

Discrimination and Disrespect

OXFORD PHILOSOPHICAL MONOGRAPHS

Editorial Committee

WILLIAM CHILD R. S. CRISP A. W. MOORE
STEPHEN MULHALL CHRISTOPHER G. TIMPSON

OTHER TITLES IN THIS SERIES INCLUDE

Knowing Better: Virtue, Deliberation, and Normative Ethics

Daniel Star

Potentiality and Possibility: A Dispositional Account of Metaphysical Modality

Barbara Vetter

Moral Reason

Julia Markovits

Category Mistakes

Ofra Magidor

The Critical Imagination

James Grant

From Morality to Metaphysics: The Theistic Implications of our Ethical

Commitments

Angus Ritchie

Aquinas on Friendship

Daniel Schwartz

The Brute Within: Appetitive Desire in Plato and Aristotle

Hendrik Lorenz

Plato and Aristotle in Agreement? Platonists on Aristotle from Antiochus

to Porphyry

George E. Karamanolis

Of Liberty and Necessity: The Free Will Debate in Eighteenth-Century

British Philosophy

James A. Harris

Discrimination and Disrespect

Benjamin Eidelson

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

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

© Benjamin Eidelson 2015

The moral rights of the author have been asserted

First Edition published in 2015

Impression: 1

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 licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries 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

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

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number: 2015940444

ISBN 978-0-19-873287-7

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Contents

Introduction	1
Part I. Clearing the Ground	
1. The Concept of Discrimination	13
2. The Puzzle of Indirect Discrimination	39
Part II. Intrinsically Wrongful Discrimination	
3. The Role of Disrespect in Wrongful Discrimination	71
4. Respect for Equality	95
5. Respect for Individuality	127
Part III. Contingently Wrongful Discrimination	
6. A Case Study: Racial Profiling	173
Conclusion	223
Acknowledgments	227
<i>Endnotes</i>	229
<i>References</i>	253
<i>Index</i>	263

Introduction

Hardly anyone disputes that discrimination can be a grave moral wrong. Yet this consensus masks fundamental disagreements about what *makes* something discrimination, as well as precisely why (and hence when) acts of discrimination are wrong. This book aims to develop systematic answers to those two questions. It offers a philosophical account of what discrimination is, and a moral theory to explain what is characteristically wrong with it.

That project has two central motivations. The first is theoretical: philosophical reflection, I hope, can help us to attain a clearer understanding of our own moral convictions in this vexed area of social and political life. The second motivation, however, is intensely practical. To confront hard cases in a thoughtful way, we have to understand when and why discrimination is wrong in the first place.

Consider, for example, discrimination based on race or sex that isn't arbitrary or irrational, but rather is grounded in observable patterns and conducted with care and apparent respect. This is often called "statistical" discrimination. Is such differential treatment morally suspect? If so, is it problematic for the same reasons that "classic" race or sex discrimination is objectionable, or for different reasons? And what are these different families of reasons? How do the material effects of discrimination, the attitudes or judgments to which it gives effect, and the social context in which it takes place combine to determine its moral status?

"Statistical" discrimination poses a hard case because it incorporates some features we ordinarily associate with wrongful discrimination but eschews others. There are many other hybrid categories of

this kind. Take, as another example, discrimination on the basis of traits that are not relevant to a person's suitability for some benefit or position, but that also are not sites of widespread subordination or disadvantage. What should we make, for instance, of the British law that bars employment discrimination based on "philosophical beliefs"—including, according to some courts, a person's views about climate change or animal cruelty?¹ How does so-called lookism—targeted in some American jurisdictions by bans on "personal appearance discrimination"—compare to racism or sexism?² More broadly, how can we tell when claims of wrongful discrimination have tapped the moral root of our convictions about racism, sexism, and the like, and when they have misappropriated this normative heritage for fundamentally different ends?

We could ask similar questions about unwitting discrimination that gives effect to implicit biases, or about apparently neutral practices that give rise to disparate effects. Because the morality of discrimination is rarely thought through from first principles, the debates over all of these non-paradigm cases tend to devolve into contests among unmotivated analogies to (and distinctions from) what are, by acclamation, focal cases. To resolve these various predicaments, however, we need a *general* theory of discrimination—one that captures our convictions about paradigm cases, explains the normative grounds on which they rest, and thereby guides us in determining how far those grounds extend.

Laying such a foundation represents the most concrete "payoff" of this book. But the more fundamental goal of the book is not simply to offer practical prescriptions about some particular kinds of discrimination, but to help fill the underlying lacuna in our sense of our own ethical commitments. The fact is that many people who are certain that discrimination is a serious moral concern remain, if we're honest with ourselves, unsure just *what* we are concerned about. We know, of course, that it has to do with principles of equality, fairness, and justice. And, more specifically, we suspect it has to do with the fact that certain traits are beyond people's control, irrelevant to the decisions that are being made about them, or associated with

pervasive, structural disadvantage. We think the concept of wrongful discrimination might also be connected to the principle that people should be treated as individuals, rather than subjected to group stereotypes. And we recognize that charges of discrimination sometimes target malignant attitudes or intentions, while other times they express concern about particular patterns of outcomes. But, still, we have little sense of what these different considerations have to do with one another or how (and if) they combine to form a coherent whole. This book is an effort to make some philosophical progress from that basic starting point, and, although it leaves much unresolved, I hope it will be helpful to others who find themselves starting from a similar position.

* * *

Questions about the morality of discrimination have gone oddly neglected by moral philosophers thus far. Since the idea of wrongful discrimination is expected to do so much normative work in our social and political lives, it is something of a philosophical embarrassment that so little has been done to explore or account for it.

There is, of course, a vast literature on the value of equality, particularly as it concerns distributive justice. One might root a theory of wrongful discrimination in one of the more general views defended there. But that transplant is not straightforward, and it may not be adequate. It is far from clear, after all, that discrimination is wrong principally (let alone only) because of its distributive consequences. Much the same could be said of the literatures on values such as fairness and respect. Put another way, the central moral question about discrimination is precisely how and where these more basic moral considerations figure in. It is striking, therefore, that relatively little work, much of it quite recent, has tackled the conceptual and moral significance of discrimination as such.³

By contrast, foundational questions about discrimination are much more familiar to legal scholars. Indeed, there has been a renaissance of interest in philosophical questions about antidiscrimination law in the past few years.⁴ This book aims to contribute to that conversation,

but it attacks the problems from a different direction than legal theorists commonly do—from the bottom up, we might say, rather than the top down. My central question, in other words, is not when discrimination should be prohibited, but why it is sometimes morally wrong to begin with.

There is reason to doubt how far theories pitched at explaining or justifying institutional rules *about* discrimination can carry us in this more basic inquiry. Because laws shape behavior as well as judge it, the reasons for having certain laws are often rooted in the collective effects of compliance throughout a society. There is, for example, a popular and intuitive view that racial discrimination in employment ought to be prohibited because this prohibition contributes to equalizing opportunity and breaking down a persistent caste-like social system.⁵ It does not follow, however, that individual acts of employment discrimination must themselves be wrong, when they are, for the same reason.

Suppose, for instance, that a racist employer privately discriminates against a black candidate whom he knows will get a better job someplace else. Then, by hypothesis, the employer doesn't significantly constrict the candidate's opportunities, let alone entrench the second-class status of a whole group of people. If the act is wrong, as it certainly appears to be, it is for some other reason. More broadly, whereas antidiscrimination law has the power to open many opportunities to a person, and indeed to many people, individuals often confront the distinct choice whether to open or close one. Because those choice situations are so different, they will almost inevitably turn on different—although hardly unrelated—moral considerations. This disjuncture will be even more fundamental in domains of social life where antidiscrimination law is reluctant to tread at all.

While legal understandings of discrimination serve as a central foil throughout the book, therefore, I mean to leave the vital questions of what discriminatory conduct should be prohibited, and why, for another day. This distinction is all the more important because the law has a tendency to occupy a field of normative discourse and claim relevant concepts as its own, tacitly equating lawfulness with

moral permissibility. When the legal construction of a concept—discrimination, causation, promissory obligation, or whatever—is shaped by concerns specific to the project of systematic and coercive enforcement, we have to guard against the risk that these considerations will come to distort our thinking about the underlying moral concepts as well.⁶

* * *

My effort to address the central philosophical questions about discrimination can be broken down into three parts—one devoted to conceptual analysis, one to moral theory, and one to an extended case study in applied ethics. At the risk of spoiling any surprise, it may help to summarize the central arguments of each component in advance.

Part I of the book clears the ground and delimits the subject by developing a formal analysis of the concept of discrimination, understood in non-evaluative terms. In particular, Chapter 1 argues that discrimination can be understood as a form of action that is, in itself, morally neutral, and which carries no particular connection to socially salient groups. What distinguishes discrimination from mere differential treatment, I suggest, is its special explanatory connection to some differential ascription of a property to the discriminatees. I explore how this explanatory connection should be understood, and I offer an interpretation of practices that discriminate between people without discriminating *against* any of them.

Because of the rhetorical power of “discrimination,” there have been many efforts to extend the concept to embrace practices that seem on their face quite distinct from the core activity identified in Chapter 1. Hence the proliferation of prefixed notions of discrimination—*structural* discrimination, *institutional* discrimination, and so forth. Chapter 2 grapples with the most important of these proposed extensions, which is known in Europe as *indirect* discrimination and in the United States as discrimination based on “disparate impact.” What is the relationship between this activity and the central concept of discrimination elaborated in Chapter 1?

I argue that the connection is in fact a distant one, and that unifying these two phenomena therefore should not be a desideratum for a theory of discrimination (or of what makes it wrongful). Many cases of indirect discrimination are, upon inspection, simply ordinary, direct discrimination of a certain characteristic kind, and are best understood as such. Others are not reducible in this way, but these may not be best thought of as *discrimination* at all. Rather, prohibiting “indirect discrimination” in such cases may be a roundabout means of equalizing opportunity for members of disadvantaged social classes. Conceiving of failure to go along with this program as itself a kind of discrimination—albeit an “indirect” kind—does not bring out underlying similarities between the phenomena as much as it invites equivocation between them. Chapter 2 concludes by applying the conceptual apparatus of the first two chapters to show why the conflict that some have posited between prohibitions on direct and indirect discrimination is illusory.

With this conceptual groundwork in place, the remainder of the book turns squarely to the morality of discrimination. Part II argues that the normative root of our convictions about core cases of wrongful discrimination is the recognition that these acts manifest disrespect for the discriminatees as persons. This general idea—that wrongful discrimination is connected to moral disrespect—is a very familiar one. Indeed, drawing on Kant, Martin Luther King, Jr. wrote half a century ago that segregation “debases personality,” threatening to “depersonalize the potential person and desecrate what he is.”⁷ A disrespect-based account of wrongful discrimination thus begins from a potent moral intuition. The core task of the book is to develop that intuition into a theory—to explore how the general obligation to recognize certain features of other persons *qua* persons, such as the intrinsic value of their well-being or the character of their individual autonomy, can be employed both to explain and to critically assess our judgments about familiar and unfamiliar variants of discrimination. The first step in this project is to get clear on the *kind* of disrespect to which we should be appealing. So, in Chapter 3, I elaborate a working account of respect and disrespect for persons

as these notions bear on the morality of discrimination. On this view, the relevant sort of respect—recognition respect for someone’s standing as a person—imposes an affirmative obligation that one take to heart the various normative upshots of her being a person and regulate one’s action accordingly. What is of fundamental importance to determining whether an action is disrespectful, then, is the set of reasons for which an agent acts, rather than the effects or social meaning of her action. But her deliberation is marred not simply by the presence of some form of animus or a defamatory belief, as some have suggested, but by the *absence* of appropriate responsiveness to someone’s standing as a person, whether that absence comes about willfully or by neglect.

How then do we explain the important role of social conventions in fixing what is disrespectful? The answer lies in a distinction, also developed in Chapter 3, between *basic* and *conventional* disrespect. An act is basically disrespectful of X as a Y when it manifests a certain kind of deliberative failure—a failure to give appropriate weight to X’s standing as a Y in deciding how to act. An act is conventionally disrespectful, by contrast, when it is *understood* to be basically disrespectful by the lights of prevailing social meanings.

What is of basic moral importance, I argue, is basic disrespect. But social conventions matter a great deal because it is usually *basically* disrespectful to ignore them in view of the psychological and material harms that come with being *conventionally* disrespected. By taking this approach, I argue, we can reconcile three commitments: the “expressivist” insight that the moral status of an act, including an act of discrimination, often turns on its social meaning; the conviction that the wrongness of discrimination does not rest on its contingent effects alone; and the commonsense view that social conventions do not have the power to determine directly what is right and wrong.

Once we have settled on a view of the general conceptual structure of respect for persons, we have to ask just *what* we are obligated to recognize about other persons, and how this obligation casts light on the various forms of wrongful discrimination. Different sorts of

wrongful discrimination, I suggest, manifest a failure of recognition with respect to different attributes of personhood. Whereas discrimination on the basis of differential concern often fails to respect some people as beings of equal intrinsic value, discrimination on the basis of statistical generalizations often fails to respect them as autonomous individuals. The intuitive unity of our thought about these disparate cases—reflected, for instance, in the inclination to call each of these different sorts of discrimination at least potentially “racist” or “sexist”—is rooted in the unity of the personhood that, in different ways, the different practices demean or deny.

Each of these moral requirements must be specified carefully. In Chapter 4, I consider discrimination that manifests a failure of recognition for the equal *value* or *worth* of persons as such. I suggest that respect for a person’s equal value is best understood as grounding a defeasible presumption of equal consideration—akin to the Principle of Insufficient Reason, which purports to govern probability judgments in the absence of relevant evidence. This presumption can be overridden, without disrespect, by genuinely relevant differences between persons’ claims to one’s consideration. But it is respect for people’s equal value that sets a baseline of equal concern and requires a special reason for weighing some people’s needs more heavily than others’. One paradigmatic form of wrongful discrimination stems from a failure to respect this presumption of equal regard. Chapter 4 also introduces an important distinction between ordinary disrespect and *contempt*, understood as willful defiance of normative authority; this distinction helps to explain what is and is not morally similar about classic bigotry and implicit bias. Finally, having analyzed the demands of respect for people’s equal value in this way, Chapter 4 concludes by considering some forms of partiality that the requirement of recognition for a person’s equality may *not* reach—including discrimination based on preferences that govern the way one “spends” the weight properly afforded to one’s own interests, rather than how one weighs the interests of others.

Chapter 5 considers the complementary idea that people have a moral claim to be treated as individuals out of respect for their

autonomy. That notion is often criticized as incompatible with the widespread and mundane use of statistical evidence in a host of contexts. But in fact, I suggest, respecting people as individuals is not a matter of forswearing the use of any sort of evidence about them. What respect for a person's individuality requires is not that we *exclude* informative patterns about those who share some of her traits, but rather that we at least *include* certain kinds of information where it is reasonably available—information that manifests her own efforts, as an autonomous agent, to give shape to her life. In this and in other particulars, the disrespect theory connects our understanding of wrongful discrimination to our grasp of moral personhood more broadly—to the question what constraints a person's standing as a person imposes on the way we should relate to her.

The disrespect-based theory of Part II aspires to explain what is characteristically and intrinsically wrong about core cases of wrongful discrimination. But one of the significant features of the account is that it does not condemn all of the discrimination that typically elicits moral concern. In Part III, I consider at some length one example of the kind of discrimination that is not necessarily ruled out by the moral demands of respect for persons: racial profiling in law enforcement. Such profiling, I suggest, exemplifies the challenges that arise in assessing discrimination that exploits statistical relationships among socially salient characteristics in pursuit of valuable social ends. I argue that this practice, and others like it, may be troubling on grounds that are fundamentally contingent and hence very different in character. What makes racial profiling morally wrong may not be that it is distinctly unfair or that it is necessarily racist, but rather that it contributes to a set of *conventional* social understandings that do unjustified harm.

Despite its title, therefore, the book ultimately offers a pluralist view of what makes discrimination wrong. On the one hand, the moral core of wrongful discrimination—what earns it its apparent place among the cardinal sins of modern social life—is the failure of recognition that it characteristically shows for the equal personhood of the discriminatees. This core is made up of the practices rightly

described as *racist*, *sexist*, and the like. Different cases merit inclusion in this class for different reasons, reflecting the different morally significant facets of persons. On the other hand, discriminatory practices that do not manifest an attitude of disrespect for persons are not thereby excused from moral assessment. Many such practices are bad simply because the harm they do—including through the disrespect that they may foster or be felt to evince—is unjustified. I take racial profiling, under certain conditions, to pose an important and misunderstood paradigm of this sort of case.

PART I

Clearing the Ground

1

The Concept of Discrimination

What is discrimination? Relatively little philosophical attention has been devoted to the questions when and why discrimination is wrong. But philosophers have taken even less interest in this more basic question: what *is* it?

Perhaps this question has attracted little discussion because the answer simply seems obvious. We are all familiar with many cases of discrimination—in employment, in public accommodations, in private relationships—and the interesting question about these seems to be what makes at least some of them morally objectionable.

To a certain extent that is true: the main reason to be interested in the concept of discrimination, I think, is that discrimination is sometimes a serious moral wrong. Discrimination is probably not an especially important concept from the detached standpoint of philosophy of action, in other words. But that does not mean that we can or should skip right to the moral question. Rather, our interest in the morality of discrimination gives us a reason to get clear on what discrimination is in the first place—which may not itself be primarily a moral question.

It is also not as straightforward an inquiry as it may appear. To discriminate is a verb, something one *does*; but it involves a complex constellation of relations to one or more other individuals or groups, as well as to some manner of treatment and some dimension of difference. Each of these relations can be characterized in different ways. In this stage-setting chapter, I offer a working account of key

aspects of this conceptual structure, with eyes to both ordinary usage and theoretical utility.

Of course, in everyday conversation, arguments about what constitutes discrimination are often normatively charged. As David Wasserman says, “to call discrimination ‘wrongful’ is merely to add emphasis to a morally-laden term.”¹ One possibility, then, is that the question what discrimination is and the question what is wrong with it are really one and the same, for “discrimination” simply names a certain kind of wrong.²

I think we should resist that approach and instead take discrimination as a phenomenon with no built-in moral status. This too is consistent with important facets of ordinary usage. Many who support efforts at “positive discrimination,” for instance, call them by this name. More broadly, it seems to be widely accepted that one can intelligibly ask whether or not discrimination of a certain kind is wrong, which suggests that it is not always so by nature. (One cannot do the same with a moralized concept such as murder.) But the most important reason for proceeding this way is not that it better reflects ordinary usage—which is undeniably inconsistent on this point—but that it structures our inquiry in a more fruitful way.

Here is what I mean. There is plainly something common to (1) wrongful discrimination (whether we understand this phrase as redundant or not) and (2) what *would* be wrongful discrimination if only it were wrongful—that is, conduct that is just like wrongful discrimination, except that it lacks the former’s wrong-making features. Whichever construal of “discrimination” you favor, you cannot adequately understand the essential characteristics of discrimination without an account of the features that characterize (1) and (2) alike. For even if you think discrimination is necessarily wrong, you will want to know what properties it distinctively has besides those that make it wrong. Similarly, an account of murder presupposes an account of killing. So we need an analysis of the category that subsumes both and only the phenomena enumerated as (1) and (2). And easily the most natural name for that category, it seems, is

“discrimination,” since we can simply add the modifier “wrongful” to refer to (1) in particular.

Perhaps another contender for this role is the more general notion of “differential treatment.” But that is not a viable alternative if there are significant features that are common to the two classes I’ve described that also are *not* common to all cases of differential treatment. I think there are indeed such features, and by offering a conceptual analysis of discrimination—understood without a necessary moral status—I mean in part to detail what they are.

So understood, the concept of discrimination can be decomposed into two fundamental parts. First, to discriminate against someone, one must treat him differently in some dimension than others, and this differential treatment must be comparatively disadvantageous to him in some respect. To be sure, this schematic formulation leaves various issues unsettled. It says nothing about what interests a person has or what constitutes a form of treatment, for instance, although both of these will necessarily contribute to determining whether something constitutes discrimination according to this definition.

Without denying that there are interesting issues here, I want to leave this first component of the concept—the part to do with disadvantageous differential treatment—largely unanalyzed for now. It suffices to say that discriminating against (or in favor of) someone requires treating her worse (or better) than others in some respect. At the end of the chapter, I will consider certain aspects of this condition more closely.

My primary focus, however, will be on a second aspect of the concept. What is distinctive of discrimination, as compared to differential treatment generally, is that it is necessarily responsive to a difference among people along some dimension. Even the broad dictionary definition of discrimination as a “process by which two stimuli differing in some aspect are responded to differently” builds in this feature: one discriminates by *responding* differently to some difference in the features *of the stimuli*.³ For this reason, one can discriminate on the basis of religion, or on the basis of eye color, and

so on, but one cannot discriminate on no basis at all. To discriminate on no basis is simply to not discriminate. An analysis of discrimination must specify, then, what has to be true of the disadvantageous differential treatment described above in order for it to constitute discrimination on the basis of a given trait (and hence also for it to constitute discrimination at all).

I proceed as follows. In §1.1, I offer a general definition of discrimination on the basis of some trait, and I highlight some of the consequences of this analysis. In §1.2, I consider and reject a proposal to limit the ambit of discrimination to differentiation on the basis of socially salient traits. Finally, in §1.3, I consider what is required for differential treatment with the appropriate etiology to qualify as discrimination *against* or *in favor of* someone, and what we should make of discrimination between people that is not discrimination against or in favor of anyone. With these conceptual clarifications behind us, we will be equipped to turn to the perplexing notion of “indirect” discrimination in the next chapter, and then to the morality of discrimination generally.

1.1 A Definition

If you treat various people differently out of pure whimsicality, you engage in differential treatment but not discrimination. This is why “discriminating” taste is a virtue: it suggests not just that a person responds differently to some objects than to others, but that her variable responses are sensitive to differences in the objects.

The sense of “discrimination” that concerns us does differ from this older sense of the word in two respects, however. First, to say that someone has discriminating taste is normally also to say that she is sensitive to *relevant* differences in the objects, an implication that is not shared by the use of “discrimination” that we are interested in. (Hence Christopher Hitchens’ remark, exploiting the gap between these two meanings, that “the racist precisely shows himself incapable of discrimination.”)⁴ The laudatory use of “discriminating” only makes sense, however, because of the more basic feature—responsiveness to

qualities of the objects—that the two conceptions share, and which distinguishes them both from mere differential treatment.

Second, to say that a person discriminates in the sense that concerns us is not necessarily to say that the differences he perceives in the discriminatees are genuine. Consider, for instance, the research showing that American employers often consider identical résumés less favorably when they carry stereotypically African-American names (e.g., Jamal, Lakisha).⁵ If the mechanism of this differential treatment is that the employer ascribes different racial classifications to such candidates, and is less disposed to interview those she takes to be African-American, then I think it would be accurate, by the lights of prevailing usage, to describe this as a case of discrimination on the basis of race (or “racial discrimination” for short). In characterizing the case this way, we do not appeal to the fact (if it is a fact) that an applicant named Jamal, who is disfavored in this way, is *actually* African-American. We do not even presuppose a view of what “actually” being a member of a particular racial group involves. Thus to discriminate on a given basis is to respond differentially to what one *takes* to be a difference between the discriminatees in respect of that trait.

With these considerations in mind, I will adopt the following working definition of discrimination:

X (directly) discriminates against Y in dimension W on the basis of P if and only if:

(Differential Treatment Condition) X treats Y less favorably in respect of W than X treats some actual or counterfactual other, Z, in respect of W; and

(Explanatory Condition) a difference in how X regards Y P-wise and how X regards or would regard Z P-wise figures in the explanation of this differential treatment.⁶

Before turning to these conditions individually, let me make some initial points about the formula generally.

First, I mean for the variable “X” to range over agents only—be they people, corporations, institutions, or groups of these. So when we say that some act or enactment discriminates in such-and-such a way, as

we often do, I take us to mean that it is an *instance of* discrimination, by some unspecified or implied agent—not that the enactment can itself be substituted for “X” in the definition and satisfy it.⁷

Second, “Y” and “Z” also range over individuals (be they persons or institutions) or groups of individuals. But it is important in this context that a group of individuals be understood as just that: a set of particulars. The problem is that often we specify the set of people that is discriminated against in a given case by reference to some trait that they share. This invites us to confuse whom X discriminates against with the basis on which X does so. If some act discriminates “against immigrants,” for instance, Y is made up of people who are immigrants. But sometimes we read “against immigrants” with an implicit “as such” attached, allowing this one phrase to fix both who is discriminated against *and* on what basis (in this case, on the basis of being an immigrant). These are very different issues, however. Many immigrants are discriminated against on the basis of race, for example; and in some contexts it is important both that it is immigrants who are discriminated against and that it is on the basis of race that they are discriminated against. So the Differential Treatment Condition should be read without any implicit “as such” following “Y”. Even if Y is a set specified by reference to a trait shared by its members, we should not infer anything about the *kind* of discrimination X commits.

Third, “W” ranges over dimensions of treatment. A dimension of treatment is a set of exclusive options regarding how one treats someone. You engage in differential treatment in the dimension of hiring when you hire some people and not others, for instance. The exclusivity that defines this dimension of treatment is not across individuals, however, but within them. That is, what makes “hiring” a dimension of treatment is not that you cannot hire both Y and Z, but that for each of Y and Z, you cannot both hire and not hire them.

Finally, I have cast the definition as an account of “(direct) discrimination.” I put it that way because, if we are pressed to distinguish between “direct” and “indirect” forms of discrimination, mine is a definition of what is generally described as the “direct” variety. In fact, however, I think “indirect” discrimination is not usefully thought of

as a distinct form of discrimination at all, except as a piece of legal jargon—a position I will defend in the next chapter. So, relative to my preferred typology, the definition above is also a definition of discrimination generally.

With these clarifications behind us, let us consider the Explanatory Condition more closely. It seems obvious that one person's less favorable treatment of another must have some explanatory connection to the latter's possession of a given trait to qualify as discrimination on the basis of that trait. But not just any explanatory connection is sufficient. If it were, you would always discriminate on the basis of any traits that explain a discriminatee's possession of a trait on the basis of which you discriminate, which is not true.

Suppose, for instance, that the operator of a roller coaster turns away Sally, a seven-year-old girl, because he measures her height and determines that she is less than four feet tall. In so doing, he discriminates against her on the basis of height.⁸ Plausibly, however, Sally's being less than four feet tall is explained by her age. If her age partially explains her height, and her height partially explains her being turned away, then her age partially explains her being turned away. But the operator obviously does not discriminate against Sally on the basis of age, at least not directly.⁹

The Explanatory Condition supports and explains this judgment by identifying a discriminator's *perception* of a trait, rather than the trait itself, as essential to discrimination on that basis. Sally's age may explain her height, but it is the operator's perception of her height—or, as I put it in formulating the Explanatory Condition, how he regards her height-wise—rather than his perception of her age that explains the way he treats her. So she is discriminated against on the basis of height, not age. Metaphorically, we might say that the grounds on which a person discriminates are marked by the links on a chain of explanation that go through her mind. The connection between Sally's age and her height simply precedes that point on the explanatory chain.¹⁰

By contrast, suppose that rather than measuring Sally's height, the operator explicitly or tacitly estimates her age, infers from this

estimate that she is probably less than four feet tall, and then excludes her for this reason. In this case the Explanatory Condition implies that he discriminates *both* on the basis of age and on the basis of height. For how he regards her in each of these respects figures in the explanation of the fact that he turns her away (when he would not turn away others). A parallel analysis could be given in the case of employers who disfavor applicants whose résumés bear stereotypically African-American names: they discriminate both on the basis of name and on the basis of race, since they perceive one by way of the other, and a difference in respect of either perception could therefore have resulted in different treatment.¹¹ These cases make clear that what matters for the purposes of the Explanatory Condition is not the *rationale* for X's action, or the explanation X would adduce, but rather which of his perceptions actually contribute to explaining the way he acts.¹²

The possibility that a single instance of differential treatment may constitute discrimination on multiple bases in the manner just described is sometimes overlooked, so it is worth considering it more closely. Of course, we are familiar with the notion that a person can be disfavored on two grounds simultaneously—race and sex, for instance. Such “intersectional” discrimination is one kind of discrimination on multiple bases.¹³ The Explanatory Condition recognizes this, but it also allows for multiple-basis discrimination of another kind. To take another example, suppose that the director of an adoption agency disavors single applicants because, having reviewed the empirical evidence, she concludes that single people prove less effective on average than two-parent couples as adoptive parents. Suppose that the director is then accused of discriminating on the basis of relationship status. She might concede that she does this, but deny that discrimination of this kind is wrong. Alternatively, however, she might respond that if she discriminates, it is simply on the basis of *likely effectiveness*. That merely has an adverse effect on single people—a fact that is traceable to the state of the evidence about effectiveness, not discrimination on the basis of relationship status on her part.

The Explanatory Condition disallows this second sort of argument. To be sure, the director *does* discriminate on the basis of likely effectiveness. Suppose that Mary is a single applicant and John and Michelle are a couple. Part of the explanation of the director's differential treatment of Mary relative to John and Michelle is that (i) the director regards Mary as less likely to be effective, and (ii) the director is more disposed to grant applications from people she regards as more likely to be effective. But it is no less true that she discriminates against Mary on the basis of being a single applicant. For the fact invoked in clause (i) is itself explained in part by the fact that she regards Mary as a single applicant, whereas she takes John and Michelle to be a couple. That means that the latter fact is part of the explanation of the differential treatment as well—which therefore satisfies the Explanatory Condition for discrimination on the basis of relationship status.

Put another way, the Explanatory Condition takes no notice of whether the explanatory force of X's P-perceptions, with respect to X's differential treatment of Y and Z, is due to their relevance to an operative premise or an auxiliary premise of X's reasoning.¹⁴ In this case, relationship status figures only in an auxiliary premise of the adoption director's reasoning—her operative aim being simply to distribute children to promising parents—but that makes no difference, according to my account, to whether she discriminates on the basis of relationship status.

Indeed, by posing a question of explanation rather than rationale, this account allows that one may discriminate on the basis of a trait without really *reasoning* about it at all. The Explanatory Condition is satisfied whenever differential treatment is explained by "how X *regards* or would regard Y P-wise." All of the cases we have considered thus far involve simple *beliefs* that a person falls under one or another description—belief being the simplest form of regarding someone a certain way with respect to a certain dimension of difference. Suppose, though, that a driver stuck in traffic feels an impulse to lock his doors whenever a black pedestrian enters his visual field, but not when other people do. The driver may not even

be consciously aware of the approaching person or of perceiving anything about him; he may only be aware of feeling vaguely unsafe, and moved to lock the doors as a result. In such a case, it seems a stretch to say that the driver has any beliefs about the race of the approaching man. But it still seems possible that the driver discriminates—unbeknownst to himself, or even despite himself—on the basis of race.

The right way to characterize the mental state out of which this driver acts—a state that seems to fall, as Tamar Gendler says, “between reason and reflex”¹⁵—touches on subtle and actively debated issues in philosophy of mind. Without getting bogged down in the details, it seems that in the case I described, certain representational content is essential to accounting for the driver’s action, but this content is “processed in a relatively shallow way,” triggering certain affective responses that in turn conduce to a behavioral reaction.¹⁶ When that content can be understood to represent a person as falling under a description—here, it seems that it must represent the man as black—the agent *regards* the person as being that way in the sense that I intend.*

The Explanatory Condition also implies, a fortiori, that even if one does consciously perceive differences in respect of some trait P, one need not be aware of the role of one’s P-perceptions in explaining one’s differential treatment (much less intend that they play this role) for one to directly discriminate on the basis of P. Instead, as we have seen, the definition rests weight only on whether these perceptions contribute to *explaining* one’s differential treatment. It thus resembles Linda Hamilton Krieger’s revisionist understanding of “disparate treatment” liability under U.S. employment discrimination law, according to which “[t]he critical inquiry would be whether . . . group status ‘made a difference’ in the employer’s action, not whether the decisionmaker intended that it make a difference.”¹⁷

* The “regarded as” language itself is borrowed from the Americans with Disabilities Act of 1990, which extends protection to those “regarded as” disabled. See 42 U.S.C. § 12102(1)(C).

These views put me at odds with some others. Andrew Altman, for instance, contends that “direct discrimination is intentional discrimination.”¹⁸ But I think this proposed equivalence can’t be true. We can begin from the observation that the act-type delimited by the definition of discrimination I have given, whatever we call that type, is such that acts of this type could be committed either intentionally or unintentionally. I think that much is unarguable. In choosing among raised hands, for instance, a teacher may call on male students more frequently than female students.¹⁹ Suppose that this discrepancy is not a simple statistical accident or a consequence of some correlation between sex and other features, but rather explained by a conjunction of the sex-perceptions of the teacher and certain sex-sensitive dispositions of his. Holding this much constant, it is clearly possible that the teacher does not *intend* to treat male and female students differently. Alternatively, the teacher could intend to act this way—perhaps because he thinks male students are more likely to make comments that advance the class discussion.

Now there is some common thing that the teacher is doing in each of those two variants of the example, unintentionally in one case and intentionally in the other. I think we ought to call that thing discrimination on the basis of sex—and “direct” discrimination if pressed to distinguish it from what is called “indirect” discrimination. Altman would apparently say that there is direct discrimination only in the intentional case. But then what is there in the unintentional case?

It seems absurd to say that there is no sex discrimination at all, so this must be a case of “indirect” discrimination. Without getting too far ahead of ourselves, Altman defines “indirect” discrimination in the typical way, as encompassing acts or policies with disproportionate negative effects on members of some group that fail some special standard of justification. But assimilating the teacher’s unintentional differential treatment of men and women to that category would elide something essential about it. The teacher does not merely do something that is disproportionately to the disadvantage of women (and do so without sufficiently good reason)—as he might if he were simply more disposed to call on taller students, whatever their sex,

and this led to his calling on women less often than men. Rather, he is less likely to call on women precisely because of the impression that their being women makes on him. There is nothing “indirect” about that; it is simply unintended. What he does should be counted as direct sex discrimination, therefore, and we should consequently decline to equate direct discrimination with intentional discrimination.

In an instructive article, Patrick Shin argues that dominant legal regimes for policing discrimination assume a “justificatory,” rather than “causal,” conception of discrimination.²⁰ The justificatory view understands the question whether a person discriminates on a given ground as fundamentally concerned with certain features of the agent’s reasons for action or “putative justificatory rationale.” As Shin notes, a turn toward causal explanation raises important issues concerning the nature of moral responsibility for wrongful discrimination (and related questions concerning legal liability), which we will consider in somewhat greater depth later on. But these questions are separable from questions about what constitutes discrimination in the first place. Even if we took the extreme view that a person should *never* be morally or legally liable for discrimination that is unconscious or unintentional, it would still be appropriate to describe him as engaged in unwitting or unintended discrimination on the basis of the relevant trait.

1.1.1 *Structural discrimination*

To conclude our discussion of the Explanatory Condition, consider Sally and the roller coaster once again. In the original version of that case, we concluded that Sally is discriminated against only on the basis of her height, and not on the basis of her age, because age is not what the operator is responding to—even if her height, to which he is responding, is itself largely explained by her age. There is nevertheless an important pattern of relations that obtains in this case between (a) the operator, (b) Sally’s exclusion, and (c) Sally’s age. It will help to have a name for this pattern of connections.

In fact, I think it is appropriate to think of this as a case of *structural* discrimination. On this understanding, an act or policy is

structural discrimination on the basis of P if it is comparatively disadvantageous to one group defined in terms of P relative to another one, and its being so is partly explained by their difference in respect of P. The essential difference between ordinary and structural discrimination, according to this account, is that in the latter case the explanatory connection between the comparative disadvantage and the trait in question need not be mediated by anything to do with how the (structural) discriminator perceives or reacts to the trait.

This yields the result that Sally's height-based exclusion from the roller coaster is an instance of structural age discrimination against her committed by the operator.²¹ Similarly, awarding scholarships to children on the basis of academic achievement may amount to structural discrimination on the basis of family wealth, if family wealth partly explains which students achieve highly. So too if banks grant loans to black applicants less often than whites, and this is because blacks have genuinely worse credit history than whites on average, but that fact is itself traceable to their being black—for instance, because blacks are targeted for predatory loans that damage their credit. (Here we have an instance of another common pattern, where the standing of one practice as structural discrimination on the basis of some trait is traceable to a distinct instance of ordinary discrimination on the basis of that trait in a different dimension of treatment.)

Acts that are structural discrimination on some basis need not be, or be connected to, an instance of ordinary discrimination on the same or another basis. If an agency offers a vaccine which, unbeknownst to it, only works for men, it may not engage in any ordinary discrimination at all. But this act arguably constitutes structural sex discrimination against women: it is comparatively disadvantageous to them, and its being so is traceable to their being women.

As I hope these examples demonstrate, my definition of structural discrimination captures the diffuse sense that this concept concerns the imposition of burdens on certain classes of people through formally neutral procedures—procedures that adversely affect these classes only because of underlying physical or social dynamics

(or “structures”) which, for instance, connect age with height, or race with credit history. At the same time, my proposal preserves a distinction between structural discrimination and situations in which burdens fall disproportionately on members of a given group but not in any way *because* they are members of that group.²²

Still, because the scholarly and popular usage of “structural discrimination” is inconsistent, I hesitate to cast my definition as more than stipulative.²³ On some interpretations, moreover, the very point of recognizing “structural” discrimination as a distinct phenomenon is that it is not a discrete act committed by an agent; for better or worse, my definition does not allow the concept to serve that function.²⁴ And, importantly, by describing this category as structural discrimination, I don’t mean to imply that it really is a kind of *discrimination*. Prefixed concepts are often ambiguous in this way: they can either mark out a subspecies of a phenomenon, or they can qualify the claim that the thing in question falls under the concept at all. Given my definition of discrimination, my definition of structural discrimination is of the latter kind.

We can set this category aside for now. In the next chapter, I will appeal to it in the course of considering possible interpretations of the idea of indirect discrimination.

1.2 Discrimination and Social Salience

We now have the fundamentals of a conceptual analysis of discrimination on the table. As it stands, the definition imposes no limits on the bases on which people can and do discriminate—that is, on the kinds of traits that can be substituted for P and satisfy the definition. We can discriminate on the basis of race or sex, but we can equally discriminate on the basis of handedness, position in a queue, trustworthiness, and indefinitely many other qualities.

Some have urged that the concept of discrimination relevant to moral analysis be hemmed in. Kasper Lippert-Rasmussen, for instance, says that to be discriminatory, differential treatment of Y and Z must be “suitably explained by Y’s and Z’s being (or [being] believed by X to be)

(members of) different, socially salient groups.”²⁵ He further specifies that “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.”²⁶ Lippert-Rasmussen acknowledges that this leaves it “somewhat unclear when a group is socially salient,” but suggests, reasonably enough, that this “simply reflects the fact that the contours of our concept of discrimination are somewhat fuzzy.” I will now argue that this conceptual constraint, which we can call the Social Salience Requirement, should be rejected.

In fact, the Social Salience Requirement is ambiguous between two significantly different positions. First, we might require that X’s differential treatment of Y and Z be suitably explained by Y and Z’s belonging to different groups, one or each of which has the feature of being socially salient. Second, we might require that X’s differential treatment of Y and Z be suitably explained by Y and Z’s *belonging to different socially salient groups*—that is, with the social salience of the groups itself contributing to the explanation of the differential treatment.

This ambiguity is not unique to Lippert-Rasmussen’s explanation of the Social Salience Requirement. Altman, for instance, writes: “The wrongfulness of . . . discrimination is tied to the fact that the discriminatory act is based on the victim’s membership in a salient social group.”²⁷ Again, either reading seems to fit this sentence equally well: what distinguishes wrongful discrimination could be (1) that (a) the act is “based on” the victim’s membership in a group and (b) this group has the feature of being socially salient, or it could be (2) that the act is “based on” the victim’s membership in a salient social group in such a way that, if the group were not socially salient, the act’s basis would be removed.

Either interpretation of the Social Salience Requirement yields some unattractive results. On the first reading, whether a person discriminates turns on whether the basis on which he differentiates is in fact a socially salient cleavage, even if that social significance has no connection whatsoever to his action. As a thought experiment, then, suppose that an odd individual is born with an innate tendency

to favor one kind of person over another: men over women, dark-skinned people over light-skinned people, left-handed people over right-handed people, or something like this. It would be very strange, I think, to say that whether *he* discriminates hinges on whether his idiosyncratic partiality happens to align with a distinction that is socially salient in his society at the time, although for reasons wholly unrelated to him or his partiality. That will bear on the consequences of his action in many cases, of course, in ways that are potentially important to its moral status. But the idea that it determines whether he discriminates conflicts with what seems to me a basic “internalist” intuition about the concept. If we know everything about how X treats Y and Z and why he treats them each as he does, we should know whether he is discriminating on any given basis in so doing.

So perhaps the second reading is better: whether a person discriminates depends on whether the social salience of the groups he differentiates between is *itself* part of the explanation for his differential treatment. Certainly this is closer to the intuitive idea reflected in paradigm cases of wrongful discrimination, in which a person discriminates on the basis of race or sex in part because of the social meanings of these categories. But this reading makes the Social Salience Requirement implausibly strong. Suppose, for instance, that the manager of a factory prefers to hire men rather than women because of a non-spurious generalization that men possess greater upper-body strength on average. Regardless of whether this policy is objectionable, it is obviously an instance of sex discrimination. But just as clearly, the manager’s preferential treatment of men need not be *explained by* the social salience or significance of sex. Rather, it is explained by the manager’s generalization about upper-body strength, which he could just as well have derived from a biology textbook as from his social milieu. Then the differential treatment would not be recognized as discrimination by the definition we are considering.

So I think the Social Salience Requirement is unsustainable, on either reading. The best way to reconcile our intuitions about the two cases I’ve described is simply to disconnect social salience from whether an act constitutes discrimination, and hence to recognize

both of them as discriminatory. This seems to yield the right results in other cases as well. If a firm were to give a preference in hiring to blond-haired applicants, for instance, it would be perfectly sensible to say that it is discriminating on the basis of hair color, even if hair color bears little significance in social life writ large.

Lippert-Rasmussen suggests that restricting the concept of discrimination to socially salient groups has the virtue of “explain[ing] why we do not talk about discrimination against non-family-members, unqualified applicants or the undeserving.”²⁸ A more straightforward explanation of this fact seems more plausible, however. Adverse treatment of non-family-members, unqualified applicants, and the undeserving is normally accepted as morally unproblematic, and as I observed at the outset, people often use “discrimination” in a moralized sense. If we simply shifted to a context where disfavoring non-family-members *would* normally be considered objectionable (such as in hiring for public offices), I suspect many people would be prepared to speak of a preference for one’s own children as discriminatory.

In other words, if there is a theme to ordinary usage of “discrimination,” it is not that the targeted groups are socially salient, but that the differential treatment is taken to be at least potentially objectionable. Witness, for instance, the popular idea that insurers should not “discriminate” against people with pre-existing medical conditions in the United States,²⁹ or the legal protections against discrimination on the basis of trade union membership or “philosophical belief” in the United Kingdom.³⁰ Similarly, the *New York Times* recently printed four critical letters to the editor regarding its reporting on hospitals that refuse to hire smokers; three of the letters described the practice as “discrimination.”³¹ Surely what is meant by invoking discrimination in these contexts is not that the policies in question injure people on account of their membership in groups that are important to the structure of social relations. Rather, it is that they objectionably employ these differences between people as grounds for adverse treatment.³²

There is certainly a mismatch between my account and ordinary usage here, since wrongfulness is no part of my proposed concept of discrimination. But both Lippert-Rasmussen and I have already bitten

that bullet by defining discrimination as, in his phrase, “not necessarily morally objectionable.”³³ What explains the intuitive appeal of the Social Salience Requirement, I think, is that it partially and indirectly patches this failure of extensional fit, by introducing a feature that is intuitively related to the wrongfulness we have decided to forgo including in the concept. But I have argued that there are good reasons for defining discrimination in non-evaluative terms, and regardless, splitting the difference in this way only invites confusion.

It should be clear, then, that by rejecting the Social Salience Requirement I do not mean to deny that discrimination on the basis of socially salient traits may be a morally significant category. If it is, however, this should emerge organically from comparing different kinds of discrimination, rather than by way of a flat stipulation that only one of them is discrimination in the first place. Indeed, this is the source of the most powerful reason for rejecting the Social Salience Requirement: it obscures what perhaps ought to be the default position about the moral significance of paradigm cases of discrimination, which do involve socially salient groups. It could be that rejecting a job applicant because of race or sex is bad in just the same way that rejecting her because of her eye or hair color would be, and that the former kinds of discrimination only command more of our attention because they are more widespread (which of course is itself largely due to the fact that these distinctions are socially salient).³⁴ Part of the job description for a moral analysis of discrimination is to determine whether there is more to the story than that—and, if so, what. But this fundamental question is necessarily marginalized when we build social salience into the relevant concept of discrimination.³⁵

1.3 Discrimination Against, In Favor Of, and Among

Thus far we have focused on how differential treatment must be explained in order to constitute discrimination on a particular basis, and on which potential bases count as possible grounds of

discrimination. But what makes a given instance of discrimination qualify as discrimination *against* or *in favor of* someone? This is a complicated issue, and it has received little direct discussion in the philosophical literature.³⁶

Plainly enough, to say that X discriminates against Y is to say that X treats Y not only differently but *worse* than X treats (or would treat) Z. The Differential Treatment Condition thus specifies that X must treat Y *less favorably*, in respect of the dimension W, than Z. To say that X discriminates against Y is therefore necessarily an evaluative statement—evaluative not of the morality of what X does, but of the comparative value of the two treatments in question. Importantly, however, these judgments of comparative value are sometimes made in terms of overall value, and sometimes in respect of the satisfaction of some specific interests, which often go unidentified.

Suppose, for example, that the admissions committee at The Girls School automatically rejects all male applicants. Moreover, the committee does this in the correct belief that any male student would be disserved by attending The Girls School rather than some other school. It follows that the school's policy is not comparatively disadvantageous to male applicants in terms of their overall well-being; by hypothesis, they are better off having their applications handled this way than having them assessed in the way that female applicants' applications are. Nor does the committee value the well-being of male applicants any less than that of female applicants, or believe that male applicants would make worse students or contribute less to the school than female applicants. Indeed, we can imagine that the committee would count male students as a greater asset to the school, and would *prefer* to admit them, but excludes them out of an overriding paternalistic concern for the applicants' own well-being.

Despite all of this, it would be typical to describe the school as discriminating against male applicants, on the basis of sex, in admissions. After all, it denies their applications for admission, and it does so simply because they are male. (Discriminating in this way in this situation may or may not be *wrong*, of course, but that is a separate question.) The best way to make sense of this fact about our use of the

concept, I think, is to say that being discriminated against must be comparatively disadvantageous to a person in some specific respect, but being comparatively disadvantaged in that respect need not be globally comparatively disadvantageous to him (and may even be globally comparatively advantageous to him). Here the specific interest at issue is in being admitted to the school.³⁷

Unfortunately, I do not have a general theory about the specific interests people possess in the sense relevant here. But it seems that when people supplicate for some opportunity—whether by applying for a job or admission to a school, or by raising their hands in a class, or whatever—they thereby come to possess a specific interest in being granted it, whether or not they had such an interest prior to applying, and whether or not having the opportunity in question will make their lives go better overall. Many paradigm cases of discrimination involve differential satisfaction of such an interest in being granted an opportunity one has sought.

Besides specific interests in being granted the various things we have asked for, I think we also have specific interests in being dealt with in certain ways. I take for granted that people have an interest in being treated with respect, for example. James Woodward suggests a useful example: the racially discriminatory denial of an airline ticket to Smith harms Smith, even if the plane goes on to crash, because “the action violates a certain highly specific interest of Smith’s in avoiding racial stigmatization and unfair or inequitable treatment.”³⁸ Applying the framework developed thus far, we could describe this case as follows. If the denial of the ticket is motivated by racism, then Smith is treated comparatively worse, relative to those who are given tickets, in respect of two different interests of his: in getting a ticket (even if he would be better off without one), and also in being treated with respect. Either of these alone would suffice to make the denial of the ticket, on the grounds that it was in fact denied to him, an instance of discrimination against him.³⁹

Consider, then, a rather different case. Many serial killers have some kind of preferred victim-type. Suppose that a particular killer targets white women. He discriminates against white women, then,

insofar as he treats them comparatively worse than men and non-white women. But there also seems to be an intuitive sense in which his selection procedure could be described as disfavoring men and non-white women, although this is surely to their benefit.

When The Girls School disfavors male applicants, we say that this is discrimination against them, even though it is to their benefit. Should we similarly say that the serial killer discriminates against men and non-white women, then, by not targeting them for killing, even though this is to their benefit? I think we should not. According to my theory, that is explained by the fact that people do not have a specific interest in being targeted for killing. Fittingly, if a non-white woman did come to the killer requesting to be killed, and he rejected her because she was not white, I think we *would* be inclined to say that he discriminated against her on the basis of race, although we might still maintain that this was to her comparative benefit overall. That is because she would have created (or revealed) a specific interest in being killed by him.*

The example of the serial killer raises another important point. It seems odd, in a way, to say that the killer discriminates *against* white women. After all, he prefers them, seeks them out, and so on. But this appearance is due to the fact that what he does falls under what we might call the *non-interest-regarding* concept of discrimination, as well as the *interest-regarding* concept with which we are centrally concerned. To say that a person discriminates “in favor of” or “against” some class of objects, in the non-interest-regarding sense, is simply to say that he prefers or disprefers it. People discriminate among wines in this way, for instance. The killer does approach the field of possible targets in a manner akin to a selection among wines, and so in the non-interest-regarding sense, it would be wrong to say

* Suppose that he avoids killing non-white women out of a kind of racial contempt for them. Then *perhaps* he treats them worse than white women in terms of their specific interest in being respected. If so, it would not be entirely wrong to say he discriminates against non-white women even if none request to be killed. But the non-white women are clearly also treated better in terms of their interest in not being killed. I discuss cases of conflict among interests of this kind below.

that he discriminates against the targets he prefers. Whether he discriminates against them in the interest-regarding sense is a separate matter, however. In particular, whether something constitutes discrimination “against” or “in favor of” someone in this sense has to do with the interests of the discriminatees in being treated one way or another—interests that wines do not have. We should simply do our best to tune out the verbal influence of the non-interest-regarding concept on our judgments about what constitutes discrimination in the sense with which we are concerned.

As it turns out, then, the Differential Treatment Condition actually has two dimensional aspects. First, the differential treatment is itself understood by reference to a dimension of treatment—a set of exclusive options regarding how one treats someone. If X’s discrimination against Y consists in hiring Z and not hiring Y, for instance, X discriminates in the dimension of hiring. Second, as we have seen, the comparative disadvantage to Y is *itself* found in a dimension or a respect—the specific interest of Y’s at issue.

These two dimensional aspects of the concept are easily run together. To say that The Girls School discriminates against boys “in admissions,” for instance, could mean either that admissions decisions form the dimension of treatment in which the school discriminates against boys, or that it is in respect of their interest in being admitted that the boys are treated comparatively worse. In this case, it likely means both. But although these features often converge, they are not the same.

This rather technical point is potentially important to understanding the moral structure of some discrimination. Suppose, for example, that an author accidentally submits a paper to the *British Journal of Political Science*, when she meant to submit it to the *British Journal for the Philosophy of Science*. I don’t think she has an interest in having the paper accepted for publication by the political science journal, assuming she would not want that to occur. Nonetheless, if the journal’s editors reject her paper out of sexism, they discriminate against her on the basis of sex. What’s more, they discriminate against

her in the dimension of paper acceptance. The interest of hers that *makes* their action in that dimension discrimination against her, however, is not an interest in having the paper accepted. It is rather an interest in being treated with respect.⁴⁰

Two final notes about usage are in order. First, to say that X discriminates against Y in some dimension, without separately specifying an affected specific interest, can be to say *either* that the way Y is treated in that dimension is comparatively disadvantageous to Y in respect of *at least one* specific interest, or that it is comparatively disadvantageous to Y *overall*. This ambiguity is another recurring source of confusion. Consider, for example, the exclusion of African Americans from jury service in the United States. This exclusion constituted discrimination against African Americans in respect of various specific interests of theirs—in social recognition and prestige, in decision-making power, in protection from biased verdicts, and so on—but also discriminated in favor of African Americans in respect of a far less important, but nonetheless genuine, interest. Insofar as jury service is involuntary and burdensome, people are benefited, in one respect, by being exempted from it. So were African Americans discriminated against or in favor of in the dimension of social decisions about whom to include in jury service? It is not wrong to say *both*, since each is true in a respect. But it is also not wrong—and in this case it is pragmatically more appropriate—to say that they were discriminated against *rather than* in favor of, because they were so clearly comparatively disadvantaged overall.

Second, we might wonder how the Differential Treatment Condition handles cases of intentional but inefficacious discrimination. Sometimes, that is, the appropriateness of calling something discrimination seems to be settled simply by the intentions of the agent—independent of any actual effects on anyone's actual interests. I think such cases are handled well by the account we have developed. Some of them actually invoke the non-interest-regarding concept of discrimination: we simply mean that the agent prefers people of one description or another, not that she *treats* them better or worse.

In other cases we do mean that she treats some better than others, but we are supposing that is true simply by virtue of her differential aims: we are supposing that they have interests in her pursuing certain aims with respect to them. That possibility can be accommodated by the skeletal account of specific interests we have developed. Moreover, as I will suggest in the next chapter, recognizing this point is important to understanding the structure of some cases of “indirect” discrimination.

1.3.1 *Discrimination among and “separate but equal” treatment*

Our very general conceptual discussion to this point has particularly significant implications for how we should think about one distinctive class of cases. Specifically, it bears on cases in which differential treatment is explained in the way that discrimination must be, yet does not treat one party comparatively worse than another. Some “separate but equal” arrangements seem to have this character. What should we make of them?

Based on the foregoing analysis, we should distinguish between instances of differential treatment in which neither party is treated less favorably in respect of *any* interest, and those in which each party is treated less favorably in respect of a *different* interest, where these disadvantages have equal net effects on overall well-being. In a moment I will suggest that most schemes that promote or enforce “separate but equal” roles for different types of people can only claim, at best, to fall in the latter category. But the former category is not empty. It is populated by acts and policies that do not treat people more or less favorably in respect of any of their interests—and hence are not instances of discrimination in the interest-regarding sense—but which can nevertheless be described as discriminating among people, in the non-interest-regarding sense. Suppose, for instance, that I invite Alfred but not Betty to an event that I know neither of them has any interest in being invited to. And suppose that I treat them differently in this way simply because Alfred’s name comes first alphabetically. I discriminate on the basis of name, then, but only in the non-interest-regarding sense.⁴¹

Attention to the role of specific interests in the concept of discrimination should lead us to sharply distinguish this kind of case from acts or policies that assign distinct but equally valuable opportunities to different classes of people. These *do* affect their various specific interests differentially; the discriminator may simply manage to do this in a manner that keeps the effects on the overall well-being of the parties in equipoise (though that would seem a remarkable feat in most cases).

For example, some claim that traditional religious doctrines prescribing distinct social and familial roles for men and women do not amount to sex discrimination, because the roles are of equal value and founded on a principle of mutual complementarity.⁴² Of course, if what is meant by this claim is that the prescriptions do not call for *wrongful* or *unjust* discrimination, the conceptual analysis I have offered does not speak to that question.⁴³ But one might think that, insofar as nobody is treated worse than anybody else, these practices do not constitute discrimination against or in favor of anyone at all—not even discrimination of the possibly permissible kind.

According to the view I have outlined, that is mistaken. If an institution makes it more difficult for men than women to be homemakers, and more difficult for women than men to be breadwinners, it very likely thereby discriminates against (and in favor of) both men and women on the basis of sex. For I take it that many men have an interest in the opportunity to be homemakers, and many women have an interest in the opportunity to be breadwinners, even if the roles are in some more global sense equally valuable. Indeed, perhaps everyone has an interest in at least having these two *opportunities*. Even if that were not so, however, it would mean only that the sex discrimination in question is discrimination against only some men, and only some women—in each case, those who want to occupy the role for which they are disfavored. That does not make the practice any less an instance of sex discrimination, because not all sex discrimination is discrimination in favor of or against *all* people of a given sex. This is one reason I suggested above that it is important to distinguish in the Differential Treatment Condition between the sets of people

who compose Y and Z—sets which may be defined in part by their members' possession of some trait—and the trait P, on the basis of which they are discriminated against.

1.4 Conclusion

The conceptual analysis of discrimination I have offered aims to bring out the features that distinguish the way we use this idea from the way we use its conceptual neighbors, like differential treatment, while also maintaining a healthy skepticism about certain regularities in usage—such as the tendency to employ “discrimination” as an epithet—that seem to confuse more than they clarify. It gives us some rules for when to describe something as discriminating *against* someone, and explains how this can be squared with its making him better off. Perhaps most importantly, the account specifies how differential treatment must be explained in order to constitute discrimination on the basis of a particular trait P (and therefore to constitute discrimination at all). In particular, the differential treatment simply has to be due, in part, to the fact that the agent regards the discriminatees differently in respect of that trait—whether or not she intends to be sensitive to this apparent difference, and whether or not she is sensitive to it only as a proxy for some other trait with which she is concerned. So, going forward, when we say that something is or is not discriminatory—or is or is not race or sex discrimination, and so on—we will know what we mean.

One influential idea that I have discussed only obliquely, however, is that of “indirect” discrimination. I turn to how we should understand this concept, and its practical significance, in the next chapter.