of what looks to be sound moral judgments and arguments in moral philosophy. In fact, it underlies many of the cases Temkin appeals to in order to show the attraction of his Slogan—a slogan, which, of course, he rejects.[65] For instance, Nozick's celebrated Wilt Chamberlain argument for holding that no one can "complain on grounds of justice" about voluntary exchanges draws on the assumption that there is no one for whom such exchanges are worse.[66]

Even if we ultimately reject the Deontic Slogan, it is still useful to bring it up here. For, first, to claim that discrimination is wrong per se one must deny the Deontic Slogan. So, perhaps, the Deontic Slogan brings out a crucial underlying

[63] Parfit (1998), "Equality and Priority."

[64] Temkin (1993), *Inequality*, p. 248.

[65] See Temkin (1993), *Inequality*, pp. 249–255 and Kasper Lippert-Rasmussen (2006), "The Insignificance of the Distinction between Telic and Deontic Egalitarianism," in NilsHoltug and KasperLippert-Rasmussen (eds.), *Egalitarianism: New Essays on the Nature and Value of Equality* (Oxford: Oxford University Press), pp. 101–124.

[66] Robert Nozick (1974), *Anarchy, State, and Utopia* (New York: Basic Books), p. 161.

disagreement between those who think discrimination is wrong per se and those who do not.

Second, some people, such as desert prioritarians, who are persuaded by the leveling-down argument that equality is not wrong per se consider most cases of inequality morally wrong. Similarly, it is possible to accept the Deontic Slogan, deny that discrimination is wrong per se, and yet maintain that most actual cases of discrimination are morally wrong. For instance, it is open to someone who accepts the Deontic Slogan to say that, actually, discrimination is often wrong, not because it is wrong as such, but because it often involves treating people badly, e.g., by not giving them adequate opportunities or by humiliating them.

Appendix: Moreau on Deliberative Freedom and Discrimination

In a recent, highly interesting article, Sophia Moreau offers "an account of what discrimination involves and why it is unjust that takes seriously a common, but not often discussed, feature of antidiscrimination laws."[67] Given her focus on making sense of discrimination law, it is not her aim as such to explain what makes discrimination morally wrong, but, at the least, her article implicitly provides some such account, to wit, that a discriminator reduces the deliberative freedom of the discriminatee and thereby "injures" her "interests."[68] To this extent, her account can be seen as a particular version of a harm-based account of the wrongness of discrimination. However, I doubt that it provides a satisfactory explanation of the wrongness of discrimination and in this appendix I explain why.

Deliberative freedoms are "freedoms to have our decisions about how to live insulated from the effects of normatively extraneous features of us, such as our skin color or gender," where normatively extraneous features are "traits whose costs we ought not to have to factor into our decisions in" the particular contexts where we enjoy deliberative freedom.[69] Antidiscrimination law seeks to give us deliberative freedoms "by preventing our employers, service providers, landlords, and others from acting in ways that deny us opportunities because of [normatively extraneous] traits, so that when we deliberate about such things as where to work and where to live, we do not have to think about these traits as costs."[70] We do not, Moreau thinks, have an entitlement to such freedoms in all spheres of life, e.g., we do not have it in our love lives, but we have it in the

[67] Moreau (2010), "What Is Discrimination?" p. 145.

[68] Moreau (2010), "What Is Discrimination?" p. 147. Presumably, she thinks that her aim of "developing a plausible account of discrimination laws" (p. 152)—is achieved, if she shows that it, and its core features, are morally justified.

[69] Moreau (2010), "What Is Discrimination?" p. 155.

[70] Moreau (2010), "What Is Discrimination?" p. 147.

spheres to which anti-discrimination law applies, e.g., "the provision of goods and services, rental housing, and employment."[71] Moreover, the relevant entitlement is an entitlement that we have independently of whether others enjoy the relevant deliberative freedoms.

A deliberative freedom has a subjective as well as an objective dimension: "in order for me to have a particular deliberative freedom with respect to a certain decision, it has to be true not just that I believe I can make that decision without having to worry about pressures from a certain extraneous trait, but that I really am free from those pressures."[72]

I agree with Moreau that, generally, it is very desirable, morally speaking, that we are able "to deliberate about and decide how to live in a way that is insulated from" certain "pressures stemming from extraneous traits of ours."[73] Clearly, it can make people's lives worse if they know that when applying for a certain job in their deliberations they will have to face the fact that they have a disfavored race or gender, or they simply do not have the option of getting the job for which they would like to apply. This much is common ground between Moreau's account and standard well-being-focused, harm-based accounts. However, Moreau seems to think that she goes beyond such an account and, to the extent that she does so, I find her account less convincing.

First, there are cases of wrongful discrimination that do not involve any restriction of deliberative freedoms as the present and the following two objections show. The notion of "costs" is ambiguous between costs incurred through choosing one option rather than another—opportunity costs, as economists

[71] Moreau (2010), "What Is Discrimination?" p. 143. Moreau thinks that many different kinds of considerations bear on which traits are normatively extraneous, in which contexts they are so, and how much moral weight they have when compared to conflicting interests. Accordingly, she does not offer "a single reductive explanation of the wrong of discrimination—that is, an explanation that traces the wrong of discrimination to some further single kind of normative fact that is operative in all cases" (p. 157). While I do not think this incompleteness makes her theory impossible to assess—desert prioritarianism as expounded in this chapter does not amount to a complete theory of the wrongness of discrimination—obviously, it makes it harder to tell when one faces a counterexample (or one that supports her account). For instance, if we do not think that a 67-year-old is wrongfully discriminated against when his application for a job as a commercial airline pilot is rejected on grounds of age, or when, due to affirmative action, an Asian applicant is turned down by an Ivy League university but would have been admitted had she been African American, we cannot readily tell if these are counterexamples to Moreau's account—she mentions age and race as normatively extraneous features in some contexts—or just show that in the particular contexts at stake they are not normatively extraneous features or that while they are, the relevant deliberative freedoms are outweighed by other relevant concerns, e.g., safety or social cohesion.

[72] Moreau (2010), "What Is Discrimination?" p. 149. A minor quibble here is whether the subjective dimension is better captured in terms of what one has reason to believe (as opposed to what one actually believes), because, arguably, my entitlement to deliberative freedom is not violated by someone, who conveys certain job-relevant information to me, and I, because of *my* unreasonability, take it to imply that I must consider my age or nationality a cost, when applying for the relevant job.

[73] Moreau (2010), "What Is Discrimination?" p. 147.

would put it—and between costs incurred whatever one chooses. It is the former kind of costs that are relevant to deliberative freedom, but, *pace* Moreau's account, the latter are not irrelevant to the moral wrongness of discrimination. Suppose 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. There is a sense in which women "do not have to think about [their sex as a cost]." For it is not the case that if they make one choice rather than another, they will be penalized because of gender, e.g., they will avoid informal social sanctions if, and only if, they choose to stay home and look after the kids rather than pursue a career. In the example I imagine, *whatever* choice they make they are penalized because of their gender. Hence, they can make decisions in "a way that is insulated from the pressures or burdens caused by certain extraneous traits."[74] Not just might it be true that agents believe that they "can make that decision without having to worry about pressures from a certain extraneous trait," but it is also true that, objectively speaking, they are "free from those pressures" since as a matter of fact they can do nothing to avoid them.[75] They are subjectively as well as objectively deliberatively irrelevant. Precisely for this reason, it is misleading to say that they are denied any *deliberative* freedom on account of their option set being worse than that of men. Their disadvantage simply consists in their set of opportunities being worse. Hence, Moreau's account is unable to explain the wrongness of the discrimination in the case imagined.[76]

Second, Moreau thinks that an important advantage of her account is that it is able to explain why discrimination harms individual discriminatees, and not just the groups to which they belong. However, there are possible cases of which this is not true. Suppose an employer directly discriminates against women in hiring, but also indirectly discriminates against men in a way that exactly counterbalances her direct discrimination against women such that, all in all, applicants will never in the relevant sense be accepted or rejected *because* of their gender. These facts are unknown to the employer. Suppose also that the 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. This employer unlawfully, directly discriminates and, on some accounts,

[74] Moreau (2010), "What Is Discrimination?" p. 149.

[75] Moreau (2010), "What Is Discrimination?" p. 149.

[76] Another way to show that Moreau overemphasizes agency and deliberative freedoms in her account of the wrongness of discrimination is the following: suppose a number of patients are hospitalized in a permanently paralyzed condition and, let us assume, as a result have no deliberative freedom. They do, however, suffer pain and the doctors can affect the level of pain they endure. Sexist doctors do so giving priority to men over women. Presumably, this would amount to wrongful discrimination and yet this cannot be accounted for in terms of how doctors restrict anyone's deliberative freedom.

engages in direct, possibly wrongful discrimination, even if she restricts no one's deliberative freedom.

In view of this objection, one might suggest that we are entitled to a more general deliberative freedom than the one Moreau ascribes to us: say, the freedom "to have our decisions about how to live insulated from the effects, as well as counterbalanced effects, of normatively extraneous features of us, such as our skin color or gender." While one accommodates the present counterexample if one revises Moreau's account in this way, this move comes with a cost, namely, that it is less plausible that we have an entitlement to deliberative freedom as extensive as that.[77]

Third, suppose a certain employer is an incompetent sexist—he wants to exclude women, but, as all potential applicants are aware of, invariably he fails to do so, e.g., in his job adverts he writes that only men need apply and nevertheless he ends up hiring women, because often he mistakes the name of a man with the name of a woman and vice versa. Presumably, such an employer might engage in unlawful as well as morally wrongful, direct discrimination—suppose the outcome would be better if he did not try to exclude women—even if he does not restrict anyone's deliberative freedom. Admittedly, he does have a "low opinion" of women, but in itself—i.e., independently of how this opinion worsens the opportunities available to female applicants—that does not violate any deliberative freedom entitlement of ours. The mere fact that we correctly believe that someone with whom we deal is motivated by a "low opinion . . . of our race or gender" does not suffice for it not to be the case that "we can decide where to work and where to live without having to think about the low opinion others may have of our race or our gender."[78]

Fourth, there are cases of discrimination that restrict people's deliberative freedoms and yet appear not to be wrongful. Suppose that the only way to give everyone better options is through reducing people's deliberative freedom, e.g., either neither men nor women have to consider their gender when applying for a job, in which case everyone faces two options yielding, say, 10 and 12, respectively, or they will have to do so, in which case they face two options yielding 15 and 17, respectively, with men facing a lesser chance of enjoying the best outcome in the latter case due to sex discrimination. On Moreau's view, ensuring that we face the latter set of options deprives us of our deliberative freedom and, thus, amounts to wrongful discrimination. Yet, this would be a case where everyone benefits overall from discrimination. To many, but not all—recall the discussion in chapter 6.9—this suggests that Moreau's account is problematic.

---

[77] Similarly, it is much less plausible to think that we have a deliberative freedom, which is violated in cases of indirect discrimination. As Moreau (2010), "What Is Discrimination?" p. 159, acknowledges, 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.

[78] Moreau (2010), "What Is Discrimination?" p. 155.

Fifth, Moreau's account does not explain how discrimination differs from nepotism or from idiosyncratic disadvantageous differential treatment. If an employer openly treats nonfamily applicants or applicants with green eyes disadvantageously, e.g., will not hire such an applicant if she is not at least 20% better qualified than the best family or non-green-eyed applicant, the employer would restrict these people's deliberative freedoms exactly as much as if the employer treated women or African Americans equally disadvantageously, e.g., by not being willing to hire such an applicant unless she is at least 20% better qualified than the best male or non–African American applicant. Yet, the latter cases seem more wrong for reasons not captured by Moreau's account.

Consider finally an employer who actually gives men and women the same opportunities, but who either gives the impression that she does not, or fails to counteract the independently generated and false impression that she does not. This employer canvasses a reduction of the deliberative freedom of applicants no less than an employer who sex discriminates, but either gives the false impression that she does not sex discriminate, or fails to counteract the independently generated and false impression that she does not. Suppose the same applicants apply for jobs in the two cases. Since both employers canvass the reduction of the deliberative freedom of applicants—the first by not making sure that the objective condition of deliberative freedom is satisfied, and the second by not making sure that the subjective condition is satisfied—Moreau's account is unable to explain why the latter employer acts more wrongly. Admittedly, like the previous objection, this last one does not show that her account is false, but it shows that it stands in need of being supplemented.

This page intentionally left blank

{ PART III }

Neutralizing Discrimination

This page intentionally left blank

{ 7 }

Discrimination and the Aim of Proportional Representation

1. Introduction

The first two parts of the book looked at the concept of discrimination and the moral qualities of discrimination in general. This last part focuses on a more applied question, namely, what we should do about discrimination, i.e., how we ought to seek to eliminate it or to remedy its effects. Obviously, our answer to the question of what makes discrimination wrong will affect our answer to the question about what should be done to eliminate discrimination or its effects, whenever such actions should be taken. However, there is by no means any direct inference from claims about what makes discrimination wrong to what we should do about it because, as we shall see in this chapter as well the next one, it is far from clear what eliminating the effects of discrimination amounts to. The mere fact that it is wrong to discriminate does not imply that others are morally required or for that matter even morally permitted to prevent others from discriminating.

Many organizations, private companies, and public institutions are committed to certain representational aims as regards the composition of their workforce. Usually, the aim is to ensure that the composition of the relevant workforce reflects the composition of the population of the area from which it is drawn. For instance, in its Equal Opportunity in Employment Policy Statement, Birmingham City Council writes of its intention that "its workforce, at all levels, should reflect the composition of the City's population." It appears from the statement that the compositional variation motivating this intention includes "gender, race, disability, color, ethnic and national origin, nationality, sexuality, marital status, responsibility for dependants, religion, trade union activity and age (up to 65)."[1]

---

[1] http://www.birmingham.gov.uk/GenerateContent?CONTENT_ITEM_ID=1199&CONTENT_ITEM_TYPE=0&MENU_ID=0 (accessed June 30, 2005).

This kind of proportional representation has been defended both as an instrumental and as a noninstrumental good. Anne Phillips, for instance, believes that in the case of gender, the absence of significant differences between the sexes in "income or time or roles" is an intrinsically desirable goal.[2] Moreover, it is "a safe bet" that any such differences will reflect inequality in initial opportunities.[3] In a spirit similar to that of Phillips's latter claim, Iris Marion Young believes that where "all positions of high status, income, and decision-making power" are not "distributed in comparable numbers to women and men," there is a strong prima facie case for the existence of injustice against (or oppression of) the underrepresented sex.[4] Thomas Pogge is critical of the anonymity condition on principles of distributive justice—the condition says that changes in which persons occupy which social positions makes no difference justice-wise—and argues that a "morally plausible criterion [of justice] would have to take account of" cases in which "salient groups" are overrepresented among those suffering hardship, especially if social factors play a large role in the genesis of this overrepresentation.[5] Like Phillips, Pogge ascribes noninstrumental significance to proportional representation.

Many different motivations lie behind these representational aims. In this chapter, I want to focus exclusively on the motivational idea that proportional representation is a way of nullifying the repressive effects of discrimination

---

[2] This contrasts with the view on political representation expounded in Phillips's excellent study of the politics of presence. Here she contends that "serious claims to guaranteed representation have to be grounded in something more than a statistical mismatch between voters and those who represent them. There has to be an additional analysis of existing structures of exclusion" (Anne Phillips [1995], *The Politics of Presence* (Oxford: Clarendon Press), p. 171; p. 64).

[3] Anne Phillips (2004), "Defending Equality of Outcome," *Journal of Political Philosophy* 12, 1–19, p. 13.

[4] Iris Marion Young (1990), *Justice and the Politics of Difference* (Princeton, NJ: Princeton University Press), p. 29. On Barry's reading of Young, she claims that disproportionate representation in itself "*constitutes* discrimination"; see Brian Barry (2001), *Culture and Equality* (Cambridge: Polity Press), p. 93. I am unable to find good textual evidence in favor of this interpretation. In a later piece, Young writes that if we find that a certain group is worse off on a large number of parameters, then "we are entitled to say that members of this group *probably* [my emphasis] suffer injustice" provided that "we can tell a *plausible structural story* that accounts for the production of the patterns"; Iris Marion Young (2001), "Equality of Whom? Social Groups and Judgements of Injustice," *Journal of Political Philosophy* 9, 1–18, p. 16. Assuming she thinks discrimination is unjust—but see appendix 1, chapter 1—this shows that by 2001 Young did not hold the view ascribed to her by Barry in *Justice*. Barry, of course, will still disagree that the mere existence of a structural explanation, as opposed to a structural explanation of a certain sort, of why a group ends up worse off should entitle us to conclude that, e.g., underrepresentation in attractive job categories constitutes an injustice.

[5] Thomas Pogge (2002), *World Poverty and Human Rights* (Polity Press: Cambridge University Press), p. 44. Pogge's point presupposes that social positions are not defined, wholly or in part, by membership of socially salient groups.

on underrepresented groups.[6] This assumption is detectable in the statement of Birmingham City Council's equal opportunity policy above. It says that the council sees the elimination of discrimination and the reduction of the effects of past discrimination as a means of ensuring that its workforce reflects "the composition of the City's population." This makes best sense if we take the council to believe that discrimination, whether past or present, is what causes the city's workforce to be unreflective of its population at present.

My primary aim is to arrive at a formulation of representational aims that best meets the widely recognized need to nullify the effects of discrimination. Unsurprisingly, matters turn out to be much more complicated than is usually assumed in public policy statements. This is partly because policy statements about representational aims must reflect a wide range of concerns—notably, pragmatic concerns about the feasibility of, and costs involved in, implementing different policy aims—as well as views about which of many different, potentially conflicting aims should enjoy political priority. For present purposes, I want to disregard such concerns about implementation in order to ask the prior question: What does the desire to nullify the effects of discrimination involve? Although my analysis will certainly have a bearing on policy aims, it is not meant to state them. So, quite compatibly with my main line of argument, the views about representation, which I criticize below as failing to capture our desire to nullify the effects of discrimination, may be acceptable, in some contexts, when construed as stating policy aims. Accordingly, I see this chapter, not as a criticism of policies involving representational aims, but first and foremost as a reflection on one of the putative values underlying these statements whose exact bearing on the issue is insufficiently articulated at present.

My main claims are, first, that the achievement of representational aims, in its most plausible form, is reducible to a concern to eliminate the effects of wrongful discrimination on individuals; and, second, that this very important concern is compatible with the representation of various groups diverging widely from their share of the overall population. Obviously, these claims are

[6] Cf. Anderson (2010), *The Imperative of Integration* (Princeton, NJ: Princeton University Press), p. 136; Sabbagh (2007), *Equality and Transparency: A Strategic Perspective on Affirmative Action in American Law* (New York: Palgrave Macmillan), pp. 49–85. This focus excludes the following alternative motivations (there may be others): to reduce the level of overall discrimination, to compensate for past discrimination, to ensure efficient use of talents, to send a symbolic message of inclusion, to ensure representation of all viewpoints, and to increase the social cohesion of society. While these aims are often thought to motivate proportional representation, this may not be the best way to construe them, e.g., the aim of compensation might be fulfilled even in the absence of proportional representation for the reasons given in this chapter.

neutral with regard to whether the proportional representation of groups, numerically speaking, serves other desiderata and the latter, thus, is abundantly justified for this reason. However, even if this were the case, we would still want to know how the aim to nullify the effects of discrimination shapes representational aims.

In the three previous chapters, I reviewed three different accounts of what makes discrimination wrong. For present purposes, however, we can set aside differences between these accounts. The problem I address here arises for all views about wrongful discrimination: namely, how we should think of representational aims as a way of nullifying the effects of wrongful discrimination. Presumably, people could agree on this issue even if they disagree over the conditions under which discrimination is wrongful. Similarly, philosophers of law may agree on what it is for a crime to be nullified through punishment and compensation even if they disagree over the kinds of act that ought to be prohibited by law.

Section 2 rejects the simple view that suitable representation of a group is representation, which is proportional to the group's share of the relevant total population. The simple view is insensitive to nondiscrimination-related causes of under- or overrepresentation, e.g., innocent differences in ambitions across groups. Sections 3 and 4 explore a number of counterfactual views according to which suitable representation depends on how well the relevant group would have been represented in the absence of discrimination. However, different relevant counterfactual worlds—e.g., worlds that differ in terms of whether only the relevant group or all groups have not been subjected to discrimination, or in terms of which people exist—are offered as relevant baselines and it is suggested that different purposes will lead us to adopt different baselines. Sections 5 and 6 contrast group-focused and individualized representational aims arguing in favor of the latter view, which, however, is quite different from the view subscribed to many who are concerned with suitable representation.

2. The Simple View and Ambition-Sensitivity

What sort of representation nullifies the effects of discrimination? The *Simple View* says that a group is suitably represented in a particular category of employment (or, in short-hand, is suitably represented) if, and only if, the proportion of people from this group employed in the relevant category is the same as the proportion of these people in the total population. I do not mean to suggest that the Simple View is most people's considered view on the present matter. This view is, however, a natural starting point for our inquiry. It underpins many official statements of representational aims. For instance, Birmingham City Council presumably does not think that the composition of its

workforce reflects the composition of Birmingham's population unless 50% are women.[7]

To apply the Simple View, one needs to address the following questions. (a) What groups are to be suitably represented? (b) What are the categories of employment within which there should be suitable representation? (c) What is to be taken as the total population? Regarding (a), lists such as the one offered by Birmingham City Council (see section 1) cannot provide a satisfactory answer alone. For we can ask, why are the items on the list included, and why are other items, such as social and economic background, ignored? Given that we are concerned to nullify the effects of discrimination, the groups to be included are those that are either subjected to discrimination or discriminate or benefit from discrimination.[8] In practical policy-making, it may of course be sensible to focus on large groups that suffer from the worst forms of discrimination and to ignore the less urgent claims of other groups.

Turning to (b), many would want to employ multilevel as well as criss-crossing accounts of the categories of employment calling for suitable representation. Take the case of women. It might be held that women should constitute roughly half the workforce. But, obviously, few defenders of the Simple View will consider this representational aim sufficient. If half the workforce is made up of women, but women are heavily overrepresented in jobs that involve nursing, care, and children, few defenders of the Simple View will be content.[9] So here we have a case of a multilevel representational aim. This aim is bound to involve criss-crossing dimensions as well. For suppose that we have equal representation of men and women across different professions, e.g., those that involve nursing, care, and children and those that do not, but that women are less well represented in well-paid jobs within the different professions. Surely this will not satisfy many defenders of the Simple View. In fact, then, many different kinds of employment categories are relevant to the Simple View. Let us call this approach to job individuation the "fine-grained" approach. Alternatively, we might construct an index comprising the benefits involved in jobs—e.g., income, prestige, power, freedom, self-respect, or well-being—and

[7] Sometimes representational aims are formulated in terms that differ slightly from the Simple View. This occurs, for instance, when representation is defined in such a way that it is sensitive to the different age profiles or qualification profiles of distinct ethnic groups. So, for instance, if we want to know if African Americans are suitably represented among New York City doctors, we might think that this depends on whether the proportion of African American doctors corresponds to the proportion of people in New York City with the relevant age and a degree from medical school who are African American (http://www.nyc.gov/html/dcas/html/eeopol.html [accessed November 20, 2004]). Since the criticism of the Simple View that I shall put forward below applies to these slightly different representational aims as well, I shall, to avoid unnecessary diversions, stay with the Simple View alone.

[8] I ignore here the complication involved in promiscuous views about groups; see George Sher (1999), "Diversity," *Philosophy and Public Affairs* 29, 85–104, pp. 89–93.

[9] For a feminist critique of feminized, welfare-state care work, see Nancy Fraser (1994), "After the Family Wage: Gender Equity and the Welfare State," *Political Theory* 22, 591–618.

then claim that, for jobs ranked equally on this index, suitable representation requires groups to be suitably represented (or, if they are not for some values of the index, counter-balancing by under- or overrepresentation in other categories of employment). Call this the "coarse-grained" approach to job individuation. My discussion below applies to fine- as well as coarse-grained approaches, although it appears more persuasive when directed at fine-grained approaches.

Consider finally (c): what is the total population? This question introduces many complications, too. One question is whether the relevant boundary is geographical. Should we, for example, include all and only people living within the jurisdiction of the council, or should we focus on people affected by what the council decides? The latter might not include all people living in Birmingham and might include some people living outside. Whether the criterion is geographical or not, we also need to decide whether one is a member of the total population irrespective of one's nationality. So, should the composition reflect the segment of non-British nationals living, temporarily or not, in Birmingham? How we answer these and other similar questions helps to determine whether a certain group is under-represented. As with (b), the answer to (c) might be quite complex.

Complex representational aims are likely to be harder to achieve than simple aims. For one thing, in complex cases we need to know how to rank the representational aims. However, (a)–(c) appear not to point to any problem with the Simple View that is decisive or distinctive. To some extent, the problems that arise are located at the level of implementation. Hence, I shall set them aside and instead focus on an apparently decisive objection to the Simple View.

Suppose we want our theory of representation to be ambition-sensitive as well as endowment-insensitive, so that representational aims respect the notion that "[p]eople's fate should depend on their ambitions (in the broad sense of goals and projects about life), but should not depend on their natural and social endowment (the circumstances in which they pursue their ambitions)."[10] In this view, if two people with equal natural and social endowments choose differently when it comes to how much they work or what kinds of work they do, the resulting differences between them should not be nullified. Suppose, moreover, that people can differ systematically with regard to such decisions for reasons unconnected with discrimination. This seems to be possible. After all, different groups of middle-class, heterosexual, white males systematically make very different choices, yet they may not be differentially affected by discrimination. Based on these two suppositions, we should reject the Simple View. Suitable

[10] Here I borrow from work, especially by Ronald Dworkin, on equality. Many egalitarians nowadays think that a theory of equality should be both ambition- and endowment-sensitive. See Will Kymlicka (2002), *Contemporary Political Philosophy* (Oxford: Oxford University Press), p. 74; Ronald Dworkin (2000), *Sovereign Virtue* (Cambridge, MA: Harvard University Press), p. 89. Barry similarly rejects the Simple View, because he thinks it is incompatible with his interpretation of liberal equality of opportunity; see Barry (2001), *Culture and Equality*, pp. 90–98.

representation should be insensitive to discrimination, but it should also be sensitive to nondiscrimination-related differences in people's ambitions. If different groups have different ambitions for such reasons, their proportion in a certain job category, or for that matter in the workforce as a whole, may diverge from their proportion of the total population. So the main conclusion of this section is that a theory of suitable representation should be ambition-sensitive, and that since the Simple View obviously is not, it should be rejected.[11]

It might be replied that it is one thing for justice to accommodate differences in the ambitions of different individuals, but quite another for it to accommodate the systematic differences between the ambitions of different groups, e.g., those dividing African Americans and European Americans or men and women. In the latter kind of case there is bound to be a structural explanation of the difference, something implying that the underrepresentation of a group in attractive job categories is the upshot of wrongful discrimination.[12]

Underrepresentation of the kinds that group defenders of proportional representation care about typically reflect discrimination or some other kind of injustice, but in my view the mere fact that systematic differences in ambition have a structural explanation does not render them unjust. First, one can offer examples in which minorities are underrepresented, and intuitively this does not reflect wrongful discrimination against this group by nonmembers. Consider the Amish. Their commitment to simple ways of life, and to religious as opposed to secular education, seems sufficient to explain why as a group they are underrepresented in high-paying jobs. The mere fact that the systematic differences between Amish and non-Amish Americans have a structural explanation—e.g., in terms of differential expectations embodied in different roles or norms and reproduced through their differential upbringing—seems irrelevant to the question whether the underrepresentation of Amish people in high-paying jobs is due to wrongful discrimination against them by non-Amish people.[13] Different groups may have different social and cultural structures, and these structural differences may explain why, at an aggregate level, their members make very different choices. If these groups have had, say, little interaction with one another, or if they have interacted but neither group dominates the other, nor did so in the past, the mere fact that there is a structural explanation

[11] Cf. Sabbagh (2007) *Equality and Transparency*, pp. 28–30.

[12] Obviously, injustice need not involve wrongful discrimination. We can ignore this for present purposes, but I want to register my belief that, for reasons that appear below, systematic differences in ambition need not reflect injustices of any kind.

[13] Similarly, in Israel people sometimes say that the student body in universities should mirror Israeli society minus the ultra-orthodox who prefer studying in Yeshivas rather than universities. Also, recall my example of Indians and Chinese in chapter 2. Admittedly, this view and the view regarding the underrepresentation of Amish people presuppose that there is no requirement to accommodate the Jewish ultra-orthodox and the Amish ways of life. Presumably, many will think these cases differ from the case of disabled people, where such a requirement of accommodation does apply.

of the groups' different patterns of choice at an aggregate level surely leaves it an open question whether this reflects discrimination of one group against the other. With these crucial qualifications, we should accept Nathan Glazer's claim that "racial and ethnic communities have expressed themselves in occupations and work groups. Distinctive histories have channeled ethnic and racial groups into one kind of work or another, and this is the origin of many of the 'unrepresentative' work distributions we see."[14]

Second, setting aside the details of cases, one can see that in many theories of justice, structurally induced differences in ambition are not unjust as such. Suppose, for instance, that we endorse the egalitarian ideal of equality of welfare and, hence, reject the view that differences in choice can justify overall inequalities in welfare. It will remain true that achieving equality of welfare may require respecting differences between the choices of different groups with regard to work versus leisure and career. Consider next Rawls's principle of fair equality of opportunity. In this view, justice requires that people who are equally talented and equally motivated to reach a certain position should have an equal chance of reaching that position whatever their luck in the social lottery.[15] The principle does not exclude systematic differences across groups in respect of the ambition to become a high flier.[16] Consider, finally, the view of a critic of Dworkinian ambition-sensitivity to gender differences, Andrew Mason.[17] Despite his critique of blanket ambition-sensitivity, Mason believes that it may well be the case that women should bear the cost of their autonomous and authentic choices, even if their choices differ systematically from those of men, provided that this results from differences in socialization not involving any sexist norms.[18] One might, of course, disagree with Mason over details, but since autonomy and authenticity are not comparative concepts—i.e., whether one's choice is autonomous and authentic depends on its intrinsic qualities, not on how it affects one relative to the way in which comparable choices made by others affect them—there is no reason to deny that an upbringing not involving sexist norms could lead men and women to make

[14] Nathan Glazer (1975), *Affirmative Action* (New York: Basic Books), p. 203. Ezorsky points out that Glazer's point may have less or no relevance in relation to African Americans, but she accepts Glazer's general claim: see Gertrude Ezorsky (1991), *Racism and Affirmative Action: The Case for Affirmative Action* (Ithaca, NY: Cornell University Press), pp. 56–57.

[15] John Rawls (1971), *A Theory of Justice* (Oxford: Oxford University Press), p. 73.

[16] Rawls's principle of fair equality of opportunity fails to condemn cases in which the ambitions of some groups are stunted as the result of a bigoted and unfair process of socialization: see Richard Arneson (1999), "Against Rawlsian Equality of Opportunity," *Philosophical Studies* 93: 77–112, pp. 78–79. Neither Arneson nor Mason (see below) concludes that any process of socialization that leads to variation in ambitions across groups is unjust.

[17] Andrew Mason (2000), "Equality, Personal Responsibility, and Gender Socialisation," *Proceedings of the Aristotelian Society* 100, pp. 227–246.

[18] Mason (2000), "Equality," p. 243.

systematically different autonomous and authentic choices.[19] Hence, I conclude that systematic differences in ambitions across groups resulting from structural causes that were not generated by wrongful discrimination (or any other kind of injustice) could exist; indeed they may well do so. Later—see the last four paragraphs of section 4—I endorse a broader notion of ambition-sensitivity. This notion allows for cases in which suitable representation needs to reflect differences between the ambitions of members of different groups, even if these were caused by unjust discrimination, whether of a structural kind or not.

Before describing a way in which the Simple View can be revised to accommodate ambition-sensitivity, I want to respond briefly to the allegation that my criticism of the Simple View is politically dangerous (see also the introduction to this book). I allow that a minority may be numerically underrepresented in attractive segments of the labor market as result of something other than wrongful discrimination. It might be said that I thereby provide assistance to those who wish to treat the underrepresentation of some minorities as just rather than wrongfully discriminatory. This allegation is not a challenge to the truth of my claims; but because I challenge a premise thought to license the inference from numerical underrepresentation to unsuitable underrepresentation, the danger must be admitted. However, those who wish to make the inference will have to explain (among other things) why their view does not encourage a similarly erroneous tendency to assume that if a minority is over-represented in an attractive segment of the labor market, this manifests discrimination against the numerically underrepresented nonmembers of this minority. And indeed by rejecting the view that it is desirable for representation to be proportionate to the relevant group's share of the total population, one can accommodate a possibility that surely needs to be acknowledged, i.e., that a minority is numerically overrepresented and yet subjected to discrimination negatively affecting its representation. The Simple View eliminates this possibility. Numerical underrepresentation of one group logically implies numerical overrepresentation of the group of individuals who are not members of the underrepresented group, and whatever explains the underrepresentation of the former group also explains the overrepresentation of the latter. Hence, the overrepresentation of the latter, in the view I criticize here, seems necessarily to be due to discrimination against the underrepresented group. So by way of illustration, consider the case of numerical overrepresentation of Asian Americans among US graduate students, especially in the sciences. Various

[19] Men and women do not constitute two distinct cultural groups. Hence, it is less likely that, in the absence of wrongful discrimination, they will be differently represented in different attractive employment categories than it is that, in the absence of wrongful discrimination, different cultural groups will be differently represented in different attractive employment categories. This observation is entirely consistent with what I am saying here.

explanations of this have been suggested (e.g., that many post–World War II Asian immigrants have been professionals, and that many Asians attach honorific significance to children's educational achievement); but whatever is the correct explanation, surely, we should not conclude, from the mere fact that Asian Americans are not numerically underrepresented, that their representation is not negatively affected by discrimination against them; nor should we conclude that a suitable level of representation for Asian Americans (suitable, that is, from the point of view of eliminating the effects of discrimination) would be lower than the actual one.[20]

3. The Counterfactual, Holistic View

These problems with the Simple View are serious enough to prompt exploration of the quite different approach embodied in the *Counterfactual View*: a group is suitably represented if, and only if, the proportion of people from this group employed in the relevant category is identical to what it would have been had its representation not been affected by discrimination against or in favor of it.[21] In this view, how large a proportion of the total population a certain group comprises is *in itself* irrelevant to determining whether it is suitably represented: what matters in itself is the proportion in which it would have been represented had it not suffered representation-affecting discrimination. Of course, if we have reason to think that, in the absence of discrimination against a certain group, it would have been represented in the same proportion as it has in the total population, the latter will be relevant—not in itself, but as an indicator of suitable representation.

The Counterfactual View is consistent with the possibility that a group is subjected to discrimination without being objectionably underrepresented—something that would happen if the discrimination did not affect this group's representation within occupational categories and hence did not prevent it from being suitably represented. This accommodation seems desirable. Note

[20] Some would even argue, rightly or not, that any measures taken to reduce the numerical overrepresentation of Asian American students would constitute wrongful discrimination against them: see L. Ling-Chi Wang (1988), "Meritocracy and Diversity in Higher Education: Discrimination Against Asian Americans in the Post-*Bakke* Era," *The Urban Review* 20, pp. 189–209. A similar point could be made in relation to the "overrepresentation" of UK Hindus "in the professional and managerial categories"; see Sandra Fredman (2011), *Discrimination Law*, 2nd ed. (Oxford: Oxford University Press), p. 81.

[21] Being discriminated against is not desirable. Yet, it is logically possible for someone to have the ambition to be discriminated against, where that ambition is not itself a causal result of prior discrimination. I rejected the Simple View because it is insensitive to ambitions. Does it not follow that I must reject the Counterfactual View? While I concede the force of this challenge, I take the ambition to be discriminated against to be sufficiently outlandish for me to ignore it below.

also that the Simple View can be derived from the Counterfactual View given the additional, and as I have argued, implausible premise that if the representation of any group diverges from its proportion of the total population, then this divergence is due to discrimination for or against this group.

Unfortunately the Counterfactual View is flawed. To see why, suppose that overt discrimination against African Americans no longer involves status discrimination, but rather takes the form of trait discrimination. Thus it is not that employers are unwilling to hire African Americans as such. The situation is, instead, that they are unwilling to hire applicants whom they feel have a "belligerent, unassimilated" style and they are much more likely to assume that an African American man has a belligerent, unassimilated style than they are to make the same assumption about a European American or an African American woman.[22] Suppose that, as a result of non-African American employers acting on the basis of this differential perception, European Americans and African American women often get jobs that, in the absence of trait discrimination, would have gone to African American men. In this case, African American women might be better represented in the relevant employment categories (which, in this example, are unlikely to collect well-paid, prestigious, or powerful jobs) even though they are not directly affected in a negative way by this kind of discrimination. Generalizing from this example, we can see that in principle a group can be overrepresented in a certain category of employment, including categories of attractive work, even if its proportion has not been affected by discrimination (against or for it) provided that its representation has been boosted by discrimination against other groups.

The problem with the Counterfactual View seems to be that it focuses on one group's representation in isolation. It does not ask how discrimination against or in favor of other groups affects that group's representation. This suggests that we should revise the Counterfactual View as follows: a group is suitably represented if, and only if, the proportion of people from this group employed in the relevant category is what it would have been had the representation of this group not been affected by discrimination against or in favor of it *or any other group*. This view—the *Holistic, Counterfactual View*—meets the objection to its predecessor raised above, because in this objection I imagined that the representation of one group was affected positively by discrimination against another group.

Three complications should be mentioned. First, it is possible that a certain group is better represented in the category of attractive jobs as a result of discrimination against other groups, but that it is less well represented as a result of other, merely unfortunate (i.e., non-discrimination-related) circumstances.

[22] Joleen Kirschenman and Kathryn Neckerman (1991), " 'We'd Love to Hire Them, But . . . ': The Meaning of Race for Employers," in C. Jencks and P. Peterson (eds.), *The Urban Underclass* (Washington, DC: Brookings Institution Press), pp. 203–234.

In such a case discrimination serves as a compensation for unjust differences in circumstances. I can ignore this complication here. It simply reminds us that what interests us is suitable representation from the point of view of nondiscrimination, not suitable representation from the broader point of view of justice.

Second, I set aside here the issue of whether the phrase "all other groups" in the Holistic, Counterfactual View means all other groups within the relevant population (e.g., nationals), or all other groups in the population as a whole, and thus everyone living anywhere in the world. As attested by the history of emigration to the United States of members of religious groups suffering discrimination, or worse, at the hands of the majority populations in their homelands, the proportion of the total population that belongs to a particular religious group in a country either tolerant, or less intolerant, of it may be affected by discrimination against this religious group in other countries. So whether a group is appropriately represented in a country, as it were, nationally speaking, and whether it is appropriately represented, again, in that country, globally speaking, are separate matters.

Third, discrimination and individual ambition may combine to bring it about that a group suffers from discrimination, is underrepresented within a certain category, numerically speaking, and yet is overrepresented according to the Holistic, Counterfactual View. Suppose discrimination against a certain group makes it harder for its members to get a job in all categories of employment, but that the degree to which access for members of this group is reduced by discrimination varies between categories. If access to one particular category of employment, which members of this group would rarely apply for anyway (i.e., in the absence of discrimination), is less affected by discrimination than access to other categories, members of this group may be underrepresented in this particular category of employment, numerically speaking, and yet overrepresented according to the Holistic, Counterfactual View, because their representation in this category would have been even lower in the absence of discrimination.

4. Which Counterfactual Scenario?

The Holistic, Counterfactual View is attractive, but it needs to be clarified. However, as I shall show in this section, identifying the relevant baseline counterfactual situation can be tricky.

Discrimination can affect not just the situation of members of a group wanting to apply for a certain job, but also—positively or negatively—how many people belong to the discriminated group. Indeed, it may help to determine what social groups and individuals exist at all. Suppose women suffer from discrimination, and that many prospective parents use artificial fertilization

techniques and selective abortion to ensure that they have male offspring. This is what is presently taking place in some Asian countries: "On average, women round the world give birth to 106 baby boys for every 100 girls; more boys than girls are then lost in childhood. In China, the sex ratio for first births matches that average . . . ; but for every subsequent birth, the surplus of boys increases. This imbalance has been growing rapidly since 1979. In 1982, there were 107 boys aged under five for every 100 girls; in 1990, 110; in 1995, 118. 'Cumulatively,' says William Lavely, a demographer at the University of Washington, '8.7m females are missing . . . or roughly 5%. But from those born in the 1990s, approximately 10% are missing.' China is by no means the only Asian country where the ratio of boy to girl babies is on the rise: the same is true in South Korea and Taiwan. The rise has gathered pace since the mid-1980s, roughly the point when new technologies . . . capable of predicting the sex of a fetus with reasonable accuracy became available."[23] Suppose that, as a result of the application of such techniques, 55% of China's population will be male in the not-so-distant future. What would be a suitable representation of women in a certain category of employment in a case like this?

One answer is 50%. If women had never suffered discrimination, 50% of the total population would have been female and they would, we can suppose, have been equally interested in, as well as equally qualified for, the relevant jobs. To endorse a smaller representation would be to entrench the sexist structure of society, which led to the imbalance between the number of men and women in the first place. In this view—call it the *Global View*—the counterfactual scenario determining whether a certain group is suitably represented is, roughly, one in which no group ever suffers from discrimination which affects its representation, whether by affecting the relative sizes of the relevant groups or by affecting how well those who are members of these groups are represented. If the actual representation of the group being considered is what it would be in the counterfactual scenario, then, and only then, it is suitably represented. (The Global View, along with its five successors described below, is stated in a slightly more formal way in the endnotes.)[24] In short, then: suitable representation of a group is the representation it would have had in a world history without representation-affecting discrimination.

While the Global View has the attractions alluded to above, one may object to it on the grounds that, given that only 45% of the members of a generation

---

[23] *The Economist*, December 17, 1998, http://economist.com/displayStory.cfm?story_id=179826 (accessed June 30, 2005).

[24] The Global View: a group G is suitably represented in a given category of employment E at time *t* if the proportion of G-people in E at *t* is what it is at *t* in the closest possible world in which G and other groups (whether the latter actually exist or do not actually exist but would exist in the relevant counterfactual situation of nondiscrimination) do not suffer from discrimination affecting their representation, whether by affecting the relative sizes of these groups or by affecting how well those who are members of these groups are represented, *at any time*.

are women, desirable elimination of the effects of discrimination cannot require women to occupy 50% of well-paid jobs. This point remains forceful even if 50% is the proportion of well-paid jobs that women would have occupied had the proportion of women in the relevant generation been what it would have been in the absence of past, sexist use of IVF technologies.

Would discrimination, in this situation, be properly redressed by a 45% representation of women in well-paid jobs? The belief that it would fits neatly with the *Constant People View*. On this view the counterfactual scenario determining whether a certain group is suitably represented at a particular time is, roughly, one in which the same people exist as exist in the actual world at that time and none of these people have suffered from discrimination affecting their representation at any time in their lives.[25] So, in this view, unlike the Global View, determining suitable representation does not require us to imagine a world history free from discrimination from the dawn of mankind onward (a tall order by any standard). It requires us to imagine only alternative, recent histories in which the same people exist as those who in fact exist; we must imagine that these people's representation in different job categories is unaffected by discrimination. Presumably, then, the scenario the Constant People View requires us to imagine will differ less from the world, as we know it, than the one pressed into service by the Global View. To return to the case of suitable representation given sexist use of IVF treatment: had the women who now exist never suffered from discrimination (unlike their unborn sisters, metaphorically speaking), and had their career ambition been no different from that of men, then, it seems plausible to assume, 45% of the people in well-paid jobs would have been women.

The reason I only go so far as to say that the Constant People View *fits neatly* with the 45% answer is that sex is a special case in that very few individuals change sex after coming into existence. Other socially salient groups differ in this respect. To take one example, a significant number of people change their religious beliefs during their lives, and discrimination may affect not only how well members of a particular religion are represented, but also how many people are members of that religion. Bearing this in mind, we can see that the idea that discrimination, in the situation we have been considering, would be properly redressed by a 45% representation of women in well-paid jobs could also be grounded in the *Constant Groups and People View*. In this view the relevant counterfactual scenario is, roughly, the most realistic counterfactual scenario satisfying the conditions stipulated in the Constant People View plus

[25] The Constant People View: a group G is suitably represented in a given category of employment E at time *t*, if the proportion of G-people in E at *t* is what it is at *t* in the closest possible world in which: (i) the same people exist at *t* as exist in the actual world; and (ii) none of these people suffer from discrimination affecting their representation, whichever groups they are members of in this possible world, *from the time of their coming into existence until t*.

the following further requirement: no discrimination has taken place during the life-time of presently existing individuals affecting the relative sizes of the relevant groups or which groups exist in the first place.[26]

These three views have conflicting implications. However, by way of mitigating the conflict, it might be argued that the Global View serves a purpose different from those served by the other two views introduced in this Section. If we take the Global View to be a principle telling us what to aim for now and here, obvious problems arise. For suppose that members of a certain group suffer discrimination. Suppose, moreover, that this group would never have existed in the first place, had no one ever suffered discrimination, e.g., because the group's emergence was a response to discrimination. According to the Global View, suitable representation of this group is 0%. This may be right for situations in which a group is not only formed in response to discrimination but will also immediately disappear with the cessation of discrimination. In the future, in a situation like this, we might bring about a state in which the relevant group has no members—and hence a situation in which it is plausible to say that suitable representation of the relevant group is 0%. But the 0% figure cannot be right for situations in which the group will survive the elimination of discrimination against it, as may be the case where, for example, a sense of shared history survives the demise of the discrimination that made the history shared.

Suppose, alternatively, that while discrimination does not occur now, it occurred in the past. Suppose, moreover, that had discrimination never occurred, everyone would have had the religion that nowadays only a small minority adheres to. In the Global View, suitable representation of this group is 100%. Again, this seems highly counterintuitive.

Finally, suppose that in the actual setting two groups exist. In the counterfactual setting one additional group exists, each of the three groups occupying a third of the relevant labor-market segment. On the view entertained here, the two actual groups should each have a representation of 33%, which gives a total of 66%! Once again, this cannot be right if the Global View tells what suitable representation is for a certain group here and now. One explanation of the failings of the Global View here is this: suitable representation of a group depends on its numbers, and the effect of past discrimination on the present size of groups is not something that ought to be rectified. Thus, for example, if there are fewer Catholics in Ireland today than there would have been had no religious discrimination ever occurred in Ireland, we will not now improve matters,

[26] The Constant Groups and People View: a group G is suitably represented in a given category of employment E at time *t*, if the proportion of G-people in E at *t* is what it is at *t* in the closest possible world in which: (i) the same people exist at *t* as exist in the actual world; (ii) none of these people suffers from discrimination affecting his or her representation, whichever groups they are members of, in this possible world, from the time of their coming into existence until *t*; and (iii) *no discrimination takes place during the life-time of presently individuals existing at t of the sort that affects the relative sizes of the relevant groups or which groups exist in the first place*.

from the point of view of suitable representation, by boosting Catholic numbers. This does not prevent the Global View from revealing suitable levels of representation in a world, which is ideal from the point of view of nondiscrimination. The contention is only that, if we are interested in what to strive for given past discrimination, we will have to return to the Constant People View and the Constant Groups and People View.

To see the difference between these two views, consider a religious minority that suffers from discrimination, albeit not discrimination of the kind ensuring the minority is underrepresented in attractive job categories relative to how many members it has at any given moment. (Perhaps an affirmative-action program nullifies the effect of discrimination on jobs.) Discrimination against members of this minority does, however, induce some of them to convert to other religions. Hence, there are now fewer members of the minority than there are in the closest possible world in which the same individuals exist; but none of these members suffered discrimination affecting their religious beliefs from the time of their coming into existence. If we now ask what is now a suitable representation for this religious minority, the Constant People View and the Constant Groups and People View will give different answers—in particular, the former view might well imply lower levels of suitable representation than the latter. In the world which, according to the Constant Groups and People View, is relevant for the purpose of determining suitable representation, there will be more members of the relevant religious minority: no discrimination that makes members of the religious minority convert to other religions will have occurred in some limited, relevant part of the past, and for this reason, presumably, the proportion of people belonging to the minority in the relevant attractive job category might well be higher. This means that, when it is construed as a view about what is now and here suitable representation now or in the future, the Constant Groups and People View is vulnerable to the objections I have just rehearsed against the Global View. For, like the Global View, it implies that suitable representation is sensitive to the way discrimination in the past may have affected the present size of the relevant group, negatively or positively.[27] In this respect, the Constant Groups and People View is less plausible than the Constant People View.

Suppose we ask ourselves what now would be a suitable representation of Buddhists ten generations from now. According to the Constant People View, that depends partly on what religion people who actually will exist ten generations from now will acquire given that they experience no discrimination

[27] Unlike the Global View, it is not the case that according to the Constant Groups and People View any piece of past discrimination that affected the size of a certain group may affect suitable representation for that group, since only discrimination that affected the size of the relevant group after its present, actual members came into existence affects suitable representation for this group according to the Constant Groups and People View.

from the time of their coming into existence. Discrimination occurring prior to that moment is irrelevant according to the Constant People View. This answer seems wrong. Suppose that Buddhists will suffer from discrimination from now until ten generations later and that as a result their numbers will be reduced at each consecutive generation. Surely what is now the suitable representation of Buddhists ten generations into the future should be sensitive to this.

In the light of the problems faced by the Constant People View and the Constant Groups and People View, I want to propose the *Future-Oriented View*. On this view, the relevant scenario determining whether now a certain group is suitably represented at some particular time (now or in the future) is, roughly, one in which no group suffers from discrimination affecting its representation or its relative size now or any time hereafter (if this is relevant).[28] Unlike the three views previously presented in this Section, the Future-Oriented View is purely future-oriented because it implies that past discrimination makes no difference to suitable representation now or in the future beyond fixing the initial starting point from which we extrapolate a nondiscriminatory future.[29] (I disregard issues of compensation.) To see this, consider a case in which two groups have acquired ambitions other than they would have acquired had it not been the case that the former group was discriminated against and the latter discriminated in favor of. If at some particular time we permanently eliminate all such discrimination, we can imagine either that the relevant difference in ambitions is entirely robust with regard to the elimination of discrimination, i.e., that although the difference in ambitions arose due to discrimination, once in place its existence is independent of this discrimination; or that the difference in ambitions is fully responsive to the elimination of discrimination, i.e., that once discrimination is eliminated the difference in ambitions vanishes immediately. (We could imagine something between these scenarios.) The Future-Oriented View implies that what constitutes suitable representation for

[28] The Future-Oriented View: at time t_x, a group G is suitably represented at time t_y (where t_y is later than, or simultaneous with, t_x) in a given category of employment E if the proportion of G-people in E at t_y is what it would have been at t_y had G and all other groups (whether groups that actually exist or groups that do not actually exist but would exist in the relevant counterfactual situation of nondiscrimination) not suffered from discrimination affecting their representation or their relative sizes *at t_x or any time thereafter*.

[29] The Future-Oriented View is time-relative in the sense that it implies that what constitutes suitable representation for a certain group at a certain time depends on when we ask the question. Consider a religious minority whose numbers are decreasing rapidly as a result of discrimination. If we now ask what constitutes suitable representation for them at some later time, we may find that this differs from what later constitutes a suitable representation for them at that time. This might be considered a problem if the Future-Oriented View is regarded as a principle that tells us what a perfect world would look like from the point of view of nondiscrimination. However, it is not a problem if the Future-Oriented View tells us what, now, is a suitable representation for a certain group now or later.

these two groups differs. This implication sits well with the argument so far. Remember that we rejected the view that suitable representation of groups whose size results from past discrimination, but whose present size is robust with regard to the present elimination of discrimination, should be sensitive to the size that the group would have had in the absence of past discrimination. But then it is hard to see why one should not say something similar here: that differences in ambitions which are robust with regard to the removal of discrimination despite arising through discrimination may help determine what constitutes suitable representation of a group.

The following case nicely illustrates this feature of the Future-Oriented View. Suppose, for the sake of argument, that if we eliminate discrimination against women now, the differences between men's and women's ambitions will narrow in such a way that in this generation 33% of all top managers will be women, in the next generation 40% will be women, and two generations from now 50% will be women.[30] According to the Future-Oriented View, a now suitable representation of women among this generation's top-managers will be 33%, a now suitable representation in the next generation will be 40%, and a two-generations-from-now suitable representation will be 50%! This apparently counterintuitive outcome is bound to alarm some observers, who will object that, given that the difference between men and women's ambitions is wholly due to discrimination against women, surely now—not two generations hence—the suitable representation of women is 50%.

In my view, the problem here is easily exaggerated. First, with regard to the example given, we should not assume that men and women differ in terms of equality of opportunity (or in terms of equality of outcome). It is simply that they have different ambitions (as an innocent result of past discrimination), and thus live different lives, statistically speaking. Hence, I am not here addressing a case in which past discrimination made some people worse off, all things considered, than they otherwise would have been; the case is simply one in which past discrimination made people choose differently than they otherwise would. Moreover, the explanation of why they have different ambitions is not that some of them are subjected to discrimination, but that women in the past suffered from discrimination.

Second, to my mind, at least, the implication of the Future-Oriented View mentioned above ceases to be counter-intuitive when we consider the reverse case. Consider the present trend in some Western countries toward numerical

[30] For the purpose of illustration, I assume here that there is a precise, determinate answer to how well a certain group will be represented at some particular time in a discrimination-free future. No doubt, often, there will be no such answer. In my view, this does not show that the Future-Oriented View is flawed. This feature of the Future-Oriented View is a quite general feature of normative views appealing to counterfactuals (that are not extremely straightforward) and it can hardly be said that such views are in general objectionable.

overrepresentation of female students at universities.[31] Let us suppose that, following the elimination or reduction of various forms of sex discrimination, 67% of all students in the present generation will be women, and that this figure will fall to 60% in the next generation and then 50% two generations from now. This may be the case if, as a result of their history of discrimination, women strive for what previous generations of women were prevented from achieving. This numerical overrepresentation of women, which is the result of their greater educational ambition, which in turn is the result of past discrimination, is in no way objectionable. Yet if a suitable representation is the representation a group would have had in the absence of past discrimination, we are committed, implausibly, to the view that this case involves objectionable overrepresentation of women. As my second example shows, the Future-Oriented View is not implausible because in some way it displays a greater tolerance of male privileges owed to past discrimination (as my initial example may have mislead some to conclude).

5. Is Absence of Discrimination Necessary for Suitable Representation?

Suppose that a certain population is divided into a number of groups. Each group has a claim to a suitable representation, and no other groups have this claim, although of course the population can be divided into further subgroups. Obviously, it will be hard to defend a view as to which division satisfies this condition, but I shall assume that some such division can be identified and defended. Suppose, finally, that although discrimination is not eliminated, it so happens that each group is suitably represented according to the Future-Oriented View, i.e., for every group, representation corresponds to what it would have been in the absence of discrimination in the relevant way. This could be the case because different kinds of discrimination neutralize the effects of each other. For example, there might be discrimination against members of the Catholic Church qua Catholics, but discrimination in favor of people with a higher education qua possessors of higher education, and it might so happen that people with a higher level of education tend to be Catholics. Generalizing, it is possible for the same individuals to be discriminated both against and for simultaneously in such a way that these forms of discrimination have no overall effect on group representation.

[31] At the largest Danish university, the University of Copenhagen, the proportion of female students among those admitted in 2005 was 60%: see http://www.ku.dk/nyheder/?content=%2Fnyheder%2Fvis%5Findlaeg%2Easp%3Ftype%3D3%26id%3D486%26ref%3D31 (acccessed June 30, 2005). I realize that the causal background to the trend toward the numerical overrepresentation of female students is complex and has to be seen in the wider context of, e.g., numerical overrepresentation of male faculty members. However, these further facts are irrelevant to the point I want to make here.

Mutual neutralization of the effects of different forms of discrimination could occur in two ways. In an unproblematic version, suitable representation will be achieved because each individual ends up having exactly the same job that this person would otherwise have got. In a more problematic version, neutralization will occur only at the aggregate level and not at the level of each individual. This latter possibility may be realized if discrimination increases the representation of Catholics with higher education relative to Catholics with no higher education and increases the representation of non-Catholics with higher education relative to non-Catholics with no higher education. Since we have stipulated that the representation of these groups does not matter, this is not bad on any group-focused view. However, we should notice that in this situation different individuals will end up in attractive positions, and we may think this is bad. Just as retributivists want guilty people to be punished and do not mind whether the proportion of punished people belonging to a certain group (e.g., immigrants, men, or young people), is identical to what it would have been in the absence of innocents being punished, some may think that the concern with nondiscrimination takes a similar individualized form. Perhaps, then, we should endorse the *Individualized View*. On this view the relevant counterfactual scenario determining whether, now, a certain group is suitably represented in a certain category of employment is, roughly, the most realistic counterfactual scenario satisfying the conditions stipulated in the Future-Oriented View plus the following further requirement: all the individual members of the relevant group that are in that category of employment would have been in that category in the absence of representation or group size affecting discrimination.[32] If the Individualized View is satisfied, so is the Future-Oriented View. Given, however, that the latter is an individualized

[32] The Individualized View: at time t_x, a group G is suitably represented at time t_y (where t_y is later than or simultaneous with t_x) in a given category of employment E if: (i) the proportion of G-people in E at t_y is what it would have been at t_y had G and other groups (whether groups that actually exist or groups that do not actually exist but would exist in the relevant counterfactual situation of nondiscrimination) not suffered from discrimination affecting their representation or their relative sizes *at* t_x *or any time thereafter*; and (ii) the G-people in E at t_y would have been in E at t_y had G and other groups (whether groups that actually exist or groups that do not actually exist but would exist in the relevant counterfactual situation of nondiscrimination) not suffered from discrimination affecting their representation or their sizes *at* t_x *or any time thereafter*. Suppose the only effect of discrimination from now and onwards would be to affect which individuals come into existence. It might then be asked why it should be considered desirable now to neutralize the effects of discrimination as the Individualized View implies. This question has a bearing on the so-called nonidentity problem, i.e., the problem of accounting for the wrongness (or rightness) of actions that do not affect individuals who exist independently of whether these actions are performed, but determine which individuals come to exist in the first place. I shall set aside this issue, since it is huge. Moreover, the concern of the present chapter is to articulate the concern to nullify the effects of discrimination; undoubtedly, who comes into existence is such an effect.

theory of suitable representation, the Future-Oriented View might be satisfied when the Individualized View is not.[33]

In my view, the Individualized View is the best account of suitable representation available. However, it raises a question. Suppose we think that discrimination is wrong even where there is suitable representation of different groups. And suppose that if discrimination is eliminated now and henceforth, then representations of different groups will be from now on unobjectionable. If this were so, the concern with suitable representation in its most plausible form would be reducible to a concern that the effects of discrimination on individuals are nullified. For we tend to think discrimination is bad even if it does result in deviations from suitable representations of groups. Moreover, suitable representation according to the Individualized View is defined in such a way that if discrimination is eliminated, it follows that suitable representation obtains. Although the Individualized View is formulated in terms of suitable representation, there is no account of what that is which does not amount to "whatever representation would obtain in the absence of discrimination." Hence, the motivation for representational aims under scrutiny in this chapter really is best construed as fundamentally a concern to nullify a certain kind of effect of discrimination on individuals.

6. Second-Best Representational Aims

Suppose we agree that the Individualized View is the best view about suitable representation. Still, it might be urged that the Individualized View often fails to address relevant questions about representational aims. For often we take it to be unavoidable that certain groups will continue to suffer discrimination in the foreseeable future—"unavoidable," not in the sense that if we all acted differently, no discrimination against members of this group would occur, but in the sense that it is certain that some will continue to discriminate against members of this group (say, because they want to, or because although they would like to stop, doing so would be costly or difficult). We might think that, in such cases, our representational aims should reflect the presumption that discrimination is unpreventable. This thought is not a criticism of the Individualized View. However, it motivates the following, *Partial Compliance View*: in this view, the

[33] The Individualized View does not imply that it serves no useful purpose to compare the representation of different groups in different categories of employment. Given a plausible theory of the social background, this kind of comparison may serve as a reliable indicator of the presence or absence of structural discrimination, and therefore of the presence or absence of suitable representation as defined by the Individualized View. Accordingly, I can agree with much of what Young says about why it is relevant, from the point of view of justice, to compare the position of different groups, i.e., that we need to do so in order to be able to identify structural injustices (one form of which I take to be structural discrimination), see Iris Marion Young (2001), "Equality of Whom? Social Groups and Judgements of Injustice," *Journal of Political Philosophy* 9, 1–18, p. 6.

relevant counterfactual scenario determining whether, now, a certain group is suitably represented at some particular time is, roughly, the sort of scenario deployed in the Individualized View except that it does not involve the absence of *any* kind of representation or group size affecting discrimination, but simply the absence of *preventable* representation or group size affecting discrimination.[34] Suitable representation in the Partial Compliance View may well diverge from suitable representation in the Individualized View.[35] In principle such representation could be higher or lower in the Partial Compliance View than it is in the Individualized View, although it is more likely to be the latter. If, for instance, in the relevant scenario of reduced discrimination, the discrimination dramatically increases the desire in members of the group being discriminated against to excel and only presents small obstacles to excelling, suitable discrimination for this group might be higher in the relevant nonideal future than it is in an ideal future of no discrimination.[36]

Nonideal futures, from the point of view of discrimination, raise an interesting question. For suppose that we can realize either of the following two scenarios of reduced discrimination. In the first, very few individuals who would have got the relevant jobs in a nondiscriminatory future actually get them. However, more individuals from the relevant group as a whole get the relevant kind of job. In the second, it is just the reverse. The representation of the group as such is further away from the group's representation in a nondiscriminatory future than in the first scenario. However, those individuals from this group who gets the jobs in the second scenario are individuals who would have gotten the jobs in a nondiscriminatory future. Hence, group-wise, the representation is better in the first scenario relative to the relevant representational aims than it is in the second. Individual-wise, however, the first scenario is worse than the second relative to the relevant representational aims. The question then arises whether one should always prefer the

[34] The Partial Compliance View: at time t_x, a group G is suitably represented at time t_y (where t_y is later than or simultaneous with t_x) in a given category of employment E if: (i) the proportion of G-people in E at t_y is what it would have been at t_y had G and other groups (whether groups that actually exist or groups that do not actually exist but would exist in the relevant counterfactual situation of *reduced discrimination*) not suffered from *preventable* discrimination affecting their representation or their relative sizes at t_x or any time thereafter; and (ii) the G- people in E at t_y would have been in E at t_y had G and other groups (whether groups that actually exist or groups that do not actually exist but would exist in the relevant counterfactual situation of reduced discrimination or nondiscrimination), not suffered from preventable discrimination (or from discrimination) affecting their representation or their sizes at t_x or any time thereafter.

[35] A similar point applies to the previously stated counterfactual views concerning suitable representation.

[36] While it may often be appropriate to ask how well a group would be represented if some but not all discrimination were to stop, it should also be noted that the more discrimination we take for granted when determining suitable representation for a certain group the less interesting, morally speaking, the resulting norms of suitable representation become. For the more discrimination we take for granted, the further away we are from a just social world.

scenario in which the number of individuals who get the job that they would have got in the relevant no- or less-discrimination scenario is maximized, or whether independent weight should attach to the fact that the relevant group's representation approaches what it would have been under the relevant no- or less-discrimination scenario, all other things being equal.

When answering this question it is important to set aside certain confusing factors. Thus the answer should not be based on the assumption that a level of representation that comes closer, group-wise, to the suitable representation of this group indicated by the Individualized View involves less discrimination against this group. This assumption implies that not all other things have been kept equal, and that our comparison is based not simply on the trade-off between better group representation and better individual representation. Similarly we should disregard the concerns set aside in section 1, such as the expressive concern to send messages of inclusion. If the present irrelevance of these factors is kept in mind, the assertion that only better individual representation matters when it comes to the concern to nullify the effects of discrimination looks far from implausible.

7. Conclusion

Accounts of suitable representation tend to be formulated as what Robert Nozick has called "end-state principles,"[37] i.e., as principles given which whether the representation of different groups is suitable depends merely on whether a structural principle of representation is met (an example of this is the Simple View). This implies that the suitableness of a representation does not depend on how it came (or will come) about; nor is it affected by which individuals occupy which positions, as long as changes of position do not affect the structural properties of the outcome being evaluated. I have argued that both of these features matter if our basic concern is to nullify the effects of discrimination. Accordingly, while the Individualized View may seem like an end-state principle, it is not: it tells us that the representation of different groups is suitable if it corresponds to what it would have been if its genesis had satisfied certain constraints.[38] On reflection, this result should not be too surprising. Discrimination is first and foremost a matter of how people are treated differently, not a matter of their having more or less of something, and how differential treatment results in differential outcomes is a contingent matter.

Should we, then, give up representational aims and focus on setting up mechanisms that eliminate discrimination? Should we strive merely to ensure

[37] Robert Nozick (1974), *Anarchy, State, and Utopia* (Oxford: Basil Blackwell), pp. 153–155.

[38] Unlike Nozick's libertarian principles of justice, however, the constraints on genesis imposed by the Individualized View are not backward looking.

that formal complaints about suspected discrimination can be made easily, put the burden of proof on employers, and use heavy fines when discrimination is found to have occurred? It depends on the circumstances. For all I have said here, in many situations suitable representational aims may function as a mechanism of this sort. Indeed my hunch is that, in many situations, suitable representational aims will be effective instruments to use in eliminating the effects of discrimination (or in reducing discrimination in the longer run). What is the use of this enquiry then? It helps us set straight the argumentative setting for determining whether representational aims are justified on the basis of the concern to nullify the effects of discrimination. In particular, it shows us that we should ensure that a preoccupation with groups' numbers in relation to the population as a whole does not lead us away from our real aim—e.g., because we have become blind to the effects of discrimination against numerically overrepresented groups, or overlook the innocently different ambitions of some numerically underrepresented groups, or forget that our ultimate concern is to eliminate the effects of discrimination on individuals, not groups as such.

{ 8 }

Discrimination in Punishment

1. Introduction

In the previous chapter I looked at cases where the numerical underrepresentation of a group is often seen as a sure sign that—or, more strongly, as a criterion of the fact that—it is being discriminated against. In this chapter, I turn to a case where overrepresentation is seen in a similar way, i.e., the case of capital punishment. All kinds of punishment can be discriminatory, but for capital punishment the charge of discrimination has played a central role in arguments for its abolishment. David Baldus concluded in a study that "the odds of a death sentence for those who kill whites in Georgia are 4.3 times higher than the odds of a death sentence for those who kill blacks."[1] In 1987 this work was used to challenge the execution of death-row prisoner Warren McClesky in the US Supreme Court on the grounds of an "unfair racial bias in the administration of the death penalty in Georgia."[2] The Supreme Court dismissed the argument, but the debate about how discrimination bears on the justifiability of punishment has continued. In 1997 the American Bar Association called for a moratorium on the death penalty partly because of "the continuing problem of racial bias in the administration" of it.[3]

Very many forms of discrimination are wrong, and discrimination in punishment seems particularly objectionable. First, a fundamental legal value is that of equality before the law. Because discrimination, by its very nature, involves unequal treatment, discriminatory punishment unavoidably clashes with the equal legal status of all.

---

[1] M. L. Radelet and M. J. Borg (2000), "The Changing Nature of Death Penalty Debates," *Annual Review of Sociology* 26, 43–61, p. 48. But see P. Cassell (2004), "In Defense of the Death Penalty" in H. Bedau and P. Cassell (eds.), *Debating the Death Penalty: Should America Have Capital Punishment?* (Oxford: Oxford University Press), 183–217, pp. 203–205.

[2] Radelet and Borg (2000), "Changing Nature," p. 48.

[3] Radelet and Borg (2000), "Changing Nature," p. 48.

Second, punishment has a different status than discrimination by private individuals. It is carried out in the name of the state and is thus authorized by all of us. Hence, discriminatory punishments have a communicative and symbolic dimension that many forms of private discrimination do not. This seems particularly obvious in the case of severe punishments such as capital punishment. Noting that the punishment that one receives in the United States for "murdering an African-American is often significantly less than the [punishment one receives for] murdering anyone else," Michael Cholbi infers that "[i]mplicitly, then, the lives of African-Americans are treated as less valuable than the lives of others."[4] In Deborah Hellman's account, wrongful discrimination is discrimination that demeans the victim, and to demean someone is to treat this person "as not fully human or not of equal moral worth."[5] If Cholbi is right, discriminatory capital punishment appears to be one of the most demeaning and, thus, based on the objective meaning account of the wrongness of discrimination discussed in chapter 5, most wrongful forms of discrimination.

To set the stage for an exploration of the moral qualities of punitive discrimination, section 2 of this chapter distinguishes between different loci of legal discrimination. Focusing on the third locus, section 3 discusses what constitutes an adequate criterion for the existence of discrimination in sentencing. This question is complicated, partly because of how considerations about legal protection of potential victims of crime interact with considerations about equal treatment of perpetrators of crime. Section 4 discusses the question of what, if anything, is wrong with discriminatory punishment when we set aside all those morally problematic features with which discriminatory punishment might be contingently correlated. Section 5 confronts an influential response to the claim that discriminatory punishments cannot be unjust, because provided those who are as a matter of fact punished deserve their punishment they are in no position to complain. But not being in a position to complain about a certain treatment may result from factors other than those that bear on the permissibility of this treatment. Thus, suppose that a murderer, who might have received a lighter sentence had he had a different race or sex, deserves, non-comparatively speaking, capital punishment and for that reason is in no position to complain

[4] M. Cholbi (2006), "Race, Capital Punishment, and the Cost of Murder," *Philosophical Studies* 127, 255–282, p. 268. See also R. L. Kennedy (1988) "*McClesky V. Kemp*: Race, Capital Punishment, and the Supreme Court," *Harvard Law Review* 101, 1388–1443, p. 1391. According to Radelet and Borg, most post-*Furman* US Supreme Court decision—a decision reducing discretion in the imposition of the death penalty intended to restrict the influence of biases, notably racial biases, in the use of capital punishment—studies "conclude that for crimes that are comparable, the death penalty is between three or four times more likely to be imposed in cases in which the victim is white rather than black," Radelet and Borg (2000) "Changing Nature," p. 47.

[5] Hellman (2008), *When Is Discrimination Wrong?* (Cambridge, MA: Harvard University Press), p. 35.

about being sentenced to death. Yet, on the view defended here, capital punishment for this murderer might not be just.

Before proceeding, I should briefly explain why I think the topic of the present chapter relates to that of the previous one. Like in discussions of proportional representation on the job market, in discussions of proportional representation in relation to capital punishment many assume that if a certain group's representation diverges from its proportional one, this reflects discrimination and, accordingly, that to eliminate discrimination or at least some of its effect we should strive for proportional representation. As was shown in the previous chapter, this assumption is problematic. In this chapter we will see further reasons for questioning it. Unlike being underrepresented at the job market, being underrepresented on death row is a good thing and in this respect the focus of this chapter differs from that of the previous one. This is helpful in part because it helps to identify what may be problems in the ideal of proportional representation motivated by adherence to (worthy) causes. For instance, while women are underrepresented on the labor market they are also underrepresented on death row. Presumably, we want an explanation of the wrongfulness of the former that does not compel us to think of the latter underrepresentation as being wrongful to men.

2. Loci of Legal Discrimination

Discrimination in punishment has several *loci*. First, laws may discriminate. As previously observed, it is common to distinguish between direct and indirect discrimination and claim that, although most laws no longer discriminate directly, many do so indirectly. The official EU homepage on discrimination defines the two species of discrimination as follows:

> Direct discrimination occurs when a person is treated less favourably than another in a comparable situation because of their racial or ethnic origin, religion or belief, disability, age or sexual orientation. An example of direct discrimination is a job advert, which says "no disabled people need apply." . . . Indirect discrimination occurs when an apparently neutral provision, criterion or practice would disadvantage people on the grounds of racial or ethnic origin, religion or belief, disability, age, gender or sexual orientation unless the practice can be objectively justified by a legitimate aim. An example of indirect discrimination is requiring all people who apply for a certain job to sit a test in a particular language, even though that language is not necessary for the job.[6]

[6] http://ec.europa.eu/employment_social/fdad/cms/stopdiscrimination/resources/glossary/?langid=en#I (accessed February 1, 2010). The indirect discrimination bit of this text passage was also quoted and discussed in appendix 1 in chapter 2.

Based on this definition, presumably, a law that stipulates that the level of punishment for crimes involving crack cocaine is significantly higher than punishments for crimes involving powdered cocaine, might—see below—indirectly discriminates against blacks provided that blacks primarily use crack cocaine and whites primarily use powdered cocaine.[7] This is so even if the law is apparently neutral—it does not distinguish for the purpose of meting out punishment between blacks and whites—and provided that there is no "legitimate aim" that can "objectively justify" the differential treatment.

Second, the law can be enforced in a discriminatory way. For instance, if the police are more inclined to stop black motorists and search them for drugs than to stop and search white motorists, they might arrest and prosecute more black people for possession of illegal drugs even if the underlying real crime rate in relation to possession of such drugs does not differ between the two groups.[8] The distinction between direct and indirect discrimination applies to enforcement of laws as well. It is one thing that police tactics are openly motivated by racial animosity—in that case the practice is not even "apparently neutral"; it is another if it is formulated in entirely neutral form, e.g., the police make it a top priority to combat certain kinds of crimes where a particular group is overrepresented, even though there is no justification for this priority.

Third, the way lawbreakers are punished can be discriminatory. This occurs whenever members of two groups do not receive equally severe punishments for crimes that, legally speaking, are not relevantly different. This kind of discrimination may not be intentional. Judges and jurors may reach the verdicts they do in ways that reflect their intention to punish some groups of criminals harder than others, or because they believe that certain crimes are graver when committed by members of some groups against members of others. But more likely discrimination in punishment takes a non-intentional form. That is, judges and jurors intend to treat members of different groups equally; they think this is what they are doing, and yet there exists a pattern of differential punishment, which is best explained by the role the legally irrelevant group membership plays (see chapter 1.8 on discrimination of which the discriminator is unaware).

Fourth, the way punishments are implemented can be discriminatory. Many convicts do not serve the full sentence but are released on parole, and biases may determine who gets it and when.

---

[7] S. L. Myers Jr. (1993), "Racial disparities in sentencing: Can sentencing reforms reduce discrimination in punishment?" *University of Colorado Law Review* 64, 781–808, p. 790; Anonymous (2001), "The Rhetoric of Difference and the Legitimacy of Capital Punishment," *Harvard Law Review* 114.5, 1599–1622, p. 1619. But see R. Kennedy (1994), "The State, Criminal Law, and Racial Discrimination: A Comment," *Harvard Law Review* 107, 1255–1278, pp. 1267–1270.

[8] See the discussion of "'disproportionate' investigation" in Matthias Risse and Richard Zeckhauser (2004), "Racial Profiling," *Philosophy and Public Affairs* 32.2, 131–170, pp. 140–142.

Below I focus exclusively on the third kind of discrimination. It is worth stressing that even if no discrimination takes place in terms of sentencing this does not imply that discrimination takes place in none of the other loci.[9] In fact, discrimination in one locus may cause greater discrimination in one of the others. For instance, John P. Pittman believes that discrimination against African Americans in policing strategies causes more African Americans to be apprehended and convicted. The result is statistics showing higher crime rates for African Americans, which again causes jurors to have a "suspicion of guilt" and thus to be more inclined to convict African American defendants.[10]

An interesting possibility that I shall disregard is that members of a certain group are discriminated against in terms of one particular crime and discriminated in favor of in terms of others, e.g., that women get harsher sentences for child molestation than men do, but lighter sentences for murdering a member of the opposite sex than men do. Here judges discriminate *between* men and women, but they may not discriminate *against* either group on the whole.

3. Criteria vs. Indicators of Discrimination

Initially, it may seem easy to determine whether discrimination in punishments exists. According to the simple view, if a group receives a greater share of the total sum of punishment meted out by courts than its share of the total population, its members suffer sentencing discrimination. On this view, which is very similar to the Simple View on representation on the job market discussed in chapter 7, the fact that African Americans are strongly overrepresented among those who receive capital punishment strongly suggests that they are subjected to capital punishment discrimination. Since overrepresentation is even more massive in the case of men in the US, the simple view implies that men are massively discriminated against.[11] Few accept the latter inference and, thus, few—none, in fact—should accept the former. Independently of whether African Americans suffer sentencing discrimination, we should reject the simple view.[12]

---

[9] Discrimination taking place outside these *loci* could also bear on the moral status of legal discrimination. For instance, Reiman believes that "unjust discrimination in the recruitment of murderers undermines the justice of applying the penalty under foreseeable conditions in the United States," J. H. Reiman (1985), "Justice, Civilization, and the Death Penalty," *Philosophy & Public Affairs* 14.2, 115–148, p. 133n22.

[10] J. P. Pittman (1997), "Punishment and Race," *Utilitas* 9.1, 115–130, pp. 117–118. See also B. Stevenson (2004), "Close to Death: Reflections on Race and Capital Punishment in America" in Bedau and Cassell (eds.), *Debating the Death Penalty*, p. 87, and my remarks on self-fulfilling statistics in chapter 11.

[11] "Between 1977 . . . and 1996, 301 men and only one woman were executed in the United States," P. E. Devine (2000), "Capital Punishment and the Sanctity of Life," *Midwestern Studies in Philosophy* 24, 229–243, p. 241n23.

[12] W. J. Bowers and G. L. Pierce (1980), "Arbitrariness and Discrimination under Post-Furman Capital Statutes," *Crime & Delinquency* 26, 563–632, p. 575.

The obvious flaw with the simple view is that to the extent that crime rates differ across groups, in the absence of sentencing discrimination some groups *should* receive a larger proportion of the total punishment than their proportion of the total population warrants. This is the principal reason we do not infer from the fact that men receive a greater proportion of the total punishment than women do that men are being subjected to sentencing discrimination. No doubt, different crime rates may reflect extralegal discrimination. If a group faces discrimination on the job market, more of its members may on average be unemployed and thus be more criminal, but this is irrelevant here given our focus on punishment. So what I shall call the complex view says that a group suffers sentencing discrimination if, and only if, its share of the total sum of crime differs from its share of the total sum of punishment.[13] This view allows that a group might receive a disproportionately large part of total punishment and yet not suffer sentencing discrimination. Indeed, consistently with this state of affairs, it may even *benefit* from sentencing discrimination in its favor.

While the complex view is less flawed than the simple view, still it is not quite right. First, suppose a certain group suffers sentencing discrimination. Presumably, this means that, on average, its members will receive harsher punishments for crimes for which members of other groups will be punished more leniently. But this could be the case, even if some individual members of this group are sentenced in a way where no discriminatory biases influence the verdict they receive.[14] Presumably, this was part of the reason why the US Supreme Court did not overturn Warren McClesky's death penalty. Moreover, some individual members of the group experiencing sentencing discrimination in their favor might be discriminated against and, as a result, receive harsher sentences than they would have received had biases not influenced their trials. Given what is the case on average, such cases are not likely to be particularly common, but the fact remains that a group might suffer sentencing discrimination even if some of its members are actually discriminated in favor of and some members of other groups are discriminated against. Hence, if two groups' shares of the total crime correspond exactly to their shares of the total punishment, we cannot infer that none of their members suffered or enjoyed sentencing discrimination. Such discrimination might be going on in individual cases even though its effects cancel each other out such that no differences show up at an aggregate level.

---

[13] The notions of total punishment and total crime are rather complicated entities, but I disregard the complications that this raises for the purpose of exploring the basic thought. One complication is recidivism. Suppose we think that the second murder is worse than the first and thus deserves harsher punishment. In that case, the second murder should count for more in relation to the measure of total crime than the first.

[14] In my view, it is possible for the group of women to be discriminated against even if it is true of some individual women that they are never subjected to gender discrimination.

The second problem is that the complex view ignores, which groups victims of crime belong to. Suppose that killers of whites tend to get harsher sentences than killers of blacks; that black killers of whites are punished more harshly than white killers of whites and white killers of blacks; and that most, but not all, murders are intraracial. In that case, the complex view might imply that no sentencing discrimination exists even though discrimination might well explain differential sentencing depending on the race of the victim.[15]

The third problem with the complex view concerns how we identify the victim of sentencing discrimination. Suppose a certain group receives a greater share of the total punishment than its share of the total crime warrants. It is not clear that this implies that members of this group are victims of sentencing discrimination overall. Consider the following thought experiment: blacks receive only ten years in prison for murder, whereas whites receive twenty and murder is the only crime committed. All murders are intraracial. Suppose that the more years in prison a race-identical murderer serves, the greater the deterrent effect of the punishment is and the fewer members of the relevant community will be murdered; suppose also that the murder rates within the black and the white communities vary with the level of deterrence only.[16] Suppose, finally, that the level of punishment does not alter the fact that blacks kill blacks only and whites kill whites only, so potential black murder victims only have reason to care about the punishment meted out to black murderers at least as far as deterrence goes and the analogous point is true about potential white murder victims. Obviously, these assumptions are unrealistic, but this is beside the present conceptual point.[17] Initially, we might think that this scenario involves sentencing discrimination against whites, because the punishment whites receive is twice as severe as the punishment blacks receive for exactly the same crime, i.e., murder.[18] But this is only part of the picture. For in addition to the severity of the punishment that criminals from different groups receive, we should also consider the degree to which the legal system offers protection to

[15] W. J. Bowers and G. L. Pierce (1980), "Arbitrariness and Discrimination under Post-Furman Capital Statutes," *Crime & Delinquency* 26: 563–632, pp. 573, 577.

[16] It is unclear whether the deterrent effect of capital punishment is greater than that of long-term imprisonment; see Radelet and Borg (2000), "The Changing," p. 47. If capital punishment does not have a greater deterrent effect than long-term imprisonment, an argument in favor of the former is undermined, but it is not in itself an argument *against* capital punishment.

[17] They are not wholly unrealistic, however: "About 80 percent of violence occurs among persons of the same race"; see R. L. Kennedy (1994), "*McClesky V. Kemp*: Race, Capital Punishment, and the Supreme Court," *Harvard Law Review* 101: 1388–1443, p. 1255n2.

[18] Van den Haag notes that since "most black murderers kill blacks, black murderers are spared the death penalty more often than are white murderers . . . The motivation behind unequal distribution of death penalty may well have been to discriminate against blacks, but the result has favored them." E. van den Haag (1986), "The Ultimate Punishment: A Defense," *Harvard Law Review* 99.7, 1662–1669, p. 1664; see also E. van den Haag (1985), "Refuting Reiman and Nathanson," *Philosophy & Public Affairs* 14.2, 165–176, p. 173, where he argues that "[t]he practice invidiously discriminates against black victims of murder."

different groups. In the hypothetical scenario the legal system offers much less protection to blacks against the risk of being murdered than it does to whites. So rather than saying that sentencing discriminates against whites, perhaps we should say that it discriminates *against* blacks as such or blacks who are potential murder victims (because it offers them less protection than it does to white, potential murder victims) and that it discriminates *in favor of* black murderers (because it punishes them less than white murderers). Possibly one would not want to say the latter, e.g., because black murderers are not a socially salient group and, thus, for conceptual reasons not a group that can be the object of discrimination (see chapter 1), or because the causes of black murderers receiving less severe punishments have nothing to do with racial prejudice against whites etc. However, these observations are beside the point: that, unlike what the complex view implies, the consequences of a certain sentencing practice for potential victims of the crimes also bear on the question of whether group differences in terms of the severity of punishment reflect discrimination.[19]

The conclusion in the previous sentence stands even if we relax the assumption that all crimes are "intra-group" crimes, but I will not pursue this issue further for it is irrelevant to establishing this section's main claim: it is complicated to show that discrimination in punishments exists and, of course, to show that it does not exist. In what follows I set aside these complications and assume that discrimination in punishment has been shown to exist.

4. The Pure Discrimination Case

I now turn to the injustice of sentencing discrimination. I am interested in a rather narrow question: what, if anything, makes sentencing discrimination unjust or morally wrong in itself? Obviously, sentencing discrimination may be morally wrong for a lot of other reasons than because of what it is in itself. For instance, sentencing discrimination is likely to be believed to exist, when it does, and this may lead to a loss of trust in general and trust in the legal system in particular. Moreover, it may cause a breakdown of relations between different communities and, in extreme cases such as the acquittal of the four police officers, who beat up Rodney King, to rioting. Also, it is hurtful for people to know that they, or people like them, are singled out for discriminatory treatment.[20] These are all flaws that, empirically speaking, accompany

[19] Randall Kennedy stresses this concern (1988), "*McClesky V. Kemp*," pp. 1422, 1425. He thinks that "the main problem confronting black communities in the United States is not excessive policing and invidious punishment but rather a failure of the state to provide black communities with the equal protection of the laws"; see R. L. Kennedy (1994), "The State, Criminal Law, and Racial Discrimination: A Comment," *Harvard Law Review* 107: 1255–1278, pp. 1256, 1259.

[20] Joel Feinberg (1974), "Noncomparative Justice," *Philosophical Review* 83.3, 297–338, p. 318.

sentencing discrimination, but they are not flaws that are necessarily true of sentencing discrimination and, thus, not flaws that could explain why sentencing discrimination is morally wrong *in itself*. It is possible, even if unlikely, for sentencing discrimination to exist without these flaws. Yet, discriminatory punishment may seem unjust even if none of its victims, say, believe they are being subjected to discriminatory treatment and, thus, none of them are hurt by what they see as their being singled out for discriminatory treatment.

To focus on this rather narrow question, suppose that for any crime there is some level of punishment such that the perpetrator deserves exactly this level of punishment and suppose that justice in punishment simply is that all criminals receive exactly the level of punishment they deserve. On this view justice in punishment is retributivist and entirely noncomparative, to use a terminology coined by Joel Feinberg.[21] It is unjust if, say, men receive 10 units of punishment less than they deserve and women receive 10 units more. Given the stated assumption, this discriminatory situation is no less unjust than one in which they all receive 10 units less (or more) and in which, consequently, let us suppose, no sentencing discrimination exists. On the assumptions made, the question of whether sentencing discrimination is unjust per se translates into the question of whether sentencing discrimination entails that some people do not get the punishment they deserve, noncomparatively speaking.

There is no doubt that sentencing discrimination often involves noncomparative injustice. So take the direct legal discrimination of the 1816 Georgia statute explicitly requiring "the death penalty for rape or attempted rape if the crime was committed by a black against a white" (and in such cases only).[22] On the impeccable assumption that there is no difference as such in punishment deserved for raping black and white people, there is no way that this sentencing discrimination could exist consistently with all rapists receiving the level of punishment they deserve. But sentencing discrimination is consistent with all those who are punished receiving exactly the punishment they deserve. This is very unlikely, but it is logically possible. Even if judges and jurors are systematically biased in the assessment of evidence and systematically intend to impose harsher sentences on members of some communities than members of others, they might fail to carry out their intentions and their biased assessment of evidence may be neutralized by counteracting factors, e.g., that members of one community generally commit crimes in a way that makes it harder to prove that they committed them. Insofar as we find sentencing discrimination unjust even in this scenario, justice in punishment must consist of something more than all criminals receiving exactly the punishment that they deserve, noncomparatively speaking, since ex hypothesis this is what all criminals do and yet the situation is unjust.

---

[21] Feinberg (1974) believes that retributive punishment has both a comparative and a noncomparative element, "Noncomparative," p. 311.

[22] Bowers and Pierce (1980), "Arbitrariness," p. 757.

Some defenders of noncomparative justice might bite the bullet and say that the scenario I imagine is perfectly just. Surely, judges and jurors can be criticized for the way they reach their verdicts. But unjust reasoning is different from reaching an unjust decision.[23] After all, in my scenario everyone receives a punishment that perfectly matches the one they deserve.

This response strikes me as having considerable force, even though I do not find it decisive. First, a legal system that operates in a biased way even though through sheer luck it manages always to reach the right verdicts fails to reflect fully the value of equality before the law and, thus, conflicts with an ideal of a society in which we relate to one another as equals.[24] It sends an objectionable message despite the fitting punishments it imposes and may for that reason be one that justice requires to be reformed. Second, criminals are arguably wronged even though they receive exactly the punishment they deserve and cannot complain that they are being punished more harshly than others who have committed comparable crimes.

Setting aside the way sentences are reached, I now turn to comparative justice. To test the relevance thereof, consider a slightly different scenario. Suppose we have a group of murderers all of whom deserve the death penalty. Some, but not all of them, receive it. Those who do not receive the death penalty receive some lesser punishment, say, life imprisonment.[25] Ex hypothesis, all of those who are put to death deserve to be put to death and, thus, suffer no noncomparative injustice.[26] But it is also the case that some who deserve to be put to death are not put to death. Perhaps those being put to death can complain of comparative injustice: how come they are being put to death when others who have committed comparable crimes are not? I want to ask three questions with regard to this scenario: (1) Does it involve injustice in punishment? (2) If so, is the relevant injustice greater for reflecting discrimination as opposed to pure arbitrariness? (3) If the relevant injustice is greater for reflecting discrimination, what should be done about it to reduce injustice?

Starting with the first question, it is obvious that the situation involves injustice. Or, at least, insofar as there is such a thing as comparative justice, this is one of the clearest candidates for comparative injustice. In an exchange between Ernest van den Haag and Stephen Nathanson, the former denies that, in

---

[23] Thomas Scanlon (2008), *Moral Dimensions: Permissibility, Meaning, and Blame* (Cambridge, MA: Belknap Press of Harvard University Press), p. 27.

[24] Elizabeth Anderson (1999), "What is the Point of Equality?" *Ethics* 109, 287–337; Samuel Scheffler (2003), "What is Egalitarianism?" *Philosophy & Public Affairs* 31.1, 5–39; Jonathan Wolf (1997), "Fairness, Respect, and the Egalitarian Ethos," *Philosophy & Public Affairs*, 27, 97–122.

[25] For the sake of argument, I assume that, for some crimes, it is true that the criminal deserves the death penalty and that the death penalty is always a harsher punishment than life imprisonment.

[26] Feinberg discusses some comparable examples in Feinberg (1974), "Noncomparative," pp. 313–318. He thinks that they involve unfairness and comparative injustice, but that, generally speaking, claims of noncomparative justice are superior to claims of comparative justice.

the scenario imagined, any execution is unjust, whereas the latter thinks they all are. Van den Haag, who favors the death penalty, says: "[I]f the death penalty is morally just, however discriminatorily applied to only some of the guilty, it does remain just in each case in which it is applied."[27] Nathanson and van den Haag both suppose that if this claim is true, the discrimination objection to capital punishment and, thus, an important abolitionist argument fails.[28] Accordingly, Nathanson, who opposes capital punishment, rejects the claim. In support of its rejection, he offers two (relevantly similar) counterexamples to van den Haag's claim, one of which is the following: "I tell my class that anyone who plagiarizes will fail the course. Three students plagiarize papers, but only one receives a failing grade. The other two, in describing their motivation, win my sympathy, and I give them passing grades."[29] In this case, Nathanson submits, "the justice of giving [the plagiarizing students] what they deserve appears to be affected by the treatment of others," despite what Nathanson admits to be van den Haag's forceful rhetorical question: "How can it possibly be unjust to punish someone if he deserve the punishment?"[30] More generally, Nathanson claims that "the treatment of classes of people is relevant to determining the justice of" treatment of, including punishments for, individuals.[31]

Both van den Haag and Nathanson fail to consider a third view, which can concede the force of van den Haag's rhetorical question and yet is consistent with denying the justice of a discriminatory but fitting punishment. On this view, we should be more careful than van den Haag and Nathanson in our individuation of the state of affairs that we deem to be just. Specifically, the state of affairs, "X is executed and Y receives a lighter punishment," is different from another state of affairs, "X is executed," and, given that X and Y are equally deserving, the fact that the latter state of affairs is just does not make the former state of affairs just too. It is entirely open to an opponent of capital punishment

[27] E. Van den Haag (1978), "The Collapse of the Case against Capital Punishment," *National Review* 31, p. 397; Quoted from S. Nathanson (1985), "Does It Matter if the Death Penalty is Arbitrarily Administered?" *Philosophy & Public Affairs* 14.2, 149–164, p. 151. Elsewhere, van den Haag distinguishes between equal justice (all get what they deserve), unequal justice (some, but not all, get what they deserve) and equal injustice (no one gets what they deserve), arguing that the former is best and that unequal justice is "certainly better than equal injustice," van den Haag (1985), "Refuting Reiman and Nathanson," p. 174. It is unclear if "better" here means "better in terms of justice."

[28] "Discrimination or capricious distribution thus could not justify abolition of the death penalty," van den Haag (1986), "The Ultimate Punishment," pp. 1662–1663.

[29] Nathanson (1985), "Does It Matter?" p. 156.

[30] Nathanson (1985), "Does It Matter?" p. 157, p. 156.

[31] Nathanson (1985), "Does It Matter?" p. 160. Sidney Morgenbesser famously quipped that it was unjust for the police to beat him up at a demonstration, but that they acted fairly, because they beat up everybody equally. In the case of a discriminatorily sentenced murderer, it is the other way round. The murderer can complain about unfairness but not about injustice. Moreover, it might be held that such a fairness complaint does not reflect a right on part of the claimant not to receive capital punishment. I am grateful to Shlomi Segall at this point.

to concede of each individual state of affairs of a deserving criminal receiving the punishment he deserves that it is just (and it is open to Nathanson to deny that it is unjust that the plagiarizing student fails), while denying the justice of the more complex states of affairs that some deserving criminals receive capital punishment and others do not. (The state of affairs of my offering to carry an elderly, fragile person's bag involves a manifestation of politeness. But the state of affairs of my offering to carry one elderly person's bag and not offering to carry another elderly, fragile person's bag does not, when they stand next to each other, I could easily carry both bags, and there is no reason why I should carry one person's bag but not the other's.) If the latter injustice is sufficiently weighty, one might answer van den Haag's rhetorical question the way he thinks the question should be answered and still oppose capital punishment. So, in sum, Nathanson might fail in showing that discriminatory sentencing undermines the justice of individual punishment, but van den Haag and Nathanson err in thinking that to oppose capital punishment, discriminatorily administered, one must deny the justice of individual cases of capital punishment.[32]

So let us turn to the second question by comparing two situations: one in which it is entirely arbitrary who receives capital punishment and who does not—say, judges simply flip a coin whenever someone is found guilty of a crime that carries the death penalty—and one in which some kind of discriminatory bias determines who gets it.[33] Arguably, the lottery situation in which any guilty person runs an equal risk of receiving capital punishment is more just (or, if you follow Morgenbesser's distinction, more fair) than one in which, due to bias, some people do not. Both situations are unjust to the extent that some people receive lighter sentences than they deserve, but the latter involves an additional unfairness to offenders from those groups that are discriminated against: namely, that they have a smaller chance of enjoying the benefit of a lighter punishment than they deserve. Of course, it ought not to be the case that anyone receives a lighter sentence than she deserves, but given that some, but not all, do, it is fairer that everyone has an equal chance of receiving a lighter sentence.[34] This explains why the situation in which some, but not all, receive

[32] For a different response to van den Haag, see J. Reiman and L. P. Pojman (1998), *The Death Penalty: For and Against* (Lanham, MD: Rowman & Littlefield), pp. 118–121.

[33] In cases involving significant racial bias, McDermot would consider the punishing agent illegitimate and, thus, object to the assumption that those being punished receive a legitimate punishment even if, by sheer coincidence, the punishment is identical to that which they deserve from a legitimate punishing agent, see D. McDermott (2001), "A Retributivist Argument against Capital Punishment," *Journal of Social Philosophy* 32.3, 317–333, p. 326, p. 328.

[34] Note that here I disregard unfairness to potential victims. Arguably, members of groups who face a higher risk of being murdered might argue that this is unfair. Assume, moreover, that the only way to eliminate this unfairness is to impose harsher sentences on people who murder members of this group. In that case, it might be less unfair, all things considered, that some groups of defendants receive harsher punishments than others for crimes that are identical except for the extrinsic features that some group of potential victims face a higher risk of being subjected to this crime than others.

lighter sentences as a result of idiosyncratic biases is fairer than one in which some, but not all, receive lighter sentences as a result of discrimination, since, given the idiosyncratic nature of the bias, there is a sense in which all defendants bear an equal risk of facing judges and jurors who are idiosyncratically biased against them.

Turning now to the third question: what should we do about the injustice involved in the scenario that I have imagined? We accept that discriminatory punishment in the scenario that I have described is unjust and that it is more unjust, because more unfair, than the punishment lottery where some get lighter sentences than they deserve but all have an equal chance of getting a lighter sentence. So the question is how we eliminate or reduce this injustice. Basically, there are three options to consider.[35] First, we might impose mandatory capital punishment on all. This seems to be the position taken by van den Haag.[36] Second, we might refrain from imposing the death penalty on anyone. This seems to be the position taken by Nathanson.[37] Finally, we might try to neutralize the systematic bias in the legal system without ensuring a perfect match between the punishment received and the punishment deserved.

Consider van den Haag's position. We might endorse it given the assumptions on which this discussion has been conducted. Yet, once we address the plausibility of these assumptions, we might be more critical. For provided that biases have a deep influence on how discretion in capital punishment is exercised, there is reason to believe that such biases will influence who among deserving as well as non-deserving defendants gets capital punishment, even if discretion is eliminated. For instance, biases may influence who is prosecuted and convicted for crimes that carry mandatory capital punishment.

Now consider Nathanson's position. Suppose we agree that discriminatory punishment is unjust. Does it follow that we should stop imposing it?[38] Note, first, that Nathanson is not suggesting that we refrain from punishment altogether. His proposal is that we continue to punish murderers with long term imprisonment rather than impose capital punishment on anyone. If we canvass the criticism of van den Haag that, in practice, biases will influence sentencing even if discretion is eliminated, it would seem difficult to argue that

---

[35] Indisputably, a legal system in which all criminals receive exactly the punishment they deserve is best. The interesting question concerns situations in which a legal system cleansed of all discriminatory biases is unfeasible.

[36] Philip Devine agrees that this is what the discrimination concern motivates even if he thinks capital punishment is hardly ever justified. Devine (2000), "Capital Punishment," p. 232.

[37] Cholbi thinks a moratorium is warranted until the discriminatory bias has been eliminated, Cholbi (2006) "Race, Capital Punishment," pp. 270–277.

[38] It does not follow from the fact that some penal practice is not perfectly just, nor from the fact that it is unjust because discriminatory, that we should eliminate it altogether. The consequences of having no penal practice as opposed to a flawed one might be disastrous.

biases will not influence sentences lighter than capital punishment as well.[39] If so, Nathanson's proposal will not eliminate the injustice of discriminatory sentencing and—again: given the assumptions made for the purposes of the enquiry in this chapter—it will result in the additional injustice that more people will be punished more lightly than they deserve to be punished. The question is whether it will reduce the injustice of discriminatory justice in such a way that it more than outweighs the additional noncomparative injustice of some people not receiving the capital punishment they deserve.

With regard to the last question, it should be noted that in practice the death penalty is imposed on very few people relative to how many people are convicted of crimes that might in principle be punished by death. This suggests that the additional noncomparative injustice will be relatively small.

With regard to the former question, it might be argued that the injustice of discrimination depends on the importance of the benefits and harms involved in the discriminatory practice.[40] So discriminatory ticketing for parking offences is less unjust than discriminatory capital punishment, even if the two discriminatory practices are identical in all other respects than the severity of the harm involved in the punishment, i.e., a small fine versus death. So if capital punishment involves a much greater harm than, say, a lifetime in jail, the injustice involved in discriminatory capital punishment is much greater than the injustice involved in discriminatory lifetime imprisonment.[41] Still, if we argue that the noncomparative justice involved in eliminating capital punishment is small because few people receive capital punishment, it seems we are also committed to saying that the reduction of the injustice of going from discriminatory capital punishment to discriminatory lifetime imprisonment is small because of the relatively few people involved. A further issue is whether there are alternatives to eliminating capital punishment that eliminate discriminatory sentencing, but do not reduce the number of people deserving capital punishment who receive it. This is where the third alternative comes into the picture.

Suppose, for instance, that rather than receiving a death penalty, defendants receive a death penalty lottery ticket where the risk of losing varies between the two groups that are being discriminated between such that this differential risk eliminates the overrepresentation of one of these groups within the relevant penal category. So to take the case of death penalty in the US: if you are a black defendant you do not receive a straightforward capital punishment verdict, but rather a lottery ticket that gives you, say, a 25% chance of capital punishment.

---

[39] Unless we suspect that the death penalty "strikes at deep-seated racial prejudices in a way that milder penalties do not," Reiman (1985), "Justice," p. 134n22; Bowers and Pierce (1980), "Arbitrariness," p. 574.

[40] E.g., J. Greenberg (1986), "Against the American System of Capital Punishment," *Harvard Law Review* 99.7, 1670–1680, p. 1678n42.

[41] " . . . Death is a much more severe punishment than imprisonment," Nathanson (1985), "Does It Matter?" p. 161.

Or if you are a white defendant charged with a crime for which death penalty could be imposed you do not get a life-term imprisonment sentence. Rather, you get a lottery ticket that gives you a certain risk of receiving the death penalty. I am not putting this forward as a serious proposal. But note that given the assumptions on which the discussion is based, it is hard to see why it is not better than Nathanson's proposal. True, it involves an arbitrary unfairness in terms of who among the pool of comparable defendants receives the death penalty, but as I argued above, this unfairness seems smaller than the unfairness involved in differential capital punishment rates reflecting discrimination against certain groups. Also, unlike Nathanson's proposal the capital punishment lottery does not leave people who deserve capital punishment off the hook.

I conclude that given the assumptions made and insofar as discriminatory capital punishment constitutes an injustice it is not clear that the best alternative to the present practice is to abolish capital punishment altogether. It is worth stressing that this conclusion is different from a defense of capital punishment. First, a number of considerations other than justice bear on whether capital punishment should be abolished.[42] Second, the discussion has been conducted on the basis of a number of assumptions that might be rejected, e.g., the retributivist assumption that murderers generally deserve capital punishment, and that discriminatory biases affect only whom among those who deserve the death penalty actually receive it.

5. The No-Complaint Argument

In the final section I address an argument—the "no-complaint argument"—that seems to show that discriminatory punishment may be just:

(1) It is true of all those who are punished either that they deserve the punishment they get or that the punishment they get is comparable to the harm they have imposed on the victim.
(2) If one deserves the punishment one gets or when it is comparable to the harm one has imposed on one's victims, one cannot complain about one's punishment.
(3) Hence, none of those punished can complain about being punished.
(4) If no one can complain about being punished, no one is punished unjustly.
(5) Hence, none of those being punished is punished unjustly.
(6) If none is being punished unjustly, the scheme of punishment is just.
(7) Hence, the scheme of punishment is just.

[42] See, for instance, Reiman in Reiman and Pojman (1998), *The Death Penalty*, pp. 107–118.

So, by way of illustration, suppose that all who are sentenced to death are murderers who deserve capital punishment and for whom the harm is comparable to what they imposed on one or several of their victims, then (1) is true in their case. It follows from this and (2) that these people cannot complain about being sentenced to death. Since they are not in a position to complain about their punishment, none of them is being punished unjustly (by (4)), and accordingly the scheme of capital punishment is just (by (6)). If this argument is sound, discriminatory sentencing is just provided that no one receives a punishment they do not deserve. This conclusion is favorable to van den Haag's position.

Other versions of this argument are suggested by remarks by Jeffrey Reiman (who ultimately rejects the death penalty) and Saul Smilansky (who discusses punishment in general and, thus, takes no stand on capital punishment in the work quoted). Reiman entertains a Kantian argument in favor of capital punishment which he takes to imply that "a rational being cannot complain of being treated in the way he has treated others, and where there is no valid complaint, there is no injustice, and where there is no injustice, others have acted within their rights."[43] Similarly, Saul Smilansky writes that: "morally, a person cannot complain when others treat him or her in ways similar to those in which the complainer freely treats others."[44]

I concede the first two premises of the argument. There is a clear difference between a complaint of being subjected to capital punishment pressed by a convicted serial killer on his own behalf and a complaint from a human rights activist with no criminal record on behalf of the convicted serial killer. Suppose the killer says: "Human life is sacred and therefore you should spare mine." It seems right to dismiss the complaint—"Look who's talking"—on the ground that the serial killer is not in a position to complain given his own violation of the very norm to which he appeals.[45] In a wide range of cases, there is a sense in which one is not in a position to complain about a certain treatment when that complaint is based on an appeal to the very norm that one has violated to an equal or greater degree.

My primary quarrel with the argument is with premise (4). I want to put forward two objections. First, it is important to be precise about what the punishees are not in a position to complain about. While it may be true that they

[43] Reiman (1985), "Justice," p. 124. Cf. I. Kant (2006 [1798]) *Metaphysics of Morals* in P. Kleingeld, *Toward Perpetual Peace and Other Writings Politics, Peace, and History* (New Haven: Yale University Press), 6:334.

[44] Saul Smilansky (2007), "The Paradox of Moral Complaint," in his *Ten Moral Paradoxes* (Oxford: Blackwell Publishing), 90–99, p. 91.

[45] G. A. Cohen (2006), "Casting the First Stone: Who Can, and Who Can't, Condemn the Terrorists?" *Royal Institute of Philosophy Supplement* 58, pp. 113–136; Kasper Lippert-Rasmusen (2013), "Who Can I Blame?" in Michael Kühler and Nadja Jelinek (eds.), *Autonomy and the Self* (Dordrecht: Springer).

are not in a position to complain about being punished, it does not follow that they are not in a position to complain about the following: that they are being punished, while others who are equally deserving of punishment are not. This complaint could be put forward even by a murderer (not guilty of similar differential treatment of his victims). So one might say: "True, I am not in a position to complain about being sentenced to death given that I am murderer, but I am in a position to complain about being sentenced to death when others equally deserving murderers get off much more lightly."[46] Accordingly, (4) is false: one might be treated unjustly even if one is not in a position to complain about one's punishment as such provided one is in a position to complain about some conjunctive fact of which this is an essential part.

Second, even if we set aside the objection of the previous paragraph there seems to be a wide range of situations when someone is not in a position to complain about how he is being treated where this does not show that he is not being treated unjustly.[47] For instance, it seems an unjustly harsh punishment to torture someone for torturing others even if a torturer is not in a position to complain about being tortured. Similarly, a person who believes that there is no such thing as a right to freedom of speech might not be in a position to complain about the police preventing him from voicing his opinion—at least not on the ground that they violate his right to freedom of speech.[48] Yet, those of us who believe in the right not to be tortured and a right to freedom of speech might agree that he cannot complain about the way he is being treated and yet also agree that it would be unjust to torture or silence him.

For these reasons, one can be treated unjustly even if one is not in a position to complain about being treated that way per se. Hence, even if it is true that those being sentenced under a discriminatory system of sentencing deserve the punishment they receive—which is highly unlikely in any case—and, for that or some other reason, is not in a position to complain about their punishments, this treatment might be unjust. Accordingly, the no-complaint argument fails to establish that discriminatory sentencing may not be unjust.

---

[46] One might deny this on the ground that one is not in a position to complain about a certain wrong, when one has committed a wrong that, albeit of a different kind, is much greater, e.g., a rapist cannot indignantly complain about being sexually harassed even though rape and sexual harassment are different wrongs along a number of important dimensions.

[47] Admittedly, there is a sense of "moral complaint" in which someone has a moral complaint if she is treated unjustly. However, if the term is used in that sense in the complaint argument premise (4) is a tautology and asserting premise (2) begs the question against those who think discriminatory punishment can be unjust, even when the punishee gets the punishment she deserves. Also, there is another sense of "moral complaint" in which someone may not have a moral complaint even if she is treated unjustly. This sense is the one employed in the Rawls quotation in the next endnote.

[48] "A person's right to complain is limited to violations of principles he acknowledges himself," John Rawls (2000), *A Theory of Justice*, rev. ed. (Cambridge, MA: Harvard University Press), p. 190.

6. Conclusion

It is complicated to establish the absence of sentencing discrimination and it is hard to determine its overall magnitude due to how discrimination against groups of lawbreakers interacts with discrimination against groups of potential victims of lawbreaking. To the extent that it exists, it is unjust if there is such a thing as comparative justice. Given retributivist premises, it may in some cases of discriminatory sentencing be more just to let the unjust sentencing system continue than to abolish it. Finally, it may be true of a criminal who receives the punishment he deserves through discriminatory sentencing that he is in no position to complain about his punishment even if it is false that his being punished involves no injustice.

{ 9 }

Reaction Qualifications

1. Introduction

In this chapter, I return to issues related to discrimination and the job-market and other similar contexts. More specifically, I will discuss the ideal of meritocracy, the realization of which has been seen by many as a bulwark against discrimination. Meritocracy plays an important part in the political self-understanding of most contemporary Western societies. We rightly take pride in the fact that, nowadays, differences between people's life prospects are determined less by accidents of birth such as race or sex than they were in the past. Many of us think that justice requires, among other things, that "there should be a fair competition among individuals for unequal positions in society," and that in this competition successes (and failures) ought to be determined by qualifications only.[1] It is testimony to the power of the ideal of meritocracy that, under the label "Fair Equality of Opportunity," it plays a prominent role in John Rawls's influential theory of justice.[2]

The ideal of meritocracy is often regarded as an extension of the ideal of nondiscrimination (and vice versa). This equation can be detected, for instance, in Lyndon B. Johnson's image of opening up "the gates of opportunity" by eliminating racial discrimination thereby ensuring a competition that is "completely fair."[3] In this chapter I shall take the ideal of meritocracy for granted with an eye to exploring the complex relationship it has with nondiscrimination (and, to a lesser extent, other forms of differential treatment such as nepotism). Shortly I say a bit more about the specific complexity I have in mind, but before this I need to clarify the *basic meritocratic norm*.

---

[1] James Fishkin (1983), *Justice, Equal Opportunity, and the Family* (New Haven, CT: Yale University Press), p. 19.

[2] John Rawls (1971), *A Theory of Justice* (Cambridge, MA: Harvard University Press), p. 72,

[3] Lyndon B. Johnson's commencement address at Howard University 1965; quoted from Steven M. Cahn (ed.) (1995), *The Affirmative Action Debate* (New York: Routledge), p. xii.

The basic meritocratic norm says that each selector ought to fill the position she is charged to fill with someone who is no less well qualified for that position than any other candidate.[4] Throughout this chapter I mean by "selector" an individual (or a body of individuals, e.g., a selection panel) that determines who will occupy a certain position to which the meritocratic norm applies. By "candidate" I mean someone who can potentially occupy such a position. Employers, applicants, and jobs constitute paradigmatic selectors, candidates, and positions, respectively. However, the scope of the meritocratic norm extends beyond the job market—e.g., a prison warder who must assign a prisoner to fill a position that requires special skills and offers special privileges should, meritocratically speaking, pick the best-qualified prisoner rather than a personal favorite. Finally, just who is "no less well qualified" depends on whether we determine a candidate's qualifications at a micro- or macrolevel. Doing the former, we determine which candidate is best qualified, disregarding the way in which our selection of the candidate will indirectly affect the overall performance of the larger group of which the candidate will become a member if selected. Doing the latter, we take such effects into account. Having flagged this important complication, I shall ignore it in what follows, as it does not bear directly on the main issue of this chapter, i.e., reaction-qualifications.[5]

It is tempting to assume that where selectors comply with the basic meritocratic norm the resulting selection process will be innocent—that is to say, free of discriminatory attitudes and nepotistic biases. Each candidate's chance of obtaining the position appears then to depend wholly upon his or her abilities.[6] I shall now explain why things are not that simple.

---

[4] Fishkin (1983), *Justice*, esp. pp. 19–30; Anthony Flew (1981), *The Politics of Procrustes* (London: Temple Smith), pp. 45–58, 111–116; 61–92; Andrew Mason (2006), *Levelling the Playing Field* (Oxford: Oxford University Press), pp. 39–67; Norman Daniels (1978), "Merit and Meritocracy," *Philosophy & Public Affairs* 7, pp. 206–223; David Miller (2001), *Principles of Social Justice* (Cambridge, MA: Harvard University Press), pp. 131–202; George Sher (1988), "Qualifications, Fairness, and Desert," in Norman E. Bowie (ed.), *Equal Opportunity* (Boulder, CO: Westview Press), 113–127; Alan H. Goldman (1979), *Justice and Reverse Discrimination* (Princeton, NJ: Princeton University), esp. pp. 22–64; Michael Walzer (1983), *Spheres of Justice* (Oxford: Basil Blackwell), pp. 129–164; Shlomi Segall (2012), "Should the Best Qualified Be Appointed?," *Journal of Moral Philosophy* 9, pp. 31–54.

[5] Norman Daniels (1978) discusses micro- versus macrolevel qualifications in his "Merit and Meritocracy," *Philosophy & Public Affairs* 7, pp. 206–223.

[6] Suppose, implausibly, that everyone is equally qualified for every position. In that case, the basic meritocratic norm is compatible with all attractive positions being filled discriminatorily on the basis of ascriptive characteristics, in which case candidates' chances would not depend only upon their talents. In response to a related problem, David Miller (2001) remarks that the merit principle "works in tandem with the principle of equality of opportunity" (*Principles*, p. 314n22). This would rule out systematically selecting candidates on the basis of ascriptive characteristics. However, even this constellation of principles does not condemn a scenario in which everyone is equally qualified and half of the employers sex-discriminate against men, and the other half against women, thereby preserving equal opportunities. In what follows, I disregard these rather unusual circumstances under which meritocracy and equality of opportunity may be satisfied even though selectors discriminate.

In an influential study Alan Wertheimer draws a distinction between *technical* and *reaction* qualifications.[7] A candidate has good reaction qualifications when she has "abilities or characteristics which contribute to job effectiveness by causing or serving as the basis of the appropriate reaction in the recipients."[8] Following Wertheimer, I mean by "recipients" individuals whom a candidate interacts with in virtue of filling the relevant position. Customers and co-workers are paradigmatic "recipients," and paradigmatic "appropriate reactions" include buying the goods the candidate sells and uncomplainingly complying with orders issued by the candidate. A salesperson who can sell more cars than her colleagues because her winning personality encourages customers to buy has better reaction qualifications. Any qualification that is not a reaction qualification is a technical qualification. A software programmer who can write more error-free lines per hour than other programmers has better technical qualifications than those others.

Wertheimer argues brilliantly that the characterization of an innocent selection process sketched above becomes problematic when qualifications include reaction qualifications. Suppose such qualifications are important for a wide range of jobs, and that the majority of people dislike a certain minority whose members have worse reaction qualifications as a result. In that case, hiring the best qualified will systematically favor the majority. This puts the tight connection between meritocracy and antidiscrimination under pressure in a way that challenges simple conceptions of meritocracy. Andrew Mason puts the point as follows: "It would surely be quite out of keeping with the ideal of meritocracy to allow that a white candidate is best qualified for a job as a salesperson because most customers are racially prejudiced."[9]

When, in seeking to preserve a tight connection between the antidiscrimination norm and the ideal of meritocracy, we look at the options, the first position that springs to mind is the view that reaction qualifications do not count at all for the purpose of determining which candidate is best qualified.[10] However,

---

[7] Alan Wertheimer (1983), "Jobs, Qualifications, and Preferences," *Ethics* 94, pp. 99–112.

[8] Wertheimer (1983), "Jobs," p. 100. If the "serving as the basis of" component is necessary, the relevant relation must be noncausal; otherwise it is already covered by the "causing" component. However, I am not sure what relation Wertheimer has in mind.

[9] Mason (2006), *Levelling*, p. 34. See also Segall (2012), "Should the Best?," pp. 47–49 and Miller (2001), *Principles*, p. 314n22.

[10] The design of institutions, work practices, and the like may substantially determine employees' reaction qualifications. For instance, a sexist company's dress code may lead to greater differences in reaction qualifications across gender than a unisex code. For present purposes, I largely set aside issues raised by the way in which institutional contexts promote or inhibit reactions relevant to candidates' qualifications. The end-state norm that I introduce in section 3 has implications vis-à-vis the duties selectors have to design jobs in such a way that objectionable responses underlying reaction qualifications are minimized. Also, I set aside difficult issues regarding how positions should be defined, e.g., is it part of the job as an airline steward or stewardess to provide mild sexual titillation to passengers to ease the boredom of flying or is the job more narrowly and

this cannot be correct. In a wide range of jobs the aim is precisely to elicit a certain reaction in people, and the reactions in question, therefore, determine in good part the employee's effectiveness.[11] This is true, for instance, of jobs involving teaching, modeling, providing medical advice, policing, selling, and managing. One's performance in these positions depends crucially on one's reaction qualifications, and it would make little sense asking which candidate is best qualified if we disregarded these.

Seeing this, most theorists tackling the reaction qualification problem have pursued a different option. One such option—the second option—is to concede that reaction qualifications count from the point of view of meritocracy, but urge that moral concerns other than that of ensuring that the best qualified candidate is selected constrain applications of the meritocratic norm. For instance, broader concerns about distributive justice might imply that, all things considered, selectors ought not to select the best qualified where their doing so will disadvantage a minority via reaction qualifications rooted in discriminatory attitudes. So while, on this view, it might be true in some contexts that a candidate from a minority group is less well qualified as a police officer because she has worse reaction qualifications, the meritocratic norm may be trumped by a broader concern about distributive justice. Where it trumps the meritocratic norm, this concern will entail that, all things considered, selectors should deviate from the meritocratic norm by hiring a less qualified minority applicant.

The third option is to adopt a moralized notion of qualification in response to the reaction qualification problem.[12] The core idea here on the way in which reaction-qualifications bear on merit is that some, but not all, reaction qualifications play a role in determining which candidate is best qualified. Accordingly, the person best qualified when all the factors relevant to job performance have been taken into consideration, i.e., factors including those that should be disregarded on the present view, may not be best qualified in terms of moralized qualifications. On this view, the minority candidate from my previous example may well have the best reaction qualifications given a suitably moralized account of those qualifications. Accordingly, in selecting this person one complies with the meritocratic norm on its present interpretation. While the second

---

conventionally defined as selling beverages, making sure that safety protocols are complied with etc., see Kimberly A. Yuracko (2006), "Sameness, Subordination, and Perfectionism," *San Diego Law Review* 43, 857–897, pp. 872–875. The issue of reaction qualifications, which is the focus in this chapter, arises even if we adopt a narrow and conventional understanding of positions, e.g., as in the example above that the job of a salesperson is to maximize profits.

[11] Wertheimer (1983), "Jobs," p. 101.

[12] Cf. Wayne Sumner (1987), "Positive Sexism," *Social Philosophy and Policy* 5.1, 204–222, p. 207. Similarly, it is possible to adopt notions of job performance and of the aims of the company etc. relative to which characteristics are deemed to be qualifications that are moralized in the relevant dimension.

and the third option contain a good deal of truth, I shall argue here that neither is quite right.

The main aim of this chapter is to improve on our understanding of how we should think of the meritocratic norm in the light of reaction qualifications, especially problematic reactions qualifications. In section 2 I survey the three leading accounts of reaction qualifications: Wertheimer's, Mason's, and David Miller's. All of these are variants of the last two options briefly described in the previous two paragraphs. I then propose an account of my own which preserves symmetry between negative evaluations of antimeritocratic bases of selection and negative evaluations of qualifications rooted in comparable antimeritocratic reactions. If, say, employers should not select among applicants on the basis of their (the employers') racial preferences, customers should not respond to services offered by the applicant actually hired on the basis of their (the customers') racial preferences. One advantage of the symmetric account is that it provides a more coherent explanation of our judgment that counting certain reaction qualifications is undesirable, meritocratically speaking, than that offered by existing accounts. For existing accounts tend not to link acceptable bases for selection with acceptable bases for qualifications. Partly in the light of cases involving objectionable reaction qualifications, section 3 considers how we should decompose the meritocratic ideal into four separate norms. One important upshot of this move is that the meritocratic ideal applies not just to selectors, but to all of us. Here my analysis suggests that, in the light of problematic reaction qualifications, meritocratic norms can be seen to have a wider scope than is normally assumed. Section 4 takes a critical look at the assumption that qualifications rooted in preferences that are problematic from the point of view of merit, such as racist preferences, should be discounted by selectors across the board. I defend a view, more permissive of such qualifications than the one currently prevalent in the literature, according to which the ideal of meritocracy has a looser connection with the antidiscrimination norm than Mason and Miller, among others, believe. Specifically, there are cases where taking these preferences into account does not unfairly disadvantage anyone, including the targets of the objectionable preferences; and in these cases the resulting reaction qualifications need not be discounted by selectors even though such situations appear to be in other ways inimical to meritocracy. This underlying rationale, the less restrictive (selector-wise) view motivated by it, and the symmetric view expounded in section 2 represent the chapter's main contributions to the literature on reaction qualifications. In section 5, I rebut an objection to my view of reaction qualifications rooted in objectionable preferences, which grounds the norm of meritocracy in respect for agency. I argue that it is unclear that my view conflicts with the respect rationale, and that it is also not clear that this rationale really offers the best account of the ideal of meritocracy. Section 6 restates my conclusions.

2. Discounting Qualifications Based on Illegitimate Preferences

Wertheimer is the first theorist to address the issue of reaction qualifications at length. He submits that "counting [reaction qualifications] seems legitimate in some cases and (more or less) illegitimate in others."[13] For instance, he thinks that counting reaction qualifications rooted in racist views is illegitimate when doing so will systematically disadvantage members of the disfavored minority, whereas counting reaction qualifications rooted in patients' preference for a same-sex doctor is legitimate. However, he thinks that it is impossible to offer a neat test showing which reaction qualifications can be legitimately counted. Instead, he offers a list of four factors which will sort the acceptable from the unacceptable in most cases. First, "the nature of the preference upon which the reaction qualification is based" matters: the more "illegitimate" the preferences upon which the reaction qualification is based, the less it counts when we apply the meritocratic norm.[14] The following four characteristics make a preference less legitimate: the preference concerns characteristics that "are not acquired through effort or ascribed (in some way) to the agent's free action" (e.g., the candidate's race); "it is based on hierarchical judgments about social groups" (e.g., a sexist view favoring the subordination of women); it is for members of a dominant group and is held by persons who are members of that (or another dominant?) group (e.g., a preference by members of a dominant religious group for other members of that group); or it is controllable or not "deeply rooted in developmental patterns" (e.g., to the extent that it were a "general fact of child development that six-year-olds respond better to female teachers" it would be permissible to count this positive reaction qualification of female applicants when recruiting a teacher).[15]

Wertheimer mentions three additional factors that do not concern the nature of the preference in which the reaction qualification is rooted: in jobs typically involving relatively intimate interaction, such as the job of a doctor, "something akin to a right of association, may exercise some moral pull"; the greater the loss in social utility involved in ignoring reaction qualifications, the more legitimate it is to count them; and, finally, counting certain reaction qualifications based on preferences that are innocent in themselves may conflict with social justice when counting them would systematically disadvantage members of particular social groups.[16]

While I agree with Wertheimer that, at some level, all of the four factors mentioned—the nature of preferences, the quasi-right of association, utility, and distributive justice—affect the legitimacy of reaction qualifications, I want

[13] Wertheimer (1983), "Jobs," p. 106.
[14] Wertheimer (1983), "Jobs," p. 107.
[15] Wertheimer (1983), "Jobs," p. 107.
[16] Wertheimer (1983), "Jobs," pp. 108–112.

to set aside the last three factors for present purposes.[17] First, unlike the nature of preferences, the last three factors do not bear specifically on the issue of reaction qualifications. They are also relevant to determining the legitimacy of *technical* qualifications. For example, if physical strength is a technical qualification for many jobs, counting this qualification may unjustly disadvantage women. Second, the quasi-right of association will at most establish that in a rather narrow range of cases people have a quasi-right to act in ways that conflict with meritocracy. Also, unlike the appeals to utility and distributive justice, the appeal to the quasi-right of association cannot show that it is morally acceptable, all things considered, to act in a way that conflicts with meritocracy. It can only show that people have a right to act in this morally unacceptable way. Such a conclusion is of little interest to an enquiry concerning the nature of the meritocratic ideal.

Let us take a closer look at the four characteristics listed in the first paragraph of the present section. Wertheimer takes these to affect the legitimacy of preferences and, derivatively, the legitimacy of reaction qualifications rooted in these preferences.[18] I shall start with the first characteristic (whether the preferences concern characteristics reflecting effort or the agent's free action) and the last (the controllability of the preferences). In these cases, I believe the issue is not really the legitimacy of the preferences themselves but the legitimacy of *counting* them. Unlike, say, a preference for male bosses rooted in sexist hierarchical judgments, there does not seem to be anything illegitimate as such about a preference for characteristics that its possessor did not acquire through effort. It is not as if the more of a person's preferences for personal attributes that concern characteristics "acquired through effort or ascribed (in some way) to the agent's free action," the more just this person is. Also, if equality of opportunity obtains, and if some people have, say, worse reaction qualifications for some jobs on account of characteristics they did not acquire through their own effort or on account of other people having preferences under their full control, presumably, this disadvantage must be offset by their having better (reaction) qualifications for other jobs. In this situation, we would not consider accommodating these preferences illegitimate. This suggests that the reason we deem it illegitimate to count them in normal cases has to do with the consequences of doing so—e.g., that accommodating them leads to inequality of opportunity or is harmful to people.

---

[17] For instance, I agree that when sufficient amounts of welfare are at stake, all things considered, it may be immoral to disregard qualifications rooted in immoral reactions; see Larry Alexander (1992-1993), "What Makes Wrongful Discrimination Wrong?" *University of Pennsylvania Law Review* 141, 149–219, p. 176, and chapter 4. I insert the qualification "at some level" because, from the perspective of the harm-based account (see chapter 6), acting on preferences based on the view that some people are inferior to others are often humiliating and for that reason more harmful than discriminatory acts that are comparable in all other respects.

[18] Wertheimer uses the term "legitimate," but I take it that he means by the term more or less the same as what I mean by "not wrong."

What about reaction qualifications rooted in preferences for (other) members of the dominant group? There are three cases to consider. In the first, the preference involves a hierarchical judgment about the dominant and the dominated group. Here Wertheimer's second characteristic is what really lies behind the illegitimacy of the preference. In the second, the preference is not so grounded and it is entirely contingent that it is a preference had by a member of the dominant group for other members of that group. For example, it might be a preference had by one member of a linguistic community for other people who have the same mother tongue, where this linguistic community happens to be the dominant one. In this case, the preference in itself is hardly illegitimate—it is not a preference for other members of the dominant group *qua members of this group*—even if counting reaction qualifications rooted in it may be (e.g., because doing this systematically disadvantages members of the dominated group in a way that conflicts with social justice). In the third case, the preference is a brute, intrinsic preference in favor of dominant groups. Such a preference seems no less objectionable from the point of view of merit than an otherwise comparable preference rooted in a value judgment to the effect that the dominant group is superior. In fact, although this is out of keeping with Wertheimer's account, such a preference seems objectionable even when had by people who are not themselves members of dominant groups.

Wertheimer's conception of the way in which hierarchical judgments delegitimize derivative qualifications is not entirely satisfactory. It does not address the issue of hierarchical judgments about people who do not make up social groups. Wertheimer would accept that not all groups of people are social groups (he surely must do so), but he offers little by way of explanation of what characterizes a social group as opposed to a mere "category" of people.[19] One suggestion here is that members of social groups share certain characteristics that are significant in a wide range of social settings, and that these characteristics are important to people's self-conception (see chapter 1). This would go a long way towards explaining why, generally speaking, the counting of reaction qualifications rooted in hierarchical judgments about social groups is more problematic than it would be were the reaction qualifications rooted in other hierarchical judgments. For such judgments are more likely to contribute to the stigmatization of certain groups; they are also more likely to be widespread, and hence it is more likely that counting them will disadvantage social groups judged to be inferior (see chapter 6). However, this is neither here nor there, since Wertheimer's view is that these preferences are illegitimate independently of any consequences of counting qualifications rooted in them.

[19] Wertheimer (1983) mentions people who are nice, or sensitive, or intelligent, as collections of people who do not form social groups, "Jobs," p. 107n19.

Moving now to Mason's account, we note that, like Wertheimer, Mason thinks that some, but not all, reaction qualifications count.[20] In fact, he thinks that any plausible version of the meritocratic ideal rests on a "moralized account of qualifications" appealing to "some broader conception of equality of opportunity (or justice more generally)."[21] Mason does not set out exactly what this broader conception of distributive justice says. Nor does he explain its implications vis–à–vis the question of what constitutes a qualification. However, it is clear that he takes the broader conception, plausibly elaborated, to "accommodate the idea that prejudice against members of a group cannot legitimately make it the case that not being a member of that group counts as a qualification for a job."[22]

Mason's approach differs from Wertheimer's in two important respects. First, it involves the idea that qualifications cannot depend on prejudice against social groups, whereas none of the four factors that Wertheimer mentions, and which determine whether a qualification is rooted in preferences that are by their very nature illegitimate, concerns prejudice. To be sure, Wertheimer and Mason appeal to the same kinds of example to illustrate the distortion of merit involved in counting certain reaction qualifications, and certainly hierarchical judgments about social groups may well be highly correlated with prejudice against social groups that are judged to be inferior and vice versa. However, while it is difficult, psychologically speaking, to consider members of an entire social group to be inferior without the assistance of prejudice, this is not psychologically impossible; and certainly it is not conceptually impossible.[23] Members of one group can be prejudiced against members of another group that they do not consider inferior, e.g., members of an oppressed group might harbor various prejudices against members of an oppressing group even if they do not consider them to be inferior. Conversely, members of one group can judge a certain group to be inferior without being prejudiced against it.

On reflection, it is not clear why Mason focuses on prejudice in particular. Suppose people generally hold false, but reasonable (on the available evidence) beliefs about a certain minority group. Suppose also that this negatively affects the way in which they respond to members of this group. Mason does not appear to think that reaction qualifications cannot be so grounded. Yet, if the reason prejudices do not count towards qualifications is that it would be unjust

---

[20] Mason (2006), *Levelling*, p. 33.

[21] Andrew Mason (2001), "Equality of Opportunity," *Ethics* 111.4: 760–781, p. 778; Mason (2006), *Levelling*, pp. 32, 38; see also David Miller (1992), "Deserving Jobs," *Philosophical Quarterly* 42, 161–181, pp. 175–176.

[22] Mason (2006), *Levelling*, pp. 34–35. This claim would appear to apply even to the case where members of a certain group have worse reaction qualifications because they are prejudiced against other members of their own group.

[23] The entry in the *New Shorter Oxford English Dictionary* for "prejudice" states: "Preconceived opinion not based on actual experience; bias, partiality . . . A prior judgement; *esp*. a judgement formed hastily or before due consideration."

for them to do so, it is not clear why one cannot similarly argue that it is unjust to members of a minority that people respond to them on the basis of false beliefs about them. And if that is so, it is not so much the fact that people's prejudices involve motivated beliefs not supported by their evidence as the fact that these beliefs are false that explains why the relevant reaction qualifications should not be counted. Moreover, there are cases involving, not prejudice or, for that matter, false beliefs, but simply dislike or even hostility where, presumably, the qualifications cannot be rooted in such noncognitive responses.[24] Finally, there are the sorts of hierarchical judgment which Wertheimer focuses on and which, logically speaking, can be prized apart from prejudice.

Probably the best way to understand Mason is to assume that he mentions the case of prejudice simply as an example of the way in which the notion of a qualification must be unpacked within a broader theory of justice, leaving open the possibility that there are other examples. Certainly, he says nothing that has the effect of ruling out this interpretation.[25] Moreover, he appeals to the notion of disrespect for persons to explain the notion of a qualification: we "can allow that attitudes of customers, clients, and members of the workforce may crucially affect what counts as a qualification but insist that these attitudes can do so only if they do not display a failure to respect others."[26] But then reactions rooted in prejudices will be permitted to count provided they are not disrespectful. Equally, reactions not rooted in prejudices should not count if they are disrespectful. Qualifications rooted in hatred of another group, for instance, should not count even if they are not mixed up with prejudice provided that the hate involves disrespect.

The second way in which the approaches of Wertheimer and Mason differ is this. Wertheimer wants to identify the conditions under which it is legitimate to count reaction qualifications. Accordingly, his factors are relevant as factors that limit the meritocratic norm. Setting aside other differences between Wertheimer and Mason, Mason would say that the conditions Wertheimer identifies are those that determine, in part, who in the relevant sense really is best qualified. So Mason's considerations regarding a moralized notion of qualification do not operate to constrain the meritocratic norm. On the contrary, they are intended to determine what would count as the correct application of this norm. Now, the view that selectors should choose candidates who are best

[24] In one place Mason mentions in passing that it is unjust for selectors to choose among candidates on the basis of their dislike of the groups of which they are members (Mason [2006], *Levelling*, p. 57). Perhaps he has a wider sense of "prejudice" in mind than that provided by the dictionary.

[25] Mason says that his "considerations suggest that any adequate account of 'qualifications' has to meet two conditions," one of which is to allow for the possibility of reaction qualifications and the other of which is to discount prejudice when determining qualifications (Mason [2006], *Levelling*, p. 34). This does not rule out the possibility that Mason would insist that there are further conditions to be met. However, he does not indicate a commitment to such further conditions.

[26] Mason (2006), *Levelling*, p. 63.

qualified on the relevant moralized conception of qualification may, all other things being equal, be only nominally different from the view that, within the constraints set by certain qualification-external moral concerns, selectors should choose candidates who are best qualified. At least this is so when the constraints within which selectors should choose the best qualified on the latter view are the very same moral concerns, which, on the former view, determine who counts as best qualified. In this case, these views might lead to the selection of the same persons. Still, to avoid confusion this nominal difference should be kept in mind.

Consider, finally, David Miller's brief discussion of reaction qualifications. Like Wertheimer and Mason, Miller thinks that some, but not all, reaction qualifications count.[27] However, he stresses that it is difficult to draw the line between reactions that are relevant to merit and those that are not. Negative reactions to candidates based on expectations that are "of a discriminatory or antimeritocratic kind," including "the belief that there is a whole category of persons who are unfit to take on work of a certain kind," clearly are not relevant to merit, Miller submits.[28] But there are intermediate cases—the situation, for example, where employees are expected to "retire to the pub at the end of the working day, and the camaraderie thus developed eases working relationships," and where people with family responsibilities may find it hard to take part.[29]

Taking our cue from Miller's focus on expectations, or (better) attitudes, that are of an "antimeritocratic kind"—I prefer "attitudes" to "expectations" because, unlike expectations, attitudes can be both cognitive and noncognitive in character—I now want to present my own account of the reaction qualifications that count:

> *The symmetric view*: (1) An attitude is antimeritocratic in a non-individualized way if, and only if, (a) its object is a group of people and (b) it is such that if a selector decides between candidates influenced by it, the decision is not based solely on merit.[30] (2) A reaction qualification

---

[27] Miller (1999), *Principles*, p. 192.

[28] Since it is not obviously antimeritocratic to hold that blind people are unfit to take on work as commercial airline pilots, Miller must have in mind a certain subset of person categories.

[29] Miller (1999), *Principles*, pp. 192, 314n25.

[30] By a "group of people" I mean any group of people, not just groups of people that qualify as "social groups" on Wertheimer's account. Note also that it is not uninformative, because circular, to say that an attitude is antimeritocratic if, and only if, "(b) it is such that if a selector chooses between candidates influenced by it, the selector decides in a way that amounts to not choosing solely on the basis of merit." I define antimeritocratic attitudes through antimeritocratic choices (and not through antimeritocratic attitudes themselves, which would indeed by uninformative because circular. True, I do not define merit, but then that is not the notion my definition is supposed to capture.) Third, my account allows that attitudes in people other than selectors, i.e., recipients, may be antimeritocratic. Note, finally, that (1) allows an attitude to be antimeritocratic (in an individualized way) even if its object is not a group of people, as when a selector dislikes the best qualified candidate for no particular reason and therefore selects someone else.

> counts from the point of merit, i.e., makes a difference to who is best qualified in the relevant sense, if, and only if, it is not rooted in non-individualized, antimeritocratic attitudes.

By *an attitude that is nonindividualized* I mean an attitude whose object is a group of people, or one that is rooted in some such attitude. If a male person dislikes being given orders by his female superior because he dislikes women being in superior positions, his attitude is nonindividualized. If he dislikes being given orders by her simply because he dislikes her style of leadership, although he does not conceive of that style as typical of female superiors, or indeed of members of some other (e.g., racial or religious) group the properties of which he takes his superior to instantiate, his attitude is individualized.

By *an attitude that is not based solely on merit* I mean an attitude that is not based exclusively on an unbiased estimate of how well the candidate will perform in the position for which she is being selected. Suppose that the relevant position is that of a car salesperson, and assume that the only task of the appointee is simply to sell as many cars as possible. In this case, any appointment influenced by a factor that is irrelevant to this person's likely sales performance, such as whether this person belongs to a particular religious minority, or by a biased estimate of a relevant factor, will not be based solely on merit. I readily concede that questions about the factors that will in point of fact determine how well a person performs in a certain position are often open to discussion, partly because it is often open to discussion what functions a certain position will involve. By way of illustration: does the fact that a certain candidate has a more attractive personality and is therefore more pleasant to work with imply that she performs better in her job, all other things being equal? For present purposes we can set this issue aside and merely note that my definition can be variously applied depending on the factors that determine the quality of a candidate's performance.

On the symmetric view we discount reaction qualifications rooted in the nonindividualized, antimeritocratic attitudes we want to prevent from influencing the selector's decisions, e.g., racial preferences. Hence we know which reaction qualifications count and which do not once we know which attitudes are antimeritocratic in a nonindividualized way. In a sense, then, the issue of which reaction qualifications count cannot be raised without also asking what, precisely, meritocracy requires of selectors when they choose among candidates. Here my account differs from those offered by other theorists in that the latter do not explicitly tie permissible reaction qualification to permissible grounds for selection.

Several considerations support the symmetric view. First, the issue of reaction qualifications is thought to be intellectually intriguing because the very factors that we want selectors to ignore, such as the sex of the applicant, may nevertheless help to determine who, given a nonmoralized account of

qualifications, is best qualified—and thus indirectly influence selectors' choice in a way that distorts the ideal of merit. But this picture presupposes that if selectors should not be influenced by a certain nonindividualized attitude, nor should recipients' reactions to candidates be influenced by that attitude thereby affecting candidates' reaction qualifications.[31]

Second, given the rather uncontroversial view that racist, sexist and other discriminatory attitudes are of the nonindividualized, antimeritocratic kind, the symmetric view accommodates the examples of reaction qualifications that clearly should not count offered by Wertheimer, Mason, and Miller. Thus it deals well with the candidate who has worse reaction qualifications due to preferences rooted in supremacist judgments, prejudice, or the expectation that people in a particular category are as such unfit for a certain kind of job.

Third, as we have seen neither Wertheimer nor Mason provides an account of the full range of problematic reaction qualifications. The symmetric account explains why: neither of their accounts reflects the full range of nonindividualized, antimeritocratic attitudes that we want to prevent from influencing selectorial choice.

Finally, the symmetric account handles seemingly powerful counterexamples that can be raised against it well. For instance, it is acceptable for recipients to react positively to whether candidates are charming, but selector's decisions should not be directly influenced by whether candidates are charming. Hence, there is an asymmetry between individualized attitudes that can determine reaction qualifications and individualized attitudes that can determine selectors' choice in ways that are compatible with merit. However, because the symmetric view applies only to nonindividualized attitudes it accommodates these claims. In the standard case, a selector might overestimate the qualifications of a charming candidate, or she might take the fact that she finds the candidate charming to be good reason to think that others may do so as well—something that may count as a qualification. But in neither of these cases does the selector have an attitude to the group of charming people. She is free of the nonindividualized, antimeritocratic attitude even though her feeling that a person is charming might influence her to choose that candidate rather than another candidate who might be better qualified.

Generally speaking, one responds to a charming person on the basis of his or her distinctive behavior, not because one has an attitude to a certain group (the group of charming people) and takes the person to be a member of that group. However, cases of the latter kind are possible. Here the symmetric view says that a selector whose decision is influenced by such an attitude is

[31] Like the second and the third consideration below, this one supports the "only if" claim, but not the "if" claim, in clause (2) of the symmetric view. The fourth consideration supports the "if" claim.

influenced by antimeritocratic views. But this does not represent a problem for the symmetric view, because reaction qualifications rooted in responses to a candidate qua member of the group of charming people seem to distort merit in exactly the same way that reaction qualifications rooted in preferences for other groups do. For instance, compare the normally not so charming waiter to whom customers respond less well, even though he for once serves the coffee when it is hot and manages to be cheerful and amiable, with a waiter with a reputation for being charming who is, on this occasion, grumpy and serves the coffee lukewarm. The former might complain on this occasion, if customers prefer the services of the latter, that he is not being judged on his merits.[32]

One suggestion as to why individualized, antimeritocratic attitudes generally bear differently on the selection process and reaction-qualifications is this. Selectors typically act as agents or representatives of a collective body, such as a university, whereas recipients typically act on their own behalf. Agents who act in the first capacity, unlike those who act on their own behalf, ought to base their decisions on the goals of the collective entity as a whole—these, not the idiosyncratic attitudes of selectors, are what define qualifications. This explanation implies that when selectors act on their own behalf they are permitted to base their decisions on their own individualized attitudes; and that when recipients act as agents of a larger, collective body they ought to disregard their own individualized attitudes and base their decisions on the relevant body's goals. So, to illustrate the first point, an employer who hires a person for her own shop may count the fact that she likes a certain candidate as a qualification. After all, within limits she may define what aims the shop pursues and, thus, which characteristics count as qualifications. To illustrate the second point, it is problematic from the point of view of merit if customers who represent collective entities base their decisions on individualized attitudes to salespersons. So, suppose a charming, unimaginative architect is hired instead of a reclusive, highly innovative architect; and that because the appointing officials base their decisions to a large extent on personal preferences, the more reclusive applicant is less likely than his charming competitor to win contracts for the company. In this situation, the recluse may complain that she is not being judged on merits. I do not propose this explanation with great confidence, and this is why my symmetric view only ranges over nonindividualized, antimeritocratic attitudes. But note that if it is correct, it implies that a symmetry claim

[32] Arguably, Wertheimer's example of the university health service that prefers a female to a male gynecologist because she has better reaction qualifications can be accommodated in the same way (Wertheimer [1983] "Jobs," p. 101). For instance, a brilliant, male gynecologist, who is superb in making his patients feel comfortable during medical examinations, can complain that he is not being judged on his merits when patients prefer his female colleague simply because of her sex, even if, as a result of something like Wertheimer's quasi-right-of-association, patients are within their rights to act on their preference for a female gynecologist.

analogous to that expressed by the symmetric view applies also in the case of individualized, antimeritocratic attitudes.[33]

3. Refining Meritocracy

Consider now reaction qualifications that derive from antimeritocratic attitudes. According to Wertheimer, Mason, and Miller, and according to the symmetric view proposed in the previous section, such qualifications do not count from the point of merit. This suggests that in determining selectors' compliance with the basic meritocratic norm, we should simply disregard reaction qualifications rooted in antimeritocratic attitudes. To test this suggestion, compare the following three employers:

Racist Customers: The first, profit-maximizing employer has racially prejudiced customers, while she herself is entirely innocent of racism. She hires a white salesperson, because she has been told that her iPod, with all her favorite music, will be destroyed unless she hires this particular person. The person she hires will sell twice as much as the best black candidate. Had the customers been racially neutral the candidate who sold most would have been a black person. Let us furthermore suppose that under racial neutrality this person would have sold twice as much as the candidate actually hired sells under the actual circumstances of racial partiality. While the employer caters to recipients' illegitimate preferences in hiring the person with the best reaction qualifications in the sense that she acts in such a way that these preferences will be satisfied, neither her knowledge of these preferences, nor their existence, motivates or otherwise explains her choice among candidates. Rather, she selects the white salesperson to save her iPod. The employer is aware of all these facts.

No Racists: Neither the second employer nor her customers are racially biased. However, the employer hires a white salesperson, because she has been receiving a threat identical to the one received by the first employer. The person hired is able to sell half as much as the best qualified candidate, who happens to be black. Again, the employer is aware of all the facts.

Racist Employer: The third employer's customers are not racially biased. However, the employer is a racist and while she has been threatened that her iPod will be destroyed unless she hires a particular white candidate,

[33] The unrestricted symmetry claims says: (1) an attitude is antimeritocratic if, and only if, it is such that if a selector chooses between candidates influenced by it, the selector decides in a way that amounts to not choosing solely on the basis of merit. (2) A reaction qualification counts from the point of merit if, and only if, it is not rooted in antimeritocratic attitudes.

> she in fact hires this person on racial grounds alone. The person hired is able to sell half as much as the best qualified candidate, who is black. Again, the employer is aware of all the facts.

Let us suppose that qualifications are wholly determined by how much a salesperson can sell, adjusted for illegitimate factors determining that. On this assumption each of the three employers hires a candidate who is only half as qualified as the best qualified candidate, in which case all three employers' violate the basic meritocratic norm to the same degree. (The threat is not so grave as to make it permissible to infringe the meritocratic norm.)

But it does not seem right that the racist employer acts in a way that clashes no more with a concern for merit than do the first two employers. Also, the black candidate in the second example, and even more so in the third example, seems to have a complaint against the employer, which the black candidate in the first example lacks (which is not to say that the latter candidate has no complaint at all).[34] Unlike in the last two cases, the employer in Racist Customers can adopt an excusatory stance. That is, she can concede that while she was wrong to care so much about her iPod at least this disproportionate concern of hers did not lead her to hire a salesperson with less good selling abilities. Finally, in No Racists illegitimate reaction qualifications are not an issue. Hence, it would seem that, from the point of view of merit, the employer would have to stand firm and hire the salesperson who will actually sell most in the face of threats graver than those the employer in Racist Customers would have to resist in order to hire the salesperson who would sell most if customers were not racially biased.[35] This suggests that reaction qualifications rooted in objectionable preferences cannot simply be irrelevant from the point of merit even if they are not unproblematic either.

One way to explain these intuitions is to supplement the basic meritocratic norm with three additional norms. The first—the *recipient norm*—enjoins recipients not to respond to candidates on antimeritocratic grounds. Like the basic meritocratic norm, this is an agent-relative norm giving each recipient the aim that she does not respond to candidates on the basis of, say, sex or race, but on the basis of qualifications only. So, by way of illustration, as a customer one ought, meritocratically speaking, to suppress (or, better, eliminate) whatever biases one may have when dealing with the representatives of a company so that one does not respond differently to financial advice from a same-race bank representative and similar advise from a different-race bank representative.

---

[34] I assume that the duty of selectors to hire the best qualified candidate corresponds to a claim-right of the best qualified candidate to be hired: see, e.g., Sher (1988), "Qualifications," p. 113.

[35] This is not to deny that there may be reasons other than meritocratic ones to stand firm in face of graver threats in Racist Customers. It may be known, for example, that giving in to the threat in the latter case is more likely to harm someone who is worse off.

Meritocracy is almost always described as a view that constrains the way in which selectors are permitted to choose between candidates. The case of reaction qualifications shows why this focus is too narrow.

The second additional norm—the *noncollaboration norm*—enjoins selectors to offer positions to candidates who are best qualified when reactions rooted in antimeritocratic attitudes are disregarded. Arguably, this agent-relative duty is less stringent than the duty a selector has not to select on the basis of her own illegitimate preferences.[36] At least, it seems that selectors have to bear greater costs to avoid selecting candidates on this basis than they have to bear to avoid selecting candidates on a basis that includes reaction qualifications rooted in the illegitimate preferences of recipients.[37] On the present proposal, the basic meritocratic norm should be interpreted in the light of the noncollaboration norm. When objectionable reaction qualifications enter the picture selectors ought to choose the candidate who is best qualified when objectionable reaction qualifications are discounted—so in that sense reaction qualifications rooted in antimeritocratic reactions do not count from the point of view of merit. However, this duty is less stringent than the duty to hire the best qualified candidate in an otherwise comparable case involving no objectionable reaction qualifications. Thus, in this respect, even reaction qualifications rooted in antimeritocratic reactions do count from the point of view of merit. Correspondingly, consider cases where no problematic reaction qualifications bear on the position in question. Obviously, in these cases, from the point of view of merit, the candidates who are best qualified on the basis of nonproblematic reaction qualifications can make a stringent claim on the selectors to be offered the relevant positions. Moreover, this claim is more stringent than the claim the best qualified candidates have in cases where problematic reaction qualifications do bear on the position in question and these are discounted to determine who is best qualified, all other things being equal.

The third additional norm—the *end-state norm*—says that each of us ought to do our bit to bring about a state of affairs in which, first, selectors fill positions with someone who is no less well qualified for that position than any other candidate, and second, recipients do not respond to candidates on

---

[36] By saying that one duty is more stringent than another, I mean that for it to be morally permissible for one to fail to do the former duty it takes worse consequences of fulfilling it than for it to be morally permissible for one to permissibly fail to fulfill the latter. By way of illustration: my duty of not torturing an innocent is more stringent than my duty to return a corrected student essay on time, because there is a wide range of bad consequences of fulfilling the latter duty, e.g., that fulfilling it requires that I could not take my child to the hospital, that would render it morally permissible not to fulfill it, but not render it morally permissible not to fulfill the former duty.

[37] Arguably, a respect-based account of meritocracy can account for that by appeal to the claim that it is more disrespectful to reject candidates in response to one's own illegitimate preferences than it is to reject them because other people's illegitimate preferences affect one's pursuit of legitimate aims, such as the aim to make a living.

antimeritocratic grounds. Unlike the three other meritocratic norms, this norm is agent-neutral and gives all of us the same aim.

The conjunction of the four meritocratic norms enables us to replace the somewhat vague submission, made by Mason, Wertheimer, and Miller, that "it is" illegitimate to count reaction qualifications (illegitimate for *whom* to count? Presumably, they have in mind selectors) with more refined talk about *who* violates *which* meritocratic norm. Note first that in all cases where people are selected on the basis of objectionable reaction qualifications, recipients violate the recipient norm. In most such cases, selectors violate the noncollaboration and the end-state norms by catering to recipients' illegitimate preferences, thereby presumably entrenching them further.

It is true that outlying cases can be imagined in which selectors do not violate the end-state norm. Suppose, oddly enough, that by catering to recipients' illegitimate preferences a selector will actually promote improvement from the point of view of merit—e.g., because candidates with the preferred gender (or whatever it is) have such bad technical qualifications that recipients will come to reconsider their sexist preferences. In such cases, a selector who hires on the basis of illegitimate reaction qualifications will comply with, not violate, the end-state norm. However, this seems perfectly acceptable. Consider the unlikely eventuality of the selector knowing that by hiring the person who will sell most (given her illegitimate reaction qualifications and despite her lousy technical qualifications) she will actually foster a situation in which candidates' meritocratically unimpeachable qualifications are the only determinants of their ability to sell (and hence to be chosen by selectors). In this case no violation of the end-state norm would occur. We might approve of the selection decision, all things considered.

Second, the noncollaboration and end-state norms explain why selectors and recipients cannot complain about laws criminalizing selections and recipient responses based on illegitimate preferences. Since people have a duty not to cater to recipients' objectionable norms, and to promote a state in which a person's qualifications are judged by his or her genuine qualifications only, they are presumably required to bear some costs should these be necessary to bring about this state.[38] Accordingly, the noncollaboration and the end-state norms

[38] In cases where employers do not control whether their recipients have objectionable norms employers, who through no responsibility of their own, have many recipients with objectionable norms may complain that it is unjust that they will must bear greater burdens in not catering for their recipients' objectionable norms than employers who are fortunate enough to have fewer recipients with objectionable norms. I agree with this observation, but believe that while it may motivate some refinement of the view stated in the main text, the core claim about the permissibility of the state forbidding employers from hiring on the basis of illegitimate reaction qualifications even though not doing so will eat into the employer's profits is not affected. In any case, I think employers may have a duty to bear some costs of the relevant sort even if the state does not ensure that costs are justly shared between different employers. Thanks to Shlomi Segall for pointing out the need to address this complication.

explain why it is permissible for the state to forbid employers from hiring on the basis of illegitimate reaction qualifications even though not doing so will eat into the employer's profits. And they explain why it is permissible for the state to make it illegal for customers to, say, order a cab with a white driver even if compliance with that law may oblige some customers to accept a ride which they will be less comfortable about.

Finally, the present conjunction of norms explains why we evaluate the three employers differently. The racist employer violates the basic meritocratic norm by choosing a candidate on an antimeritocratic basis. This norm imposes a duty on employers more stringent than those imposed by the norms violated by the two other employers. Specifically, to compare the first and the third employer, selectors have a right to disregard recipients' illegitimate preferences, and even some duty to do so. However, it would seem that the duty they violate in not disregarding recipients' illegitimate preferences is less stringent than the duty they violate if they do not disregard their own illegitimate preferences (as the racist employer does not). From this point of view, there appears to be a relevant difference between the first, profit-maximizing employer and the more problematic second employer. When the latter fails to hire the best qualified candidate no objectionable reaction-qualifications enter into the picture, whereas the first employer only fails to hire the best qualified candidate once we discount reaction qualifications rooted in the antimeritocratic responses of the recipients. But, all other things being equal, the duty not to cater to recipients' illegitimate preferences seems less stringent than the duty one has to hire the best qualified when no questions arise about the objectionable reaction qualifications rooted in recipients' illegitimate preferences.

4. Illegitimate Preferences Not Disadvantaging Targeted Groups

So far, like other theorists in the field, I have assumed that all reaction qualifications grounded in antimeritocratic attitudes are problematic,[39] but I would like now to nuance this assumption. More specifically, I want to argue that, in a range of cases, the basic meritocratic and the noncollaboration norms do not require selectors to disregard reaction qualifications rooted in prejudice (and so on). Consideration of this range of cases teaches us something about what makes it morally wrong not to disregard such qualifications in standard cases. In the relevant kinds of case normally entertained here a certain group is subjected to antimeritocratic attitudes. Members of that group therefore come to have poor reaction qualifications, and as a result of this they end up worse off. Consider a case where the opposite happens:

[39] See the quotation from Mason in the final sentence of the sixth paragraph, section 1.

> *The Ironic Case*: Whites are prejudiced against blacks, and accordingly they are less likely to make purchases when dealing with a telephone salesperson they believe to be black. Suppose that white salespersons do not come across as distinctively white on the phone, whereas black salespersons do (through their deliberate efforts). As a result black salespersons outsell their white colleagues. Suppose that an employer hires a black salesperson and that, in the absence of racial prejudice, the best white applicant would sell more than the person actually hired.[40]

This situation is not ideal from the point of view of meritocracy, since there is widespread violation of what I have called the recipient norm. Also, the end-state norm enjoins all of us to put an end to this. However, let us set these norms aside and focus on the basic meritocratic norm and the noncollaboration norm. Should they be construed in such a way that a selector who hires the best-selling black candidates violates them? That is, do they entail that the white candidate—who, on a suitably moralized conception of qualifications, is the best qualified candidate—should be selected?

This—the claim that the white candidate who in the absence of racial biases would be most successful at selling should be selected over the black candidate who, given the presence of racial biases, as a matter of fact is most successful at selling—would not be the most plausible way to construe the two norms. Accordingly, the Ironic Case offers further support for the view defended in the previous section that it is not true, across the board, that reaction qualifications based on illegitimate preferences are wholly irrelevant from the point of view of merit. In applying the basic meritocratic norm and the noncollaboration norm, reaction qualifications that persons have as an ironic result of prejudice against them do count. If so, the question becomes why it is that some reaction qualifications based on illegitimate preferences should be discounted while others should not.

One suggestion is that what underpins our judgment here is simply a more general concern about distributive justice. After all, the Ironic Case and the standard case of reaction qualifications reflecting unjust prejudice and hostility differ in the respect that, in one case, members of the underprivileged minority benefit while in the other they are harmed. So perhaps we resist counting reaction qualifications in the standard case because doing so worsens the situation of the badly off; and perhaps we do not resist in this way in the Ironic Case, because counting them improves the situation of the badly off minority.[41] Possibly, when Mason claims that "prejudice against members of a group

[40] For the experiment which served as inspiration for the Ironic Case, see Samuel Gaertner and Leonard Bickman (1971), "Effects of Race on the Elicitation of Helping Behavior: The Wrong Number Technique," *Journal of Personality and Social Psychology* 20.2, pp. 218–222.

[41] I take it that the view that we should discount bigoted attitudes, no matter in favor of whom they work, is implausible.

cannot legitimately make it the case that not being a member of that group counts as a qualification for a job" he has in mind only the standard situation.[42] On the present suggestion we can agree that in the standard case it would be impermissible for selectors to count reaction qualifications based on illegitimate preferences. However, the reason this is so has nothing to do with the illegitimacy of reaction qualifications based on illegitimate preferences *as such*. Rather, this reason derives entirely from the violation of wider principles of distributive justice that counting those qualifications involves in the standard case.

This explanation does not capture our intuitions fully satisfactorily. Consider, first, a version of the standard case where it turns out that, for some reason, counting illegitimate reaction qualifications does not harm the worse off minority targeted by the relevant objectionable preferences. True, some members are not hired for jobs for which they are better qualified when illegitimate reaction qualifications are discounted, but let us suppose that this negative effect is somehow offset by compensating factors. Hence, the overall situation of the relevant minority is not affected for the worse. This situation appears problematic from the point of view of merit. Consider next a version of the Ironic Case where counting the relevant reaction qualifications does not affect the overall situation of the worse off minority. True, some of them are hired for jobs for which they would not have been best qualified if no one had harbored illegitimate preferences against them, but this beneficial effect is somehow offset by negative counterbalancing factors. Again, this would not eliminate entirely our inclination to say that the black salespersons really are the best qualified in the Ironic Case. To the extent that we share these intuitions, we should resist the view that the illegitimacy of reaction qualification based on illegitimate preferences derives wholly from the violation of wider principles of distributive justice such as equality of opportunity.

What makes the Ironic Case different from the standard case is that in the standard case counting reaction qualifications unfairly disadvantages people who are targeted by the illegitimate attitudes. This is not so in the Ironic Case. In this case it is the people that harbor the illegitimate preferences against

[42] Mason, *Levelling*, p. 35. Mason might respond that in the Ironic Case what counts as a qualification is not *being white*, but *appearing to be white*. He might say that his view leaves it an open question whether the latter reaction qualification legitimately counts. But in making this response he would render his position vulnerable to a reinterpretation of standard cases where we would not want to count race as a reaction qualification: e.g., it might be said that the relevant qualification of a white applicant to a position as a shop assistant is not that she is white but that she appears to be white. Alternatively, Mason might respond that in these particular circumstances, and contrary to what some of his claims suggest (e.g., the one quoted in the sentence to which this note is attached), coming across as white *is* a qualification, since it is a performance-related characteristic, and that it is not unjust to select black applicants on this basis. However, this response assumes that the case is unproblematic from a meritocratic point of view, which is doubtful. Given that the recipient norm is transgressed, all of us ought to put an end to this.

members of other groups that end up being disadvantaged, and *they* cannot complain about the unfairness of their ending up worse off for *that* reason. In saying this I neither want to tie fairness or justice to the requirement that no relevant person is in a position to complain, nor intend to rely on any general theory of fairness. All I want to claim is that if someone treats someone else unfairly and ends up worse off as a result, it may—depending on the relevant unfair treatment and the nature of the disadvantage—not be unfair that she suffers this disadvantage.[43] At least, this looks right if we assume that the unfair treatment is blameworthy, and that the disadvantage suffered is not out of proportion with the unfairness of the relevant treatment (as it is not in cases of just punishment). My hope is that this specific claim about when disadvantages are not unfair is compatible with a broad range of theories of fairness, widely acceptable, and clearly applicable to the present case. Accordingly, on the view proposed here, what really lies behind our initial inclination to discount all illegitimate reaction qualifications is a concern that no one is unfairly disadvantaged in the competition for positions. We initially conjectured that being targeted by illegitimate preferences that give rise to reaction qualifications creates an unfair disadvantage, but the Ironic Case shows that this is not necessarily so.

In further support of the relevance of unfair disadvantage here consider a different version of the Ironic Case in which half the minority members come across as majority members on the phone and the other half do not. In this case, counting illegitimate reaction qualifications will benefit some minority members and unfairly disadvantage others. Here our view about counting reaction qualifications tracks our views on unfair disadvantage. That is, if we only consider the situation of the advantaged minority members and the majority, counting reaction qualifications is unproblematic. But if we contrast the situation of advantaged minority members with that of the unfairly disadvantaged minority members, we are much less inclined to think that selectors are permitted to count reaction qualifications rooted in illegitimate preferences.

Consider next a case where the white candidates—who, recall, do not come across as being distinctively white on the phone—unlike other whites, do not harbor any objectionable attitudes to blacks. In this case, they may complain about being unfairly disadvantaged in the competition for jobs, since they are being deselected on the basis of illegitimate biases they do not share. I am inclined to think that, at some level, this gives the selectors a reason not to count the reaction qualifications rooted in illegitimate attitudes. If that is right, we

[43] This claim is compatible with my claim in chapter 8 that capital punishment might be unjust (or unfair) even if a murderer is in no position to complain about the punishment. I am not contending that it is the fact that the agent is in no position to complain about the relevant disadvantage resulting from her unfair treatment of others that makes the disadvantage that she suffers not unfair. Rather, it is the fact that the disadvantage results from her unfair treatment of others that, possibly, makes it not unfair.

have further support for the following claim: the basic meritocratic norm enjoins selection of the best qualified, where this involves disregarding reaction qualifications rooted in illegitimate attitudes only when taking such qualifications into account would unfairly disadvantage some applicants (though not necessarily the applicants who are the targets of the illegitimate attitudes).

In the Ironic Case blacks exercised skill to come across as whites on the phone (not because this requires skill in itself, of course, but because it requires skill to come across differently from how you would come across unthinkingly). In acquiring this reaction qualification they have trained, or in other ways exercised agency, in imaginative ways. True, this would not have been necessary had people harbored no illegitimate attitudes, but in the example they do, and, in general, there is nothing unfair about people acquiring better qualifications—including those relating to appearance—through training and the imaginative exercise of agency. Arguably, things might be different if black applicants came across as whites on the phone effortlessly. In that case, they might enjoy an unfair advantage over white applicants who harbor no illegitimate attitudes to blacks and who do not come across as white on the phone. However, since it is not clear that reaction qualifications really should be counted in this case, these scenarios do not speak against the account I have offered here.[44]

5. Respect and Reaction Qualifications

George Sher and Andrew Mason think that the ideal of meritocracy is based on, or shaped by, the notion that we should respect individuals as agents. Sher submits that meritocracy is rooted in our duty to treat people as "agents whose purposeful acts are capable of making a difference in the world" and not as "passive links in causal chains."[45] Mason, who accepts Sher's views on this matter, albeit with some significant reservations, holds that the best argument in favor of meritocracy appeals to the fact that, generally speaking, failing to select those best qualified "is disrespectful to the candidates."[46] This rationale might seem to clash with my reinterpretation of the basic meritocratic norm. This reinterpretation involves a distinction between cases where counting reaction qualifications unfairly disadvantage some and cases where it does not.

---

[44] To accommodate the claims made in this section we should revise (2) in the symmetric view so that it says: "A reaction qualification counts from the point of merit . . . if, and only if, it is not rooted in nonindividualized antimeritocratic attitudes and counting it unfairly disadvantages someone."

[45] Sher (1988), "Qualifications," pp. 119–120, 123. If making an effort counts as "making a difference in the world," and if, as is surely the case, effort and qualifications do not always correspond, hiring on the basis of expected effort to do the job well, as opposed to expected performance in the job, may well be a case of respecting applicants as agents.

[46] Mason (2006), *Levelling*, p. 64; see also p. 60. Unlike Sher, Mason does not tie respect and desert together.

But this distinction is not necessarily aligned to a distinction between treating people as "agents whose purposeful acts are capable of making a difference" and not so treating them. A minority member who effortlessly comes across as a majority member simply through the tone of her voice is not being unfairly disadvantaged even if she is not gaining an advantage as an agent "whose purposeful acts are capable of making a difference in the world."

In response to this I offer two considerations. First, it is not clear that my reinterpretation really clashes with the respect for agency account. In the Ironic Case I stipulated that black salespersons had acquired their ability to come across as whites through training and imagination. Arguably, it would be more disrespectful of their agency to ignore qualifications so acquired than to take them into account.

Still, there could be cases where, through no control of their own, the salespersons simply came across as white, and in those cases counting the reaction qualifications does not unfairly disadvantage anyone: those who are disadvantaged by the counting here are not disadvantaged *unfairly*, since it is they who harbor the illegitimate attitudes. Here I want to say that the notion of agency includes more than those aspects of what we do that are under our control. So when I speak with a certain dialect that is still an aspect of my agency in a broader sense: these are things I do even if I do not control my dialect, not things that simply occur to me. While Sher and Mason may think that it is only aspects of our agency that are under our control that merit respect, they have not (yet) argued for this. Moreover, a rationale in terms of respect for agency in a narrow sense would be in conflict with many judgments about qualifications that we find unproblematic—e.g., that the charming actor really is the best qualified applicant for the role even though the aspect of her agency that makes her charming is not under her control.

My second response is that the motivation of meritocracy in terms of respect for agency is problematic in ways that threaten the notion that it would be a decisive objection to my reinterpretation of the meritocratic norm that it conflicts with this rationale. First, it is not clear that respect for someone's agency will align suitably with the practice of hiring on the basis of qualifications. Sher submits that "when we select among applicants for reasons other than their ability to perform the tasks that define the positions they seek, we treat them as passive recipients of largesse or links in causal chains."[47] Similarly, Mason writes that "respect for persons as agents requires that we consider only the qualifications of the candidates to do the job as opposed to (say) their membership in some group or their connections to others."[48] But the contrasts they draw here are misleading. Suppose the relevant group is the group of marathon

[47] Sher (1988), "Qualifications," pp. 123–124.
[48] Mason (2001), "Equality of Opportunity," pp. 769–770.

runners who have won gold medals at the Olympics recently; and imagine that, partly as a result of her membership, a member of this group is hired as a headmaster instead of a number of other and better qualified candidates.[49] It is far from obvious that the wrong done to the better qualified candidates consists in their agency not being respected—that rejected candidates are treated "as mere bearers of needs and claims, as passive links in causal chains, or as interchangeable specimens of larger groups or classes."[50] To win a marathon one has to exercise agency in ways of which very few people are capable. For example, one has to engage in meticulous long-term planning and to stick with it despite strong temptations to give up. In general, there are other bases for selecting candidates than qualifications for the position in question such that selecting on these bases count as respecting candidates as agents.

One might take this objection to motivate, not wholesale rejection of a respect-based approach, but *refinement* of it. (I am not endorsing this move, as will become clear in the next paragraph. I just want to explore its credibility.) In this refined view, meritocracy would be rooted, not in a requirement to respect the agency of candidates, but in the more specific requirement to respect those agential powers of candidates that are relevant to the specific range of tasks involved in the position being appointed to. In the case of the less well qualified Olympic gold medal winner this more specific requirement of respect is flouted. For we can assume that the agential powers that are relevant to the job of headmaster (e.g., administrative skills, and the ability to inspire and motivate teachers, staff and students) are very different from the powers relevant to a marathon runner (e.g., the ability to persist despite exhaustion and physical pain). Hence, the best qualified candidate for the position as a headmaster can complain that selectors fail to respect those of her agential powers that are specifically relevant to the position at hand even if they do fail to respect anyone's agential powers as such. Whether anyone's agential powers as such are respected is simply relevant on the present account, since the content of the relevant requirement of respect is specifically tailored to the position at hand.

This move strikes me as unpromising. For one thing, it is unclear that the specific kind of respect involved is mandatory.[51] Familiar cases appealed to in support of the respect-based approach involve violation of the more general requirement of respect (where, say, a talented female candidate is rejected solely

[49] A similar point applies to cases where candidates are selected on the basis of a reasonable expectation that their being selected for the position will lead them to perform tasks that are not related to the position brilliantly.

[50] Sher (1988), "Qualifications," p. 123; quoted approvingly by Mason (2006) *Levelling*, p. 58.

[51] A further problem is that it is often very hard to say exactly which agential powers are relevant to a specific position Alexander (1992), "What Makes?" pp. 151–153 n1; Jan Narveson (2001), *The Libertarian Idea* (Peterborough: Broadview Press), pp. 316–318. However, this problem is one that any meritocratic view faces in one form or another; see my second commentary on the symmetric view above.

because she is female), not this more specific one. Moreover, the more specific requirement comes too close to the meritocratic norm it is supposed to motivate to serve as an independent rationale for it. Those who reject this norm could reasonably refuse to countenance any attempt to justify it referring to the putative fact that failure to select the best qualified candidates disrespects their relevant agential powers simply by pointing out that whether there is such a requirement is part of the very issue that the parties are trying to resolve.[52]

The second problem I wanted to highlight is that the respect for agency account is neutral between a state that involves perfect satisfaction of the meritocratic norm and one that involves a breach of that norm but perfectly satisfies a virtuocratic norm: that is, a norm that is structurally identical with the wide meritocratic norm but in which moral virtue is substituted for qualifications so that, for example, when this norm is satisfied each position is filled with the most virtuous candidate. At least, this seems so to the extent that one becomes morally virtuous by exercising one's agency in the right way. But, presumably, a satisfactory rationale for meritocracy would provide an account of why we should not be indifferent as to whether the best qualified or the most morally virtuous candidate is selected.[53]

6. Conclusion

I have defended the symmetric view—a view that preserves symmetry between antimeritocratic bases for the selection of candidates and qualifications rooted in antimeritocratic attitudes. I have decomposed the ideal of meritocracy into four separate norms addressing different categories of person: the selector, the recipient, as well as all others, e.g., people who interact with recipients. I have argued that the basic meritocratic norm should be reinterpreted in a way that allows reaction qualifications rooted in antimeritocratic preferences to be counted provided no one is thereby unfairly disadvantaged. An important upshot of this and other arguments is that the ideal of meritocracy and the norm of nondiscrimination are less closely tied than has commonly been assumed. Finally, I have tried to show that my views on reaction qualifications do not clash with a requirement of respect in a way that undermines them.

[52] Mason might broadly agree with this point, since he does not think that we have an independent idea of respect from which we can derive the meritocratic norm (Mason (2006) *Levelling*, p. 64).

[53] Note that an appeal to legitimate expectations cannot do this job. Presumably, in a virtuocratic society, people would legitimately form the expectation that they will be rewarded on the basis of moral virtue.

{ 10 }

Discrimination in the Private Sphere

1. Introduction

In the previous three chapters I focused on discrimination outside the so-called private, as opposed to the public, sphere by collective agents such as companies or courts. Most adopt a quite critical stance in relation to discrimination in the public sphere. In particular, most would consider a state's discrimination against its citizens—say, by refusing to hire them on account of their sex, race, religion, or ethnicity—clearly wrong, something in urgent need of rectification. Yet we often take a less censorious view of discriminatory acts by private individuals who choose not to share their neighborhood with, marry, associate with, work with, befriend, trade with, or be buried in the same graveyards as, people of a different sex, race, religion, or ethnicity.[1] In all of these arenas many believe that the state should not prevent people from discriminating and, indeed, that acts that would otherwise qualify as acts of discrimination are mere cases of permissible differential treatment when performed by private

[1] The reason people choose not to do such things need not reflect any negative attitudes or value judgments on their part against people of a different sex, race, etc. as such. Instead, it may reflect something analogous to reaction qualifications on the job market (see chapter 9) and, like that form of discrimination, be rational from the point of view of the discriminator, cf. Russell Hardin (1995), *One for All: The Logic of Group Conflict* (Princeton, NJ: Princeton University Press), p. 90. By way of illustration, while I am indifferent to which faith my partner has, my close co-religionists, e.g., members of my family, care strongly about the issue and would impose severe informal sanctions on me were I to start a relationship with a person of a faith different from my own and, accordingly, to avoid such sanctions I discriminate in favor of same-religion potential partners and against different-religion potential partners. Like with reaction qualifications on the job market, this kind of discrimination is morally different from one that is rooted in noninstrumental, negative attitudes toward potential partners with a religion different from one's own, and yet persons may sometimes be morally required to bear some costs involved in refraining from discrimination rooted in the prejudiced reactions of others in the interest of eliminating discrimination. I say "some costs," because I do not think, say, that women who face the threat of honor killings or family ostracism are morally required not to discriminate against potential, different-faith partners.

individuals. For instance, in the United States, the Supreme Court takes the view that the First Amendment protects church autonomy. In accordance with this view, courts have granted religious organizations exemptions to discriminate on the basis of sex in cases where secular organizations would not have been allowed to do so.[2]

Chapter 9 addressed this asymmetry in relation to the meritocratic norm. Implicitly, I expressed skepticism about construing it asymmetrically in the present sense suggesting that, properly construed, it also pertains to, among other things, the market behavior of individuals, viz., the recipient norm (chapter 9.3). However, the asymmetry between the way in which many think about discrimination in the public and the private sphere is perhaps most salient when it comes to religion. Mainstream versions of many religions believe that priests etc. should not be women or homosexuals or atheists or people who believe in a different religion or in the same religion but in a heretic version.[3] Many think that religious communities should have a legal right to engage in discrimination against such candidates for religious positions when deemed inappropriate on religious grounds; that exercising this right is not wrong; or, even more strongly, that, for some not wholly perspicuous reason, it is not discrimination when, say, a women is refused a job on account of her gender when those who refuse to hire her do so on religious grounds and in their capacity as agents of a religious institution. In assessing these beliefs, it is worthwhile to keep in mind that, not so many years ago, many Christian churches opposed having nonwhite priests or even nonwhite members of the congregations offering religious justifications for this stance.[4] Few people today would accept racial discrimination practiced on religious grounds. Yet, for apparently no good reason, many accept sex discrimination and discrimination on grounds of sexual orientation when religiously motivated. To back up this claim consider the following: nowadays many accept that legal protection against discrimination on religious grounds should be understood in such a way that it extends to discrimination on secular grounds in certain contexts and when the secular

---

[2] *McClure v. Salvation Army*, 460 F2d 553 (5th Cir 1972); Christopher L. Eisgruber and Lawrence G. Sager (1994), "The Vulnerability of Conscience: The Constitutional Basis for Protecting Religious Conduct," *University of Chicago Law Review* 61.4: 1254–1315, p. 1275n52.

[3] To allow, inter alia, that it may not amount to discrimination when a church treats same-religion and different/no-religion applicants differently I included (v) in my definition of group discrimination in chapter 1; (v), however, does not rule out that such cases often involve discrimination, e.g., because the differential treatment is motivated by the belief that people of one's favored religion should not intermingle—worship-wise at least—with atheists or members of different religion. Also, there is the issue of specifying the recipient norm in relation to members of a certain church. Does it not require that, all other things being equal, they respond equally well to a different/no religion priest as to a same religion priest?

[4] "[T]he Dutch Reformed Church in Apartheid South Africa . . . refused to admit blacks because of a belief that God made them inferior," see Sandra Fredman (2011), *Discrimination Law* (Oxford: Oxford University Press), p. 210.

convictions in question are relevantly similar to religious beliefs in being in some sense deep.[5] So, for instance, many accept that a suitably nondiscriminatory unemployment benefit law should accommodate not only an unemployed person, who on religious grounds refuses an available job in a weapons factory, but also one who out of a deep, secular, pacifist conviction turns down such a job offer. Yet, they do not similarly think that because religious organizations should be legally permitted to discriminate on grounds of gender or sexuality, secular organizations infused with nonreligiously justified sexist or homophobic values should also be legally permitted to engage in discrimination against women or homosexuals. Why, one must ask, should religion be treated alike to secular convictions in one context and not in the other?[6]

Others take a much more critical view and claim that while religions may not by their very nature be discriminatory, as a matter of fact, they are discriminatory to the bone. Simone de Beauvoir famously complained that Christianity represented women as being the second sex. According to the Christian account of Genesis, the first woman was created from one of Adam's ribs, and this de Beauvoir took as a sign of how women were subordinated to men and largely restricted to matters of family and home. Similar objections can be raised to other religions and on that basis one might argue that there is no reason why private religious institutions should be allowed to engage in forms of discrimination that other private institutions are legally prohibited from engaging in.

On reflection, this asymmetry between public and private discrimination is puzzling: why should it not be morally wrong for private individuals to do what it would be morally wrong for the state to do?[7] First, it cannot be explained by saying that state discrimination has graver consequences involving more people than private discrimination. True, considered on their own, private discriminatory acts often have negligible consequences; but together, such acts often form patterns that are no less consequential than, say, the passing of a piece of discriminatory legislation. For instance, marital and geographical racial segregation arising from cultural norms and individual choices is central to racial discrimination.[8] Also, individual acts of discrimination by the state or a state institution likewise will often not harm many significantly. A similar question is, why it is morally right for there to be a legal right for private individuals and

[5] I set aside here that people's religious attitudes need not be deep, but might simply reflect conformism or habit.

[6] If we compare discrimination on the basis of adherence to religious dogmas by religious organizations to discrimination on the basis of adherence to, say, political dogmas by political parties etc., a similar asymmetry probably does not exist. Yet, in principle, a priest (political leader) could be very good at promoting and representing a religion (political view) that he or she does not endorse.

[7] M. Cavanagh (2002), *Against Equality of Opportunity* (Oxford: Clarendon Press), p. 169.

[8] Elizabeth S. Anderson (2010) *The Imperative of Integration* (Princeton, NJ: Princeton University Press), pp. 112–117; Fredman (2011), *Discrimination Law*, p. 68.

others to engage in discrimination, when it is not morally right for state institutions to have such a right.[9]

Second, while the absence of legal regulation of actions performed within the private sphere is sometimes justified as a way of protecting privacy, it is hard to see that matters pertaining to the private sphere always involve a concern for privacy while actions in the public sphere never do and accordingly hard to see that the distinction between public and private can be drawn on the basis of privacy.[10] For instance, while regulations forbidding ethnically or religiously segregated graveyards may be morally wrong, they do not seem to involve any violations of privacy and yet graveyards are normally considered a matter pertaining to the private sphere. Business transactions between a prostitute and a customer are normally considered to involve a concern for privacy even though business transactions are thought to belong to the public sphere.

Third, it is unclear on what ground the distinction between public and private is drawn, and sometimes the distinction seems simply to be a distinction between those spheres that the person who draws the distinction thinks should be subjected to legal regulation and those that should not be (or at least, should be so only to a significantly lower degree).

We need a better understanding of the morality of private discrimination. Specifically, according to the desert-prioritarian account proposed in chapter 6, basically, there is no intrinsically morally relevant difference between private and nonprivate discrimination. However, as we shall see, for various instrumental reasons, there might be a gap between the morality of discriminatory acts by private individuals and the morally required legal status of such acts, e.g., it might be the case that one ought, morally speaking, not engage in a certain form of discrimination, which, however, should be legally permitted. To determine how, morally speaking, discrimination ought to be legally regulated, I want to return to desert-accommodating prioritarianism (chapter 6); other

[9] Eisgruber and Sager (1994) argue that religion does not have a privileged legal status just because religious associations are legally permitted to engage in sex discrimination when hiring priests. In their view, exemptions from normal antidiscrimination laws are motivated through there being a way of ensuring that "government treat the deep, religiously inspired concerns of minority religious believers with the same regard as that enjoyed by the deep concerns of citizens generally," p. 1283, and, thus, simply reflects a desire to protect vulnerable religious groups. I find their argument unconvincing. First, even setting the aside the unclear nature of Eisgruber and Sager's core distinction between privileging and protecting—see Andrew Koppelman (2006), "Is It Fair to Give Religion Special Treatment?" *University of Illinois Law Review* 3, 571–604, pp. 581–583—their argument cannot show why nonminority religions should be exempt from anti-sex-discrimination law. Second, given that women have not been treated with equal regard, how could a concern for equal regard lie behind a legal exemption for religious associations to discriminate against women? In any case, Eisgruber and Sager's argument is to some extent orthogonal to mine, because my main concern is with what laws regulating discrimination ought to be, morally speaking, and how people ought to act, morally speaking, and not with which concerns have motivated various aspects of the law as a matter of empirical fact.

[10] Cf. Eisgruber and Sager (1994), "The Vulnerability of Conscience," p. 1245.

moral principles will lead to different conclusions. Desert-prioritarianism implies that there ought to be a law prohibiting a certain kind of wrongful private discrimination if, and only if, that law maximizes moral value. In fact, according to this view, there ought to be a law that enjoins private discrimination that is morally wrong if the existence of such a law maximizes moral value. There are six cases to consider:

	Private discrimination which there ought to be a legal duty to engage in.	**Private discrimination which there ought to be a legal right, but not a legal duty, to engage in.**	**Private discrimination which there ought to be a legal duty not to engage in.**
Private discrimination that is morally wrong.	1. (Section 2)	2. (Section 3)	3. (Section 4)
Private discrimination that is not morally wrong.	4. (Section 5)	5. (Section 5)	6. (Section 6)

The world we live in may well instantiate all six possibilities—some more often than others. The examples that follow describe some considerations relevant to the classification of different cases. Before doing so, one issue needs to be briefly addressed: how can we determine which alternatives are relevant to determining which laws maximize moral value? One question here is whether one should use the same range of alternatives for the purpose of determining which kinds of private discrimination are wrongful and for the purpose of determining which kinds of private discrimination ought to be legal. Using different sets of alternatives may render it too easy to find examples of wrongful private discrimination that ought to be legal and examples of morally permissible private discrimination that ought to be illegal.

Laws that maximize moral value in one set of alternatives may be very different from those that maximize moral value in another. Assume, for instance, that in one alternative people neither engage in morally wrong discrimination, nor engage in other kinds of unjust behavior and most private discrimination is legal. That alternative might well realize greater moral value than one in which most people engage in various kinds of morally wrong private discrimination and most kinds of private discrimination are illegal. Equally, it might realize greater moral value than an alternative in which people engage neither in morally wrong private discrimination, nor in other kinds of unjust behavior and most kinds of private discrimination are illegal (say, because of the costs involved in having laws that are never transgressed). This might show that in a world that is pretty ideal from the point of view of nondiscrimination there ought to be a right to engage in wrongful private discrimination; but it would not show that, given the nonideal worlds that we could realistically hope to realize, there ought to be a legal right to engage in such behavior. To properly

explain which kinds of wrongful private discrimination ought to be legal, one must specify the full range of alternatives relative to which one answers the question. This undertaking is beyond the scope of this chapter. Hence, the following examples assume that the relative alternatives will parallel ours, in terms of people's dispositions to behave in certain ways, but will have different laws on private discrimination. Accordingly, the present discussion deals with alternatives that are quite far away from those that bear on ideal theory. The evaluation of the alternatives considered will be rough, sketchy, and done from the armchair. At best, it identifies some relevant considerations.

2. A Legal Duty to Engage in Wrongful Private Discrimination

In cases involving widespread biases, it might emerge that there ought to be a legal duty to engage in wrongful private discrimination. Consider a country whose population is divided into a large Protestant majority and a small Catholic minority. Some members of the Protestant elite are anti-Catholic, and as a result there are certain laws that, in effect, impose on Protestants a legal duty to engage in private discrimination against Catholics. For instance, Protestants are not allowed to hire or live together with Catholics. Members of the Protestant elite and Protestants generally, care about these discriminatory laws mainly as a symbolic manifestation of Protestant supremacy. Hence, Protestants and Catholics who engage in the legally forbidden private interaction experience little trouble: few inform on lawbreakers; the police give low priority to violations of these discriminatory laws; and only a few lawbreakers experience social sanctions.

Generally speaking, it is morally wrong for Protestants to discriminate in the ways they are legally obliged to—they would bring about more moral value if they abstained from doing so. However, trying to change the laws would cause an uproar, and possibly even civil war, and hence they ought, morally speaking, to be retained for the time being. More generally, in cases involving widespread biases that focus mainly on the symbolic significance of discriminatory laws, a gap between legally required and morally permissible discriminatory acts may develop. Such situations are normally best construed as partial-compliance situations that are not optimal from the point of view of nondiscrimination.

3. A Legal Right to Engage in Wrongful Private Discrimination

There are plenty of discriminatory acts that are wrongful and yet ought to be legally permitted. First, some wrongful private discriminatory acts are likely to be so rare that little good will come from outlawing them. For a wide range of commodities people will purchase whatever is the best offer from a purely commercial point of view. Hence, if some consumers will not buy petrol

produced by companies that predominantly employ people of a certain ethnicity, they might act in a wrongful way and yet denying them a legal right to do so might not, given the rarity of this behavior, improve the situation.[11]

Second, in connection with some kinds of wrongful private discrimination, enforcement of the laws would be very difficult and costly. For instance, people who discriminate on the basis of race, ethnicity, and religion in their choice of spouse may act wrongly—especially in situations of racial, ethnic, or religious tension. But often it will be difficult, and harmful to many parties, to prove that someone has discriminated in his or her choice of spouse.

Third, as the example just provided illustrates, effective law enforcement may destroy some important goods altogether. The law can force a company to hire a higher proportion of its workforce with a minority background without seriously affecting, for better or worse, its economic viability; but a person cannot be forced by law to marry a person with a race, religion, or ethnicity different from his own without seriously affecting his love life. (Imagine: "This court finds beyond a reasonable doubt that had you not been such a bigot, you would have married Y rather than your present spouse, Z.") However, it should not be forgotten that, through indirect and nonlegal means, the state might try to reduce the amount of private discrimination of this sort. For instance, the state might seek to ensure that people mix across races and thereby to increase interracial marriage and as a long-term result reduce racial discrimination. Obviously, in these cases, laws cannot create a state of affairs that is ideal from the point of view of non-discrimination. This requires a certain kind of ethos manifested in the choices that people make within the constraints set by the law.

4. A Legal Duty Not to Engage in Wrongful Private Discrimination

The claim that this case is instantiated is much less controversial than the previous one—except, of course in cases such as the one mentioned in section 2, where there are moral reasons why there should be a legal duty to engage in wrongful discrimination—and very little shall be said about it. For instance, there ought to be laws forbidding private employers from paying women less than men for doing the same job when these are easily enforced at little cost and with little harm to employers, female employees, or consumers. One highly relevant factor is the way in which laws will affect the strength and existence of the relevant discriminatory preferences. When discriminatory preferences are a rather superficial feature of people's identity, outlawing conduct reflecting such

[11] Jobs and services may differ from many commodities, in that these involve a much greater density of social interactions. See Jan Narveson (2002), *Respecting Persons in Theory and Practice: Essays on Moral and Political Philosophy* (Lanham, Maryland: Rowman and Littlefield Publishers, Inc., 2002), pp. 203, 213.

preferences may make them almost extinct, strengthening the argument for outlawing the relevant conduct.[12]

5. A Legal Duty or Permission to Engage in Private Discrimination That Is Not Wrongful

Affirmative action programs will often be of this kind.[13] Consider first the case for a legal duty. Suppose a certain minority has suffered from seriously wrongful discrimination. As a result, its members suffer from substandard levels of education and income. Suppose, finally, that affirmative action programs at private universities and companies will either not make things worse or make them considerably better from the point of view of desert-accommodating prioritarianism—because they redistribute resources to badly off people who are in some cases more deserving, and because they lead to a society with less discrimination that makes better use of the pool of talent. In the latter case, where affirmative action programs at private universities and companies make things considerably better from the point of view of desert-accommodating prioritarianism, a legal duty to engage in affirmative action might be justified because it ensures that more private universities and companies engage in affirmative action. In the former case, where affirmative action programs at private universities and companies simply do not make things worse from the point of view of desert-accommodating prioritarianism, a legal duty might be justified because it sends a powerful message of inclusion which may have positive consequences quite independently of the consequences resulting from various agents' efforts to comply with the law.

There may be cases in which (i) the net benefits from universities and companies adopting affirmative action programs in response to a legal requirement are negative and (ii) where a mere legal permission to adopt such programs will involve roughly the same symbolic benefits as a legal requirement. Equally, there may be cases in which (i) these symbolic benefits of legal permission will be smaller than the symbolic benefits of a legal duty, (ii) but where there is a somewhat greater net loss involved in having some companies and universities adopting affirmative action programs that they would not have adopted had the law not forced them to do so, e.g., because they will see these programs as being imposed on them and accordingly act in ways that make them run less smoothly than they would have done had they been legally optional. In both situations desert-accommodating prioritarianism favors legal permission to engage in affirmative action over a legal duty to do so.

---

[12] See Larry Alexander (1992), "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies," *University of Pennsylvania Law Review* 141.1: 149–219, p. 163.

[13] Recall that I am using the concept of group discrimination, not the moralized concept.

6. A Legal Duty Not to Engage in Private Discrimination That Is Not Wrongful

Consider, finally, the case of a community that is prone to ethnic and religious tension. In cases involving disaster relief, people prefer to donate money to aid agencies providing assistance to people with the same religion and the same ethnic background as themselves. Partly as a result of this, only sectarian aid agencies exist. Suppose that each act of discriminatory donation brings about more moral value than would be brought about by not donating to the relevant aid agency and, hence, is not wrong according to desert prioritarianism.[14] It may nevertheless be the case that there ought to be laws prohibiting discriminatory donations. The pattern of sectarian donations might be such as to fuel ethnic and religious tension considerably. If this were so, a nonsectarian pattern of private donations would be a collective good from the point of view of all donors. Alternatively, legal enforcement of a pattern of nonsectarian donation might not appreciably reduce the total sum of donations.[15]

These observations strongly suggest that, for the desert-accommodating prioritarian, the wrongfulness of private discrimination (where it arises) should not, morally speaking, always be tracked by unlawfulness. Cases exist in which there should be a legal duty, and henceforth a legal permission, to avoid private discrimination even though that discrimination is not morally wrong. In these respects, wrongful private discrimination is no different from other kinds of wrongful behavior; and while some moral theories may not allow such gaps between the moral status and the morally desirable legal status of discrimination, desert-accommodating prioritarianism is not alone in allowing such a gap.

7. Conclusion

This chapter differs in its approach from the one taken in the other chapters in part 3. In a sense, it employs a top-down approach and applies a general normative principle—desert prioritarianism—to a concrete issue: the relation between the moral status of legal regulation of acts and the moral status of

[14] The sets of sectarian donations may well be wrongful, then, but I disregard this complication.

[15] Another example is organ donations, where the donor for racist reasons wants his or her organ to be given to a member of her ethnic group; see Nir Eyal (forthcoming) "Levelling Down Health" in Nir Eyal, Ole F. Norheim, Samia A. Hurst, and Daniel Wikler (eds.), *Inequalities in Health: Ethics, Measurement, and Policy* (New York: Oxford University Press). Suppose that we are talking about a live donor who wants to donate one of his kidneys, something which, let us suppose, goes beyond the call of duty. Most Western countries forbid organ donations so conditioned.

these acts themselves. The other chapters in part 3 employ a bottom-up approach where the relevant concrete issue is explored on the basis of concrete moral intuitions concerning the case at hand. More specifically, I have argued that in a desert-prioritarian account it may well be the case that legal and moral status deviate (e.g.,, sections 2, 3, and 6), e.g., it might be morally right to have a legal permission or even a legal duty to engage in wrongful discrimination.

This implication of desert prioritarianism strikes me as being independently plausible and I suspect it is one that is acceptable to many who reject desert prioritarianism. The law is one among many instruments that we may use to regulate one another's actions and like other instruments it will have to be judged on its effects. Given this perspective it is unsurprising that we sometimes achieve the most desirable effects through laws that permit, or even mandate, morally objectionable actions. Admittedly, this perspective is very different from a strong moralistic perspective on laws, i.e., one according to which one should have a legal duty not to act in a certain way if, and only if, one is morally required not to act in this way. However, that there are strong objections to this view. For instance, the moral desirability of a legal prohibition on so-called presumed offenses—i.e., actions such as selling magazines without a front page at a discount price that in themselves are not morally wrong, but where the fact that they are performed is a reliable indicator of other actions that are morally wrong though hard to detect directly, i.e., getting a refund from publishers for magazines that one falsely declares to be unsold—shows that one can have a legal duty not to do what, in the absence of the legal duty, one has no moral duty not to do.[16] The case of lying in personal matters shows that one might be morally required not to act in a certain way even if one should not have a legal duty not to act in this way.

Returning finally to the puzzle with which I opened this chapter, I note that in a desert-prioritarian account, one can see why there is a gap between public and private discrimination in the sense that the morally desirable legal status may not correspond its moral status. Hence, it might be morally justified that churches are treated differently from ordinary companies when it comes to religious discrimination in hiring, because the consequences of allowing companies to engage in religious discrimination may be very different from the consequences of allowing religious institutions to do so. However, it is important to stress what this does not imply. It neither implies that private discriminators do not discriminate, when they treat members of some socially salient

[16] F. Schauer (2003), *Profiles*, *Probabilities, and Stereotypes* (Cambridge, MA: Harvard University Press), pp. 224–250. Retailers claim refunds for unsold magazines from publishers on the basis of the number of front pages returned.

groups disadvantageously on account of their group membership, nor that do not act wrongly, morally speaking, when they do so. So in a sense private acts of discrimination are as wrongful as comparable public ones. In the last chapter I shall look at a case of a state policy that may sometimes be morally permissible, even if the social facts that make it so reflect the unjust actions of private individuals.

{ 11 }

Racial Profiling

1. Introduction

This final chapter addresses a form of discrimination that, while notoriously controversial, differs from many other forms of discrimination in that it has been cautiously defended by a number of theorists. Sometimes the police employ group identity profiling to make decisions about the use of scarce resources, guided by statistical information about crime rates among different groups of people. In the aftermath of the July 2005 terrorist bombings of the London transit system that killed 56 people and injured 700, Ian Johnston, a chief constable of the British Transport Police, said in a public statement: "We should not waste time searching for old white ladies."[1] This remark did not go down well with the public. One reason was that Johnston's remark seemed to endorse statistical discrimination, more specifically, a form of statistical discrimination called racial profiling. To that extent the response to Johnston's remarks is interesting. As I noted in chapter 3, statistical discrimination is something we all engage in. Arguably, it is something we can hardly avoid engaging in given that inductive reasoning and the tendency to act thereupon are deeply ingrained in our nature. Moreover, it is not always morally wrong.

Suppose, hypothetically, that Johnston had said publicly about the stop-and-search efforts by the British Transport Police to prevent hold-ups involving firearms that the police should not waste time searching old ladies, period. Presumably, to say so would be to endorse statistical discrimination against males and young women. Yet, I conjecture that no one, except those who mistakenly took Monty Python's sketch "Hell's Grannies" to be an alarming documentary, would take any notice.[2] This suggests that statistical discrimination is not

[1] http://www.frontpagemag.com/Articles/ReadArticle.asp?ID=19163 (accessed November 22, 2006).

[2] See for instance, http://www.youtube.com/watch?v=bIHF4rVTK4E.

intrinsically morally wrong and that we need to do some work to find out what makes it morally wrong when it is.

Racial profiling is a specific type of group identity profiling, which again is a type of statistical discrimination. It can be said to occur where, say, there is a greater likelihood that police officers will stop, search, and question people of a certain race because members of this group are believed to be more likely to possess illegal drugs than members of other groups. Hence, it is not racial profiling where the police are more likely to stop, search, and question people of a certain race when investigating a drugs-related crime because witnesses have described the perpetrator as a member of that race. In the latter case, the investigative activities are not based on a relevant sort of race-sensitive *statistical* belief. Similarly, it might be known that people who have a certain job are more likely to commit a certain type of crime, and that almost all of those who happen to have this job belong to a certain racial group. Accordingly, members of different races might be treated differently in preventive police work. But this need not amount to racial profiling (although it might involve occupational profiling). In general, differential treatment of the members of different racial groups may or may not involve racial profiling, depending on the basis of the treatment.

Racial profiling can be, and very often is, morally problematic for various incidental reasons: for its association with racial hostility, double standards, prejudice influencing the formation of statistical beliefs about crime rates in racial groups, biased conceptions of what constitutes crime, and so on.[3] But suppose we focus on what we might call an *unalloyed* case of racial profiling. By this I mean a case of the following kind. First, the police activities are based in the right way on justified, true beliefs about crime rates in different racial groups. Second, the group identity profiling enables the police to catch more perpetrators of crime, prevent more crimes, and, thus, reduce the overall incidence of crimes, i.e., it promotes the aims of the police relating to crime.[4] Third, no alternative, and equally or more effective, way of doing this is feasible.[5]

[3] Recall the features in chapter 3, which I argued might be contingently, but not necessarily, connected with statistical discrimination.

[4] Many assume that the relevant statistics here are statistics regarding the differential crime rate across different groups. However, given the stated aims, the "relevant statistics" here refers to statistics regarding the marginal effects of the use of additional police efforts across different groups. Accordingly, it may, say, be instrumentally rational to profile a racial group that has a *lower* crime rate but is more responsive to the relevant police efforts than that of another group which has a higher crime rate which, however, is largely unaffected by extra police efforts, e.g., because members of the relevant group have more to gain from crime and much worse job opportunities.

[5] For some considerations about when reliance on statistical information will be less relevant, see Laurence Thomas (1994), "Statistical Badness," *Journal of Social Philosophy* 25, pp. 37–40; David Wasserman (1996), "Racial Discrimination and Police Discretion," in JohnKleinig (ed.), *Handled With Discretion: Ethical Issues in Police Decision Making* (New York: Rowman & Littlefield), p. 123.

Fourth, the profiling is appropriately deployed, i.e., applied in all and only in those situations where this will lead to better crime-related results.[6] Finally, useful information concerning the crime rates in groups other than racial groups is treated no differently than equally useful information concerning the crime rates in different racial groups. Given the satisfaction of these conditions and only given that, we have an interesting moral question: Is unalloyed racial profiling morally justified?[7] (Henceforth, I mean by "racial profiling" unalloyed racial profiling.)

Before trying to answer this question, we should acknowledge that, in any account, racial profiling, like discrimination in general, can be more or less obviously wrong. First, racial profiling *in favor of* minorities is often taken to be less morally wrong, e.g., few would complain if at Christmas time the police gave priority to making car drivers of a non-Middle Eastern appearance take alcohol tests, and it is unlikely that drivers of Middle Eastern appearance would complain of being treated on the basis of a favorable, but, as they see it, nevertheless humiliating stereotype. Also, in some cases, profiling in itself will not harm those profiled, will be employed in fighting serious crime, and will be largely effective in achieving this aim. The police making advantageous use of race-sensitive statistical information when running names through a computer in an attempt to catch terrorists might be like this. In other cases profiling will harm those subjected to it, and/or will target less serious crime, and/or will be ineffective. Racially biased but in other ways random searches conducted in an attempt to crack down on illegal possession of soft drugs that humiliate those searched would doubtless fall into this latter category. Obviously, differences like these are relevant to our moral assessment of racial profiling, and later I shall make some brief comments about their relevance. However, for the most part, in what follows I can set aside these features and simply discuss the moral justifiability of racial profiling as such, having in mind mostly decisions as to whom to stop, search, and question. First, I discuss three objections to racial profiling, all of which I find to be nondecisive in various respects (sections 2 to 4). Second, I explore an interesting way in which racial profiling might be incompletely morally unjustified—namely, when the statistical facts to which defenders of racial profiling appeal reflect an unjust social situation created by those very individuals who prefer that the police employ racial profiling (sections 5 to 8).

[6] For the relevance of this condition, see Thomas (1994), "Statistical," p. 31; J. Angelo Corlett (1993), "Racism and Affirmative Action," *Journal of Social Philosophy* 24, pp. 164–165.

[7] The reason we do not have an interesting moral question, if we do not make these assumptions, is that, in that case, racial profiling is morally wrong for obvious reasons.

2. A Right to Be Treated as an Individual

As noted in chapter 3, many think that statistical discrimination is morally problematic as such. One important objection to racial profiling—indeed to statistical discrimination as a whole—is that it is morally impermissible to treat an individual on the basis of statistical information about socially salient groups of which he is a member.[8] All individuals have a right to be treated on the basis of their "individual properties" and not to be treated simply as statistical entities.[9] David Miller, for instance, writes that to make the assumption about a particular female applicant "that she is liable to perform at a lower level because of [a] decision to have children" on the basis of "information that relates to the whole group or class" to which she belongs, i.e., the group of women, is "to fail to treat her respectfully as an individual, and potentially to commit an injustice."[10] In a similar spirit, Peter Singer entertains, though ultimately rejects, the complaint that statistical discrimination fails to recognize us as "the unique individuals that we see ourselves as."[11]

Before this addressing this complaint, I want to explain why it is not the case that one fails to treat someone as an individual if, and only if, one engages in statistical discrimination against this person. First, one may treat someone as an individual if one treats this person on the basis of *all* available evidence, i.e., statistical as well as nonstatistical evidence, including statistical evidence about the like errors of perception, etc. Second, one may fail to treat someone as an individual even if one does not engage in discrimination. So suppose a policeman stops anyone carrying a Rolex watch to check if it is stolen. Presumably, this policeman does not treat those whom he stops as individuals—e.g., he

---

[8] E.g., Corlett (1993), "Racism," p. 166; Elizabeth S. Anderson (2010), *The Imperative of Integration* (Princeton, NJ: Princeton University Press), p. 165. To treat an individual on the basis of statistical information about nonsocially salient groups of which he is a member—e.g., college graduates who work in the education sector, have two or more kids, have been married to the same spouse for more than 10 years, have an income between £25,000 and £50,000 etc.—does not qualify as statistical *discrimination* as I have defined it in chapter 3. However, if one thinks that people have a moral right not to be treated on the basis of statistical information about the groups of which they are members, one will object to this kind of statistically motivated differential treatment as well.

[9] Ronald Dworkin (2003), "Bakke's Case: Are Quotas Unfair," in B. Boxill (ed.), *Race and Racism* (Oxford: Oxford University Press), p. 302, thinks that arguments based on an appeal to the right to be treated as an individual is "frequently made on editorial pages." In relation to admission procedures at universities, and more generally, he dismisses such arguments because "[a]ny admissions procedure must rely on generalizations about groups that are justified only statistically" (p. 303).

[10] David Miller (1999), *Principles of Social Justice* (Cambridge, MA: Harvard University Press), pp. 168–169. Miller thinks that this particular kind of injustice would not arise in a case where a woman says that "she is quite likely to want to take a career break in a few years' time" to have a child (p. 309).

[11] Peter Singer (1983), "Is Racial Discrimination Arbitrary?" in Jan Narveson (ed.), *Moral Matters* (New York: Oxford University Press), p. 316.

would even stop someone whom he could see was the CEO of Rolex Watch Inc.—and yet he does not engage in statistical discrimination as I have defined the notion, since presumably people carrying Rolex watches do not form a socially salient group.[12] These considerations show that not treating someone as an individual and statistical discrimination are more loosely connected than one might initially think.

Still, surely, some cases of statistical discrimination involve a failure to treat discriminatees as individuals, so let us return to the cogency of this moral complaint. In response, I first want to note that few reflective people can put it forward in good faith. This is so because very few people think that it is never the case that members of other groups should not be treated on the basis of statistical information about the groups to which they belong. For instance, most think that it is morally permissible to make retirement for airline pilots at the age of 60 compulsory. But if so, they do not really think that it necessarily violates a right to be treated as an individual. Or at least, they would then have to explain why this case of statistical discrimination is different from others. Similarly, often people think that they should be treated on the basis of statistical information about the group to which they belong, when the relevant information is positive. For instance, when some people complain that Muslims should not be judged on the basis of what a tiny minority of Islamist terrorists do, presumably, this complaint assumes that had it, absurdly, been the case that most Muslims were Islamist terrorists it would have been permissible to treat Muslims on the basis of this generalization. While these observations are interesting, they do not take us to the core of the issue. The mere fact that few reflective people are in a position to put forward in good faith a complaint about the way in which they are being treated on the basis of statistical information about groups to which they belong, does not show that they are not being treat impermissibly or unjustly. As noted in chapter 8, it can be true that someone is unjustly treated in a way that she is not in a position to complain about.

Setting aside who is in a position to complain about being treated on the basis of statistical information, I now turn to the correctness of the complaint. This pertinent position, when understood as involving a claim about moral impermissibility (which may not be how Miller understands his position), strikes me as hopelessly strong. Suppose we have reliable statistical information that perpetrators of racially motivated hate crimes are almost always from a

[12] While people who wear Rolex watches do not form a socially salient group, wearing one certainly is an indicator of one's belonging to one such group, if, as I have argued, rich people form a socially salient group. For this reason, one might think that the case at hand may involve some sort of derivative form of discrimination. Allowing this possibility implies a need for revising my definition of discrimination slightly. I say "slightly" because membership of socially salient groups will still play a central role even on the revised definition.

race (however defined) different from that of the victim of the crime. It would seem obviously permissible for the police to base its investigation of such crimes on this statistical information and, thus, to be more inclined to interview people of a race different from that of the victim who were in the vicinity of the scene of the crime than people of the same race as the victim who were in the vicinity of the scene of the crime. The complaint "You can't do that—you must treat me as an individual" strikes one as shrill here. It might be replied that my case differs from Miller's in that in my case the question is how likely it is that a certain individual belongs to a particular, socially salient group given that he has the feature that expresses our underlying concern (is guilty of the crime), whereas in Miller's case the situation is the opposite: the question there is how likely it is that a certain individual has the feature that expresses our underlying concern (e.g., will apply for parental leave) given that she belongs to particular, socially salient group. While this asymmetry certainly obtains, it is hard to see what moral difference it should make from Miller's point of view, since in neither case is the target of our action treated on the basis of individual properties.[13] Anyway, there are cases that are like Miller's in the relevant respect and where it is clearly morally permissible to base decisions on statistical information about groups, e.g., it would be permissible not to hire an alcoholic as a pilot given statistical information that most alcoholics from time to time fail to keep sober on the job or to have a law regarding the issuing of drivers' licenses that uses age as a proxy for driving abilities.

A weaker position says that while it is not morally impermissible to treat an individual on the basis of statistical information about socially salient groups of which he is a member, it is always morally preferable to treat individuals on the basis of their "individual" properties rather than on the basis of statistical information about the relevant socially salient groups of which they are members. This weaker claim also seems false, albeit less obviously so. To say that people ought to be treated on the basis of their individual properties is based on the assumption that when people are being treated on the basis of the individual properties that they are reasonably believed to have, it is more likely that they receive the treatment that they in fact ought to receive in virtue of the individual properties they in fact have (recall the discussion in chapter 3). But suppose now that exactly the opposite is true, e.g., I am an incompetent, street-unwise policeman, and if I treat potential suspects in accordance with their individual properties, as I perceive them, rather than on the basis of statistical information about the socially salient groups of which they are

[13] It might be urged on Miller's behalf that what really matters is not the failure to treat people on the basis of their individual properties, but that in his case, unlike in mine, the target is treated on the basis on statistical properties of the socially salient group to which she belongs and this is unjust. However, this suggestion amounts to no more than a statement of the moral proposition for which an account is needed.

members, I will be more likely to stop and search the innocent. I readily grant that this is not how things work out usually, but it is not my point to deny this. Rather my point is that we ought to treat people on the basis of the individual properties that they in fact have. We should then favor those practices that ensure that we do so to the highest degree possible. Normally, this is best achieved if people are treated on the basis of their individual properties, as we perceive them. However, when this is not the case, we should not treat people on the basis of individual properties, but as a morally problematic second-best on the basis of proxies. So, at most, in some cases it is a derived truth that we ought not to treat people on the basis of proxies and that is not enough to show that statistical discrimination is never morally permissible, all things considered.[14]

3. Unequal Treatment

It is a definitional property of a statistically discriminatory policy that it gives members of the target group worse comparative prospects of achieving the relevant benefit or of avoiding the relevant harm that the policy involves. By way of illustration suppose Nordic-Americans "as a group are especially gifted at cold-weather sports, and employers who need some minimal level of cold-weather sports ability in their employees find it most efficient simply to give preference to [Nordic-Americans] in their hiring practices."[15] As a result Chinese-Americans and Nordic-Americans with identical winter sports potential talent, and with equally strong cold weather sports employment aspirations "have different prospects of success in competition for these positions of advantage."[16] Arguably, the fact that they enjoy differential opportunities constitutes unequal treatment and is, thus, intrinsically morally wrong.

Despite the immediate plausibility of this objection, it fails to show that statistical discrimination necessarily constitutes morally problematic unequal treatment. From a Dworkinian perspective, unequal treatment is morally problematic only when it constitutes a failure to treat as equals and sometimes

[14] See also Richard Arneson (2006), "What Is Wrongful Discrimination?," *San Diego Law Review* 43, p. 787; Michael Levin (1992), "Responses to Race Differences in Crime," *Journal of Social Philosophy* 23, p. 23; Matt Cavanagh (2002), *Against Equality of Opportunity* (Oxford: Clarendon Press), pp. 190–193. For further discussion of the complaint against not being treated as an individual, see Kasper Lippert-Rasmussen (2011), " 'We are all Different': Statistical Discrimination and the Right to be Treated as an Individual," *Journal of Ethics* 15, pp. 47–59.

[15] Arneson (1999), "Against Rawlsian Equality of Opportunity," *Philosophical Studies* 93.1: 77–93, p. 105.

[16] Arneson (1999), "Against Rawlsian," p. 106. As Arneson points out, this situation violates Rawls's principle of fair equality of opportunity.

treating as equals requires unequal, i.e., nonidentical, treatment.[17] First, statistical discrimination against (in favor of) simply worsens (improves) the relative comparative prospects of the targeted group relative to what they would otherwise have been. Hence, it may lead to a situation, unlike the one in Arneson's cold weather sports example, where statistical discrimination serves to promote equal opportunities such that the relevant unequal treatment does not strike one as morally problematic. Suppose statistics show that Nordic-American professional skiers are more likely to be involved in illegal doping than Chinese-Americans. However, for some reason it is harder for the police to present evidence of illegal doping to the courts leading to the conviction of a Nordic American professional skier than it is for them to do the same with a Chinese American. Suppose the police decide on the basis of the statistical overrepresentation of Nordic Americans among professional skiers involved in doping to use a disproportionately larger share of its resources to investigate suspected case of illegal doping involving Nordic American skiers. As a result, professional skiers, who illegally seek to enhance their performance through doping, face an equal risk of being punished, whether they are Nordic Americans or Chinese Americans.

Second, we do not in general think that it is morally problematic if a policy worsens the comparative prospects of a group in one area unless this group finds itself in a similar position in a large number of other areas such that the overall prospects of members of this groups end up being worse than the prospects of members of other groups.[18] For instance, we do not reject a new tax as morally problematic simply on the ground that it worsens the comparative prospects of a group and in that sense involves unequal treatment. But then, on pain of inconsistency, we cannot consider statistical discrimination inherently morally problematic for this reason. Racial profiling cannot be impermissible simply because it involves unequal treatment (which undeniably it does).

4. Unfairness

Whenever statistical discrimination involves statistical indicators for features, which are under people's own control, it might be said that innocent people, i.e., those who refrained from acquiring the relevant feature, are subjected to statistical discrimination as a result of what other people chose to do. Since

[17] Ronald Dworkin (1985), *A Matter of Principle* (Oxford: Clarendon Press), p. 190. According to Joel Feinberg (1973), the basic principle of comparative justice is that people should be treated alike provided that there are no relevant differences between them, *Social Philosophy* (Englewood Cliffs, NJ: Prentice Hall), pp. 98–99. From this perspective, my claim is that the unequal treatment involved in statistical discrimination may map unto relevant differences.

[18] For instance, see Shlomi Segall (2012), "What's so Bad about Discrimination?" *Utilitas* 24.1: 82–100.

being statistically discriminated against involves various costs, it might be argued that statistical discrimination is unjust, provided that the object of statistical discrimination is not responsible for what others choose to do, because it involves imposing costs on innocents on account of what others chose to do.[19]

A similar objection says that statistical discrimination is unfair because it involves subjecting people to disadvantageous treatment on the basis of something that does not reflect their own choice.[20] To see why this objection differs from the first variant of the unfairness objection, consider a case where, in connection with a World Cup Match, the police prevent me from entering a certain part of the town, because I wear a typical hooligan outfit. Here it might be said that in a sense I am being treated disadvantageously on the basis of something that does reflect my choice, i.e., my choice of dress. Hence, the second variant of the fairness objection does not apply. Still, the first variant might be said to apply because supposing that I am a law abiding and peaceful person I might complain that I am being disadvantageously treated because of the unruly behavior of other people, who wear the same kind of dress as I do.[21]

Note finally that both variants of the fairness objection are different from the individual treatment complaint. If others are responsible for an individual property of mine and I am being treated disadvantageously on this basis the

---

[19] One particularly strong version of this objection says that it is unfair that the objects of statistical discrimination are being punished for what others do. When "punishment" is understood literally, the objection is obviously flawed, since statistical discrimination need not involve any expression of an authoritative judgment that the targeted individuals have violated rules good citizens are supposed to respect—see Thomas Scanlon (1999), *What We Owe Each Other* (Cambridge, MA: Harvard University Press), p. 266. For the objection not to be obviously flawed "punishment" must simply mean something like "the imposition of costs" in which case the objection really is no different from the one discussed in the main text.

[20] Note that insofar as one is a member of the relevant group of one's own choice (and the relevant group is identified partly or wholly through having the relevant target property), then the complaint that it is unfair that one is being treated of the basis of one's membership of a certain group and not on the basis of one's individual merits strikes one as weak if not forceless. This might suggest that choice and treatment on the basis of group membership interacts in such a way that the latter is morally objectionable only when group membership is not a result of choice. Note also that in the case of some groups the contrast between "being assessed on one's own merits" and being assessed "merely as a member of a group" is misleading. Suppose that a research body could give a research grant to a Nobel Prize winner or to someone else. The complaint that the research body merely assessed the applicants on the basis of their group membership and not on the basis of their individual merits seems strange in this case. Assessing researchers on the basis of their membership the group of Nobel Prize winners is one way, albeit not perfectly reliable, of assessing their individual merits.

[21] If I am not a law abiding, peaceful person but a hooligan, it is still true that I am being treated on the basis of what others choose to do. Yet, given my disposition to attack supporters of other clubs simply because they are supporters of others clubs, I am in no position to complain about statistically motivated, precautionary police measures and, arguably, such measures are not unfair to me. For a discussion of moral complaints relevant to the present case, see G. A. Cohen (2006), "Casting the First Stone: Who Can, and Who Can't, Condemn the Terrorists?" *Royal Institute of Philosophy Supplement* 58: 113–136, and Saul Smilansky (2006), "The Paradox of Moral Complaint," *Utilitas* 18, pp. 284–290.

individual treatment complaint does not apply, yet both of the fairness objections may. Conversely, in the unlikely case that I am responsible for the statistical properties of a group of which I am a member and on the basis of whose properties I am being treated disadvantageously, the individual treatment complaint applies, but neither of the fairness complaints do.[22]

I think both variants of the fairness complaint fail. Suppose we have a package of different kinds of statistical discrimination whose aggregate effects on each person are positive and positive to an equal degree, and suppose the only two feasible states are one with and one without the package of statistically discriminatory practices. So while everyone is disadvantaged by some kinds of statistical discrimination, everyone benefits from other kinds of statistical discrimination. In this scenario everyone is being treated on the basis of what others do and not on the basis of what they themselves do. However, since everyone benefits, all things considered, it is not clear that this is in any way unfair. Disregarding a complication to be introduced below in section 7, it is not unfair to ask people to bear some costs in a particular dimension, if at the same time they benefit in other dimensions such that everyone is made better off to an equal degree. Accordingly, if it is unfair to impose uncompensated burdens on people on account of what others choose to do and not on the basis of what they themselves do, statistical discrimination is not in itself unfair. At most it follows that those kinds of statistical discrimination that impose burdens on some without benefiting everyone to an equal degree are unfair. But even that weaker conclusion is warranted only to the degree that it is always unfair to impose uncompensated burdens on people on account of what others choose to do and not on the basis of what they themselves do. In fact, I think this principle is false.

Consider the last part of this claim. Suppose we have an initial equal distribution of food. A natural disaster ruins the food stock of some people. Surely, there is nothing unfair about transferring some food from the better off to the worse off if this is necessary for their survival, although the better off are being deprived of some of their food on account of something, which does not reflect any choice of theirs.

Consider next the former part of the initial claim. Suppose again that we start with an equal distribution. The population is divided into two groups,

[22] The distinction I make between the fairness complaint and the individual treatment complaint is related to, but different from, the distinction Cavanagh (2000) makes between a complaint based on the responsibility principle, i.e., "the idea that it is wrong to treat people differently according to differences for which they are not responsible," and a complaint based on "the principle that a person's prospects should depend on facts about him rather than on facts about some group of which he happens to be a member"; see *Against Equality of Opportunity*, pp. 186–187. The former complaint on Cavanagh's account concerns what "we are allowed to regard as ultimately relevant" whereas the latter concerns "ways we can pursue the things we regard as relevant," p. 187. As I have construed the complaints both concern the latter.

vegetarians and nonvegetarians. A subgroup within the nonvegetarians begins to raise freely roaming cattle, thereby depriving the vegetarians, who are all farmers, of their source of food and income, but leaving the other nonvegetarian subgroup, the hunters, unaffected. Suppose the only way to prevent the farmers from becoming very badly off is to impose certain restrictions on movement on all nonvegetarians, something which will marginally harm members of the hunter subgroup. Surely, it would not be unreasonable for the hunters to accept such a marginal, uncompensated loss to prevent a much greater loss to others and it would not be unreasonable of the farmers to refuse to accept a great loss to avoid marginal harmful restrictions on the innocent hunters on account of other people's choice to raise cattle.[23] This suggests that it may not always be unfair to impose uncompensated burdens on people on account of what others choose to do.[24]

Schauer seems to agree with the conclusion defended here: "[I]f there were something necessarily unjust about relying on a generalization] then it would be as wrong to generalize about the capacities of fourteen-year-old voters and sixteen-year-old drinkers as it is to generalize about the capacities of elderly pilots."[25] While I agree with Schauer's conclusion, I am unsure I follow his implicit argument here. Surely, the two cases mentioned may differ in respects other than that they involve relying on generalizations and these other respects might be morally relevant in such a way that one is more wrong than the other. Schauer also notes that "the total avoidance of generalization" is not only "impossible," but are often "unwise."[26] Both of these claims may be true, but it is not obvious that they imply that generalizations are not bound to involve injustice. Presumably, if race discrimination were hardwired into our genes, we might then find it excusable, but I do not think that we for this reason would not find it unjust. Also, the fact that an action might be unwise does not show that it is not the just action. The person who hands over his valet to the armed robber does not think that he is doing what is just simply because it would be unwise to do otherwise.

---

[23] Arthur Isak Applbaum (1996), "Response: Racial Generalization, Police Discretion, and Bayesian Contractualism," in Kleinig (ed.), *Handled*, 145–157, pp. 150–151; Thomas C. Scanlon (1982), "Contractualism and Utilitarianism," in A. Sen and B. Williams (eds.), *Utilitarianism and Beyond* (Cambridge: Cambridge University Press), 103–128, pp. 122–123.

[24] It might be replied that while the imposition of a marginal burden on the hunters is, all things considered, fair, it is still in one respect unfair that they are made to bear a burden, however small, on account of what other nonvegetarians choose to do. To accommodate this concern, assume that the nonvegetarians who choose to raise cattle have no acceptable alternative to doing so.

[25] F. Schauer (2003), *Profiles, Probabilities, and Stereotypes* (Cambridge, MA: Harvard University Press), p. 132.

[26] Schauer (2003), *Profiles*, p. 299.

5. The Making of Statistical Facts and the Justifiability of Statistical Discrimination

I have argued that statistical discrimination is not wrong in itself, albeit it may often be morally wrong for various contingent reasons (see chapter 3 and the sections 2 to 4 in the present chapter). Still, racial profiling seems problematic in a way that many other kinds of group profiling are not. In the remaining sections of this chapter I will propose an account thereof. I suggest it, not so much as my account of why racial profiling is morally wrong in a way that many other kinds of statistical discrimination are not wrong, as an account of what makes many people hold this belief (whether or not, ultimately, this is a good reason for holding this belief). More specially, I will suggest that the social facts on which racial profiling is based are themselves the result of morally wrongful patterns of behavior, often on part of many of those who benefit from racial profiling and encourages the racial profiling of others. In order to present this argument I want to take my point of departure in a qualified defense of racial profiling offered by Mathias Risse and Richard Zeckhauser.

Risse and Zeckhauser tentatively argue that "in a range of plausible cases" certain kinds of racial profiling by the police are morally justified.[27] This is so, according to Risse and Zeckhauser, whether we presuppose a utilitarian or a deontological moral theory.[28] In the case of utilitarianism, Risse and Zeckhauser point to the fact that crime prevention increases the quality of life for almost everyone. They also suggest that "the primary contributor" to the harm involved in racial profiling—i.e., the chief cause of "resentment, hurt, and loss of trust among" groups subjected to increased police attention as a result of such profiling—"is likely to be underlying racism or underlying socio-economic disadvantages, rather than profiling as such."[29] While these harms may be great, the incremental harm involved in racial profiling is comparatively small. In the case of deontological theories, the extra burdens imposed on some people by racial profiling are neither unfair, nor rights-violating. They are not rights-violating, because "profiling can be justified on legitimate grounds," i.e., the need to curb crime;[30] and in any case, unalloyed profiling as such "does not contribute to [an] oppressive relationship" involving pejorative discrimination against one racial group by another. The extra burdens are not unfair, because whether "functioning reciprocity" obtains or not, racial profiling is not unfair

[27] Mathias Risse and Richard Zeckhauser (2004), "Racial Profiling," *Philosophy and Public Affairs* 32.2, p. 133.

[28] The argument they offer from a utilitarian point of view applies, *mutatis mutandis*, from a desert-prioritarian view as well.

[29] Risse and Zeckhauser (2004), "Racial Profiling," p. 144. See, however, Annabelle Lever (2005), "Why Racial Profiling is Hard to Justify," *Philosophy and Public Affairs* 33.1, 94–110.

[30] Racial profiling might still be rights-infringing even if it is morally justified.

to those subjected to it. A society characterized by functioning reciprocity is one in which

> Differential burdens are imposed if five conditions are met: (1) It is required by the nature of the public good. (2) The imposition occurs through a suitable democratic process that gives proper weight to the interests of all involved. (3) The social importance of the good is widely acknowledged. (4) State interference with individual lives entailed by the imposition of differential burdens is within the limits of what citizens can reasonably be expected to bear—in particular the imposition of the good does not involve the violation of widely acknowledged rights of individuals. (5) It is widely acknowledged that, similarly, differential burdens are imposed for other public goods as well, and that this differential imposition of burdens for different good works, by and large, to society's advantage, although it may not work to every individual's advantage.[31]

Where functioning reciprocity obtains, to deny that racial profiling is a fair means of achieving the collective good of reduced crime on the grounds that it imposes differential burdens on those subjected to it would, in effect, be to imply that collective goods can hardly ever be brought about in a way that is fair. But this is implausible. Even if functioning reciprocity does not obtain, racial profiling is fair to those subjected to it, because even they (or at least almost all of them: those who are law-abiding and perhaps even some of those who are not) will be net-beneficiaries of racial profiling. At any rate, this is likely to be the case when most crimes committed by members of a particular racial group are against other members of the same group.[32]

These results, if sound, are surprising. Many observers are inclined to dismiss racial profiling, in a range of cases, out of hand. But if Risse and Zeckhauser are right, such profiling can be justified within either of two of the main contemporary moral theories.

While Risse and Zeckhauer's primary and sensible aim is not to produce a bottom-line conclusion, but to delineate the shape of the moral debate about [racial] profiling, they omit one very important consideration from the overall map. It is crucial to their argument that overall racial profiling benefits the people subjected to it. If these people were on balance harmed by racial profiling, the utilitarian argument would look much less firm and the fairness argument would not work in the absence of functioning reciprocity. For the sake of argument I shall assume that people being subjected to racial profiling are net-beneficiaries relative to a state in which no racial profiling takes place; for I want to ask in what way, exactly, the net benefit justifies racial profiling. Note that this is no more than an assumption: I am not claiming that racial profiling

[31] Risse and Zeckhauser (2004), "Racial Profiling," p. 157.
[32] Risse and Zeckhauser (2004), "Racial Profiling," p. 163.

is in fact beneficial in this way in real life. Of course, if it is not, all sorts of additional problems in justifying racial profiling will arise. The line of argument I shall explore is this: insofar as racial profiling is beneficial, this is not a statistical fact that obtains independently of what we choose to do, but a fact that obtains exactly *because of* what we choose to do. This means that we cannot simply appeal to this statistical fact in a discussion of whether racial profiling is justified. Rather, we must ask whether we can justify acting in such a way that this fact obtains.

Risse and Zeckhauser's argument focuses on the use of racial profiling in the United States. Based on US crime figures, much of their discussion concerns the racial profiling of African Americans. Whether these statistics are reliable is a complex issue different from the one I shall pursue here. I shall assume along with Risse and Zeckhauser that they are reliable to a significant extent. While the concept of unalloyed racial profiling is, like the presentation in this section, color-blind, I shall follow Risse and Zeckhauser in discussing cases where African Americans are the objects of racial profiling in order to engage with their conclusions about the United States as directly as possible. The conclusions I draw are, of course, easily generalized across hypothetical cases in which members of different racial groups (however these are delineated) switch history and circumstances, and to actual, non-American cases.

6. Putting the Argument to the Interpersonal Test

As Risse and Zeckhauser rightly point out, their argument in support of profiling "in a range of plausible cases" is "consistent with" (their) support for far-reaching measures to reduce "racial inequities and inequality."[33] However, a stronger claim is warranted. Plausible justifications for racial profiling suggest not merely that we may, but that we should, adopt measures to reduce racial inequities and inequality.

Let me explain. Suppose (i) that African Americans are more likely to commit certain crimes than European Americans solely as a result of the deprivation resulting from discrimination and unjust, racial inequality. Hence, if discrimination and unjust, racial inequality were eliminated, the crime rates of European Americans and African Americans would converge over time.[34]

[33] Risse and Zeckhauser (2004), "Racial Profiling," pp. 134, 161.

[34] Here I focus on overall crime rates. I disregard the possibility that different types of crime have different causes. I also ignore the fact that, even if differential overall crime rates of European Americans and African Americans reflect inequalities and racial discrimination, there might be particular types of crime for which differential crime rates do not reflect, or reflect to a lesser degree, inequalities and racial discrimination. Finally, I ignore that not even all differences in overall crime rates across groups need reflect injustice and discrimination (recall the discussion

Suppose, next, (ii) that all European Americans could choose to act so that, in the long run at least, African Americans would no longer suffer unjustly from discrimination and racial inequality. Suppose (iii) that given the existing discrimination and racial inequality, racial profiling will benefit African Americans as well as European Americans. It will benefit African Americans because, although African Americans will have to bear the costs of racial profiling, they will also enjoy the lion's share of the benefits in the form of reduced crime, since African Americans are more likely than European Americans to be victims of crime. Suppose, finally, (iv) that relative to a state in which there is neither discrimination nor racial inequality, European Americans generally benefit from discrimination and racial inequality in their favor, while African Americans are generally harmed (a supposition that is consistent with the idea that in some respects European Americans as well as African Americans may benefit from the cessation of discrimination and racial inequality).

Before I proceed, I need to clarify these assumptions as well as the role they play in my argument. First, all four assumptions, and especially (i) and (ii), are controversial, but they are made quite commonly. Consider, for instance, Bernard Boxill's view of (i): "A typical argument [. . .], is that the mainstream would rob and steal as readily as the black poor if it had to do these things in order to achieve a decent living. Now, I think that there is much to be said for this argument; but it cannot explain why crime among the poor is disproportionately black. I suggest that we should not overlook

---

of proportional representation in chapter 7). In a response to an article that contains much of the material, including the previous part of this note, presented here in chapter 11, Mathias Risse (2007), "Racial Profiling: A Reply to Two Critics," *Criminal Justice Ethics* 26.4, 4–19, p. 7, notes that he finds the "mono-causality" involved in (i) "implausible" and offers a convincing reason to doubt it. He seems to accept that a weaker version of (i) might be true, but thinks that it is "too weak for Lippert-Rasmussen's purposes," p. 7. However, my primary aim there, as well as here, is in part to map the contours of the debate about racial profiling, and, for the reasons stated, assuming (i), even if false, is helpful in achieving this aim. Indeed, in the article, as well as here, I myself consider how the argument presented below is affected if a weaker version of (i) is true. In fairness to Risse, I should say that he notes that my article "may well be read in such a way that [Kasper Lippert-Rasmussen] agrees with [Risse's critique of (i)]," p. 18*n*15. I agree. Risse also notes that I use "(i) to trigger the idea that it the whites' fault (or at the very least have made it true) that blacks have higher crime rates and that this, under the circumstances, diminishes their claim to the protection which profiling might contribute," p. 7. It is unclear what Risse means by saying that I "trigger" this idea. At any rate, I do not claim (nor deny) that whites are at "fault" or to "blame," so the disjunction is misleading; I explicitly state the discriminatees cannot refuse "to bear any cost, however small, if bearing that cost will prevent the violations of" discriminators' right to life and liberty; and I distance myself from the claim that African Americans "have no influence whatsoever on their own situation." Similarly, when Risse thinks that he quibbles with (i) when he points to "the differential involvement of whites in racist practices" and notes that this renders it "implausible that no, or few, whites now have the claim to the kind of protection toward which the use of race in police tactics is supposed to contribute," p. 7, this is seriously misleading. First, I explicitly noted (and note below in section 7) the differential involvement to which Risse refers. Second, (i) does not say anything about the kind of protection that whites have a claim to.

the part that racism probably plays in this."[35] I have no settled opinion as to whether (i)–(iv) are true. Below I shall consider what follows if (i) and (ii) are false. The primary argumentative aim in sections 5 to 8 is to map the structure of the philosophical debate about racial profiling, i.e., to draw conditional conclusions of the sort claiming: to the extent that *these* empirical assumptions obtain, racial profiling *is* or *is not* justified. For this purpose, I need not actually make any empirical claims. However, for what it is worth, I should say that, in my view, (i)–(iv) are not wildly inaccurate. Thus, if all European Americans acted differently in certain ways, there would indeed, before long, be a significant reduction in the extent to which African Americans suffer from unjust discrimination and racial inequality—and as a consequence a significant reduction in the differences in crime rates across racial groups. I suggest, then, that the considerations below do point to a real problem for Risse and Zeckhauser's contention that racial profiling is justified.

Second, I suspect that readers who are sympathetic to assumptions (i)–(iv) will tend to agree with the criticism of Risse and Zeckhauser's argument in favor of racial profiling I shall make shortly, and that those who are unsympathetic will tend to reject the criticism although they might concede that if these factual assumptions were true, the criticism would hit its target. The fact that this is so, if it is—the fact, in other words, that one's tendency to agree with Risse and Zeckhauser's argument is affected by one's view of what leads racial profiling to provide a net benefit to African Americans—is evidence that I have achieved my primary aim of charting the debate about racial profiling.

Third, Risse and Zeckhauser seem willing to grant suitably weakened versions of assumptions (i) and (ii).[36] For they write: "Plausibly, the disproportionate tendency of minorities to engage in criminal activity is, to some extent, a symptom of discrimination."[37] And they offer a list of suggestions as to how political initiatives may reduce discrimination and racial inequality that could presumably be implemented with some success if a majority of the American population were to support them.[38]

Fourth, by assuming that European Americans could act differently in such a way that discrimination and racial inequality would disappear, I mean just

---

[35] Bernard Boxill (1991), "Wilson on the Truly Disadvantaged," *Ethics* 101, p. 588.

[36] They themselves argue for the (iii); their article is neutral on (iv).

[37] Risse and Zeckhauser (2004), "Racial Profiling," p. 152.

[38] In part for that reason, and in part because of my explicit reservations regarding (i)-(iv), it is somewhat surprising that Risse (2007), "Racial Profiling: A Reply," p. 8, claims that "there is no politically feasible sense in which we now 'could' [bring about a nonracist society]." Again: I have not committed myself to (ii); I introduce it to map the debate about racial profiling; and, like Risse, I do not think (ii) is "wildly inaccurate" because I think we could, in the political feasibility-sense of "could," take steps to reduce racial discrimination, so the argument presented in this section might cast some critical light on Risse's and Zeckhauser's defense of racial profiling.

that. I do not mean to suggest that European Americans are racially prejudiced to a higher degree than African Americans are; that African Americans, unlike European Americans who enjoy perfect control, have no influence whatsoever on their own situation; that racial discrimination explains more, exactly as much as, or less of the differential crime rates than does the worse social and economic situation of African Americans; or that no serious progress has been made since the start of the civil rights movement to reduce discrimination and racial inequality. These and similar claims play no part in my argument.

With assumptions (i)–(iv) granted, I now want to put the defense of racial profiling appealing to benefits to African Americans to what G. A. Cohen calls the "interpersonal test."[39] Cohen uses this test to criticize an interpretation of Rawls's difference principle. According to this interpretation, inequality that is necessary to make the worst off as well off as possible because, in the absence of inequality-inducing incentives being offered to talented people, these people will choose to be less productive than they choose to be when offered ample incentives, is justified. Hence, the difference principle so construed makes room for incentive-based inequality. Cohen shows that in a just, Rawlsian society, where people are guided in their daily lives by the Rawlsian principles of justice, talented people do not through their economic choices bring it about that incentives are necessary to make the worse off better off than they would be under equality. The failure of the Rawlsian incentive argument for inequality is perspicuously displayed when, instead of a third-person setting for the argument, we imagine an interpersonal setting in which talented people appeal to the difference principle in order to convince nontalented people that since they, the talented people, will insist on incentives or otherwise be less productive, nontalented people should endorse inequality-producing incentives. Thus presented, the argument has an unattractive ring to it, mainly because the incoherence of the professed acceptance on the part of talented people of the difference principle and their insistence on incentives in the absence of which worse off people could be better off becomes more vivid. Arguments in favor of racial profiling that appeal to the benefits for African Americans produced by racial profiling suffer from an analogous flaw, given (i)–(iv) at any rate.

Compare the following two statements of an argument appealing to the net benefits to African Americans of racial profiling. The first is impersonal, the second interpersonal. To illustrate this difference, let us imagine that the impersonal version of the net benefit argument is offered by a group of Russians to another group of Russians arguing about American police tactics, while the interpersonal version is offered by all European Americans to an

[39] G. A. Cohen (1992) "Incentives, Inequality and Community," in G. Petersen (ed.), *The Tanner Lectures on Human Values* 13 (Salt Lake City: University of Utah Press).

audience of only African Americans in an all-American, Rousseau-inspired, republican-assembly-like context:

Impersonal

(1) What will provide African Americans with a net benefit should be implemented.
(2) Implementing racial profiling will provide African Americans with a net benefit.
(3) Thus, racial profiling should be implemented.

Interpersonal

(4) We, African Americans and European Americans, should implement what will provide you, African Americans, with a net benefit.
(5) Implementing racial profiling will provide you with a net benefit.
(6) Thus, we should implement racial profiling.

Since these two sets of statements are versions of the same argument, they differ neither in respect of the truth of their premises nor in respect of the validity of their inferences. Still, they come across quite differently. In the interpersonal version, the people putting the argument make the factual premise true *ex hypothesi*. Were European Americans to act differently, discrimination against African Americans and racial inequality would cease to exist. As a result, racial profiling would no longer be a net benefit to African Americans, since the crime rates of European Americans and African Americans would then converge ensuring that racial profiling is ineffective in reducing crime. Hence, according to the normative premise of the argument—which presumably European Americans accept, since otherwise they would not put forward the argument—European Americans should act differently. In the interpersonal context it makes sense for the addressee of the argument to point out to proponents of the argument that, by the very same standards that they appeal to, they should not make the factual premise of the argument true; and the proponents of the argument cannot simply dismiss this criticism by denying that the truth of the factual premise has anything to do with them. The same cannot be said of the Russian proponents of the impersonal version of the argument. Hence, their appeal to it cannot appropriately be met by confronting *them* with their failure to conform to the normative premise—a premise, notice, to which they appeal to in the very same argument.

The asymmetry between the impersonal and the interpersonal versions of the benefit argument shows three things. First, the argument is not just consistent with the view that discrimination and racial inequality are morally undesirable. Given (i)–(iv), the normative principle appealed to implies that nondiscrimination should be implemented and, disregarding genuine leveling-down situations, racial inequality should be eliminated.

Second, if European Americans put forward this argument, it might reflect badly on them.[40] For if they believe, or should believe, (i)–(iv), they will refrain from reviewing their own conduct in the light of their own normative principles and, in proposing the argument to African Americans, they will invite African Americans to acquiesce in, in effect, a piece of hypocrisy on their part. African Americans might well respond by asking: But why do you not stop countenancing our being worse off in the first place? Why do you suddenly care about what is in our interest, when you've otherwise disregarded it? Is it because in this particular case it is to your own benefit?

Third, the proposed justification is incompatible with what Cohen calls a justificatory community, i.e., a community where a norm prevails to the effect that if the justification for a certain policy is based on a factual assumption that a certain subset of the members of the community will act in a certain way, then that way of acting is in itself justified.[41] Accordingly, any European Americans putting the argument for racial profiling expecting African Americans not to question why they were making the factual premise true would thereby demonstrate that they do not see their relation to African Americans as involving a justificatory community. So when Risse and Zeckhauser claim that racial profiling is justified if it provides a net benefit to African Americans, even in the absence of functioning reciprocity, they cannot be talking about *comprehensive* justification.

The noncomprehensive justification with which Risse and Zeckhauser might be concerned is in some ways similar to that involved in an argument offered by me to people whom I have captured and made desperately badly off that they should agree to a miserly work-for-food scheme that will make them less badly off and me even better off. Even though my addressees care about their self-interest and this self-interest would be served by the scheme, they might care more about the injustice of their being put in a situation where, in order to get

---

[40] In saying this I am not arguing *ad hominem*. Risse and Zeckhauser are not offering an argument for racial profiling on behalf of European Americans to African Americans, just as Rawls was not defending incentives-based inequality to untalented people on behalf of himself and other talented people. Rather, they are engaged in the perfectly respectable enterprise of describing and assessing a line of argument for the benefit of the academic community (and beyond).

[41] G. A. Cohen (1992), "Incentives," pp. 279, 282–283. Risse (2007), "Racial Profiling: A Reply" observes that "crime prevention measures, no matter what their character, cannot be" comprehensively justified, because they "are all based on the assumption that some people act in ways that are unjustified (which is why prevention is needed)." Strictly speaking, this is not true because of the possibility of a gap between an action's moral and legal status (see chapter 10). However, Risse is surely right to point out that most of the crimes that racial profiling might prevent are not morally justified and, accordingly, my present pedantic point makes no substantive difference here. Still, like Risse I do not think this point undermines the value of the present discussion, because we might still think that a "policy P, which imposes special burdens on group G1, is not [conclusively] justified if it is because of the [justified] behavior of G2 that G1 is in the position in which it makes sense to impose special burdens on G1 to realize the goals pursued by P," Risse (2007), "Racial Profiling: A Reply," p. 9.

food, they must accede to my outrageous demands. In response they might display uncooperative anger in declining my offer. In refusing to cooperate with me in this kind of unfairness, they show that they conceive of me as something other than an unwelcome natural event—in fact, as an agent who can be held accountable for his or her behavior and someone with whom one can reason.

Before proceeding with my main line of argument, I need to note and set aside the complication connected with transitional periods. To the extent that eliminating discrimination and racial inequality takes time, there is a sense in which racial profiling may be comprehensively justified in the period transitional to the successful elimination of discrimination and racial inequality.[42] During this period, European Americans might rightly say that they are doing all they can to eliminate discrimination and racial inequality, and that they should act no differently, in keeping with the normative premise to which they appeal. Accordingly, they can justify their making the factual premise true (to the extent they do so) during the transitional period by appeal to the normative premise of their argument. However, given that the present might have been different if European Americans had acted differently in the past, racial profiling is comprehensively unjustified even in the transitional period in a temporally global sense. After all, racial profiling is comprehensively justified in a temporally local sense only because there have been injustices in the past.

One might alter the normative premise of the benefit argument so that no incoherence is involved in the fact that the European Americans make the factual premise of the argument true despite their adherence to the normative premise. This, for instance, would be the case if the normative principle, to which appeal is made, is that one should do what would provide one with a benefit. Similarly, no incoherence would obtain, if European Americans appeal to the normative premise of the net-benefit argument, not because they themselves accept it, but because they believe that African Americans do. These possibilities merely show, however, that the claim that the relevant arguer's appeal to a certain normative principle and his own conduct are at odds with one another is different from the claim that the arguer is unable to justify his failure to make it false that racial profiling provides African Americans with a net benefit. Accordingly, changing the normative principle appealed to, or putting it forward in an *ad hominem* mode, so as to avoid incoherence, would not in itself make it acceptable that the European Americans make the factual premise of the argument true. The reason why it is unacceptable that European Americans make it the case that racial profiling provides African Americans

[42] Risse (2007) "Racial Profiling: A Reply," p. 19n24, responds to my criticism that "racial profiling should be thought of as a short- or medium-run measure, whereas [measures to change society in such a way as to make it true that race is not significantly correlated with a propensity to commit certain crimes] should be seen as long-term measures." Given that this view is hard to distinguish from the one I canvass here regarding transitional periods, I am uncertain as to why he thinks that he offers a *response* to a critic at this point.

with a net benefit has nothing to do with which principles European Americans use to defend the practice. Of course, this is not to deny that their appeal to certain principles may be relevant when it comes to the valuation of other things; e.g., it may be significant in the assessment of the moral sensitivity of European Americans that propose the argument to an African American audience failing to observe any dissonance.

I assume without argument that while the normative principle appealed to in the net-benefit argument is not a plausible *basic* moral principle—since such principles do not refer to racial groups—it may well be derived, given certain factual assumptions, from one that does have initial plausibility, such as utilitarianism, the desert-prioritarian principle, egalitarianism or the principle of fairness. And I assume, similarly without argument, that a society in which discrimination and racial inequality make the factual premise of the benefit argument true cannot be comprehensively justified if one of these prima facie plausible moral principles, from which the normative premise of the net-benefit argument is derived, is accepted.

It might be replied at this point that I have failed to address Risse and Zeckhauser's deontological case for racial profiling under functioning reciprocity. Given such reciprocity, society has "done all it can be expected to do to distribute burdens and benefits fairly."[43] Presumably, in a society with reciprocity the distribution of benefits and burdens is comprehensively justified, and so here, as long as the factual premise of the net-benefit argument is true, racial profiling can be comprehensively justified, too. In response, I would suggest that in a fair *society* that has "done all it can be expected to do to distribute burdens and benefits fairly," the crime rates of different racial groups would not vary significantly and hence the factual premise of the net-benefit argument is likely to be false. I suspect that what Risse and Zeckhauser have in mind when they speak of something that may or may not have functioning reciprocity is the *state*, not society as such.[44] However, the mere fact that the state can have functioning reciprocity does not ensure that the distribution of benefits and burdens in the society in which it is located is fair. For a society to be fair it is necessary that both the state and nonstate parts of that society, e.g., its civil society and economy, are fair. The latter requirement entails, among other things, that people's everyday, nonlegally mandated dealings with one another are noninequality sustaining.[45] For this reason one can deny that racial profiling

[43] Risse and Zeckhauser (2004), "Racial Profiling," p. 158.

[44] While "a suitable democratic process" may be one that involves society as such, as would be the case in an ideal of participatory democracy, nothing suggests that this is what Risse and Zeckhauser have in mind.

[45] Admittedly, this assumption is controversial and many Rawlsians reject it. See, for instance, Thomas Pogge (2000), "On the Site of Distributive Justice: Reflections on Cohen and Murphy," *Philosophy & Public Affairs* 29.2: 137–169; Andrew Williams (1998), "Incentives, Inequality, and Publicity," *Philosophy & Public Affairs* 27.3: 225–247.

is fair even though the state complies with norms of functioning reciprocity without being committed to the unattractive view that collective goods could never be brought about in a fair manner. Similarly, a state could be characterized by functioning reciprocity even though the distribution of benefits and burdens on men and women are unfair—e.g., because state policies are more beneficial to men than to women owing to their different gender-determined roles, although these benefit both men and women.[46] In the light of these considerations, I conclude that, given assumptions (i)–(iv), any appeal to the way in which racial profiling is a net benefit to African Americans cannot *comprehensively* justify racial profiling. This conclusion is not affected by the theoretical choice of utilitarian or deontological starting points.

7. Noncomprehensively Justified?

The conclusion of the last section raises two further questions: Is racial profiling noncomprehensively justified given assumptions (i)–(iv)? Can racial profiling be comprehensively or noncomprehensively justified if we amend one or more of (i)–(iv), and if so, which? A policy is noncomprehensively justified when its justification assumes, correctly, that a certain subset of members of the community will act in a way that is unjustified. Noncomprehensively justified policies are either such that both their implementation and their nonimplementation are permissible or such that while their implementation is permissible, failure to implement them is impermissible. It is not clear to me which of these positions, if either, Risse and Zeckhauser would take on racial profiling. Their reluctance to pursue paternalistic arguments in favor of racial profiling suggests, but does not establish, that they would take the first option. However, some of their assumptions make more sense if the second option is in play.

Let me start with the first question. I suspect that from a utilitarian point of view while a society free of discrimination and racial profiling is better than one with discrimination and racial profiling, the latter may well be better, for the reasons given by Risse and Zeckhauser, than a society in which there is discrimination but no racial profiling. This, of course, is significant if correct. However, from this it would not follow that racial profiling is noncomprehensively justified from a utilitarian point of view.[47] That conclusion would follow only if: (a) there were no alternatives to racial profiling that were both superior to such profiling and left the rest of society, including present racial

[46] For a defense of the view concerning the scope of principles of distributive justice expressed here, see G. A. Cohen (1997), "Where the Action Is: On the Site of Distributive Justice," *Philosophy and Public Affairs* 26.

[47] Risse and Zeckhauser do not discuss racial profiling from the perspective of desert prioritarianism. However, the line of argument I develop here on utilitarian premises can also be developed, mutatis mutandis, on desert-prioritarian premises.

discrimination and present racial inequality, unchanged; or (b), although there were alternatives of this kind, none was better than the option in which racial profiling is implemented along with other measures. Risse and Zeckhauser offer no reason to think that (a) and (b) are true and hence give us no reason to believe that racial profiling is noncomprehensively justified from a utilitarian point of view. It might be replied that they simply, and plausibly, assume that no genuine alternative involves a significant reduction in existing racial discrimination and inequality. But, first, this suggestion sits uneasily with Risse and Zeckhauser's own list of measures to change that situation—e.g., Head Start, affirmative action programs, and reparations for African Americans.[48] To demand that it be shown that there are no logically possible alternatives to racial profiling (apart from those that are comprehensively justified) that are at least as good would be to set the standards for utilitarian arguments for racial profiling unreasonably high, but restricting oneself to a simple comparison of the status quo with and without racial profiling is to fail to set the standards high enough.

Second, it would be odd anyway to assume that no significant reduction of existing racial discrimination and inequality is possible. For, surely, given the will, we could reduce racial discrimination and inequality. To refuse to treat this possibility as a relevant alternative simply on the grounds that many will not actively support (and may even resist) the required measures because these are costly for them, is to imply that a great many things are noncomprehensively justified; and this, in effect, merely ensures that noncomprehensive justification carries little or no recommendation. In short, the fewer possibilities we allow to be relevant, the less is the moral significance of saying of some alternative that it is noncomprehensively justified.

I now turn to the question whether, from a deontological point of view, racial profiling can be noncomprehensively justified. I shall focus on fairness. The considerations about available alternatives that were discussed in the preceding section are also relevant here, but to clarify whether racial profiling conflicts with a concern for fairness I need to shift the focus a little. Much here hangs on whether or not it would be justified for African Americans to refuse to implement racial profiling. The mere fact that it is morally impermissible for European Americans to render the factual premise of the net-benefit argument true does not show that African Americans are noncomprehensively justified in rejecting racial profiling where European Americans make the factual premise true. If somebody impermissibly and credibly threatens that unless I donate $5 to Oxfam he will shoot someone at random, then what he does is, obviously, impermissible. Yet, it may nevertheless be impermissible for me not to donate the $5 (regardless of whether it was before the threat). What is relevant in this

[48] Risse and Zeckhauser (2004), "Racial Profiling," p. 161.

example is the fact that the costs for me of compliance with the threat are slight and the consequences for a third party of my noncompliance are very bad. If Risse and Zeckhauser's argument is right, the acceptance of racial profiling has no net-costs for African Americans and its rejection has very bad consequences for third parties, i.e. for people who are the victims of serious crimes that would otherwise have been prevented. This suggests that African Americans are not justified in expressing anger at the terms of collaboration offered by European Americans by refusing racial profiling useful in reducing serious crime. It seems to show that, given (unjustified) discrimination and racial inequality, racial profiling is justified in some cases at least.

However, deontologists are likely to point to two facts with a direct bearing on the issues here. First, *ex hypothesi*, European American victims of crimes that would have been prevented had racial profiling been employed contribute to the discriminatory practices and racial inequality that leads to higher crime-rates among African Americans; and for that reason, they are partly and indirectly causally responsible for the crime from which they suffer. Just how heavily this consideration weighs may depend on how blameworthy European Americans are in failing to eliminate discrimination and racial inequality, on the extent to which crimes that are preventable through racial profiling are causally independent of the existence of discrimination and racial inequality, and on the seriousness of the types of crime racial profiling is being employed to prevent. Presumably, by contributing causally to discriminatory (or inequality-inducing) practices, one does not automatically render it permissible for one's discriminatees to refuse to bear any cost, however small, if bearing that cost will prevent the violation of one's right to life and liberty. In situations involving the self-defense of innocent individuals against blameworthy attackers, the case for holding that the innocent are morally required, in defending themselves, to bear some costs to avoid disproportionate harm to their attackers is, if anything, weaker than the case for saying that a discriminator is permitted to force a discriminatee to bear any costs involved in avoiding harm to the discriminator (and other discriminatees) as a result of aggression on part of the discriminatee where that aggression is partly a result of the discrimination. While the innocent defender can say that he is in no way responsible for the fact that he has to bear certain costs in choosing his means of self-defense in order to avoid disproportionate harm to the blameworthy aggressor, the aggressor-discriminatee cannot say that he is in no way responsible for the fact that he has to bear certain costs in order to avoid disproportionate harm to the discriminator.[49] Yet, even in the former case, such costs may have to be born. If I try to spit in

[49] I assume here that we are talking about cases of disproportionate harm to the discriminator. I am not assuming that there is no such thing as proportionate harm to the discriminator that it would be impermissible for the discriminator to protect him-or herself against. Also, the analogy is misleading in that racial profiling will affect not only discriminatees that commit crime, but

someone's face after having set his house on fire, I am partly responsible for his attempt to shoot me to prevent me from spitting in his face. However, this does not render it impermissible for me to take measures to prevent him from shooting me, including the use of physical force. More generally, we do not take the fact that A is partly to blame for B's aggression against him, e.g. because A was himself aggressive, to render it impermissible for A to take measures that harm B in the course of protecting himself against B's aggression.[50]

Second, African Americans who become victims of crimes that would not have occurred under racial profiling might themselves have rejected racial profiling. They may therefore be said to have consented, in an indirect sense, to their increased risk of exposure to crimes that racial profiling would have prevented. (Of course, this is not to say that African Americans have consented to being put in the kind of circumstances where racial profiling is an effective means of reducing crime.) Hence, while the absence of racial profiling increases the risk that one's rights will be seriously violated, everyone is either responsible for this fact or has consented to exposure to the increased risk. This suggests that if African Americans successfully resist racial profiling, no one will be in a position to complain that his rights have been violated as a result of the police refraining from using racial profiling techniques to curb crime, or that he suffers from unfair treatment as a result of this.[51] From a deontological point of view, then, given that all African Americans oppose racial profiling, it need not be true that a society with racial profiling is less unjustified than one without such profiling. For if all African Americans oppose racial profiling, no one can complain that his rights have been violated or that he has been treated unfairly as a result. If African Americans agree to racial profiling, it may be justified. But this stance will not be less unjustified than one in which African Americans successfully resist racial profiling. On the deontological view deployed here, one cannot simply set out the case for racial profiling and then present it to the potential objects of such profiling. The case for racial profiling depends on whether or not they agree to it.

---

also the great majority of discriminatees who do not. To accommodate this disanalogy, one could imagine instead a group of innocent people only a small minority of whom are threatened by aggression from blameworthy attackers. Presumably, all of these innocent people, even those who are not threatened with aggression, are morally required to bear some costs to avoid disproportionate harm to the blameworthy aggressors.

[50] This is not to say that it has no moral implications that the defender is partly responsible for the aggression against which she was justified in defending herself, e.g., she may be morally required to share some of the medical costs of making the wounded aggressor recover. I owe this point to Shlomi Segall.

[51] As noted in chapters 6 and 8, the fact that no one is in a position to complain does not entail absence of injustice or moral wrongfulness.

I now turn to the second question raised at the start of this section: Is racial profiling justified from a noncomprehensive point of view, if we amend one or more of (i)–(iv)? To continue with the line of argument pursued in the last section, let us relax assumption (ii) and instead assume, plausibly, that some, but not all, European American victims of crimes preventable by racial profiling refuse to countenance discrimination and racial inequality—perhaps some of them actively oppose them—and, thus, play no part in making the factual premise of the net-benefit argument true. Moreover, let us assume that some, but not all, African American victims of crimes preventable through racial profiling favor the implementation of racial profiling. Here we have a case in which, if racial profiling is successfully resisted, some people impose avoidable risks of harm on people who are neither responsible for these risks nor consented to by them. Of course, this will not ensure that racial profiling is noncomprehensively unjustified, but presumably from a deontological point of view the greater the proportion of these people, and the more stringent the rights, violation of which is more likely in the absence of racial profiling, the stronger becomes the deontological case for saying that it is noncomprehensively unjustified not to implement racial profiling. I am not endorsing this implication, but simply saying that it seems to follow from a deontological point of view given the stated assumption.

Let us now alter (ii) in a different way. Suppose that European Americans are not responsible for the fact that racial profiling is an efficient means of curbing crime. One might now argue that it is impermissible not to implement racial profiling. To see this, suppose racial profiling of non-Swiss passengers is used at security checks at Swiss airports to prevent hijackings. The point here is not whether this is morally justified all things considered. It is that we would think differently of the net-benefit argument when presented by Swiss people to the passengers who are subjected to racial profiling. Even if the Swiss are required to justify their actions to air travelers, they can say that the fact, let us suppose, that racial profiling will reduce the risks of hijackings has nothing to do with them. It is independent of their will—a situation they must cope with in the same way as one has to cope with bad weather one way or another.

8. Challenges

I have suggested that there are striking parallels between the incentives argument and the net-benefit argument for racial profiling; and that, accordingly, Cohen's critique of the former can be transferred to Risse and Zeckhauser's argument in favor of racial profiling. However, this view may be resisted for two reasons. First, in the incentives argument the effect of the choices of talented people does not come about through the choices of nontalented people. If

talented people choose to work less hard without incentives, nontalented people will be worse off whatever they do. However, the effects of the discriminatory and inequality-inducing or inequality-sustaining choices made by European Americans run partly through the choices of African Americans, including the choices of some African Americans to commit crimes. Hence, it can be argued that African Americans might have chosen otherwise even given discrimination and racial inequality, in which case the factual premise of the net-benefit argument would not have been true. So neither all African Americans nor all European Americans can regard it as a fact independent of their will that the factual premise of the net-benefit argument is true; and insofar as African Americans simply disregard their contribution to making the factual premise true they show that they do not see themselves and European Americans as members of the same justificatory community.

In response to this objection—and without raising fundamental doubts about free will and responsibility in a way that would be inappropriate in the present context—it must be conceded that it may well be the case that neither European Americans who discriminate or refrain from eliminating inequality, nor African Americans who commit crimes or refrain from fighting crime, can regard the truth of the factual premise of the net-benefit argument as true independently of their will.[52] However, this has no real bearing on the conclusions above. It is not as if European Americans become able to justify their making (together with African Americans) the factual premise true because the factual premise would be false if some African Americans were to act otherwise. Generally, one cannot justify one's making something bad happen simply by pointing to the fact that had others acted otherwise, it would not have happened.

Second, in the case of the incentives argument the issue is how talented people will respond to a certain policy proposal, whether incentives are offered or not; but this is not the issue here. Whether racial profiling is adopted will make, or is being assumed to make, no difference to the European Americans' behavior. What we are evaluating is not a piece of their conduct, which is conditional upon which policy is adopted. Rather, it is a piece of their behavior that has an impact on the effects of a certain policy. European Americans are not warning that they will countenance discrimination and racial inequality if racial profiling is not implemented, but talented people are warning that they will be less productive if they are not offered incentives.

[52] For this reason, people who commit crimes (or, perhaps, a disproportionately large number of crimes) and thereby help to make it true that the profiled group to which they belong has a higher crime rate than other groups cannot reasonably object that profiling of his group rests on the disrespectful belief that it commits a disproportionately large number of crimes, see Applbaum (1996), "Response: Racial Generalization," p. 153.

Again it must be conceded that there is this difference between the arguments. However, again I would want to deny that it makes any moral difference, i.e., that we should now evaluate the arguments differently. What matters here is that the factual premise of the proposed argument is true because European Americans refrain from making it false and they cannot justify their not doing so. More generally, the mere fact that I have reasons to continue to conduct myself in a certain way independently of how others choose to conduct themselves, where my avoidable way of acting partly determines the outcome of their choices, does not imply that I can take my way of acting for granted when offering these people reasons for acting one way rather than another, let alone that my conduct falls outside the scope of whatever normative principles that I appeal to when so doing. In saying this, I am not suggesting that what Bernard Williams refers to as a "strong doctrine of negative responsibility" is part of common-sense morality—or, more to the point, deontology.[53] I can allow that there is a sense in which "each of us is specially responsible for what he does, rather than for what other people do"[54] as long as this is not taken to imply rejection of a *weak* doctrine of negative responsibility, i.e. not taken to imply that we are *not at all* responsible for how our actions influence what other people do, even when what we do to them is unjust and makes the kind of conduct that some of them end up being involved in much more attractive than it would otherwise be.

9. Conclusion

Many people sometimes express themselves as if they find statistical and in particular racial discrimination per se morally objectionable. Yet, it is easily seen that people do not generally object to group profiling as such. For instance, most endorse gender profiling in police work based on the assumption that most perpetrators of violent crime are male. Similarly, few would be disturbed by racial profiling techniques used to prevent or investigate racist hate crimes against African Americans. A plausible explanation of these facts is that, for good reason, people do not think of European Americans and males (European American as well as African American) as being either unjustly worse off or more likely to commit violent crime or racist hate crimes against African Americans because African Americans and women (European- as well as African American) have made them unjustly worse off or subjected them to

[53] J. J. C. Smart and Bernard Williams (1973), *Utilitarianism: For and Against* (Cambridge, Cambridge University Press), pp. 93–100.

[54] Smart and Williams (1973), *Utilitarianism*, p. 99.

discriminatory treatment.[55] However, the case of African Americans is different and is believed by many to be different. One aim of this chapter has been to explore the way in which this difference shapes discussion of the morality of racial profiling.

Often cases of statistical discrimination are morally objectionable, not because they constitute cases of statistical discrimination, but for other reasons contingently related to statistical discrimination, e.g., biased uses of statistics or because of the use of unreasonably coarse-grained statistics (see chapter 3). When we set aside these objectionable features contingently related to statistical discrimination, it might be said that statistical discrimination is intrinsically wrong because it violates a claim to be treated as individuals, because it necessarily involves unequal treatment, or because it unfairly treats people on the basis of what others do and not on the basis of what they themselves do. Against these claims, I argued that treating people on the basis of their perceived individual properties is a mere means of treating them on the basis of the individual properties they have in fact; that the unequal treatment, admittedly necessarily involved in statistical discrimination, need not amount to morally problematic treatment as unequals; and that it is not in general unfair to have to bear uncompensated burdens even though this imposition is not based on facts about what one oneself did or based upon facts about what others did.

Finally, statistical facts are often facts about how we choose to act. Since we can morally evaluate how we choose to act, we cannot simply take statistical facts for granted when justifying policies: we need to ask the prior question of whether it can be justified that we make these statistical facts obtain. Special problems arise in cases involving self-fulfilling expectations where the practice of statistical discrimination causes the very same statistical facts which are then appealed to in order to justify the practice of statistical discrimination and in cases where the statistical facts on which the practice of statistical discrimination is based are caused in an unjustified way by those who benefit from the practice of statistical discrimination. Specifically, I have argued that racial

[55] I concede that there are other explanations of why racial profiling is morally problematic: e.g., that African Americans are disadvantaged; or that people should not be treated on the basis membership of racial groups (but see section 2, this chapter); or that racial profiling is incompatible with a community in which everyone has equal status. Note, however, in relation to the first of these, that racial profiling is, according to Risse and Zeckhauser's argument, likely to be a net-benefit to African Americans and in the view of many people this supposed fact, when considered on its own and, thus, in isolation from the sort of considerations adduced here, strengthens rather than weakens the case for racial profiling. The second explanation implies that affirmative action and racial profiling are, in one respect, equally morally problematic. On most views, this suggests that racial profiling is morally problematic in other respects as well. The account provided here might be seen as one way of spelling out the requirement of the ideal of a community of equals. A society in which some do not see the need to offer comprehensive justifications for their policy proposals is a society that fails to live up to this ideal.

profiling is not comprehensively justified. In the process, I have relied on assumptions (i)–(iv), but these, I suggest, have some plausibility and would be regarded sympathetically by many observers. Risse and Zeckhauser have not shown that racial profiling is noncomprehensively justified from a utilitarian point of view. But it may be noncomprehensively unjustified not to implement racial profiling from a deontological point of view if the type of crime that racial profiling is being used to prevent is sufficiently grave; a sufficient number of people who are not subjected to racial profiling and harmed by its absence cannot be said to have brought it about that racial profiling is an effective means of fighting crime; and a sufficient number of people belonging to the racial group subjected to racial profiling wants these measures to be taken. If these conditions are not all satisfied, then it is noncomprehensively unjustified from a deontological point of view to implement racial profiling.

In sum, we cannot dismiss racial profiling as morally wrong simply because it constitutes statistical discrimination. That a certain practice constitutes a case of statistical discrimination may indicate the presence of some of the moral problems described above, but to establish that these problems really render a certain practice of differential treatment morally impermissible we would need to look at the details in each case.

This page intentionally left blank

{ BIBLIOGRAPHY }

Adler, M. (2000) "Expressive Theories of Law: A Skeptical Overview." *University of Pennsylvania Law Review* 148.1: 1363–1501.

Alexander, Larry (1992) "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies." *University of Pennsylvania Law Review* 141.1: 149–219.

Altman, A. (2011) "Discrimination." In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, URL = http://plato.stanford.edu/archives/spr2011/entries/discrimination/ (Stanford, CA: The Metaphysics Research Lab).

Anderson, Benedict (1983) *Imagined Communities: Reflections on the Origin and Spread of Nationalism* (London: Verso).

Anderson, Elizabeth S. (1999) "What Is the Point of Equality?" *Ethics* 109.2: 287–337.

——— (2010) *The Imperative of Integration* (Princeton, NJ: Princeton University Press).

Anderson, Elizabeth S., and R. H. Pildes (2000) "Expressive Theories of Law: A General Restatement." *University of Pennsylvania Law Review* 148.5: 1503–1575.

Anonymous (2001) "The Rhetoric of Difference and the Legitimacy of Capital Punishment." *Harvard Law Review* 114.5: 1599–1622.

Arneson, Richard J. (1990) "Liberalism, Distributive Subjectivism, and Equal Opportunity for Welfare." *Philosophy & Public Affairs* 19.2: 158–194.

——— (1999) "Against Rawlsian Equality of Opportunity." *Philosophical Studies* 93.1: 77–93.

——— (1999) "Egalitarianism and Responsibility." *Journal of Ethics* 3: 225–247.

——— (2002) "Equality of Opportunity." In E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, http://plato.stanford.edu/entries/equal-opportunity/ (Stanford, CA: The Metaphysics Research Lab).

——— (2006) "What Is Wrongful Discrimination?" *San Diego Law Review* 43.4: 775–807.

Arthur, Isak Applbaum (1996), "Response: Racial Generalization, Police Discretion, and Bayesian Contractualism." In John Kleinig (ed.), *Handled With Discretion* (Lanham, MD: Rowman and Littlefield), pp. 145–158.

Baber, Harriet E. (2001) "Gender Conscious." *Journal of Applied Philosophy* 18.1: 53–63.

Bagenstos, Samuel R. (2003) " 'Rational Discrimination,' Accommodation, and the Politics of (Disability) Civil Rights." *Virginia Law Review* 89.5: 825–923.

Barry, Brian (2001) *Culture and Equality* (Cambridge, UK: Polity Press).

——— (2005) *Why Social Justice Matters* (Cambridge, UK: Polity Press).

Blum, Lawrence (2002) "*I'm Not a Racist, But . . .* " (Ithaca, NY: Cornell University Press).

Bowers, W. J., and G. L. Pierce (1980) "Arbitrariness and Discrimination under Post-Furman Capital Statutes." *Crime & Delinquency* 26: 563–632.

Boxill, Bernard (1991) "Wilson on the Truly Disadvantaged." *Ethics* 101: 579–592.

———(2001) *Race and Racism* (Oxford: Oxford University Press).

Brest, Paul (1976) "In Defense of the Antidiscrimination Principle." *Harvard Law Review* 90.1: 1–54.

Broome, John (1990) "Fairness." *Proceedings of the Aristotelian Society* 91: 87–101.

Broome, John (1994) "Fairness versus Doing the Most Good." *The Hastings Center Report* 24: 36–39.

Cahn, Steven M. (ed.) (1995) *The Affirmative Action Debate* (New York: Routledge).

Cassell, P. (2004) "In Defense of the Death Penalty." In H. Bedau and P. Cassell (eds.), *Debating the Death Penalty: Should America Have Capital Punishment?* (Oxford: Oxford University Press), pp. 183–217.

Cavanagh, M. (2002) *Against Equality of Opportunity* (Oxford: Clarendon Press).

Cholbi, Michael (2006) "Race, Capital Punishment, and the Cost of Murder." *Philosophical Studies* 127: 255–8.

Churchill, Ward (2003) "'Some People Push Back': On the Justice of Roosting Chicken." http://www.kersplebedeb.com/mystuff/s11/churchill.html (accessed June 5, 2013).

Cohen, G. A. (1992) "Incentives, Inequality and Community." In G. Petersen (ed.), *The Tanner Lectures on Human Values* 13 (Salt Lake City: University of Utah Press).

——— (1997) "Where the Action Is: On the Site of Distributive Justice." *Philosophy and Public Affairs* 26.1: 3–30.

——— (2002) *If You're an Egalitarian, How Come You're So Rich?* 34 (Cambridge, MA: Harvard University Press).

——— (2004) "Expensive Taste Rides Again." In Justine Burley (ed.), *Dworkin and His Critics* (Oxford: Blackwell Publishing), pp. 3–29.

——— (2006) "Casting the First Stone: Who Can, and Who Can't, Condemn the Terrorists?." *Royal Institute of Philosophy Supplement* 58: 113–136.

——— (2008) *Rescuing Justice and Equality* (Cambridge, MA: Harvard University Press).

——— (2011) *On the Currency of Egalitarian Justice and Others Essays in Political Philosophy* (Princeton, NJ: Princeton University Press).

Collins, Hugh (2003) "Discrimination, Equality, and Social Inclusion." *The Modern Law Review* 66.1: 16–43.

Corlett, J. Angelo (1993) "Racism and Affirmative Action." *Journal of Social Philosophy* 24.1: 163–175.

Cox, Chana Berniker (1993) "On Levin's "Responses to Race Differences in Crime." *Journal of Social Philosophy* 24.1: 155–162.

Cummisky, David (1990) "Kantian Consequentialism." *Ethics* 100.3: 586–615.

Daniels, Norman (1978) "Merit and Meritocracy." *Philosophy & Public Affairs* 7: 206–223.

Darwall, Stephen (1977) "Two Kinds of Respect." *Ethics* 88.1: 36–49.

Devine, P. E. (2000) "Capital Punishment and the Sanctity of Life." *Midwestern Studies in Philosophy* 24: 229–43.

Doyle, Oran (2007) "Direct Discrimination, Indirect Discrimination and Autonomy." *Oxford Journal of Legal Studies* 27.3: 537–553.

Dummett, Michael (2004) "The Nature of Racism." In Levine and Pataki (eds.), *Racism in Mind* (Ithaca, NY: Cornell University Press), pp. 27–34.

Dworkin, Ronald (1985) *A Matter of Principle* (Cambridge, MA: Harvard University-Press).

——— (2000) *Sovereign Virtue* (Cambridge, MA: Harvard University Press).

——— (2003) "Bakke's Case: Are Quotas Unfair?" In B. Boxill (ed.), *Race and Racism* (Oxford: Oxford University Press), pp. 297–306.

Edmonds, David (2006) *Caste Wars: A Philosophy of Discrimination* (London: Routledge).

Eidelsson, Ben (2011) *What Is Discrimination and When Is It Wrong?* D.Phil. thesis, Oxford University.

Eisgruber, Christopher L., and Lawrence G. Sager (1994) "The Vulnerability of Conscience: The Constitutional Basis for Protecting Religious Conduct." *University of Chicago Law Review* 61.4: 1254–1315.

Epstein, Richard A. (1995) *Forbidden Grounds: The Case against Employment Discrimination Laws* (Cambridge, MA Harvard University Press).

Eyal, Nir (forthcoming) "Levelling Down Health", in Nir Eyal, Ole F. Norheim, Samia A. Hurst, and Daniel Wikler (eds.), *Inequalities in Health: Ethics, Measurement, and Policy* (New York: Oxford University Press).

Ezorsky, Gertrude (1991) *Racism and Justice: The Case for Affirmative Action* (Ithaca, NY: Cornell University Press).

——— (1992) "Discrimination." In L. C. Becker (ed.), *Encyclopedia of Ethics* vol. 1 (New York: Garland Publishing): 264–267.

Feinberg, Joel (1973) *Social Philosophy* (Englewood Cliffs, NJ: Prentice Hall).

——— (1974) "Noncomparative Justice." *Philosophical Review* 83.3: 297–338.

——— (1983) *Justice, Equal Opportunity, and the Family* (New Haven, CT: Yale University Press).

——— (1984) *Harm to Others* (Oxford: Oxford University Press).

Fiss, O. M. (1976) "Groups and the Equal Protection Clause." *Philosophy & Public Affairs* 5.2: 107–77.

Fleurbaey, M. (2008) *Fairness, Responsibility, and Welfare* (Oxford: Oxford University Press).

Flew, Anthony (1981) *The Politics of Procrustes* (London: Temple Smith).

——— (1990), "Three Concepts of Racism." *Encounter* 75: 63–66.

Frankfurt, Harry G. (1971), "Freedom of the Will and the Concept of a Person." *Journal of Philosophy* 68, 5–20; reprinted in Harry Frankfurt (1988), *The Importance of What We Care About* (Cambridge: Cambridge University Press), 11–25.

——— (1999) "Equality and Respect." in Harry Frankfurt (1999) *Necessity, Volition, and Love* (Cambridge: Cambridge University Press), 146–154.

——— (1999) *Necessity, Volition, and Love* (Cambridge: Cambridge University Press).

Fraser, Nancy (1994) "After the Family Wage: Gender Equity and the Welfare State." *Political Theory* 22: 591–618.

——— (1997) *Justice Interruptus* (New York: Routledge).

Fredman, Sandra (2011) *Discrimination Law* 2nd ed. (Oxford: Oxford University Press).

Fried, Charles (1978) *Right and Wrong* (Cambridge, MA: Harvard University Press).

Fullinwider, Robert (1997) "The Life and Death of Racial Preferences." *Philosophical Studies* 85: 163–180.

Gaertner, Samuel, and Leonard Bickman (1971) "Effects of Race on the Elicitation of Helping Behavior: The Wrong Number Technique." *Journal of Personality and Social Psychology* 20.2: 218–222.

Garcia, Jorge L. A. (2001) "The Heart of Racism." In Bernard Boxill (ed.), *Race and Racism* (Oxford: Oxford University Press), 256–296.

Gardner, John (1989) "Liberals and Unlawful Discrimination." *Oxford Journal of Legal Studies* 9.1: 1–22.

——— (1996) "Discrimination as Injustice." *Oxford Journal of Legal Studies* 16.3.1: 353–367.

——— (1998) "On the Grounds of Her Sex(uality)." *Oxford Journal of Legal Studies* 18: 167–187.

Glasgow, Joshua (2009) "Racism is Disrespect." *Ethics* 120.1: 64–93.

Glazer, Nathan (1975) *Affirmative Action* (New York: Basic Books).

Goffman, Erving (1963) *Stigma: Notes on the Management of Spoiled Identity* (London: Penguin Books).

Goldman, Alan H. (1979) *Justice and Reverse Discrimination* (Princeton, NJ: Princeton University).

Greenberg, J. (1986) "Against the American System of Capital Punishment." *Harvard Law Review* 99.7: 1670–80.

Halldenius, Lena (2005) "Dissecting 'Discrimination.'" *Cambridge Quarterly of Healthcare Ethics* 14.4: 455–463.

Hampton, J. (1988) "Forgiveness, Resentment and Hatred." In J. G. Murphy and J. Hampton (eds.), *Forgiveness and Mercy* (New York: Cambridge University Press).

Hardin, Russell (1995) *One for All: The Logic of Group Conflict* (Princeton, NJ: Princeton University Press).

Hare, Richard (1952) *The Language of Morality* (Oxford: Clarendon Press).

——— (1963) *Freedom and Reason* (Oxford: Clarendon Press).

——— (1981) *Moral Thinking: Its Levels, Methods, and Point* (Oxford: Clarendon Press).

——— (1989) *Essays in Ethical Theory* (Oxford: Clarendon Press).

Harris, John (1987) "QALYfying the Value of Life." *Journal of Medical Ethics* 13.3: 117–123.

Heinrich, B. (2007) "What is Discrimination and When is it Morally Wrong?" *Jahrbuch für Wissenschaft und Ethik* 12: 97–114.

Hellman, Deborah (2003) "What Makes Genetic Discrimination Exceptional?" *American Journal of Law and Medicine* 29: 77–116.

——— (2008) *When Is Discrimination Wrong?* (Cambridge, MA: Harvard University Press).

Hill, Thomas E. Jr., (1995) "The Message of Affirmative Action." In Steven M. Cahn (ed.), *The Affirmative Action Debate* (New York: Routledge), pp. 169–191.

Holmes, E. (2005) "Anti-Discrimination Rights without Equality." *Modern Law Review* 68.2: 175–94.

Holtug, Nils (2002) "The Harm Principle." *Ethical Theory and Moral Practice* 5: 357–389.

Hooker, Brad (2000) *Ideal Code, Real World* (Oxford: Clarendon Press).

Horta, Oscar (2010) "Discrimination in Terms of Moral Exclusion." *Theoria* 76.4: 314–332.

Jackson, Frank (1987) "Group Morality." In Philip Pettit et al. (eds.), *Metaphysics and Morality* (Oxford: Blackwell Publishers).

Kagan, Shelly (1988) "The Additive Fallacy." *Ethics* 99.1: 5–31.

——— (1989) *The Limits of Morality* (Oxford: Clarendon Press).

——— (1998) *Normative Ethics* (Boulder, CO: Westview Press).

Kamm, Frances M. (2007) *Intricate Ethics* (Oxford: Oxford University Press).

Kant, Immanuel (2002[1785]) *Groundwork for the Metaphysics of Morals* (New Haven, CT: Yale University Press).

Kant, Immanuel (2006[1798]) "Metaphysics of Morals." In P. Kleingeld (ed.), *Toward Perpetual Peace and Other Writings Politics, Peace, and History* (New Haven, CT: Yale University Press).

Kappel, Klemens, and Peter Sandøe (1992) "QALYs, Age and Fairness." *Bioethics* 6.4: 297–316.

Kelman, Mark (1991) "Concepts of Discrimination in "General Ability" in Job Testing." *Harvard Law Review* 104: 1157–1247.

Kennedy, R. L. (1988) "*McClesky V. Kemp*: Race, Capital Punishment, and the Supreme Court." *Harvard Law Review* 101: 1388–1443.

Kennedy, R. (1994) "The State, Criminal Law, and Racial Discrimination: A Comment." *Harvard Law Review* 107: 1255–1278.

Kirschenman, Joleen, and Kathryn Neckerman (1991) "'We'd Love to Hire Them, But . . . ': The Meaning of Race for Employers." In C. Jencks and P. Peterson (eds.), *The Urban Underclass* (Washington, DC: Brookings Institution Press), 203–234.

Koppelman, Andrew (2006) "Is It Fair to Give Religion Special Treatment?" *University of Illinois Law Review* 3: 571–604.

Kymlicka, Will (2002) *Contemporary Political Philosophy* (Oxford: Oxford University Press).

Lever, Annabelle (2005) "Why Racial Profiling is Hard to Justify." *Philosophy and Public Affairs* 33.1: 94–110.

Levin, Michael (1992) "Responses to Race Differences in Crime." *Journal of Social Philosophy* 23: 5–29.

Lewis, David (1989) "Dispositional Theories of Value." *Proceedings of the Aristotelian Society* suppl. vol. 63: 113–137.

Lippert-Rasmussen, Kasper (2006) "Telic versus Deontic Egalitarianism." In Nils Holtug and Kasper Lippert-Rasmussen (eds.), *Egalitarianism: New Essays on the Nature and Value of Equality* (Oxford: Oxford University Press), pp. 101–124.

——— (2006) "The Badness of Discrimination." *Ethical Theory and Moral Practice* 9.2: 167–185.

——— (2007) "Private Discrimination: A Prioritarian, Desert-Accommodating Account." *San Diego Law Review* 43.4: 817–856.

——— (2009) "Reaction Qualifications Revisited." *Social Theory and Practice* 35.3: 413–439.

——— (2010) "Scanlon of the Doctrine of the Double Effect." *Social Theory and Practice* 36.4: 541–564.

——— (2011) "'We Are All Different': Statistical Discrimination and the Right to be Treated as an Individual." *Journal of Ethics* 15: 47–59.

——— (2013) "Who Can I Blame?." In Michael Kühler and Nadja Jelinek (eds.), *Autonomy and the Self* (Dordrecht: Springer), pp. 295–315.

——— (2013) "Discrimination." In Hugh LaFolette (ed.), *International Encyclopedia of Ethics* (Oxford: Wiley-Blackwell).

List, Christian, and Philip Pettit (2011) *Group Agency: The Possibility, Design, and Status of Corporate Agents* (Oxford: Oxford University Press).

Loury, Glenn C. (2002) *The Anatomy of Racial Inequality* (Cambridge, MA: Harvard University Press).

Lucas, J. R., (1985) "Discrimination and Irrelevance." *Proceedings of the Aristotelian Society* 86: 307–324.

MacIntyre, Alasdair (1984), "Is Patriotism a Virtue?" Lindley Lecture (Lawrence, KS).

Mackinnon, Catherine (1987) *Feminism Unmodified* (Cambridge, MA: Harvard University Press).

Mason, Andrew (2000) "Equality, Personal Responsibility, and Gender Socialisation." *Proceedings of the Aristotelian Society* 100: 227–246.

——— (2001) "Equality of Opportunity." *Ethics* 111.4: 760–781.

——— (2006) *Levelling the Playing Field* (Oxford: Oxford University Press).

McCrudden, C. (1982) "Institutional Discrimination." *Oxford Journal of Legal Studies* 2.3: 303–367.

McDermott, D. (2001) "A Retributivist Argument against Capital Punishment." *Journal of Social Philosophy* 32.3: 317–33.

McMahan, Jeff (2002) *The Ethics of Killing: Problems at the Margins of Life* (Oxford: Oxford University Press).

McMahan, Jeff (2009) "Intention, Permissibility, Terrorism, and War." *Philosophical Perspectives* 23.1: 345–72.

Miller, David (1992) "Deserving Jobs." *Philosophical Quarterly* 42: 161–181.

——— (1998) "Equality and Justice." In Andrew Mason (ed.), *Ideals of Equality* (Oxford: Blackwell Publishers), pp. 21–36.

——— (1999) *Principles of Social Justice* (Cambridge, MA: Harvard University Press).

Moreau, Sophia (2010) "What Is Discrimination?" *Philosophy & Public Affairs* 38.2: 143–179.

Morris, Andrew (1995) "On the Normative Foundations of Indirect Discrimination Law: Understanding the Competing Models of Discrimination Law as Aristotelian Forms of Justice." *Oxford Journal of Legal Studies* 15.2: 199–228.

Myers Jr., S. L. (1993) "Racial Disparities in Sentencing: Can Sentencing Reforms Reduce Discrimination in Punishment?" *University of Colorado Law Review* 64: 781–808.

Nagel, Thomas (1997) "Justice and Nature." *Oxford Journal of Legal Studies* 17.2: 303–321.

Narveson, Jan (1993) *Moral Matters* (Peterborough, Ontario: Broadview Press).

——— (2001) *The Libertarian Idea* (Peterborough: Broadview Press).

——— (2002) *Respecting Persons in Theory and Practice* (New York: Rowman and Littlefield).

Nathanson, S. (1985) "Does It Matter if the Death Penalty is Arbitrarily Administered?" *Philosophy & Public Affairs* 14.2: 149–64.

Nozick, Robert (1974) *Anarchy, State, and Utopia* (New York: Basic Books).

Parfit, Derek (1984) *Reasons and Persons* (Oxford: Oxford University Press).

——— (1998) "Equality and Priority." In Andrew Mason (ed.), *Ideals of Equality* (Oxford: Blackwell Publishers), pp. 1–20.

Phillips, Anne (1995) *The Politics of Presence* (Oxford: Clarendon Press).

——— (2004) "Defending Equality of Outcome." *Journal of Political Philosophy* 12: 1–19.

Pittman, J. P. (1997) "Punishment and Race." *Utilitas*, 9.1: 115–130.

Pogge, Thomas (2000) "On the Site of Distributive Justice: Reflections on Cohen and Murphy." *Philosophy & Public Affairs* 29.2: 137–169.

Pogge, Thomas (2008) *World Poverty and Human Rights*, 2nd ed. (Cambridge: Polity Press).

Radcliffe Richards, Janet (1980) *The Sceptical Feminist* (Harmondsworth: Penguin).

Radelet, M. L., and M. J. Borg (2000) "The Changing Nature of Death Penalty Debates." *Annual Review of Sociology* 26: 43–61.

Rawls, John (1971) *A Theory of Justice* (Oxford: Oxford University Press).

——— (2000) *A Theory of Justice*, rev. ed. (Cambridge, MA: Harvard University Press).

Reiman, J. H. (1985) "Justice, Civilization, and the Death Penalty." *Philosophy & Public Affairs*, 14.2: 115–148.

Reiman, J., and L. P. Pojman (1998) *The Death Penalty: For and Against* (Lanham, MD: Rowman & Littlefield Publishers).

Risse, Mathias (2007) "Racial Profiling: A Reply to Two Critics." *Criminal Justice Ethics* 26.4: 4–19.

Risse, Mathias, and Richard Zeckhauser (2004) "Racial Profiling." *Philosophy and Public Affairs* 32.2: 131–170.

Ryder, Richard D. (1975) *Victims of Science: The Use of Animals in Research* (London: Davis-Poynter).

Sabbagh, Daniel (2007) *Equality and Transparency: A Strategic Perspective on Affirmative Action in American Law* (New York: Palgrave Macmillan).

Scanlon, Thomas C. (1982) "Contractualism and Utilitarianism." In A. Sen and B. Williams (eds.), *Utilitarianism and Beyond* (Cambridge: Cambridge University Press), 103–128.

——— (1999) *What We Owe Each Other* (Cambridge, MA: Harvard University Press).

——— (2000) "Intention and Permissibility." *Proceedings of Aristotelian Society* Suppl. Vol. 74.1: 301–317.

——— (2008) *Moral Dimensions: Permissibility, Meaning, and Blame* (Cambridge, MA: Belknap Press of Harvard University Press).

Schauer, F. (2003) *Profiles, Probabilities, and Stereotypes* (Cambridge, MA: Harvard University Press).

Scheffler, Samuel (1982) *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions* (Oxford: Clarendon Press).

——— (2003) "What Is Egalitarianism?" *Philosophy and Public Affairs* 31.1: 5–39.

——— (2005) "Choice, Circumstance, and the Value of Equality." *Politics, Philosophy, and Economics* 4.4: 5–28.

Segall, Shlomi (2012) "Should the Best Qualified Be Appointed?" *Journal of Moral Philosophy* 9: 31–54.

——— (2012) "What's so Bad about Discrimination?" *Utilitas* 24.1: 82–100.

——— (forthcoming) *Equality and Opportunity* (Oxford: Oxford University Press).

Sher, George (1988) "Qualifications, Fairness, and Desert." In Norman E. Bowie (ed.), *Equal Opportunity* (Boulder, CO: Westview Press), 113–127.

Sher, George (1999) "Diversity." *Philosophy and Public Affairs* 28.2: 85–104.

Singer, Peter (1983) "Is Racial Discrimination Arbitrary?" In Jan Narveson (ed.), *Moral Matters* (New York: Oxford University Press).

——— (1993) *Practical Ethics* (Cambridge: Cambridge University Press).

Smart, J. J. C., and Bernard Williams (1973) *Utilitarianism: For and against* (Cambridge: Cambridge University Press).

Smilansky, Saul (2006) "The Paradox of Moral Complaint." *Utilitas* 18: 284–290.

——— (2007) "The Paradox of Moral Complaint." In Saul Smilansky, *Ten Moral Paradoxes* (Oxford: Blackwell Publishing), pp. 90–99.

Stevenson, B. (2004) "Close to Death: Reflections on Race and Capital Punishment in America." In Bedau and Cassell (eds.), *Debating the Death Penalty: Should America Have Capital Punishment?* (Oxford: Oxford University Press), 76–116.

Sunstein, Cass (1991) "Why Markets Won't Stop Discrimination." *Social Philosophy and Policy* 21.2: 22–37.

——— (1994) "The Anticaste Principle." *Michigan Law Review* 92.8: 2410–55.

Sverdlik, Steven (2011) *Motive and Rightness* (Oxford: Oxford University Press).

Taylor, Paul W. (1973) "Reverse Discrimination and Compensatory Justice." *Analysis* 33: 177–82.

Temkin, Larry (1993) *Inequality* (Oxford: Clarendon Press).

Thomas, Laurence (1992), "Statistical Badness." *Journal of Social Philosophy* 23.1: 30–41.

Thomsen, F. K. (unpublished, on file with author) "Discrimination and Disrespect."

Thomson, Judith Jarvis (1986) *Rights, Restitution, and Risk* (Cambridge, MA: Harvard University Press).

——— (1991) "Self-Defense." *Philosophy & Public Affairs* 20.4: 283–310.

——— (1999) "Physician-Assisted Suicide: Two Moral Arguments." *Ethics* 109.3: 497–518.

Tutu, Desmond (1999) *No Future without Forgiveness* (London: Rider).

van den Haag, E. (1978) "The Collapse of the Case against Capital Punishment." *National Review* 31, pp. 395–407.

——— (1985) "Refuting Reiman and Nathanson." *Philosophy & Public Affairs* 14.2: 165–76.

——— (1986) "The Ultimate Punishment: A Defense." *Harvard Law Review* 99.7: 1662–69.

van Parijs, Philippe (1991) "Why Surfers Should be Fed: The Liberal Case for an Unconditional Basic Income." *Philosophy & Public Affairs* 20.2: 101–131.

Walen, Alec (2006) "The Doctrine of Illicit Intentions." *Philosophy & Public Affairs* 34.1: 39–67.

Walzer, Michael (1983) *Spheres of Justice* (Oxford: Basil Blackwell).

Wang L. ing-Chi (1988) "Meritocracy and Diversity in Higher Education: Discrimination against Asian Americans in the Post-*Bakke* Era." *The Urban Review* 20.3: 189–209.

Wasserman, David (1991) "Morality of Statistical Proof and the Risk of Mistaken Liability." *Cardozo Law Review* 13: 935–976.

——— (1996) "Racial Generalizations and Police Discretion." In John Kleinig (ed.), *Handled with Discretion: Ethical Issues in Police Making Decisions* (Lanham, MD: Rowman & Littlefield), pp. 115–130.

——— (1998) "Discrimination, Concept Of." In R. Chadwick (ed.), *Encyclopedia of Ethics* (San Diego, CA: Academic Press), pp. 805–814.

Wasserstrom, R. (1977) "Racism, Sexism, and Preferential Treatment: An Approach to the Topics." *University of California Law Review* 24: 581–615.

Wertheimer, Alan (1983) "Jobs, Qualifications, and Preferences." *Ethics* 94.1: 99–112.

Williams, Andrew (1998) "Incentives, Inequality, and Publicity." *Philosophy & Public Affairs* 27.3: 225–247.

Wittgenstein, Ludwig (1953) *Philosophical Investigations* (Oxford: Basil Blackwell).

Woodruff, Paul (1995) "What's Wrong with Discrimination?" In Steven M. Cahn (ed.), *The Affirmative Action Debate* (London: Routledge), 39–42.

Young, Iris Marion (1990) *Justice and the Politics of Difference* (Princeton, NJ: Princeton University Press).

——— (2001) "Equality of Whom? Social Groups and Judgements of Injustice." *Journal of Political Philosophy* 9: 1–18.

Yuracko, Kimberly A. (2006) "Sameness, Subordination, and Perfectionism." *San Diego Law Review* 43: 857–897.

{ INDEX }

action, affirmative (see also reverse discrimination), 4, 8, 24n30, 27n41, 30, 39n64, 41n66, 43, 49–50, 52, 54, 69, 80, 159, 168, 175–176, 186, 208, 268, 294, 300n55
activism, political, 5–6, 51
ad hominem, 290–291
ageism (*see also* age discrimination), 3, 160–161
agency, respect for, 257–260
Alexander, Larry, 4, 7, 16n6, 20n20, 34–35, 60, 87n16, 103, 111, 113–128, 134, 149, 155, 167, 173n48, 241n17, 259n51, 268n12
Altman, Andrew, 7, 54–76
ambition-sensitivity, 196–202
American Bar Association, 217
Americans,
  African, 21n23, 54, 62, 71–72, 93, 107, 148, 169, 181–182, 186n71, 189, 197n7, 199–200, 203, 218, 221, 285, 286–300
  Asian, 72n47, 201–202
  European, 71–72, 148, 199, 203, 285–299
Amish, 199
Anderson, Elizabeth, 5n17, 42–43, 93n29, 96n36, 112, 179n58, 195n6, 226n24, 263n8, 275n8
animals, non-human, 19–20, 31n49, 48n79, 113, 119, 124
animosity (*see* hatred), 27–28, 46, 83, 220
anti-semitism, 32n53, 158–159
apartheid, 1, 33, 137, 171n44, 262n4
approach, bottom-up vs. top-down, 269–270
arbitrariness (*see also* irrelevance discrimination), 24n30, 226
argument, Wilt Chamberlain, 184
Arneson, Richard J., 108–109, 113, 128n62, 149–150, 166, 169n35, 175n52, 200n16, 278–279
atheists, 1, 34n58, 262
attitude,
  antimeritocratic, 245–249, 251, 253, 257n44, 260
  based solely on merit, 246
  nonindividualized, 245–248, 257n44
  propositional, 113, 122, 125
autonomy, 48, 200, 262

Baber, Harriet B., 30n48
Baldus, David, 217
Barry, Brian, 23n28, 194n4, 198n10
baseline,
  ideal account of the relevant, 158n9
  issue of, 64–65, 81, 156–161, 167, 176, 196, 204
  moralized, 64–65, 161
  no-discrimination account of the relevant, 158–160
  straightforward account of the relevant, 157–160, 167
  well-informed, unbiased agent baseline, 65
Beauvoir, Simone de, 263
belief (*see also* statistical information, statistical knowledge),
  statistical background, 83–84, 95
  statistical foreground, 84–85, 95
  statistical vs. non-statistical, 80, 82–85, 95, 273
bias, 17n11, 28, 36, 41–43, 54–55, 59–61, 65, 67–69, 72, 74, 76–77, 80, 82, 87–88, 91, 115, 117–118, 122, 124–126, 145, 149n57, 179, 217–218, 220, 222, 225–226, 228–231, 236, 243, 246, 249–250, 254, 256, 266, 273–274, 300
Bickman, Leonard, 254n40
blacks, 13, 21, 26n35, 38, 43, 79n1, 89, 13–138, 148–149, 217, 220, 223–224, 254, 256–257, 262n4, 286n34
Blum, Lawrence, 21n23, 28n42, 138n23, 171n43
Boxill, Bernard, 38n63, 44n75, 275n9, 286, 287n35
Brahmin, 147
Brest, Paul, 134n14, 169n33
British Transport Police, 272
Broome, John, 176n55
Buddhists, 91, 208–209
business necessity, 41n66

Catholics, 27–28, 61, 77, 187, 207–208, 211–212, 266
Cavanagh, Matt, 25–26, 28n44, 70n39, 106, 113, 170–171, 263n7, 278n14, 281n22

challenge,
the definition, 16, 46, 84
the disadvantage specification, 61
the discriminatee identification, 57, 88n17, 95
the indirect discrimination definition, 55
the limitation, 70, 74, 83, 95
the statistical discrimination definition, 80, 95
the trilemma, 87
character, moral, 30n48, 33, 106n5, 116, 121, 123–124, 127, 160, 163, 173
childbirth (*see also* parental leave), 37, 82
choice (*see also* choice account of the wrongness of discrimination), 38, 106–108, 187, 198–201, 245n30, 247, 249, 263, 267, 280–282, 288, 297–299
Cholbi, Michael, 218, 229n37
Christians, 25–26, 262–263
Churchill, Ward, 133
classification, 19, 31, 34, 47, 55, 62, 132, 265
Cohen, G. A., 19n15, 32n53, 90n19, 108n13, 232n45, 280n21, 288, 290, 293n46, 297
cohesion, social, 186n71, 195n6
Collins, Hugh, 1n4, 18n12, 110–111
color-blindness, 30, 37, 285
community, 175, 182, 242, 269, 290, 293, 298, 300n55
compensation, 98n37, 147, 162, 173, 178, 195–196, 204, 209, 255, 281–282, 300
competition, fair, 235
complaint,
individual treatment, 280–281
moral, 17, 46, 48, 52–53, 71, 84, 126, 135–136, 145, 159, 180, 216, 227n31, 231–233, 250–251, 275–278, 280–301
compliance (*see also* ideal vs. nonideal theory), 8, 64–65, 213–214, 249, 253, 266, 295
condition,
the causal, 71, 179, 180
the disadvantage, 56, 61–66, 75, 77
the disproportionateness, 56, 62n25, 65–69, 74–75, 76–77
the expanded no-intention, 69, 179
the individualized causal, 182
the no-intention, 56, 59–61, 75–76
the revised disadvantage, 64, 68, 179–180
the revised no-intention, 61, 68, 69
the social salience condition, 56–57, 75–76, 81, 179
consequentialism, 8, 114, 153–154, 163, 167, 175–176, 183
contempt, 70n41, 113, 119, 171n41
control, 93n27, 105, 107–108, 240–241, 252, 258, 279, 288
Copts, 113
Corlett, J. Angelo, 90n19, 274–275
crime,
hate, 138, 168n28, 276, 299
rate, 79, 220–22, 272–274, 285–289, 292, 298n52

Dalit, 147
Daniels, Norman, 236n4-5
Darwall, Stephen, 117n40
demeaningness, 3, 7, 98, 99, 128, 130–139, 150–151, 154–155, 165, 168, 176, 218
deprivation, 285
desert, moral, 8–9, 35–36, 48n80, 103n1, 105, 110n18, 121, 124, 128, 153–154, 165–168, 170–179, 181, 183, 185–186, 257n46, 264–265, 268–270, 283, 292–293
desire, brute, 42, 115, 150
difference principle, 288
dignity, 138
disadvantage,
comparative, 74
historical, 26n35
intrapersonal vs. interpersonal, 16–17
local vs. global, 62, 179–180
socioeconomic, 283
systematic, 40, 76–77, 148, 169, 237, 240, 242
discrimination,
against babies, 19, 138
against smokers, 1, 13, 107
against vs. in favor of vs. between, 14, 221
age, 1, 25, 110n19, 117
agentless, 18–19
asymmetric, 21
benevolent, 184
choice account of the wrongness of, 106–108
cognitive vs. non-cognitive, 40–43, 45, 90, 113, 115, 244, 245
comparative falsehood account of the wrongness of, 121–122
concept of, 3–4, 6–8, 13–99, 169n32, 193
criteria of vs. indicators of, 13, 221–224
definition of, 13–53
desert-prioritarian account of the wrongness of, 8, 9, 103n1, 105, 118n18, 128, 153–154, 165–186, 264–265, 269–270, 283n28, 292, 293n47
direct vs. indirect discrimination, 18n12, 36–41, 43, 50, 54–78, 83, 87–89, 95n34, 114, 131, 155, 176, 219–220
equality of opportunity accounts of the wrongness of, 109–110
essentially comparative, 16–18, 22, 81, 184
falsehood account of the wrongness of, 113–121
first- vs. second-order, 44–45
formal, 174–175

generic sense of, 14–22, 29
genetic, 33, 49, 81, 95–99
harm-based account of the wrongness of, 48, 110n18, 127–130, 134, 139, 153–189, 241n17
hierarchical, valuational vs. non-hierarchical, valuational, 44–45
in an extensional vs. intensional sense, 20, 58n12, 73
in housing, 9, 106, 108, 261
in love, 3, 26–28, 149, 185, 267
in sentencing, 217–234
in the private sphere, 261–271
income and wealth, 38–40, 46, 48, 58n15, 62, 72n46
indicator (*see* criteria for vs. indicator of discrimination)
indirect (*see* direct vs. indirect discrimination)
institutional, 76–78
irrational comparative falsehood account of the wrongness of, 121–127
irrelevance, 22–25, 30n47, 45, 91n23,105n3
irrelevance-based account of the wrongness of, 105–106
job market (*see* labor market)
labor market, 40, 52, 62, 108, 139, 142n39, 144, 146, 149, 201, 207, 219, 221–222, 235–261
legal, 22n25, 161, 174, 218, 219–222, 225
loci of, 219–221
mental-state based account of the wrongness of, 61, 103–129, 131, 164–165, 176, 183
meritocratic account of the wrongness of, 108, 171n44
moralized concept of, 24–26, 29n46, 45,183, 268n13,
morally not unjustified, 29–30, 87–89, 177–185
morally unjustified vs. unjust discrimination, 29–30
narrow concept of, 49, 53
neutralizing the effects of, 193–216
non-statistical, 81–87
objective meaning account of the wrongness of, 112n25, 114n34, 128–155, 165, 168, 176n54, 218
objectless, 18–19
paradigm cases of, 1, 2, 36, 111, 117, 163, 181
pluralistic account of the wrongness of, 183
proxy, 55, 82–83, 86, 105
racial, 2–3, 27, 43–44, 60, 77, 87–89, 107, 110–111, 117, 139–151, 166, 173n47, 176, 183, 235, 262–263, 267, 285n34, 287–288, 294, 299
racist (*see* racial discrimination)
real-life, 5–6
reflexive, 21–22
religious, 1, 27, 97, 164, 182, 207, 270
respect-based account of the wrongness of, 114–116, 119–120, 170n39, 251n37
reverse (*see also* affirmative action), 4, 29n46, 42n69, 148, 159, 162, 168, 176
self-, 6
sex, 1, 16, 20–21, 36, 46, 52, 58, 70, 82, 105, 138, 188–189, 211, 236n6, 262, 264
skeptic, 49–53
social exclusion account of the wrongness of, 110–110
statistical, 79–99, 114–115, 117, 131, 155, 176, 272–299
structural, 19, 41–42, 45, 77–78, 213n33
subjects of, 19, 38n62
subjective-meaning accounts of the wrongness of, 129
unintentional, 18n12, 60
unlawful, 187–188, 269
unsuccessful, 19n16
valuation- vs. non-valuation-based, 43–46, 49, 115, 118, 149
wrongness of, 103–189
wrongness of indirect discrimination, 177–183
disproportionateness, 56, 62n25, 65–69, 74–77
disrespect (see respect)
distribution, just, 70, 161, 194, 238, 240–241, 243, 254–255, 293n46
doctrine of double effect, 112, 124n55
dominance, 41n66, 49–53
double standards, 121, 273
Dummett, Michael, 60
Dworkin, Ronald, 5, 198n10, 200, 275n9, 278

egalitarianism, 6, 108n13, 184, 292
deontic, 184–185
telic, 74, 165
efficiency, 109, 168, 175, 195n6, 278, 297
Eisgruber, Christopher G., 262
endowment-insensitivity, 198
Epstein, Richard A., 4, 109n16
equality,
gender, 52, 58
opportunity of, 2, 26–27, 109–110, 169n35, 180n61, 198n10, 200, 210, 235–236, 241, 243, 255, 278n16
Equality and Human Rights Commission, 79
equals, moral standing as, 28, 93–94, 98, 226, 278–279, 300
equivalence, moral, 4
eugenics, 98
European Commission, 74–75

European Court of Human Rights, 54
European Union, 219
evidence, statistical, 81n6, 84–88, 90–92, 275
exclusion, 15, 39, 50, 54n4, 59, 71, 76, 93, 110–111, 143, 148, 194n2
explanation, structural, 194n4, 199
Ezorsky, Gertrude, 76, 200n14
Eyal, Nir, 269n15

facts, statistical, 84, 92, 176, 274, 283–285, 300
fairness, 25–26, 70n41, 140, 151, 163–164, 176–177, 183, 200n16, 217, 226–229, 231, 239, 255–258, 260, 279–285, 291–294, 296, 300
Feinberg, Joel, 225–226, 279n17
Flew, Anthony, 23
Frankfurt, Harry, 116
Fredman, Sandra, 18n14, 54n4, 202n20
freedom, deliberative, 179n60, 185–189
Fullinwider, Robert, 49

Gardner, John, 76
gender, 1, 6n18, 21, 23n26, 26n34, 28, 36, 43–44, 52, 55, 58, 74–75, 82, 105, 140, 185–188, 193–194, 200, 219, 222n14
generalization, probabilistic, 82
Glasgow, Joshua, 40n65, 116–117
Glazer, Nathan, 200
Goffman, Erving, 32n51
Griggs v. Duke Power Co., 41n66, 54, 71, 75
group,
  discrimination, 2n7, 26–47, 51–52, 56, 62n24, 72–74, 82, 103–105, 178, 183, 262n3, 268n13
  ethnic, 2, 33, 74, 76, 127, 140, 172, 193, 200, 219, 261, 264, 267, 269
  socially salient, 3, 26–36, 46–48, 56–58, 64, 68–70, 72–73, 76–77, 80–81, 84, 86, 88, 90, 92–93, 95–97, 104, 106, 109, 123, 125, 168–170, 174, 178–181, 194n5, 206, 224, 275–277
groups,
  inequality between, 67, 69, 179
  inequality within, 63
  relevantly contrasting, 57

Halldenius, Lena, 24, 29, 76–77
Hampton, Jean, 130
Hardin, Russell, 261n1
Hare, Richard, 29, 158n8, 160n13, 163n19, 173n49
harm
  dimensions of, 157n5, 160–162
  metric of (*see* dimensions of harm)
  symbolic, 181, 268

hatred (*see also* animosity), 18, 70n41, 123, 138n24, 164, 244
Head Start, 294
Hellman, Deborah 4, 98–99, 129–139, 149–151, 155, 218
heterosexuals, 21, 25–28, 108, 198
hierarchy, 134, 138, 139
hiring,
  preferential (*see also* affirmative action, reverse discrimination, positive discrimination), 49
  word-of-mouth, 38
history, 62, 71n45, 98, 132–133, 138–139, 205–207, 211
Holocaust, 104, 159
Holtug, Nils, 155n2, 157n6
homophobia, 31–32, 263
homosexuals, 25–28, 108, 120, 262–263
Hooker, Brad, 158n8
humiliation, 28, 39–40, 92–93, 104, 185, 241n17, 274
hypothetical cases, 4, 6, 108, 285

immigrants, 20, 43–44, 57, 62, 67, 71–72, 88n17, 109, 137, 202, 212
incentive argument, 288
inclusion (*see also* exclusion), 195n6, 215, 268
indicator, statistical (*see also* proxy discrimination), 86, 88–89, 91–92, 279
inequality, 2, 6, 52, 70, 110, 127, 131, 152, 167, 180–181, 185, 194, 241, 285–292, 294–295, 297–298
inference,
  deductive, 84
  inductive, 83, 272
information
  genetic, 33, 95–99
  statistical (*see also* statistical background vs. foreground belief, statistical knowledge), 80, 82, 84, 91–95, 272–277
insult, 132, 138, 141–145, 147, 149–151
insurance, 1, 33, 52, 80–81, 88–89, 95–99
intention,
  discriminatory, 55, 140–141, 145
  predictive significance of, 141
intuitions, 110n20, 127, 154, 168n28, 174, 179, 250, 255, 270
irrationality (*see* rationality)
irrelevance (*see* relevance; *see also* irrelevance discrimination)
isolationism, 28
issue, the baseline (*see* baseline)
IVF, 206

Jews, 17n9, 31, 117, 122, 125–126
  ultraorthodox, 117n41, 199n13
Johnson, Lyndon B., 235
justice,
  comparative, 226, 230, 234, 279n17
  distributive, 70n39, 194, 238, 240–241, 243, 254, 255, 293n46
  egalitarian, 5
  fundamental principles of, 5
  noncomparative, 226, 230
  Rawlsian principles of, 180n61, 288
  transitional, 291
justification, comprehensive vs. noncomprehensive, 290–294, 300n55

Kagan, Shelly, 144n45, 177n40
Kahnemann, Daniel, 85
Kamm, Frances M., 60
Kant, Immanuel, 112, 114, 119, 121, 124, 232
Kennedy, Randall, 224n19
King, Rodney, 224
knowledge, statistical (*see also* statistical background vs. foreground belief, statistical information), 82–83
Koppelman, Andrew, 264n9
Ku Klux Klan, 77

law,
  antidiscrimination, 110n21, 185, 264n9
  enforcement, 91n22, 267
leave, parental, 41, 64, 82–86, 277
lesbians, 21, 33, 108
libertarianism, 215n38
List, Christian, 77n58
lottery,
  fair, 140
  genetic, 33
Lucas, J. R., 15
luck, brute, 96–98

MacKinnon, Catharine A., 49–53
Malaysia, 110
Mason, Andrew, 200, 237, 239, 243–245, 247, 249, 252, 254–255, 257–260
McClesky, Warren, 217–218, 222–224
McDermott, Daniel, 228n32
McMahan, Jeff, 2n8
meaning,
  conventional, 132, 136
  descriptive 29n46, 47–48
  of an action, 129–130, 141
meritocracy, 109, 235–241, 246, 249, 251, 254, 257–260
Miller, David, 108, 236n6, 239, 245, 247, 249, 252, 275–277
Monthy Python, 272
Moreau, Sophia, 179n60, 185–189
Morgenbesser, Sidney, 227–228
Muslims, 17, 26, 44, 113, 276,

Narveson, Jan, 267n11
Nathanson, Stephen, 223n18, 226–231
nationality, 32, 175–176, 186n72, 193, 198, 204
Nazism, 18, 31, 98, 104, 117, 122, 126, 157
negligence, 130
nepotism, 23, 25, 189, 235–236
norm,
  basic meritocratic, 235–236, 249–251, 253–254, 257, 260
  end-state, 237n10, 251–252
  non-collaboration, 251, 253–254
  recipient, 250, 252, 254–255, 262
  virtuocratic, 260
Nozick, Robert, 184, 215

obesity, 13, 105
objection,
  deliberative perspective, 112
  deontological, 173
  discriminatee-focused, 172
  discrimination-tolerant, 174
  intrinsic wrongness, 175
  leveling down, 165, 184–185
  nondistinctiveness, 171
  unfairness, 176, 280–281
observational evidence, 82, 85–86
operation, discriminatory in, 38–39, 41n66
opportunity costs, 186
oppression, 49–53, 138n24, 161, 194
ostracism, 251n1

Parfit, Derek, 5n15, 170n37
paternalism, 117, 293
patriarch, 146,147
people,
  disabled, 33, 36, 39, 105, 106n6, 110n19, 176, 199n13, 219
  elderly, 33, 129
  possible, 19
permissibility, moral, 68n33, 90, 108, 112, 114, 119, 155, 158n9, 173–175, 251n36, 265–266, 271, 276–278
permission, legal, 268–270
Pettit, Philip, 77n58
Phillips, Anne, 194
philosophy,
  analytic, 4
  political, 5, 51
pilots, airline, 186n71, 245n28, 276–277, 282

Piper, Adrian, 44n75
Pittman, John P., 221
Pogge, Thomas, 70n39, 78n59, 194
police, 22–23, 79–80, 88, 91–94, 220, 224, 227n31, 233, 238, 266, 272–277, 279–280, 283, 286n34, 288, 296, 299
power, 21n23, 132, 138, 144, 197, 203
preference, legitimate, 240–256
prejudice, 21n23, 67, 87, 113, 126–127, 140, 142–145, 148, 150n57, 224, 230n39, 239, 243–244, 247, 249, 253–254
principle,
  basic moral, 136, 171, 292
  end-state, 215
  of equal moral worth, 132, 135–136
  of fair equality of opportunity, 180n61, 200, 278n16
  responsibility, 281n22
  the means, 112, 124
prioritarianism,
  desert-accommodating, 8, 9, 48m80, 103n1, 105, 110n18, 128, 153–154, 165–186, 264–265, 268–270, 283n28, 292
privacy, 264
privilege, 148, 168, 211, 236, 264n9
profiling,
  group identity, 272–273
  racial, 3, 10, 79–80, 88–89, 272–300
  statistical, 79
Protestants, 27–28, 182, 266, 281
proxy (*see also* proxy discrimination), 55, 82–83, 86, 105, 277
punishment, capital, 217–234

qualifications,
  moralized account of, 243
  reaction, 24, 28n43, 36, 43, 106, 143, 168, 235–261
  technical, 237, 241, 252

race, 3, 17, 19, 23, 26n24, 28,n43, 31, 39–40, 43–44, 46, 55–56, 70n39, 76, 89, 91n22, 99, 106–108, 116, 122, 138, 141, 148–149, 176, 186, 188, 193, 218, 223, 235, 240, 250, 255n42, 261, 267, 273–274, 277, 282, 286n34, 291n42
racism, 17, 21n23, 23, 36, 40, 43, 60, 77, 96, 116–117, 133, 138, 143, 149, 172, 176, 239–240, 247, 249–250, 253, 269n15, 283, 286n34, 287–301
rationality, 106n6, 119–120, 123, 160n13, 173n48
Rawls, John, 8, 200, 233, 235, 288, 290n40
reciprocity, functioning, 283–284, 290, 292–293
regulation, legal, 22n25, 54, 75n51, 264, 269
Reiman, Jeffrey H., 221n9, 232
religion, 13, 23, 26n34, 28, 31, 39, 42, 55–56, 74, 97, 108, 164, 184, 193, 206–208, 219, 261
representation,
  equal, 197
  over- vs. under, 198, 201
  proportional, 193–216
  suitable, 196–199, 202, 204–216
resources, 5, 17, 19, 28, 37, 74, 79, 91–92, 118, 132, 161, 167, 176
respect, 2–3, 16n7, 60–61, 98, 113–129, 132, 138, 150, 154, 157, 169–170, 244, 251n37, 257–260, 275, 298n52
responsibility, 136, 177n56, 193, 252n38, 281n22, 298–299
rights,
  of association, 240–241, 248n32
  to be treated as an individual, 275–278
Risse, Mathias, 283–301
Ryder, Richard D., 20

Sabbagh, Daniel, 39n64, 80n4
Sager, Lawrence G., 262n2, 264n9
salience, social, 30–36, 48, 56, 57, 75–76, 81, 95, 179
Scanlon, Thomas, 23n28, 112, 139–152
Schauer, Frederick, 80, 82–85, 90–91, 105, 282
Scheffler, Samuel, 68n33
secularism, 65, 199, 262–263
Segall, Shlomi, 109–110
segregation, 27n41, 263–264
self-conception, 32, 242
self-consciousness, 120, 124
self-defense, 295–296
self-esteem, 19
self-respect, 34, 150, 197
sexism, 9, 17, 20, 55, 117–118, 123–124, 129, 133, 147n51, 172, 174–175, 187–188, 200, 205–206, 237n10, 240–241, 247, 252, 263
sexuality, 32, 108, 193, 263
Shanaghan v. U.K., 54
Sher, George, 257–259
Singer, Peter, 19–20, 24n30, 107, 275
Slogan, the, 184–185
Smilansky, Saul, 232
society, as opposed to state, 292
speciecism, 19–20
status,
  legal, 9, 48n80, 217, 264, 269–270, 290n41
  moral, 15–16, 19, 98, 113, 128, 134, 137–138, 150–151, 154, 168, 221n9, 269–270
  social, 23, 132, 161
Steel Workers v. Weber, 27n41
stereotype, 21, 31, 35, 42, 80, 274

stigma, 40, 47–48, 92–94, 96–97, 99, 134, 148, 168–169, 242
Sunstein, Cass, 109n16

tastes, expensive, 5
Taylor, P. W., 162n16
Temkin, Larry, 161, 184
term, secondarily/primarily evaluative, 29
terrorism, 92, 105, 112, 272, 274, 276
test, interpersonal, 285–293
testimony, 86, 150
theory,
  ideal vs. non-ideal, 8, 266
Thomas, Laurence, 91n24
Thomson, Judith Jarvis, 112
thought experiment, 29, 150, 177, 223
tolerance, 34n58
trait, extraneous, 185–188
transsexuals, 14
treatment,
  as equals, 278–279
  idiosyncratic differential, 25
  unequal, 18, 36n61, 133, 152, 217, 278–279, 300
Tversky, Amos, 85

unfairness (*see* fairness)
US Supreme Court, 54, 75–76, 217, 222, 262
utilitarianism, 107n9, 153, 167, 283–284, 292–294, 301

values, impersonal, 66n31
van den Haag, Ernest, 223n18, 226–229, 332
veil of ignorance, 109
view of reaction qualifications, symmetric, 245–249, 257n44, 259–260
view of suitable representation,
  constant groups and people, 206–206
  constant people, 206, 208–209
  counterfactual, 202–204, 214
  future-oriented, 209–213
  global, 205–208
  holistic, counterfactual, 203–204
  individualized, 212, 214–215
  partial compliance, 213–214
  simple, 196–203, 215, 221–222

Walzer, Michael, 180n62
Wasserman, David, 25, 70–71, 79–80
Wasserstrom, Richard, 31n49
welfare, 5, 27n39, 62, 153–154, 161–162, 164, 167, 200
well-being, mental-state accounts of vs. objective-list accounts of vs. preference-based accounts of, 161–162
Wertheimer, Alan, 237, 239–245, 247–248, 252
whites, 2, 13, 26, 28n43, 33, 79n1, 89, 138, 149, 217, 220, 223–224, 254, 256–258, 286n34
will, free, 298
Williams, Bernard, 299
Wittgenstein, Ludvig, 13n1
women, 1, 9, 13–15, 20–21, 23, 26n35, 28, 33–34, 37–38, 40–43, 50–52, 55, 57–62, 64–67, 70, 75–77, 80–84, 88–89, 91, 105, 109, 117–118, 120, 124, 129, 131, 138, 147, 161, 167, 174–175, 178, 183, 187–189, 194, 197, 199–201, 203–206, 210–211, 219, 221–222, 225, 236n6, 240–241, 246, 261n1, 263–264, 267, 272, 275, 293, 299
worth, moral, 114, 117–127, 132, 134–138, 176n54, 218
wrongness,
  intrinsic, 114, 122–127, 149n57
  subjective vs. objective, 158, 167

Young, Iris Marion, 32n52, 49–53, 159n11, 194, 213n33
Yuracko, Kimberley A., 238n10

Zeckhauser, Richard, 283–301